





COSYNE 2010

7th Computational and Systems Neuroscience Meeting

Main Conference • 25–28 February 2010
Salt Lake City, Utah



Program Summary

Thursday, 25 February

- 4:00 pm **Registration opens**
- 6:00 pm **Welcome reception**
- 7:20 pm **Introductory Remarks**
- 7:30 pm **Keynote address:** R. Clay Reid, Harvard Medical School
- 8:30 pm **Poster Session I**

Friday, 26 February

- 7:30 am **Breakfast**
- 8:30 am **Morning Session** (break 10:00 – 10:30)
- 12:00 pm **Lunch break** (and last chance to see Session I posters)
- 2:00 pm **Afternoon Session** (break 3:15 – 3:45)
- 5:00 pm **Dinner break**
- 7:30 pm **Poster Session II**

Saturday, 27 February

- 7:30 am **Breakfast**
- 8:30 am **Special session honoring Horace Barlow's legacy** (break 10:00 – 10:30)
- 12:00 pm **Lunch break** (and last chance to see Session II posters)
- 2:00 pm **Afternoon Session** (break 3:15 – 3:45)
- 5:00 pm **Dinner break**
- 7:30 pm **Poster Session III**

Sunday, 28 February

- 7:30 am **Breakfast**
- 8:30 am **Morning Session** (break 10:15 – 10:45)
- 12:00 pm **Lunch break** (hotel checkout, and last chance to see Session III posters)
- 2:00 pm **Afternoon Session**
- 4:00 pm **Closing remarks**

frontiers IN NEUROSCIENCE

Editor-in-Chief: Idan Segev

ACADEMIC PUBLISHING IN THE 21ST CENTURY

Expert article review • Interactive peer-review • Democratic article tiering
High quality publishing • Rapid online publishing • Open access articles

WHY PUBLISH IN FRONTIERS?

- **Expert Editors** (over 2,600 world class editors)
- **Instantaneous Abstract Publishing** (be the first, claim your discovery date)
- **Efficient Peer-Review** (first real-time peer-review in publishing)
- **Democratic Article Tiering** (best articles selected by thousands of readers)
- **Fast Publication** (around 3 months from submission to acceptance)
- **Massive Readership** (over 225,000 neuroscience related readers)
- **Global Dissemination** (over 160 countries)
- **Real-time Impact** (online analysis of the impact of your article)
- **Archived Presence** (in Pubmed, PsycINFO and Google Scholar)

**Over 2,000 authors already have
a Frontiers article. Do you?**

FRONTIERS IN
AGING NEUROSCIENCE
Mark A. Smith and Gemma Casadesus
FRONTIERS IN
BEHAVIORAL NEUROSCIENCE
Carmen Sandi
FRONTIERS IN
CELLULAR NEUROSCIENCE
Alexander Borst
FRONTIERS IN COMPUTATIONAL
NEUROSCIENCE
Misha Tsodyks
FRONTIERS IN
ENTERIC NEUROSCIENCE
Joël C. Bornstein

FRONTIERS IN
EVOLUTIONARY NEUROSCIENCE
Steven M. Platek
FRONTIERS IN HUMAN NEUROSCIENCE
Robert T. Knight
FRONTIERS IN
INTEGRATIVE NEUROSCIENCE
Sidney A. Simon
FRONTIERS IN
MOLECULAR NEUROSCIENCE
Peter H. Seeburg
FRONTIERS IN NEURAL CIRCUITS
Rafael Yuste
FRONTIERS IN NEUROANATOMY
Javier Defelipe


FRONTIERS IN NEUROENERGETICS
Pierre J. Magistretti
FRONTIERS IN
NEUROENGINEERING
Laura Ballerini
FRONTIERS IN NEUROGENESIS
Angélique Bordey
FRONTIERS IN NEUROGENOMICS
Robert W. Williams
FRONTIERS IN
NEUROINFORMATICS
Jan G. Bjaalie
FRONTIERS IN
NEUROPHARMACOLOGY
Nicholas M. Barnes

FRONTIERS IN
NEUROPROSTHETICS
Niels Birbaumer and Eilon Vaadia
FRONTIERS IN
NEUROROBOTICS
Frederic Kaplan
FRONTIERS IN
NEUROSCIENCE METHODS
Mark H. Ellisman
FRONTIERS IN
SYNAPTIC NEUROSCIENCE
Mary B. Kennedy
FRONTIERS IN
SYSTEMS NEUROSCIENCE
Ranulfo Romo



www.frontiersin.org

Frontiers in Neuroscience • Scientific Park, PSE-A, PO Box 110 • 1015 Lausanne • Switzerland • Tel +41 (0)21 693 92 02 • Fax +41 (0)21 693 92 01



Great companies never stop moving forward

Qualcomm is a world leader in providing mobile communications, computing and network solutions that are transforming the way we live, learn, work, and play.

As a pioneering technology innovator, we are committed to continuous evolution and exploration. Our researchers and computational neuroscientists engage in a wide variety of exciting and technically challenging projects—including the application of biologically inspired learning machines to a new breed of intelligent devices.

We help you work smarter, move faster and reach further. We see success in your future. In fact, we're passionate about it.

We are futurists.

www.qualcomm.com

QUALCOMM®

COSYNE 10

iii



The Bernstein Network is a large-scale research network focused on the field of Computational Neuroscience. Computational Neuroscience combines up-to-date experimental approaches with theoretical models and computer simulations.

The German Federal Ministry of Education and Research (BMBF) initiated this funding initiative in 2004, with the aim of supporting and interconnecting scientific competence and promoting interdisciplinary teaching & training in Germany.

Participating Institutions:

- Universities
- Max Planck Institutes
- Fraunhofer Institutes
- Leibniz Institutes

Research topics:

Basic science

towards understanding the brain

- How does the brain process information?
- Sensing and perceiving, attention
- Control of movements and behavior
- Pain
- Development and reorganization of the brain
- How do thoughts arise? Can we read thoughts?
- Learning and memory
- Computer simulations of the brain


Applied Science

Neurotechnology for IT and health

- Neurological diseases, BMI, prosthetics
- Intelligent robots, new IT technologies

200 Research Groups 24 Locations

- Bernstein Award
- Bernstein Collaboration
- Bernstein Group
- Bernstein Center
- Bernstein Focus: Neurotechnology
- Bernstein Focus: Neuronal Basis of Learning
- Bernstein Coordination Site
- German INCF Node (G-Node)



What we offer:

Interdisciplinary training

- In English
- No tuition

Open Positions

- PhD students
- Postdocs
- Junior Research Groups
- Professorships

Industrial partners of the Bernstein Network:

Biomedizinische NMR Forschungs GmbH - Brain Products - certon systems - Cochlear - Daimler - Honda Research Institute - Infineon Technologies - inomed - Leica Microsystems - L-1 Identity Solutions - Magnicon - MED-EL - Multi Channel Systems MCS - neuroConn - NIRx - nisis - Otto Bock Healthcare - Robert Bosch - Schunk - Telekom - Thomas RECORDING - VITRONIC.

www.nncn.de

Innovative research at the interface
between the life sciences and the physical sciences

**Browse our content
at <http://hfspj.aip.org>**



**Open Access
for All Articles
after 6 Months**


**Immediate
Open Access
Available
for Authors**

**Online submissions
and subscription info
at <http://hfspj.aip.org>**



Brain Corporation is hiring 3 senior vision neuroscientists to join its theory group to develop a spiking model of the mammalian visual system.

Email to jobs@braincorporation.com




The Anatomy of Bias

How Neural Circuits Weigh the Options

JAN LAUWEREYNS

"Jan Lauwereyns brings together concepts that are generally treated as disparate, and traces the historical evolution of their relation to one another and to current research. The significance of this contribution will be partly as a stimulus to new ideas, as well as its achievement in situating current ideas about decision firmly in their historical intellectual milieu. *Anatomy of Bias* is the kind of book that will change people's thinking—and lives." — R. H. S. Carpenter, Cambridge University
288 pp., 32 illus., \$30 cloth




A Hole in the Head

More Tales in the History of Neuroscience

CHARLES G. GROSS

"Charles Gross is a pioneering neuroscientist with a deep sense of history and a gift for lucid, vivid story-telling. His new book, *A Hole in the Head* is enthralling, fast-paced, and as exciting as any detective story." — Oliver Sacks, author of *The Man Who Mistook His Wife for a Hat*

Visit our website to hear a podcast featuring this author
336 pp., 59 illus., \$35 cloth




Computational Modeling Methods for Neuroscientists

EDITED BY ERIK DE SCHUTTER

A guide to computational modeling methods in neuroscience, covering a range of modeling scales from molecular reactions to large neural networks.

Computational Neuroscience series
432 pp., 85 illus., \$50 cloth



Seeing

The Computational Approach to Biological Vision

Second Edition

JOHN P. FRISBY AND JAMES V. STONE

"*Seeing* is not a new edition but a completely new book, and a unique book—a carefully written, beautifully illustrated text of the computational approach to human vision that will take the reader from first principles to cutting-edge ideas about all levels of the visual process." — Oliver Braddick, University of Oxford

576 pp., 132 color illus., 399 b&w illus., \$55 paper

Brain Signal Analysis

Advances in Neuroelectric and Neuromagnetic Methods

EDITED BY TODD C. HANDY

Recent developments in the tools and techniques of data acquisition and analysis in cognitive electrophysiology.

264 pp., 12 color illus., 56 b&w illus., \$55 cloth

Learning and Inference in Computational Systems Biology

EDITED BY NEIL D. LAWRENCE, MARK GIROLAMI, MAGNUS RATTRAY, AND GUIDO SANGUINETTI

Tools and techniques for biological inference problems at scales ranging from genome-wide to pathway-specific.

Computational Molecular Biology series
400 pp., 73 illus., \$40 cloth

Control Theory and Systems Biology

EDITED BY PABLO A. IGLESIAS AND BRIAN P. INGALLS

"Historically, control theory has its roots in the analysis and understanding of physical and technological systems. However, in recent times it has revealed its wider potential as a tool for describing the complex dynamical behavior of living things. This valuable book is therefore timely, since it places control theory where it belongs—at the heart of our search for the principles by which living systems operate." — Peter Wellstead, National University of Ireland

384 pp., 138 illus., \$45 cloth

Cognitive Biology

Evolutionary and Developmental Perspectives on Mind, Brain, and Behavior

EDITED BY LUCA TOMMASI, MARY A. PETERSON, AND LYNN NADEL

"Modern progress in science depends critically on interdisciplinary endeavor, and psychology has had its share of alliances. We've had psycholinguistics, cognitive neuroscience, cognitive archeology, evolutionary psychology, not to mention the offspring spawned by the incursion of neuroimaging. None quite captures the breadth of inquiry needed to fathom the mind. This excellent volume brings together a diverse range of expertise, neatly captured by the volume's title." — Michael Corballis, University of Auckland

Vienna Series in Theoretical Biology
384 pp., 80 illus., \$50 cloth



Enabling Innovative Neurophysiology

Blackrock microsystems is recognized as a leading world-wide provider of advanced hardware and software tools that help neuroscientists, engineers, and clinicians perform cutting-edge research in areas such as fundamental neuroscience, brain machine interfaces, and neuroprosthetics.

Blackrock's products are innovative, highly-configureable neural data acquisition and analysis systems that can meet the design requirements of even the most challenging neurophysiology experiments.


System configurations support both in vivo recordings from anesthetized and awake behaving subjects (birds, mice, rats, cats, primates) and in vitro recordings from cultured microelectrode plates and brain-slice chambers.



>> The **Cerebus System** is a highly-configureable, easy-to-use multichannel data acquisition system for animal neurophysiology experiments. The Cerebus Front-End Amplifier and Neural Signal Processor captures, processes and analyzes action potentials (spikes), field potentials and other physiological signals in relation to experiment state events - in real time.



>> The **Cerebus Stim Switch** provides unprecedented electrode-switching control for neural stimulation applications by enabling high-quality neural recordings immediately after stimulation.



>> Our **Microelectrode Arrays** have become the standard for multi-channel, high-density neural recordings from large populations of neurons.

Blackrock Microsystems™ | 391 Chipeta Way, Ste. G | Salt Lake City, UT 84108
Web: www.blackrockmicro.com | Email: sales@blackrockmicro.com

The RZ2

Z-Series Processor



TDT's System 3 Z-Series processor and preamplifier deliver increased processing, throughput and input channels required to record neurophysiological data from **up to 256 channels**.

- User Configurable Real-Time Processing
- Optically Isolated Direct Digital PreAmps
- Integrated Stimulus Generation
- Powerful Software Control
- High Channel Count, High Sampling Rate System

Also available: RZ5 for low channel count recordings and RZ6 for multi-I/O and audio applications.

Streamlines
data acquisition

Minimizes
post-hoc analysis


Eliminates
data transfer bottlenecks

Increases
realizable sampling rates

Completely
System 3 compatible

Tucker-Davis Technologies • 11930 Research Circle • Alachua, FL 32615
Phone: 386.462.9622 • Fax: 386.462.5365 • E-mail: info@tdt.com

w w w . t d t . c o m



innovative
neurotechnology

toll free | (866) 324-5197
local | (801) 413-0139
fax | (801) 413-2874
email | sales@rppl.com
web | www.rppl.com

grapevine™



The **grapevine™** neural interface system offers a hardware and software platform for neuroscience research and neuroprosthesis development.

connect to



Utah microelectrode array
subdural and μ ECOG electrode grids
surface EEG & EMG electrodes
Michigan microelectrode probe
custom connections for any electrodes

Our systems are compact, portable, and optimized for real-time, closed-loop control experiments with up to 512 electrodes.

A complete system includes a neural interface processor and selected front ends for analog, digital, EMG, EEG, ECoG, LFP, and/or microelectrode array signals.



system specifications

filtering	0.01 Hz to 7.5 kHz input bandwidth additional band selection via digital filtering
input range	± 6.5 mV (± 5.0 V max allowable input)
resolution	16-bit, 0.25 μ V/bit, up to 30 ksp/s
impedance	> 250 M Ω 20 pF (input)
noise	< 2.1 μ Vrms (referred to input)



© 2010 ripple LLC, patents pending

About Cosyne

The annual Cosyne meeting provides a forum for the exchange of experimental and theoretical results in systems neuroscience. Presentations are arranged in a single track, so as to encourage interdisciplinary interactions.

The 2010 meeting consists of 15 invited talks selected by the Executive and Organizing Committees, along with 25 talks and 300 posters selected from the submitted abstracts by the Program Committee. Ten of the poster presentations will also give a four-minute "spotlight" presentation summarizing their submitted work.

The abstracts of the 2010 meeting will be published by Frontiers in Systems Neuroscience. Similar to the abstracts of the Society for Neuroscience meeting, these abstracts are citeable, but they are not full-length proceedings and therefore do not preclude further publication.

Cosyne 2010 Committees

Organizing Committee:

- *General Chair:* Maneesh Sahani (University College London)
- *Program Chairs:* Anne Churchland (University of Washington), Bartlett Mel (University of Southern California)
- *Workshop Chairs:* Adam Kohn (Yeshiva University), Mark Laubach (Yale University)
- *Communications Chair:* Byron Yu (Carnegie Mellon University)

Executive Committee:

- Anthony Zador (Cold Spring Harbor Laboratory)
- Alexandre Pouget (University of Rochester)
- Zachary Mainen (Champalimaud Neuroscience Programme)

Advisory Board:

- Matteo Carandini (University College London)
- Peter Dayan (University College London)
- Steven Lisberger (UC San Francisco and HHMI)
- Eero Simoncelli (New York University and HHMI)
- Karel Svoboda (HHMI Janelia Farms)

About Cosyne

Program Committee:

- Anne Churchland (University of Washington), co-chair
- Bartlett Mel (University of Southern California), co-chair
- Wyeth Bair (Oxford University)
- Sue Becker (McMaster University)
- Nicolas Brunel (CNRS)
- Mark Churchland (Stanford University)
- Yang Dan (University of California, Berkeley)
- Jim DiCarlo (MIT)
- Aapo Hyvarinen (University of Helsinki)
- Vivek Jayaraman (HHMI Janelia Farms)
- Adam Kepecs (Cold Spring Harbor Laboratory)
- Konrad Koerding (North Western University)
- Daniel O'Connor (HHMI Janelia Farms)
- Jonathan Pillow (University of Texas, Austin)
- Jennifer Raymond (Stanford University)
- Eric Shea-brown (University of Washington)
- Thanos Siapas (Caltech)
- Ed Vul (MIT)
- Alex Wade (Smith Kettlewell Institute)
- Mike Wehr (University of Oregon)

Reviewers:

Florin Albeanu, Justin Ales, Greg Appelbaum, Pamela Baker, Andrea Barreiro, Jeff Beck, Robert van Beers, Ulrik Beierholm, Pietro Berkes, Gunnar Blohm, Ed Boyden, Neil Burgess, Pat Byrne, Stijn Cassenaer, Matthew Chafee, Paul Cisek, Marlene Cohen, Ronald Cotton, David Cox, Poppy Crum, Kathleen Cullen, Rhodri Cusack, Peter Dayan, Mike DeWeese, Brent Doiron, Allison Doupe, Joshua Dudman, Alexander Ecker, Chris Eliasmith, Tomer Fekete, Dan Feldman, Gidon Felsen, Maria Geffen, Nicolas Heess, Judith Hirsch, Kresimir Josic, Rex Kerr, David Knill, Takaki Komiyama, John Krakauer, Nikolaus Kriegeskorte, Peter Latham, Mate Lengyel, Michael Lewicki, Chengyu Li, Kenway Louie, Eugene Lubenov, Wei Ji Ma, Christian Machens, Jamie Mazer, Ofer Mazor, Douglas McLelland, Javier Medina, Gianluigi Mongillo, Christopher Moore, Anthony Norcia, Klaus Obermayer, Bruno Olshausen, Anitha Pasupathy, Nicholas Priebe, Sridhar Raghavachari, David Reish, Alfonso Renart, Alex Reyes, Magnus Richardson, Jason Ritt, Jaime de la Rocha, Mark CW van Rossum, Alex Roxin, Nicole Rust, Philip Sabes, Adam Sanborn, Simon Schultz, Odelia Schwartz, Stephen Scott, Walter Senn, Thomas Serre, Harel Shouval, Marshall Shuler, Fritz Sommer, Dominic Standage, Robert Stewart, Karel Svoboda, Andreas S Tolias, Thomas Trappenberg, Matt Tresch, Jochen Triesch, Wilson Truccolo, Glenn Turner, Richard Turner, Thanos Tzounopoulos, Carl van Vreeswijk, Samuel Wang, Casimir Wierzynski, Ben Willmore, Haishan Yao, Angela Yu, Neta Zach

Travel Grants

Thanks to the generosity of our sponsors travel grants are available to support student and postdoc participation at Cosyne.



The recipients of the 2010 **Gatsby Cosyne Fellowships** are: Adil Khan, Adrien Peyrache, Ahna Girshick, Alexander Lerchner, Anita Schmid, Annabelle Singer, Biswa Sengupta, Chris Nolan, Elias Issa, Eran Mukamel, Greg Stephens, Hiroki Asari, Matthijs van der Meer, Michael Famulare, Neda Nategh, Nick Steinmetz, Pedro Goncalves, Pietro Berkes, Pradeep Shenoy, Ralf Häfner, Shin-ichiro Kira, Sridharan Devarajan, Steve Yaeli, Timm Lochmann, Uwe Friederich, Yongseok Yoo.



The recipients of the 2010 **Qualcomm Cosyne Fellowships** are: Annegret Falkner, Matt Nassar, Michael Vidne, Robert Wilson, Timothy Warren.



The recipients of the 2010 **Brain Corporation Cosyne Fellowships** are: Margarida Agrochao, Oren Shriki.

Conference Support

Administrative Support, Registration, Hotels:

- Denise Soudan, Conference and Events Office, University of Rochester

Online Submissions:

- Thomas Preuss, confmaster.net

Program

(Note: institutions listed in the program are the primary affiliation of the first author. For the complete list, please consult the abstracts.)

Thursday, 25 February

- 4:00 pm **Registration opens**
- 6:00 pm **Welcome reception (including buffet and cash bar)**
- 7:20 pm **Introductory Remarks**

Session 1 (Chair: Bartlett Mel)

- 7:30 pm **Keynote address:** [Towards complete structural and functional imaging of cortical circuits](#)
R. Clay Reid, Harvard Medical School 21
- 8:30 pm **Poster Session I**

Friday, 26 February

- 7:30 am **Continental breakfast**

Session 2 (Chair: Ila Fiete)

- 8:30 am [Non linear dendritic processing in cortical pyramidal neurons](#)
Jackie Schiller, Technion Medical School (**invited**) 22
- 9:15 am [Input-dependent switching of inhibitory configurations in neural networks](#)
A. D. Reyes, New York University 22
- 9:30 am [Neuronal biophysics modulate the ability of gamma oscillations to control response timing.](#)
A. Hasenstaub, S. Otte, E. M. Callaway, Crick-Jacobs Center, Salk Institute for Biological
Studies 23
- 9:45 am [Desynchronization of an electrically coupled interneuron network with excitatory synaptic
input](#)
K. Vervaeke, A. Lorincz, P. Gleeson, M. Farinella, Z. Nusser, A. Silver, University College
London 24
- 10:00 am **Refreshment break**

Session 3 (Chair: Peter Latham)

- 10:30 am [Beyond optimality to understanding individual brains: variability, homeostasis and com-
pensation in neuronal circuits](#)
Eve Marder, Brandeis University (**invited**) 24
- 11:15 am [Threshold modulation and homeostatic control of spike timing via circuit plasticity](#)
B. Doiron, Y. Zhao, T. Tzounopoulos, Dept. of Mathematics, Univ. of Pittsburgh 25

Program

11:30 am	Robust spatial working memory through inhibitory gamma synchrony D. Sridharan, S. Millner, J. Arthur, K. Boahen, Dept. of Neurobiology, Stanford University .	26
11:45 am	Influence of task-specific instructions on cross-modal sensory interactions R. Natarajan, R. Zemel, I. Murray, D. Hairston, University of Toronto	27
12:00 pm	Lunch break (and last chance to see Session I posters)	

Session 4 (Chair: Jonathan Pillow)

2:00 pm	Adaptation and inference Adrienne L. Fairhall, University of Washington (invited)	27
2:45 pm	The same neurons form a visual place code and an auditory rate code in the primate SC J. Lee, J. M. Groh, Duke University	28
3:00 pm	Visual influences on information representations in auditory cortex C. Kayser, N. Logothetis, S. Panzeri, Max-Planck-Institute for Biological Cybernetics . . .	29
3:15 pm	Refreshment break	

Session 5 (Chair: Angela Yu)

3:45 pm	Neuroethology of social attention Michael L. Platt, Duke University Medical Center (invited)	29
4:30 pm	Implications of correlated neuronal noise in decision making circuits for physiology and behavior R. M. Haefner, S. Gerwinn, J. H. Macke, M. Bethge, Max-Planck-Institute for Biological Cybernetics	30
4:45 pm	Spotlights Attention modeled as a two-dimensional neural resource D. Ballard, University of Texas at Austin	121
	A common-input model of a complete network of ganglion cells in the primate retina. M. Vidne, Y. Ahmadian, J. Shlens, J. W. Pillow, J. Kulkarni, E. P. Simoncelli, E. J. Chichilnisky, L. Paninski, Center for Theoretical Neuroscience, Columbia University	133
	Stability and competition in multi-spike models of spike-timing dependent plasticity B. Babadi, L. F. Abbott, Center for Theoretical Neuroscience, Columbia University	149
	Preparatory tuning in premotor cortex relates most closely to the population movement-epoch response M. M. Churchland, M. Kaufman, J. P. Cunningham, K. Shenoy, Stanford University	156
	Optimal neuronal tuning curves - an exact Bayesian study of dynamic adaptivity S. Yaeli, R. Meir, Department of Electrical Engineering, Technion	170
5:00 pm	Dinner break	
7:30 pm	Poster Session II	

Saturday, 27 February

7:30 am	Continental breakfast	
---------	------------------------------	--

Special session honoring Horace Barlow's legacy (Chair: Daniel Wolpert)

8:30 am	Efficiency, redundancy and sparse coding: Just a portion of Barlow's legacy David J. Field, Cornell University (invited)	30
---------	--	----

9:00 am	Predictions of visual performance from the statistical properties of natural scenes Wilson S. Geisler, University of Texas, Austin (invited)	31
9:30 am	Eyes 3, 4 and 5 would be most mystifying structures if one did not know that flies flew Simon B. Laughlin, University of Cambridge (invited)	31
10:00 am	Refreshment break	
10:30 am	A generative model of the covariance structure of images Geoffrey E. Hinton, University of Toronto (invited)	32
11:00 am	Evidence for a neural model to evaluate symmetry in V1 Horace B. Barlow, University of Cambridge (invited)	32
12:00 pm	Lunch break (and last chance to see Session II posters)	

Session 8 (Chair: Pamela Reinagel)

2:00 pm	Timing in the auditory cortex Anthony M. Zador, Cold Spring Harbor Laboratory (invited)	33
2:45 pm	Differential sensitivity of different sensory cortices to behaviorally relevant timing differences Y. Yang, A. M. Zador, Cold Spring Harbor Laboratory	33
3:00 pm	The Poisson clicks task: long time constant of neural integration of discrete packets of evidence B. W. Brunton, C. D. Brody, Princeton Neuroscience Institute	34
3:15 pm	Refreshment break	

Session 9 (Chair: Konrad Körding)

3:45 pm	Action video games as exemplary learning tools Daphne Bavelier, University of Rochester (invited)	34
4:30 pm	Pupillometric evidence for a role of locus coeruleus in dynamic belief updating M. Nassar, R. C. Wilson, R. Kalwani, B. Heasly, J. Gold, University of Pennsylvania	35
4:45 pm	Spotlights Dynamical control of eye movements in an active visual search task: theory and experiments H. He, J. Schilz, A. J. Yu, Department of Cognitive Science, University of California, San Diego	195
	Role of anterior cingulate cortex in patch-leaving foraging decisions B. Hayden, M. Platt, Duke University	199
	Ensemble activity underlying movement preparation in prearcuate cortex R. Kalmar, J. Reppas, S. Ryu, K. Shenoy, W. Newsome, Stanford University	224
	Temporal precision of the olfactory system R. Shusterman, M. Smear, T. Bozza, D. Rinberg, Janelia Farm, HHMI	243
	A normalization model of multi-sensory integration T. Ohshiro, D. E. Angelaki, G. C. DeAngelis, University of Rochester	247
5:00 pm	Dinner break	
7:30 pm	Poster Session III	

Program

Sunday, 28 February

7:30 am **Continental breakfast**

Session 10 (Chair: Nicole Rust)

8:30 am	Detection and estimation of defocus in natural images J. Burge, W. S. Geisler, Center for Perceptual Systems, University of Texas, Austin	36
8:45 am	Invariant contrast coding in photoreceptors U. Friederich, D. Coca, S. A. Billings, M. Juusola, University of Sheffield	36
9:00 am	Spike-triggered covariance and synthetic image replay reveal nonlinearities in V1 color processing G. Horwitz, University of Washington	37
9:15 am	Metamers of the ventral stream J. Freeman, E. P. Simoncelli, Center for Neural Science, NYU	38
9:30 am	The control of visual information by prefrontal dopamine Tirin Moore, Stanford University (invited)	39
10:15 am	Refreshment break	

Session 11 (Chair: Dora Angelaki)

10:45 am	High frequency entrainment of thalamic neurons by basal ganglia output in the singing bird J. H. Goldberg, M. S. Fee, Massachusetts Institute of Technology	39
11:00 am	Beside the point: Motor adaptation without feedback error correction in task-irrelevant conditions S. Y. Schaefer, I. L. Shelly, K. A. Thoroughman, Department of Physical Therapy, Washington U.	40
11:15 am	Conscious or not? How neuroscience is building a bridge to understanding recovery following severe brain injury Nicholas D. Schiff, Weill Cornell Medical College (invited)	40
12:00 pm	Lunch break (hotel checkout, and last chance to see Session III posters)	

Session 12 (Chair: Vivek Jayaraman)

2:00 pm	Hippocampal processes underlying episodic memory John Lisman, Brandeis University (invited)	41
2:45 pm	Coordinated hippocampal firing across related spatial locations develops with experience A. C. Singer, M. P. Karlsson, A. R. Nathe, M. F. Carr, L. M. Frank, Keck Center and Department of Physiology UCSF	41
3:00 pm	Temporal transformations in olfactory encoding promote rapid detection of natural odor fluctuations K. Nagel, R. Wilson, Harvard Medical School	42
3:15 pm	Experimental evolution to probe gene networks underlying cognition in <i>Drosophila</i> Josh Dubnau, Cold Spring Harbor Laboratory (invited)	43
4:00 pm	Closing remarks	

Poster Session I

8:30 pm Thursday 25th February

I-1. Intrinsic dendritic plasticity maximally increases the computational power of CA1 pyramidal neurons. R. Cazé, M. D. Humphries, B. Gutkin, INSERM U960	43
I-2. Analytical study of history dependent timescales in a generic model of ion channels D. Soudry, Y. Meir, R. , Electrical Engineering, Technion	44
I-3. Fast Kalman filtering on quasilinear dendritic trees L. Paninski, Columbia University	45
I-4. Dendritic spine plasticity can stabilize synaptic weights C. O'Donnell, M. F. Nolan, M. C. W. van Rossum, University of Edinburgh	45
I-5. Model of synaptic plasticity based on self-organization of PSD-95 molecules in spiny dendrites. D. Tsigankov, S. Eule, Max-Planck Institute for Dynamics and Self-Organization	46
I-6. Trajectory prediction combining forward models and historical knowledge J. O'Reilly, T. Behrens, FMRIB Centre, University of Oxford	47
I-7. Dynamics of fronto-parietal synchrony during working memory. N. M. Dotson, R. F. Salazar, C. M. Gray, Montana State University	47
I-8. Bayesian optimal use of visual feature cues in visual attention B. Vincent, University of Dundee	48
I-9. Neural correlates of spatial short-term memory in the rodent frontal orienting field J. C. Erlich, M. Bialek, C. D. Brody, HHMI & Princeton University	49
I-10. When to recall a memory? Epoch dependent memory trace with a power law of timescales in ACC neurons. A. Bernacchia, H. Seo, D. Lee, X.-J. Wang, Department of Neurobiology, Yale University	49
I-11. Neural encoding of decision uncertainty in prefrontal cortex R. J. Cotton, A. Laudano, A. S. Tolias, Baylor College of Medicine	50
I-12. Clocking perceptual processing speed: from chance to 75%correct in 30 milliseconds T. R. Stanford, S. Shankar, D. Massoglia, G. Costello, E. Salinas, Wake Forest University School of Medicine	51
I-13. The effect of time pressure on decision making S. Kira, M. N. Shadlen, Dept. Physiology & Biophysics, NPRC; U. WA	52
I-14. Decision-related activity in area V2 for a fine disparity discrimination task. H. Nienborg, B. G. Cumming, Salk Institute	52
I-15. Changes in functional connectivity in LIP during a free choice task A. Falkner, M. E. Goldberg, Columbia University	53
I-16. Role of secondary motor cortex in withholding impulsive action: Inhibition or competition? A. Fonseca, M. Murakami, M. Vicente, G. Costa, Z. Mainen, Champalimaud Neuroscience Program at IGC	54
I-17. Dorsomedial prefrontal cortex encodes value information during a sequential choice task C.-H. Luk, J. D. Wallis, Helen Wills Neuroscience Institute, UCB	54
I-18. Striatal activity consistent with model-based, rather than model-free prediction errors C. S. Green, P. Zhang, N. Daw, D. Kersten, S. He, P. Schrater, University of Minnesota	55
I-19. From integrate-and-fire neurons to Generalized Linear models S. Ostoic, N. Brunel, Center for Theoretical Neuroscience, Columbia University	56
I-20. Excitatory-inhibitory correlations result from opposing correlating and anticorrelating forces C. Omar, J. W. Middleton, D. J. Simons, B. Doiron, Center for the Neural Basis of Cognition, Carnegie Mellon University	57
I-21. Saliency and surround interactions via natural scene statistics: A unifying model. R. Coen-Cagli, P. Dayan, O. Schwartz, Department of Neuroscience, Albert Einstein College of Medicine	57

I-22. A three-layer model of natural image statistics A. Hyvarinen, M. Gutmann, University of Helsinki	58
I-23. Learning L_p spherical potentials for Markov Random Field models of natural images U. Koster, M. Gutmann, A. Hyvärinen, University of Helsinki	59
I-24. Grasping image statistics O. Aladini, C. A. Rothkopf, J. Triesch, Frankfurt Institute for Advanced Studies	60
I-25. Beyond magical numbers: towards a noise-based account of visual-short term memory limitations W. J. Ma, W.-C. Chou, Baylor College of Medicine	60
I-26. Identifiability of nonlinear receptive field models from sensory neurophysiology data C. DiMattina, K. Zhang, Electrical Engineering & Computer Science, Case Western Reserve University	61
I-27. Detecting a change by a single neuron H. Kim, B. J. Richmond, S. Shinomoto, Kyoto University	62
I-28. Complexity and performance in simple neuron models S. Mensi, R. Naud, M. Avermann, C. Petersen, W. Gerstner, EPFL - LCN	63
I-29. Short-term synaptic plasticity and sensory adaptation as Bayesian inference I. H. Stevenson, B. Cronin, M. Sur, K. Kording, Northwestern University	63
I-30. The speed of time M. Ahrens, M. Sahani, Cambridge University	64
I-31. Reconstruction of sparse circuits using multi-neuronal excitation (RESCUME) T. Hu, D. Chklovskii, HHMI, Janelia Farm Research Campus	65
I-32. On the connections between SIFT and biological vision K. Muralidharan, N. Vasconcelos, Statistical Visual Computing Laboratory, UCSD	65
I-33. The pairwise phase consistency: A bias-free measure of rhythmic neuronal synchronization M. Vinck, M. van Wingerden, T. Womelsdorf, P. Fries, C. Pennartz, University of Amsterdam	66
I-34. Physiology in <i>Drosophila</i> motion-sensitive neurons during walking and flight behavior J. D. Seelig, M. E. Chiappe, G. K. Lott, M. B. Reiser, V. Jayaraman, Janelia Farm Research Campus, HHMI	67
I-35. The frequency of hippocampal theta oscillations and unit firing can be manipulated by changing the t E. Pastalkova, G. Buzsáki, Janelia Farm Research Campus, HHMI	68
I-36. Compressed sensing in the brain: role of sparseness in short term and long term memory S. Ganguli, H. Sompolinsky, UCSF	68
I-37. A spike timing computational model of hippocampal-frontal dynamics underlying navigation and memory L. Jayet, P. H. Goodman, M. Quoy, Brain Computation Lab, University of Nevada Reno	69
I-38. Interaction of hippocampo-neocortical neuronal assemblies during learning and sleep A. Peyrache, F. P. Battaglia, UNIC, CNRS	70
I-39. Spontaneous activity in a self-organizing recurrent network reflects prior learning A. Lazar, G. Pipa, J. Triesch, Frankfurt Institute for Advanced Studies	70
I-40. Spike timing-dependent plasticity interacts with neural dynamics to enhance information transmission G. Hennequin, J.-P. Pfister, W. Gerstner, EPFL, SV and Brain-Mind Institute	71
I-41. Integration of new and old auditory memories in the European starling. D. Zaraza, D. Margoliash, University of Chicago	72
I-42. An avian basal ganglia circuit contributes to fast and slow components of songbird vocal learning T. Warren, E. Tumer, M. Brainard, Keck Center, UCSF	72
I-43. The BOLD response in the nucleus accumbens quantitatively represents the reward prediction error E. E. J. DeWitt, P. Glimcher, New York University	73

I-44. Internal time temporal difference model of neural valuation S. Kaveri, H. Nakahara, Lab for Int Theor Neurosci, RIKEN BSI	74
I-45. Does one simulate the other's value-based decision making by using the neural systems for his own? S. Suzuki, N. Harasawa, K. Ueno, S. Kaveri, J. Gardner, N. Ichinohe, M. Haruno, K. Cheng, H. Nakahara, Lab for Int Theor Neurosci, Riken BSI	75
I-46. Prefrontal neurons solve the temporal credit assignment problem during reinforcement learning W. F. Asaad, E. N. Eskandar, Massachusetts General Hospital	76
I-47. Changes in the response rate and response variability of area V4 neurons during saccade preparation N. Steinmetz, T. Moore, Stanford University	76
I-48. Neurophysiological evidence for basal ganglia involvement in speed-accuracy tradeoff in monkeys M. Watanabe, T. Trappenberg, D. Munoz, Queen's University	77
I-49. Modelling basal ganglia and superior colliculus in the antisaccade task T. Trappenberg, M. Watanabe, D. Munoz, Dalhousie University	78
I-50. Integration of visual and proprioceptive information for reaching in multiple parietal areas. L. McGuire, P. Sabes, University of California San Francisco	79
I-51. Sensory integration in PMd: position-dependent dynamic reweighting of vision and proprioception M. Fellows, P. Sabes, UCSF, Dept. of Physiology and Keck Center	79
I-52. A new notion of criticality: Studies in the pheromone system of the moth C. L. Buckley, T. Nowotny, CCNR, Informatics, University of Sussex	80
I-53. Cellular imaging in behaving mice reveals learning-related specificity in motor cortex circuits T. Komiya, T. R. Sato, D. H. O'Connor, Y.-X. Huber, D. Hooks, B. M. Gabitto, M. Svoboda, K. , Janelia Farm Research Campus, HHMI	81
I-54. Neural mechanisms underlying the reduction in behavioral variability during trial-and-error learning A. Dubreuil, Y. Burak, T. Otchy, B. Ölveczky, Ecole Normale Supérieure Cachan	82
I-55. Evidence for a central pattern generator built on a heteroclinic channel instead of a limit cycle K. M. Shaw, H. Lu, J. M. McManus, M. J. Cullins, H. J. Chiel, P. J. Thomas, Case Western Reserve University	83
I-56. A neural microcircuit using spike timing for novelty detection C. R. Nolan, G. Wyeth, M. Milford, J. Wiles, School of ITEE & Queensland Brain Institute, The University of Queensland	83
I-57. Using natural stimuli to estimate receptive fields in neurons that employ sparse coding G. Isely, C. Hillar, F. Sommer, Redwood Center for Theoretical Neuroscience, UC Berkeley	84
I-58. Comparison of V1 receptive fields mapped with spikes and local field potentials F. Biessmann, F. Meinecke, A. Bhattacharyya, J. Veit, R. Kretz, K.-R. Müller, G. Rainer, TU Berlin, Dept. Machine Learning	85
I-59. A novel method to estimate information transfer between two continuous signals of finite duration J. Takalo, I. Ignatova, M. Weckström, M. Vähäsöyrinki, University of Oulu	86
I-60. The firing irregularity as the firing characteristic orthogonal to the firing rate T. Shimokawa, S. Shinomoto, Kyoto University	86
I-61. Optimal information transfer in the cortex through synchronisation A. Buehlmann, G. Deco, Universitat Pompeu Fabra, Barcelona	87
I-62. Two dimensions for the price of one: the efficient encoding of vertical disparity J. Read, Institute of Neuroscience, Newcastle University	88
I-63. Analysis of subsets of higher-order correlated neurons based on marginal correlation coordinates H. Shimazaki, S. Gruen, S.-i. Amari, RIKEN Brain Science Institute	88
I-64. Spike latency code for orientation discrimination and estimation by primary visual cortical cells O. Shriki, A. Kohn, M. Shamir, Ben-Gurion University	89

I-65. Dissecting the action potential onset rapidness on the response speed of neuronal populations W. Wei, F. Wolf, MPI for Dynamics and Self-Organization	90
I-66. Diversity of efficient coding solutions for a population of noisy linear neurons E. Doi, L. Paninski, E. P. Simoncelli, New York University	90
I-67. Orientation and direction selectivity in the population code of the visual thalamus G. B. Stanley, J. Jin, Y. Wang, G. Desbordes, M. J. Black, J.-M. Alonso, Biomedical Engineering, Georgia Tech/Emory University	91
I-68. Visual hyperacuity despite fixational eye movements: a network model O. Mazor, Y. Burak, M. Meister, Harvard University	92
I-69. Sparse coding in modular networks E. Dyer, D. Johnson, R. Baraniuk, Rice University	93
I-70. Suppression of intrinsic cortical response variability is state- and stimulus-dependent B. White, L. Abbott, J. Fiser, Program in Neuroscience, Brandeis University	93
I-71. Background synaptic activity modulates spike train correlation A. L. Kumar, M. J. Chacron, B. Doiron, Center for the Neural Basis of Cognition, Carnegie Mellon	94
I-72. Dynamic population coding with recurrent networks of integrate and fire neurons M. Boerlin, S. Deneve, Group for Neural Theory, LNC, DEC, ENS Paris	95
I-73. Noise correlations in area MSTd are weakened in animals trained to perform a discrimination task Y. Gu, S. Fok, A. Sunkara, S. Liu, G. DeAngelis, D. Angelaki, Washington University School of Medicine	96
I-74. Sound texture perception via synthesis J. H. McDermott, A. J. Oxenham, E. P. Simoncelli, Center for Neural Science, New York University	96
I-75. Manipulation of sound-driven decisions by microstimulation of auditory cortex P. Znamenskiy, A. M. Zador, Watson School of Biological Sciences	97
I-76. A generalized linear model for estimating receptive fields from midbrain responses to natural sounds A. Calabrese, J. Schumacher, D. Schneider, S. Woolley, L. Paninski, Columbia University	98
I-77. Identification of excitation and inhibition in the auditory cortex using nonlinear modeling N. Schinkel-Bielefeld, S. V. David, S. A. Shamma, D. A. Butts, Department of Biology, University of Maryland	99
I-78. Top-down influences on intensity coding in primary auditory cortex L. S. Hamilton, S. Bao, University of California, Berkeley	100
I-79. Neural encoding of global statistical features of natural sounds. M. Neimark Geffen, T. Taillefumier, M. Magnasco, Rockefeller University	100
I-80. Contour representation of sound signals Y. Lim, B. G. Shinn-Cunningham, T. Gardner, Department of Cognitive and Neural Systems, Boston University	101
I-81. Neural activity as samples from a probabilistic representation: evidence from the auditory cortex P. Berkes, S. V. David, J. B. Fritz, M. Lengyel, S. A. Shamma, J. Fiser, Brandeis University	101
I-82. Two photon imaging of tactile responses during frequency discrimination in awake head fixed rats F. Haiss, J. Mayrhofer, D. Margolis, M. T. Hasan, F. Helmchen, B. Weber, University of Zurich	102
I-83. From form to function: deriving preferred stimuli from neuronal morphology J. Mulder-Rosi, G. Cummins, J. Miller, Montana State University	103
I-84. Receptive field mapping of local populations in mouse visual cortex using two-photon calcium imaging V. Bonin, M. H. Histed, R. C. Reid, Harvard Medical School	104
I-85. The encoding of fine spatial information in salamander retinal ganglion cells F. Soo, G. Schwartz, M. J. Berry II, Princeton University	105
I-86. The projective field of single bipolar cells in the retina H. Asari, M. Meister, Department of Molecular and Cellular Biology, Harvard University	105

I-87. Contribution of amacrine transmission to fast adaptation of retinal ganglion cells N. Nategh, M. Manu, S. Baccus, Stanford University	106
I-88. Perception of the reverse-phi illusion by <i>Drosophila melanogaster</i> J. C. Tuthill, M. E. Chiappe, V. Jayaraman, M. B. Reiser, Janelia Farm Research Campus, HHMI	106
I-89. The structure of spontaneous and evoked population activity in mouse visual cortex S. Hofer, B. Pichler, H. Ko, J. T. Vogelstein, N. Lesica, T. Mrsic-Flogel, University College London	107
I-90. Modelling molecular mechanisms of light adaptation in <i>Drosophila</i> photoreceptor Z. Song, D. Coca, S. Billings, M. Juusola, The University of Sheffield	108
I-91. The role of EAG K ⁺ channels in insect photoreceptors E.-V. Immonen, R. Frolov, M. Vahasoyrinki, M. Weckstrom, University of Oulu, Department of Physics	109
I-92. Memory-related activity in the PFC depends on cell type only in the absence of sensory stimulation C. Hussar, T. Pasternak, Department of Neurobiology and anatomy, University of Rochester	109
I-93. Compete globally, cooperate locally: Signal integration for behavior depends on cortical separation K. Ghose, J. H. R. Maunsell, Harvard Medical School	110
I-94. The role of inhibition in formatting visual information in the retina and LGN D. A. Butts, A. R. R. Casti, Dept. of Biology and Program in Neuroscience, University of Maryland	111
I-95. Towards large-scale, high resolution maps of object selectivity in inferior temporal cortex E. B. Issa, A. Papanastassiou, B. B. Andken, J. J. DiCarlo, McGovern Inst/Dept of Brain & Cog Sci, MIT	112
I-96. Recording a large population of retinal cells with a 252 electrode array and automated spike sorting O. Marre, D. Amodè, F. Soo, T. E. Holy, M. Berry, Princeton University	112
I-97. Models for the mechanisms of perceptual learning: linking predictions for brain and behavior M. Wenger, V. D. H. Rebecca, Department of Psychology, The Pennsylvania State University	113
I-98. Disparity tuning of the population responses in the human visual cortex: an EEG source imaging study B. R. Cottreau, A. M. Norcia, S. P. Mckee, The Smith-Kettlewell Eye Research Institute	114
I-99. Sources of response variability underlying contrast-invariant orientation tuning in visual cortex S. Sadagopan, N. Priebe, I. Finn, D. Ferster, Dept. of Neurobiology and Physiology, Northwestern University	115
I-100. Lateral Occipital cortex responsive to local correlation structure of natural images H. S. H. Steven, S. Ghebreab, A. Smeulders, V. Lamme, University of Amsterdam, Dep. of Psychology	115

Poster Session II

7:30 pm Friday 26th February

II-1. Beyond linear perturbation theory: the instantaneous response of the integrate-and-fire model	
M. Helias, M. Deger, S. Rotter, M. Diesmann, RIKEN Brain Science Institute	116
II-2. Gamma oscillations in the optic tectum in vitro represent top-down drive by a cholinergic nucleus	
C. A. Goddard, D. Sridharan, J. Huguenard, E. Knudsen, Dept of Neurobiology, Stanford University	117
II-3. Parallel channels in the OFF visual pathway emerge at the cone synapse	
C. P. Ratliff, S. H. DeVries, Northwestern University	117
II-4. The logic of cross-columnar interactions along horizontal circuits	
H. Adesnik, M. Scanziani, HHMI	118
II-5. Bang-bang optimality of energy efficient spikes in single neuron models	
B. Sengupta, M. Stemmler, J. Niven, A. Herz, S. Laughlin, University of Cambridge	119
II-6. Cognitive control recruits theta activity in anterior cingulate cortex for establishing task rules	
T. Womelsdorf, K. Johnston, M. Vinck, S. Everling, Department of Physiology & Pharmacology, University of Western Ontario	120
II-7. Neural mechanisms of competitive interaction in recurrent maps of a visual pathway	
D. Lai, R. Wessel, Dept. of Physics, Washington University	120
II-8. Attention modeled as a two-dimensional neural resource	
D. Ballard, University of Texas at Austin	121
II-9. Evidence for a race model	
D. Liston, L. Stone, San Jose State University Foundation	122
II-10. Adaptive properties of differential learning rates for positive and negative outcomes	
R. Cazé, v. d. M. Matthijs, Group for Neural Theory, ENS, Paris	122
II-11. A Bayesian model of simple value-based decision-making	
D. Ray, A. Rangel, CNS, Caltech	123
II-12. The rational control of aspiration in learning	
D. Acuna, C. S. Green, P. Schrater, Dept. Computer Science and Engineering, University of Minnesota .	124
II-13. Bonsai trees: How the Pavlovian system sculpts sequential decisions	
Q. J. M. Huys, N. Eshel, P. Dayan, J. P. Roiser, Gatsby Unit and Neuroimaging Centre, UCL	124
II-14. The temporal dynamics of human decision under risk	
L. Hunt, M. Rushworth, T. E. Behrens, FMRIB Centre, University of Oxford	125
II-15. Approaching avoidance: asymmetries in reward and punishment processing	
Q. J. M. Huys, R. Cools, M. Goelzer, E. Friedel, R. J. Dolan, A. Heinz, P. Dayan, Gatsby Unit and Neuroimaging Centre, UCL	126
II-16. Neurons in area LIP encode perceptual decisions in a perceptual, not oculomotor, frame of reference	
S. Bannur, J. Gold, University of Pennsylvania	127
II-17. Context-dependent gating of sensory signals for decision making	
V. Mante, W. T. Newsome, HHMI and Stanford University	128
II-18. The effect of value normalization and cortical variability on rational choice	
K. Louie, P. Glimcher, New York University	128
II-19. Thalamocortical changes in clinical depression probed by physiology-based modeling	
C. Kerr, A. Kemp, C. Rennie, P. Robinson, School of Physics, University of Sydney	129
II-20. Multistability as a mechanism for modulation of EEG coherences	
J. Drover, J. Victor, S. T. Williams, M. Conte, N. Schiff, Weill Medical College of Cornell University	130
II-21. A robust, bilateral line attractor model of the oculomotor system with ordinary neurons	
P. Goncalves, C. K. Machens, Group for Neural Theory, ENS, Paris	131

II-22. Cortical activity demystified: a unifying theory that explains state switching in cortex	
A. Lerchner, P. E. Latham, Gatsby Computational Neuroscience Unit, UCL	131
II-23. Pattern separation by adaptive networks: neurogenesis in olfaction	
S. F. Chow, S. D. Wick, H. Riecke, Northwestern U.	132
II-24. A common-input model of a complete network of ganglion cells in the primate retina.	
M. Vidne, Y. Ahmadian, J. Shlens, J. W. Pillow, J. Kulkarni, E. P. Simoncelli, E. J. Chichilnisky, L. Paninski, Center for Theoretical Neuroscience, Columbia University	133
II-25. Closed-form correlation-based identification of recurrent spiking networks	
M. Krumin, A. Tankus, S. Shoham, Technion - IIT	134
II-26. Neuronal variability and linear neural coding in the vestibular system	
A. Schneider, M. Chacron, K. Cullen, McGill University	134
II-27. Transition-state theory for integrate-and-fire neurons	
L. Badel, W. Gerstner, M. J. E. Richardson, Department of Statistics, Columbia University	135
II-28. Auditory textures and primitive auditory scene analysis	
R. E. Turner, M. Sahani, Computational and Biological Learning Lab	136
II-29. Bayesian Pitch	
P. Hehrmann, M. Sahani, Gatsby Computational Neuroscience Unit, UCL	136
II-30. Bayesian belief propagation and border-ownership signals in early visual cortex	
H. Hosoya, University of Tokyo	137
II-31. Orientation maps as Moiré interference of retinal ganglion cell mosaics	
S.-B. Paik, D. Ringach, Department of Neurobiology, UCLA	138
II-32. Estimation and assessment of non-Poisson neural encoding models	
J. W. Pillow, University of Texas at Austin	138
II-33. Bayesian line orientation perception: Human prior expectations match natural image statistics	
A. R. Girshick, M. S. Landy, E. P. Simoncelli, Dept of Psychology & Center for Neural Science, New York University	139
II-34. The value of lateral connectivity in visual cortex for interpreting naturalistic images	
X. Pitkow, Y. Ahmadian, K. Miller, Center for Theoretical Neuroscience, Columbia University	140
II-35. Perturbation of hippocampal cell dynamics by halorhodopsin-assisted silencing of PV interneurons	
S. Royer, B. Zemelman, A. Losonczy, J. Magee, G. Buzsaki, Janelia Farm Research Campus, HHMI	140
II-36. Behavioral state continuously modulates hippocampal information processing	
C. Kemere, F. Zhang, K. Deisseroth, L. M. Frank, UCSF	141
II-37. Hippocampal learning and cognitive maps as products of hierarchical latent variable models	
A. Johnson, Z. Varberg, P. Schrater, Bethel University	141
II-38. Computational role of theta oscillations in delayed-decision tasks	
P. Joshi, Frankfurt Institute for Advanced Studies	142
II-39. Prefrontal and hippocampal coding during long-term memory formation in monkeys	
S. L. Brincat, E. K. Miller, Picower Institute for Learning & Memory, MIT	143
II-40. How neurogenesis and modulation affect network oscillations in a large-scale dentate gyrus model.	
J. B. Aimone, F. H. Gage, Salk Institute	144
II-41. A computational approach to neurogenesis and synaptogenesis using biologically plausible neurons	
L. N. Long, A. Gupta, G. Fang, The Pennsylvania State University	144
II-42. Why is connectivity in barrel cortex different from that in visual cortex? - A plasticity model.	
C. Clopath, L. Büsing, E. Vasilaki, W. Gerstner, LCN	145
II-43. Rapid feature binding by a learning rule that integrates branch strength potentiation and STDP	
R. Legenstein, W. Maass, Graz University of Technology, IGI	146

II-44. Bimodal structural plasticity can explain the spacing effect in long-term memory tasks	
A. Knoblauch, Honda Research Institute Europe	146
II-45. The role of dopamine in long-term plasticity in the rat prefrontal cortex: a computational model	
D. Sheynikhovich, S. Otani, A. Arleo, Lab. of Neurobiology of Adaptive Processes	147
II-46. Structural plasticity improves stimulus encoding in a working memory model	
C. Savin, J. Triesch, Frankfurt Institute for Advanced Studies	148
II-47. Vigour in the face of fluctuating rates of reward: An experimental test	
U. Beierholm, M. Guitart Masip, R. Dolan, E. Duzel, P. Dayan, Gatsby Computational Neuroscience Unit, UCL	149
II-48. Stability and competition in multi-spike models of spike-timing dependent plasticity	
B. Babadi, L. F. Abbott, Center for Theoretical Neuroscience, Columbia University	149
II-49. Risk-minimization through Q-learning of the learning rate	
K. Preuschoff, P. Bossaerts, Social and Neural Systems Lab, University of Zurich	150
II-50. Model averaging as a developmental outcome of reinforcement learning.	
T. H. Weisswange, C. A. Rothkopf, T. Rodemann, J. Triesch, Frankfurt Institute for Advanced Studies	151
II-51. Learning to plan: planning as an action in simple reinforcement learning agents	
G. E. Wimmer, M. van der Meer, Columbia University	152
II-52. Dynamics of frontal eye field and cerebellar activity during smooth pursuit learning	
J. Li, J. Medina, L. Frank, S. Lisberger, UCSF	153
II-53. Idiosyncratic and systematic features of spatial representations in the macaque PRR	
S. W. C. Chang, L. H. Snyder, Duke Institute for Brain Sciences	153
II-54. High-performance continuous neural cursor control enabled by a feedback control perspective	
V. Gilja, P. Nuyujukian, C. Chestek, J. Cunningham, B. Yu, S. Ryu, K. Shenoy, Stanford University	154
II-55. The emergence of stereotyped behaviors in C. elegans	
G. Stephens, W. Ryu, W. Bialek, Lewis-Sigler Institute	155
II-56. Preparatory tuning in premotor cortex relates most closely to the population movement-epoch response	
M. M. Churchland, M. Kaufman, J. P. Cunningham, K. Shenoy, Stanford University	156
II-57. Sparse connectivity in short-term memory networks	
D. Fisher, E. Aksay, M. Goldman, Center for Neuroscience, UC Davis	156
II-58. Modeling firing-rate dynamics: From spiking to firing-rate networks	
E. S. Schaffer, L. F. Abbott, Columbia University	157
II-59. Optimal network architectures for short-term memory under different biological settings	
S. Lim, M. Goldman, Center for Neuroscience, UC Davis	158
II-60. Neuroptikon: a customizable tool for dynamic, multi-scale visualization of complex neural circuits	
F. Midgley, D. J. Olbris, D. Chklovskii, V. Jayaraman, Janelia Farm Research Campus, HHMI	159
II-61. Near exact correction of path integration errors by the grid cell-place cell system	
S. Sreenivasan, I. Fiete, Center for Learning and Memory, University of Texas at Austin	159
II-62. Deciding with single spikes: MT discharge and rapid motion detection	
B. Krause, G. Ghose, University of Wisconsin	160
II-63. Testing efficient coding: projective (not receptive) fields are the key theoretical prediction	
E. Doi, G. Field, J. Gauthier, A. Sher, M. Greschner, J. Shlens, T. Machado, L. Paninski, D. Gunning, K. Mathieson, A. Litke, E. J. Chichilnisky, E. P. Simoncelli, New York University	161
II-64. The role the retina plays in shaping predictive information in ganglion cell populations	
S. E. Palmer, M. J. Berry, W. Bialek, Princeton University	162
II-65. Odour identity is represented by the pattern of activated neurons in the drosophila mushroom body	
R. Campbell, G. C. Turner, K. Honegger, Cold Spring Harbor Laboratory	163

II-66. Efficient theta-locked population codes in olfactory cortex K. Miura, Z. Mainen, N. Uchida, JST PRESTO	163
II-67. Temporally distributed information gets optimally combined by change-based information processing R. Moazzezi, P. Dayan, Redwood Center for Theoretical Neuroscience, UC Berkeley	164
II-68. Fisher information in correlated networks D. G. T. Barrett, P. E. Latham, Gatsby Computational Neuroscience Unit, UCL	165
II-69. Positive reinforcement increases pooled population-coding efficacy in the auditory forebrain J. Jeanne, T. Sharpee, T. Gentner, UC San Diego	166
II-70. Online readout of frequency information in areas SI and SII A. Wohrer, R. Romo, C. K. Machens, Ecole Normale Supérieure	166
II-71. One-dimensional dynamics of associative representations in lateral intraparietal (LIP) area J. K. Fitzgerald, D. Freedman, A. Fanini, J. Assad, Department of Neurobiology, Harvard Medical School	167
II-72. Modelling visual crowding of complex stimuli S. C. Dakin, P. Bex, J. Greenwood, Institute of Ophthalmology, University College London	168
II-73. Decoding stimulus velocity from population responses in area MT of the macaque A. A. Stocker, N. Majaj, C. Tailby, J. A. Movshon, E. P. Simoncelli, Department of Psychology, University of Pennsylvania	168
II-74. Neural correlates of dynamic sensory cue re-weighting in macaque area MSTd C. R. Fetsch, G. C. DeAngelis, D. E. Angelaki, Washington University School of Medicine	169
II-75. Optimal neuronal tuning curves - an exact Bayesian study of dynamic adaptivity S. Yaeli, R. Meir, Department of Electrical Engineering, Technion	170
II-76. One rule to grow them all: A general theory of neuronal branching and its practical application H. Cuntz, F. Forstner, A. Borst, M. Häusser, University College London	171
II-77. Columnar transformation of neural response to time-varying sounds in auditory cortex P. Crum, X. Wang, Johns Hopkins School of Medicine	171
II-78. Is multisensory integration Hebbian? Ventriloquism aftereffect w/o simultaneous audiovisual stimuli D. Pages, J. M. Groh, Duke University	172
II-79. The kinetics of fast short-term depression are matched to spike train statistics to reduce noise W. Nesse, R. Khanbabaie, A. Longtin, L. Maler, Department of Cellular and Molecular Medicine, University of Ottawa	173
II-80. Frequency-invariant representation of interaural time differences H. Lüling, I. Siveke, B. Grothe, C. Leibold, Ludwig-Maximilians-Universität, München	173
II-81. Behavioral context in pigeons: motor output and neural substrates K. McArthur, J. D. Dickman, Department of Anatomy & Neurobiology, Washington University School of Medicine	174
II-82. Systematic analyses of receptive field of mammalian olfactory glomeruli L. Ma, S. Gradwohl, Q. Qiu, R. Alexander, W. Wiegand, R. Yu, Stowers Institute for Medical Research	175
II-83. Decoding intensity-tuned neurons in the auditory system E. N. Marongelli, P. V. Watkins, D. L. Barbour, Department of Biomedical Engineering, Washington University in Saint Louis	175
II-84. The structure of human olfactory space A. Koulakov, A. Enikolopov, D. Rinberg, Cold Spring Harbor Laboratory	176
II-85. A mechanism that governs a transition from coincidence detection to integration in olfactory network C. Assisi, M. Bazhenov, University of California, Riverside	177
II-86. Visual features evoke reliable bursts in the perigeniculate sector of the thalamic reticular nucleus V. Vaingankar, C. Soto Sanchez, X. Wang, A. Bains, F. T. Sommer, J. Hirsch, University of Southern California	177

II-87. Decoding multiple objects from populations of macaque IT neurons with and without spatial attention E. Meyers, Y. Zhang, S. Chikkerur, N. Bichot, T. Serre, T. Poggio, R. Desimone, MIT	178
II-88. Contrast dependent changes in Monkey V1 gamma frequency undermine its reliability in binding/control S. Ray, J. H. R. Maunsell, HHMI & Harvard medical school	179
II-89. Contrast suppression in human visual cortex B. Gijs Joost, D. Heeger, Center for Neural Science, New York University	180
II-90. Quantifying the difficulty of object recognition tasks via scaling of accuracy vs. training set size S. Brumby, L. M. Bettencourt, C. Rasmussen, R. Bennett, M. Ham, G. Kenyon, Los Alamos National Laboratory	181
II-91. Does the visual system use natural experience to construct size invariant object representations? N. Li, J. DiCarlo, McGovern Inst/Dept of Brain & Cog Sci, MIT	182
II-92. Stimulus timing dependent plasticity in high- and low-level vision D. B. T. McMahon, D. A. Leopold, National Institute of Mental Health	182
II-93. In vivo Ca2+imaging of neural activity throughout the zebrafish brain during visual discrimination E. A. Naumann, A. Kampff, F. Engert, Harvard University	183
II-94. Frequency dependence of the spatial spread of the local field potential D. Xing, C.-I. Yeh, S. Burns, R. Shapley, New York University, Center for Neural Sci	184
II-95. Network dynamic regime controls the structure of the V1 extra-classical receptive field D. B. Rubin, K. D. Miller, Department of Neuroscience, Columbia University	184
II-96. Modulation of speed-sensitivity of a motion-sensitive neuron during walking M. E. Chiappe, J. D. Seelig, M. B. Reiser, V. Jayaraman, Janelia Farm Research Campus, HHMI	185
II-97. Cross-correlation analysis reveals circuits and mechanisms underlying direction selectivity. P. M. Baker, W. Bair, Department of Physiology Anatomy and Genetics, Univ. of Oxford	186
II-98. The influence of pulvinar activity on corticocortical communication C. M. Ziemba, G. R. Mangun, W. M. Usrey, Center for Neuroscience, UC Davis	187
II-99. Image classification with complex cell neural networks J. Bergstra, Y. Bengio, P. Lamblin, G. Desjardins, J. Louradour, University of Montreal	187
II-100. Visual cortex unplugged: Neural recordings from rats in the wild A. Agrochao, T. Szuts, V. Fadeyev, W. Dabrowski, A. Litke, M. Meister, Harvard University	188
II-101. Developing a rodent model of selective auditory attention C. Rodgers, S. Kochik, V. D. Vu, M. R. DeWeese, UC Berkeley	189

Poster Session III

7:30 pm Saturday 27th February

III-1. Synaptic filtering of natural spike trains in central synapses: A computational study U. Kandaswamy, C. Stevens, V. Klyachko, Washington University in St Louis	189
III-2. When somatic firing undermines dendritic compartmentalization B. F. Behabadi, B. W. Mel, Biomedical Engineering Department, University of Southern California	190
III-3. The effect of inhibition on pyramidal cells: A conceptual model M. Jadi, A. Polsky, J. Schiller, B. W. Mel, Biomedical Engineering Department, University of Southern California	191
III-4. Map dynamics of rhythmically perturbed neurons E. Jan, K. Loncich, R. Mirolo, M. Hasselmo, M. Yoshida, Physics Department	191
III-5. How do temporal stimulus correlations influence the performance of population codes? K. R. Rad, L. Paninski, Columbia University	192
III-6. The dynamic routing model of visuospatial attention B. Bobier, T. Stewart, C. Eliasmith, University of Waterloo	192
III-7. Mechanisms recruited by attending to the time or frequency of sounds S. Jaramillo, A. M. Zador, Cold Spring Harbor Laboratory	193
III-8. 252-site subdural LFP recordings in monkey reveal large-scale effects of selection attention. C. A. Bosman, T. Womelsdorf, R. Oostenveld, B. Rubehn, P. de Weerd, T. Stieglitz, P. Fries, Donders Centre	194
III-9. Dynamical control of eye movements in an active visual search task: theory and experiments H. He, J. Schilz, A. J. Yu, Department of Cognitive Science, University of California, San Diego	195
III-10. Decision-making dynamics and behavior of a parietal-prefrontal loop model D. Andrieux, X.-J. Wang, Yale University School of Medicine	195
III-11. Time-varying gain modulation on neural circuit dynamics and performance in perceptual decisions R. K. Niyogi, K. Wong-Lin, Gatsby Computational Neuroscience Unit, UCL	196
III-12. An optimality framework for understanding the psychology and neurobiology of inhibitory control P. Shenoy, R. P. N. P. N. Rao, A. J. Yu, University of Washington	197
III-13. Discounting as task termination, and its implications S. Paul, C. A. Rothkopf, University of Minnesota	198
III-14. Optimal decision-making in multisensory integration J. Drugowitsch, A. Pouget, G. C. DeAngelis, D. E. Angelaki, Department of Brain & Cognitive Sciences, University of Rochester	198
III-15. Role of anterior cingulate cortex in patch-leaving foraging decisions B. Hayden, M. Platt, Duke University	199
III-16. A neuronal model for context-dependent change in preference A. Soltani, B. De Martino, A. Rangel, C. Camerer, Baylor College of Medicine	200
III-17. What, when and how of target detection in visual search. V. Navalpakkam, P. Perona, Caltech	200
III-18. Beyond the edge: Amplification and temporal integration by recurrent networks in the chaotic regime. T. Toyozumi, L. F. Abbott, Department of Neuroscience, Columbia Univ.	201
III-19. Control of persistent spiking activity by background correlations. M. Dipoppa, B. Gutkin, Group for Neural Theory, LNC, DEC, ENS	202
III-20. Multiple routes to functionally feedforward dynamics in cortical network models. E. Wallace, M. Benayoun, W. van Drongelen, J. Cowan, Dept. of Mathematics, University of Chicago	203

III-21. Visualizing classification decisions of hierarchical models of cortex	
W. A. Landecker, S. P. Brumby, M. Thomure, G. T. Kenyon, L. M. A. Bettencourt, M. Mitchell, Portland State University	203
III-22. Hidden structures detection in nonstationary spike trains	
K. Takiyama, M. Okada, The University of Tokyo.	204
III-23. Microcircuits of stochastic neurons	
S. Cardanobile, S. Rotter, BCCN Freiburg	205
III-24. Bayesian methods for intracellular recordings: electrode artifact compensation and noise removal	
Y. Yoo, J. W. Pillow, University of Texas at Austin	205
III-25. A normative theory of short-term synaptic plasticity	
J.-P. Pfister, P. Dayan, M. Lengyel, University of Cambridge, Dpt. of Engineering	206
III-26. Dynamic Bayesian network model on two opposite types of sensory adaptation	
Y. Sato, K. Aihara, Institute of Industrial Science, The University of Tokyo	207
III-27. Methods for neural circuit inference from population calcium imaging data	
J. T. Vogelstein, T. A. Machado, Y. Mishchenko, A. M. Packer, R. Yuste, L. Paninski, Johns Hopkins University	208
III-28. Coincidence detection in active neurons	
C. Rossant, R. Brette, Ecole Normale Supérieure	208
III-29. A delta-rule approximation to Bayesian inference in change-point problems	
R. C. Wilson, M. Nassar, J. Gold, Princeton University	209
III-30. Cellular mechanisms that may contribute to prefrontal dysfunction in psychosis	
V. S. sohal, K. Deisseroth, Dept. of Psychiatry and Behavioral Sciences, Stanford University	210
III-31. In vivo multi-single-unit extracellular recordings from identified neural populations in fruit flies	
M. B. Ahrens, M. Barbic, B. Barbarits, B. G. Jamieson, V. Jayaraman, Cambridge University	211
III-32. Layered sparse associative network for soft pattern classification and contextual pattern completion	
E. Ehrenberg, P. Kanerva, F. Sommer, Redwood Center for Theoretical Neuroscience, University of California, Berkeley	211
III-33. Spike-timing theory of working memory	
B. Szatmáry, E. M. Izhikevich, The Neurosciences Institute, San Diego, CA	212
III-34. Cue-based feedback enables remapping in a multiple oscillator model of place cell activity	
J. D. Monaco, K. Zhang, H. T. Blair, J. J. Knierim, Krieger Mind/Brain Institute, Johns Hopkins	213
III-35. Independent snapshot memories in hippocampus: Representation of touch- and sound-guided behavior.	
P. . M. Pavel, E. Vinnik, C. Honey, J. Schnupp, M. E. Diamond, SISSA	214
III-36. Contextual information for maintaining coherent egocentric-alloentric maps	
c. Jimenez Rezende*, D. Molter*, W. Gerstner, EPFL - LCN	215
III-37. Attention selects informative populations	
P. Verghese, A. Wade, Smith Kettlewell Eye Research Institute	216
III-38. How well does local neuronal activity predict cortical hemodynamics?	
Y. B. Sirotnin, A. Das, Columbia University, Dept of Neuroscience	216
III-39. Decrease in synaptic variance improves perceptual ability	
R. C. Froemke, M. M. Merzenich, C. E. Schreiner, University of California, San Francisco	217
III-40. Spike-based Expectation Maximization	
B. Nessler, M. Pfeiffer, W. Maass, Graz University for Technology	218
III-41. Evaluation of memories through synaptic tagging	
M. Papper, R. Kempter, C. Leibold, University of Munich	219
III-42. Optimal architectures for fast-learning, flexible networks	
V. Itskov, A. Degeratu, C. Curto, University of Nebraska-Lincoln	219

III-43. Disconnection of monkey orbitofrontal and rhinal cortex impairs assessment of motivational value A. M. Clark, S. Bouret, E. A. Murray, B. J. Richmond, Laboratory of Neuropsychology, NIMH, NIH	220
III-44. Reward-modulated spike timing-dependent plasticity requires a reward-prediction system N. Frémaux, H. Sprekeler, W. Gerstner, LCN, EPFL	221
III-45. Tagging and capture: a bridge from molecular to behavior L. Ziegler, W. Gerstner, EPFL - LCN	221
III-46. Predicting the task specificity of learning J. M. Fulvio, P. Schrater, University of Minnesota	222
III-47. Eye position modulation of visual responses in the lateral intraparietal area lags the eye movement Y. Xu, C. Karachi, M. Goldberg, Columbia University	223
III-48. Oscillatory spiking activity in primate superior colliculus is related to spatial working memory L. Lee, R. J. Krauzlis, Salk Institute for Biological Studies	223
III-49. Ensemble activity underlying movement preparation in prearcuate cortex R. Kalmar, J. Reppas, S. Ryu, K. Shenoy, W. Newsome, Stanford University	224
III-50. Network mechanisms for the modulation of gamma spike phase by stimulus strength and attention P. Tiesinga, T. Sejnowski, Radboud University Nijmegen	225
III-51. A model of vPFC neurons performing a same-different task: an alternative model of working memory J. Lee, Y. Cohen, Department of Otorhinolaryngology, University of Pennsylvania School of Medicine	226
III-52. Designing optimal stimuli to control neuronal spike timing Y. Ahmadian, A. M. Packer, R. Yuste, L. Paninski, Columbia University	227
III-53. Hidden Markov models for the stimulus-response relationships of multi-state neural systems S. Escola, L. Paninski, Center for Theoretical Neuroscience, Columbia University	227
III-54. Dopamine-modulated dynamic cell assemblies generated by the GABAergic striatal microcircuit M. Humphries, R. Wood, K. Gurney, Ecole Normale Supérieure	228
III-55. Stationary envelope synthesis (SES): A universal method for phase coding by neural oscillators H. T. Blair, A. Weldon, I. G. Shlifer, M. Bloom, K. Zhang, Psychology, UCLA	229
III-56. Representation of environmental statistics by neural populations D. Ganguli, E. P. Simoncelli, Center for Neural Science, NYU	230
III-57. Origins of contrast gain control in isolated cortical neurons: deriving the code from the dynamics M. Famulare, R. Mease, A. Fairhall, University of Washington	231
III-58. Single neuron dynamics determine the strength of chaos in the balanced state M. Monteforte, S. Löwel, F. Wolf, MPI for Dynamics and Self-Organization, BCCN	231
III-59. A non-stationary copula-based spike count model A. Onken, S. Grünewälder, M. H. J. Munk, K. Obermayer, Technische Universität Berlin	232
III-60. Multiple spike time patterns occur at bifurcation points of membrane potential dynamics J. V. Tups, J.-M. Fellous, P. J. Thomas, T. J. Sejnowski, P. H. Tiesinga, Univ. North Carolina Chapel Hill	233
III-61. Information scaling, efficiency and anatomy in the cerebellar granule cell layer. G. Billings, A. Lorincz, P. Gleeson, Z. Nusser, A. Silver, University College London	234
III-62. Dynamic changes in single cell and population activity during the acquisition of task behavior J. Swearingen, M. Reyes, C. V. Buhusi, Medical University of South Carolina	234
III-63. A stimulus-dependent maximum entropy model of the retinal population neural code E. Granot-Atedgi, G. Tkacik, R. Segev, E. Schneidman, Department of Neurobiology, Weizmann Institute of Science	235
III-64. Effects of spike-driven feedback on neural gain and pairwise correlation J. Bartels, B. Doiron, University of Pittsburgh	236

III-65. A neuronal population measure of attention predicts behavioral performance on individual trials M. R. Cohen, J. H. R. Maunsell, Harvard Medical School	237
III-66. A nearly optimal correlation-independent readout of population activity C.-L. Teng, P. Latham, J. W. Pillow, University of Virginia	237
III-67. Sampling based inference with linear probabilistic population codes J. Beck, A. Pouget, P. Latham, University College London	238
III-68. Decoding multiscale word and category-specific spatiotemporal representations from intracranial EEG A. M. Chan, E. Halgren, C. Carlson, O. Devinsky, W. Doyle, R. Kuzniecky, T. Thesen, C. Wang, D. Schomer, E. Eskandar, Harvard-MIT Health Sciences & Technology	239
III-69. Exploring the statistical structure of large-scale neural recordings using a sparse coding model A. Khosrowshahi, J. Baker, R. Herikstad, S.-C. Yen, C. J. Rozell, B. A. Olshausen, Redwood Center for Theoretical Neuroscience, University of California, Berkeley	240
III-70. Modulation of STDP by the structure of pre-post synaptic spike times G. Pipa, M. Castellano, R. Vicente, B. Scheller, MPI for Brain Research	241
III-71. Spatio-temporal credit assignment in population learning J. Friedrich, R. Urbanczik, W. Senn, Department of Physiology, University of Bern	241
III-72. Single-neuron spike timing depends on global brain dynamics C. Ryan, K. Ganguly, S. Kennerley, K. Koepsell, C. Cadieu, J. Wallis, J. Carmena, University of California, Berkeley	242
III-73. Temporal precision of the olfactory system R. Shusterman, M. Smear, T. Bozza, D. Rinberg, Janelia Farm, HHMI	243
III-74. Complementary encoding of sound features by local field potentials and spikes in auditory cortex S. V. David, N. Mesgarani, S. Atiani, S. A. Shamma, University of Maryland, College Park	243
III-75. Modeling peripheral auditory processing in the precedence effect J. Xia, B. Shinn-Cunningham, Boston University	244
III-76. Auditory cortex neuronal tuning is sensitive to statistical structure of early acoustic environment H. Koeber, Y.-T. L. Tseng, K. Gill, S. Bao, UC Berkeley	245
III-77. Up-states are rare in awake auditory cortex T. Hromadka, M. DeWeese, A. M. Zador, Cold Spring Harbor Laboratory	245
III-78. Hearing the song in noise M. Richard, P. R. Gill, F. E. Theunissen, UC Berkeley	246
III-79. A normalization model of multi-sensory integration T. Ohshiro, D. E. Angelaki, G. C. DeAngelis, University of Rochester	247
III-80. Odor trail tracking by rats in surface and air borne conditions A. Khan, U. Raheja, U. Bhalla, National Centre for Biological Sciences	247
III-81. The flow of expected and unexpected sensory information through the distributed forebrain network M. Cui, J. Fiser, D. Katz, A. Fontanini, Department of Psychology, Brandeis University	248
III-82. Sparse coding of natural stimuli in the midbrain M. J. Chacron, Department of Physiology	249
III-83. 2D encoding of concentration and concentration gradient in <i>Drosophila</i> ORNs A. J. Kim, A. A. Lazar, Y. Slutskiy, Columbia University	249
III-84. Interactions of rat whiskers with air currents: implications for flow sensing V. Gopal, M. Kim, C. Chiapetta, J. Russ, M. Meaden, M. Hartmann, Department of Physics, Elmhurst College	250
III-85. Re-testing the energy model: identifying features and nonlinearities of complex cells. T. Lochmann, J. N. Stember, T. Blanche, D. A. Butts, Department of Biology, University of Maryland	251

III-86. Human versus machine: comparing visual object recognition systems on a level playing field. N. Pinto, N. J. Majaj, Y. Barhomi, E. A. Solomon, D. D. Cox, J. J. DiCarlo, MIT	252
III-87. Model of visual target detection applied to rat behavior P. Meier, P. Reinagel, University of California, San Diego	253
III-88. Do quantal dynamics of graded synaptic signal transfer adapt to maximise the rate of information? X. Li, S. Tang, M. Juusola, Beijing Normal University	253
III-89. Bursts and visual encoding in LGN during natural state fluctuations in the unanesthetized rat E. D. Flister, P. Reinagel, UCSD	254
III-90. Velocity coding and octopamine in an identified optic flow-processing interneuron of the blowfly K. D. Longden, H. G. Krapp, Dept of Bioengineering, Imperial College London, UK	255
III-91. Exact statistical analysis of visual inference amid eye movements E. A. Mukamel, Y. Burak, M. Meister, H. Sompolinsky, Harvard University	256
III-92. Encoding stereomotion with neural populations using IOVD and CD mechanisms Q. Peng, B. E. Shi, Dept. of ECE, HKUST	256
III-93. Sparseness is not actively optimized in V1 P. Berkes, B. L. White, J. Fiser, Brandeis University	257
III-94. Motion and reverse-phi stimuli that do not drive standard Fourier or non-Fourier motion mechanisms Q. Hu, J. D. Victor, Weill Cornell Medical College	258
III-95. Temporal integration of motion and cortical normalization in macaque V1 D. McLelland, P. M. Baker, B. Ahmed, W. Bair, Department of Physiology Anatomy and Genetics, Univ. of Oxford	259
III-96. Relationship of contextual modulations in V1 and V2 revealed by nonlinear receptive field mapping A. M. Schmid, J. D. Victor, Weill Cornell Medical College	259
III-97. Black' dominance measured with different stimulus ensembles in macaque primary visual cortex V1 C.-I. Yeh, D. Xing, R. Shapley, New York University, Center for Neural Sci	260
III-98. The adapting receptive field surround of a large ON ganglion cell K. Farrow, B. Roska, Friedrich Miescher Institute	261
III-99. Functional dissection of optomotor pathways in the Drosophila optic lobe. S. E. J. de Vries, T. R. Clandinin, Stanford University	261
III-100. Retention of perceptual categorization following bilateral removal of area TE in rhesus monkeys N. Matsumoto, R. Saunders, K. Gothard, B. Richmond, AIST	262
III-101. A model of efficient change detection through the interaction of excitation and inhibition N. Bouaouli, S. Deneve, Ecole Normale Supérieure, Paris, France	263

Abstracts

Abstracts for talks appear first, in order of presentation; those for posters next, in order of poster session and board number. An index of all authors appears at the back.

DOI links will not be active until the conference.

T-1. Towards complete structural and functional imaging of cortical circuits

R. Clay Reid

CLAY_REID@HMS.HARVARD.EDU

Harvard Medical School

For the past six years, our lab has been using two-photon calcium imaging to measure the in vivo physiological activity of virtually every neuron (up to several thousand) in volumes of visual cortex spanning up to several hundred μm (Ohki et al., 2005, 2006). More recently, we have scaled up the collection of serial-section electron microscopic images so we can measure the fine anatomical structure of similarly large volumes of cortical tissue. By combining these two methods, we hope to collect data sets that provide a complete physiological and structural overview of a specific piece of the cortex. Because we are interested in large cortical volumes at high resolution, we have concentrated on very high throughput. To reconstruct the finest axons, dendrites, and synapses with electron microscopy, pixels must be on the order of $<5\text{ nm}$ and section thickness should be at most 40 nm . At this resolution, reconstruction of a $100\text{ }\mu\text{m}$ cube requires at least 1012 bytes of data (one terabyte). A $500\text{ }\mu\text{m}$ cube would require 125 terabytes. Our first large data set consists of $\sim 1,200$ neurons in a volume that spans roughly $450 \times 350 \times 50\text{ }\mu\text{m}$ of cortical tissue, which is large enough to trace many of the connections between neurons within the volume. We have also studied the physiological properties of a small fraction of the cells in this volume and solved the correspondence problem between the physiological and anatomical data sets. We can point to cell bodies, axons, and dendrites within the volume and correlate their morphology and connections with the visual physiology as determined by calcium imaging. As with many large-scale datasets, such as the genome, these anatomical data will never be fully analyzed. Instead, they will serve as a repository—an infinite slide box—of anatomical information that can answer specific well-posed questions. For any given question, only a tiny fraction of the data need be analyzed, but the completeness of the data provides a number of unique opportunities. Our lab is addressing one class of questions: are there sub-networks within the local circuit that process distinct information? If a circuit is made up of red, green, yellow and blue cells (representing a functional property such as orientation selectivity), are the red cells connected to the red cells? More generally, are intracortical connections specific, where specificity is defined as "making sense" in a functional context? This question can be posed for the multiple types of connections within a cortical circuit: between excitatory neurons within layers and between layers; and those involving the multiple types of inhibitory neurons.

[doi:](#)

T-2. Non linear dendritic processing in cortical pyramidal neurons

Jackie Schiller

JACKIE@TX.TECHNION.AC.IL

Technion Medical School

Neurons in the central nervous system typically possess an elaborated dendritic tree, which serves to receive and integrate the vast input information arriving to the neuron. Understanding the way information is processed in dendrites is crucial for comprehending the input/output transformation functions of individual CNS neurons, and in turn learning how cortical networks code and store information. Cortical pyramidal neurons, which are the major excitatory neurons in the cortical tissue, have a typical dendritic tree consisting of a large apical trunk which branches to form the oblique and tuft branches, and a basal tree branching directly from the soma. In the present work we concentrated on understanding how tuft dendrites process their incoming information. Tuft dendrites are the main target for feedback inputs innervating neocortical layer-5 pyramidal neurons but their properties remain obscure. We report the existence of NMDA-spikes in the fine distal tuft dendrites that otherwise did not support the initiation of calcium spikes. Both direct measurements and computer simulations showed that NMDA-spikes are the dominant mechanism by which distal synaptic input leads to firing of the neuron and provide the substrate for complex parallel processing of top-down input arriving at the tuft. These data lead to a new unifying view of integration in pyramidal neurons in which all fine dendrites, basal and tuft, integrate inputs locally through the recruitment of NMDA receptor channels relative to the fixed apical calcium and axo-somatic sodium integration points.

[doi:](#)

T-3. Input-dependent switching of inhibitory configurations in neural networks

Alex D. Reyes

REYES@CNS.NYU.EDU

New York University

The responses of neurons during a stimulus depend on the summed excitatory (E) and inhibitory (I) synaptic inputs. Shifts in the E-I balance cause both qualitative and quantitative changes in the neuronal firing patterns. How the E and I inputs scale with stimuli is yet unclear. An important determinant is the network architecture, broadly classified either as the lateral inhibitory network (LIN) where the I inputs to a neuron are more broadly tuned than the E inputs, or the co-tuned network (CON) where the E and I inputs co-vary for the entire stimulus range. E and I scale with input differently in each configuration to produce complementary sets of responses, suggesting that both may be needed to account for the diverse stimulus-evoked firing behavior. Here I show that a single E-I network transitions seamlessly between LIN and CON when the amplitude and spatial extent of the input to the network changes. Simulations were performed with a 2D network (10000 E cells; 2000 I cells) of Adaptive exponential integrate and fire neurons adjusted to reproduce the firing of pyramidal cells, and fast spiking and low threshold spiking interneurons. The patterns of connections between E and I cells were based on experimental data obtained from paired recordings in an in vitro slice preparation. Synaptic barrages (200 ms duration) were delivered to each cell; the number of barrages was adjusted so that the mean input was Gaussian distributed (parameterized by amplitude A and standard deviation S) in space. When the input was narrow (small S), the network was configured as LIN. E cells tended to fire tonically, and exhibited side-band inhibition. As the input broadened, the network switched to CON. The neurons tended to fire phasically and exhibited no sideband inhibition. To characterize quantitatively the variation of E-I balance and firing in the A-S space, mean field techniques were used to calculate the population activity of the neurons in the network. This yielded relations that describe how the widths and amplitude of excitatory and inhibitory inputs and the associated firing patterns change with the magnitude and spatial distribution of the input. Physiologically, the Gaussian input may represent tuning to stimuli such as tone frequency in the auditory system or disc location in the visual system; A may vary with stimulus intensity while S with bandwidth or disc diameter. The changes in firing that occur with increasing the bandwidth/diameter of auditory/visual stimulus are reproduced by tracing appropriate trajectories in the A and S space. That the LIN and CON configurations are not hardwired into the network has important implications.

The ability to switch between configurations provides potentially a mechanism for modulating the response of the network to a variety of inputs and behavioral states. Additionally, many of the heterogeneous firing and receptive field properties that had been postulated to arise from different network configurations may in fact be due to transitions within a single network that are triggered perhaps by changes in the stimulus characteristics or state of the animal.

[doi:](#)

T-4. Neuronal biophysics modulate the ability of gamma oscillations to control response timing.

Andrea Hasenstaub¹

Stephani Otte^{2,3}

Edward M. Callaway¹

ANDREA.HASENSTAUB@GMAIL.COM

OTTES01@GMAIL.COM

CALLAWAY@SALK.EDU

¹Crick-Jacobs Center, Salk Institute for Biological Studies

²Salk Institute

³UCSD Neuroscience Graduate Program

Gamma frequency oscillatory activity of inhibitory sub-networks has been hypothesized to regulate information processing in the cortex as a whole. Inhibitory neurons in these sub-networks synchronize their firing and selectively innervate the perisomatic compartments of their target neurons, generating both tonic and rapidly fluctuating inhibition that is hypothesized to enforce temporal precision and coordinate the activity of their post-synaptic targets. Indeed, in vivo and in vitro recordings have demonstrated that many neurons' firing is entrained to these oscillations, although to varying extents and at various phases. Cortical networks are composed of diverse populations of cells that differ in their chemical content, biophysical characteristics, laminar location, and connectivity. Thus, different types of neurons may vary in the amplitude and timing of the synchronized inhibition they receive, as well as in the effects of pattern of inhibitory inputs on response timing and precision. What accounts for this heterogeneity of response timing between cell types, and are these response properties fixed or flexible? To answer these questions, we use a combination of in vitro electrophysiology, dynamic clamp, and modeling to characterize the interactions between a neuron's intrinsic properties, the degree of gamma-band synchrony among its inhibitory inputs, and its spike timing. We apply these techniques to study six distinct types of cortical neurons. We find that neuron types systematically vary in the phase and precision of their spike timing relative to the peak of gamma frequency input, and the degree to which their spike time depended on changes in inhibitory synchrony. Biophysical characterizations of real neurons suggest that the membrane time constant (T_m), afterhyperpolarization amplitude and duration, and sodium channel properties are the key features governing gamma control of response timing. We confirmed these findings both in a single compartment model, and by using dynamic clamp to alter intrinsic features of neurons' physiology. Shortening neurons' time constant by adding artificial leak conductance enhanced neurons' ability to entrain their firing and phase-precise, while adding an artificial afterhyperpolarization conductance reduced neurons' ability to phase-precise. We conclude that a neuron's intrinsic physiology substantially affects the ability of gamma-synchronized inhibitory inputs to control its response timing, and that different excitatory and inhibitory neuron types systematically differ in this regard. These results suggest that the characteristic phase relationship of the discharges of neurons during gamma activity may be explained by differences in intrinsic properties rather than differences in connectivity. Further, we note that the relevant physiology is not static, but may be altered by contextual or neuromodulatory factors; for instance, a neuron's time constant decreases during intense synaptic activity, while its afterhyperpolarization is altered by neuromodulators such as noradrenaline and acetylcholine. We therefore suggest that the ability of gamma oscillations to control response timing is not fixed, but may be dynamically shaped to suit the cortex's computational requirements during attention or cognition.

[doi:](#)

T-5. Desynchronization of an electrically coupled interneuron network with excitatory synaptic input

Koen Vervaeke¹
 Andrea Lorincz²
 Pdraig Gleeson¹
 Matteo Farinella¹
 Zoltan Nusser²
 Angus Silver¹

K.VERVAEKE@UCL.AC.UK
 LORINCZ@KOKI.HU
 P.GLEESON@UCL.AC.UK
 M.FARINELLA@UCL.AC.UK
 NUSSE@KOKI.HU
 A.SILVER@UCL.AC.UK

¹University College London

²Institute of Experimental Medicine

It is well established that the electrical synapses made between interneurons contribute to their synchronized firing behaviour and to the generation of network oscillations in a number of brain regions. However, relatively little is known about how coupled interneuron networks respond to excitatory synaptic inputs that occur during a sensory stimulus. To address this we studied a network of electrically coupled interneurons in the input layer of the cerebellar cortex. This network is particularly suited to this type of investigation because inhibitory Golgi cells only form chemical synapses with granule cells and not with one another, allowing electrical signalling to be studied in isolation. We show with immunohistochemistry, electron microscopy and electrophysiology that electrical coupling between Golgi cells is mediated by Connexin 36 and that gap junctions are made predominantly between the ascending dendrites of Golgi cells. Using paired whole-cell recordings and synaptic stimulation in acute slices, we show that a sparse, temporally precise, excitatory mossy fibre input triggers a spike followed by a pause, in directly innervated Golgi cells, while a pause-only response was observed in the spontaneous firing of cells that do not receive direct excitatory synaptic inputs. Moreover, mossy fibre input also caused transient desynchronization of firing in the cell pair. Our results show that inhibition of neighbouring cells and desynchronization is caused by propagation of the action potential after-hyperpolarization through gap junctions. Biologically detailed computer simulations of Golgi cells using reconstructed cell morphologies, active conductances, measured gap junction locations and strengths, together with synaptic input distributed on the dendritic tree closely reproduced the results obtained with paired recordings. To explore the spatial properties of network dynamics we extended our two-cell model to a larger 3D network that incorporated the spatial dependence of connection probability and coupling coefficient measured from 136 paired recordings, using the software application NeuroConstruct. The 3D network model, which consisted of 45 Golgi cells (each with 4672 segments) also incorporated nonuniform input resistances, producing heterogeneous firing rates. Analysis of the network connectivity suggests that each Golgi cell is electrically connected to ~10 other Golgi cells. Simulations run on parallel NEURON predict that rhythmically active Golgi cells synchronized in the absence of correlated synaptic inputs. However, sparse synaptic excitation of the Golgi cell network caused a transient desynchronization of spiking, with cells exhibiting a mosaic of excitation-pause and pause only response. Our results suggest that several features of the sensory-evoked behaviour of Golgi cells and the rapid disappearance of granule cell oscillations observed in vivo could arise from sparse excitation of the electrically coupled Golgi cell network. Moreover, if interneuron networks with inhibitory electrical synapses are present in different regions, sparse long-range excitatory connections could potentially be used to trigger near-simultaneous desynchronization of multiple networks across the brain. Funded by BBSRC, MRC, EUSynapse and the Wellcome Trust.

[doi:](#)

T-6. Beyond optimality to understanding individual brains: variability, homeostasis and compensation in neuronal circuits

Eve Marder
 Brandeis University

MARDER@BRANDEIS.EDU

I will summarize recent theoretical and experimental work that shows that similar circuit outputs can be produced with highly variable circuit parameters. This work argues that the nervous system of each healthy individual has found a set of different solutions that give "good enough circuit performance. I will use examples from theoretical and experimental studies using the crustacean stomatogastric nervous system to argue that synaptic and intrinsic currents can vary far more than the output of the circuit in which they are found. These data have significant implications for the mechanisms that maintain stable function over the animal's lifetime, and for the kinds of changes that allow the nervous system to recover function after injury. In this kind of complex system, merely collecting mean data from many individuals can lead to significant errors, and it becomes important to measure as many individual network parameters in each individual as possible.

[doi:](#)

T-7. Threshold modulation and homeostatic control of spike timing via circuit plasticity

Brent Doiron¹
YanJun Zhao²
Thanos Tzounopoulos²

BDOIRON@PITT.EDU
YANJUNZ@PITT.EDU
THANOS@PITT.EDU

¹Dept. of Mathematics, Univ. of Pittsburgh

²Dept. of Otolaryngology, Univ. of Pittsburgh

Activity-dependent plasticity enables neuronal circuits to encode the large dynamic range of sensory stimuli. However, the known mechanisms underlying plasticity of firing rates also affect spike timing neural codes. Therefore these rate adaptation mechanisms may not be applicable to many sensory systems encoding temporally precise inputs. The auditory system is an ideal site for studying the cellular mechanisms underlying coordinated spike rate modulation and preservation of spike timing. In auditory circuits, activity-dependent spike rate adaptation is common, yet faithful encoding of the fine temporal structure of auditory inputs via an invariant temporal code is necessary for accurate stimulus identification and discrimination. We present in vitro characterization of disynaptic parallel fiber (PF) synapses in dorsal cochlear nucleus, an auditory brainstem structure. The direct PF excitation onto principal DCN neurons has a Hebbian spike timing dependent plasticity (STDP) rule, whereas excitation onto DCN interneurons has an anti-Hebbian STDP rule [1,2]. The coordinated recruitment of long-term potentiation (LTP) and long-term depression (LTD) in the PF feedforward, inhibitory circuit increases the duration of the principal cell integration window to PF-PF and PF-auditory nerve inputs. We use a simple computational model of DCN integration and spike dynamics to show that plasticity induced increase in fusiform cell integration window modulates the spike response threshold to auditory nerve inputs, while preserving a spike latency code. This latency code invariance to threshold changes permits any latency decoder to be well tuned for a broad distribution of stimulus intensities, offering a simple solution to parcel rate and temporal decoding schemes. Previous studies have established that neural circuits use homeostatic plasticity to maintain stable firing rate [3]. Nevertheless, neurons also perform time sensitive tasks and we propose that homeostatic control of temporal coding, similar to that outlined in our study, is a general feature of neural circuit design. [1] Tzounopoulos, T., Kim, Y., Oertel, D., and Trussell, L.O. (2004). Cellspecific, spike timing-dependent plasticities in the dorsal cochlear nucleus. *Nat. Neurosci.* 7, 719-725. [2] Tzounopoulos, T., Rubio, M.E., Keen, J.E. & Trussell, L.O. Coactivation of pre- and postsynaptic signaling mechanisms determines cell-specific spiketiming- dependent plasticity. *Neuron* 54, 291-301 (2007). [3] Turrigiano, G.G. The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell* 135, 422-435 (2008).

[doi:](#)

T-8. Robust spatial working memory through inhibitory gamma synchrony

Devarajan Sridharan^{1,2}

Sebastian Millner³

John Arthur²

Kwabena Boahen²

DSRIDHAR@STANFORD.EDU

SMILLNER@KIP.UNI-HEIDELBERG.DE

JARTHUR@STANFORD.EDU

BOAHEN@STANFORD.EDU

¹Dept. of Neurobiology, Stanford University

²Dept. of Bioengineering, Stanford University

³Kirchhoff Institute for Physics, University of Heidelberg

Persistent firing and gamma (30-80Hz) oscillations are known to be critically involved in working memory [1,2], yet previous models have ignored or avoided oscillations as being potentially detrimental to working memory [3]. Here, we develop a model in which these gamma oscillations have a direct role in the robust maintenance of persistent information. Current models of working memory [4,5] use a recurrent excitatory-inhibitory framework where the excitatory recurrence leads to persistence, and the inhibition prevents diffusion of the pattern after stimulus offset. However, any heterogeneities in the neurons (firing rates, thresholds etc) tend to cause drift of the stored pattern towards attractor locations (most excitable neurons). Previous attempts to create functional homogeneity in the neural firing rates have used activity-dependent synaptic scaling rules [5], whose timescale (hours to days, [6]) is not reasonable for the rapid, reversible homogeneity required for stable maintenance of information over short durations in working memory. Here we propose a simple mechanism operating at the timescale of milliseconds to seconds that can reversibly create functional homogeneity in the firing landscape: inhibitory gamma synchrony. Upon imposing an external (background) inhibitory rhythm in the gamma range, model pyramidal neurons tend to synchronize to the rhythm, and equalize their firing rates. Simply, this can be explained as follows: Each pyramidal neuron has a small window of opportunity for firing before being shunted (reset) to resting potential by the periodic inhibitory background. This prevents the cumulative expression of firing heterogeneity among neurons over time, as the membrane potentials of all neurons are "reset" at the beginning of every cycle of inhibition. The resulting firing homogeneity may permit persistent, localized firing with little drift, and could facilitate active maintenance of information in working memory. We tested this theory using a neuromorphic chip with 1024 excitatory silicon neurons arranged in a 32x32 grid (Supplementary Figure 2) with 256 interleaved inhibitory neurons. Background inhibitory gamma synchronization was achieved by global recurrence among inhibitory neurons [7]. Local recurrence among excitatory neurons permitted persistent activity. The strong background inhibitory gamma rhythm indeed facilitated homogenization of excitatory neuron firing rates (coefficient of variation was 0.59 vs. 1.06 with strong and weak synchrony, respectively, Supplementary Figure 1). The network, in the strongly synchronous case, also demonstrated stable maintenance of stored patterns with less drift of the centroid (1.64+/-0.16 grid points), compared to the weakly synchronous case (3.28+/-0.43 grid points), despite comparable levels of overall inhibition (Supplementary Figure 2). Thus, our model demonstrates a mechanism by which inhibitory gamma synchrony could facilitate robust maintenance of information in spatial working memory by rapid, reversible homogenization of firing rates. References 1. Pesaran et al, Nat. Neurosci. 5, 805 (2002). 2. Fuchs et al, Neuron, 53, 591 (2007). 3. Gutkin et al, J. Comput. Neurosci., 11, 121 (2004). 4. Compte et al, Cereb. Cortex 10, 910 (2000). 5. Renart et al, Neuron 38, 473 (2003). 6. Perez-Otano and Ehlers, Trends Neurosci. 28, 229 (2005). 7. Arthur and Boahen, IEEE Trans. Neural Netw. 18, 1815 (2007).

doi:

T-9. Influence of task-specific instructions on cross-modal sensory interactions

Rama Natarajan¹

Richard Zemel¹

Iain Murray^{1,2}

David Hairston³

RAMA@CS.TORONTO.EDU

ZEMEL@CS.TORONTO.EDU

MURRAY@CS.TORONTO.EDU

DAVEHAIRSTON@ARL.ARMY.MIL

¹University of Toronto

²University of Edinburgh

³Army Research Lab

We analyze and model subjects' behavior when localizing an auditory cue in the presence of a visual distractor. In this paradigm, the ventriloquist illusion has been reliably demonstrated: subjects are biased towards the visual stimulus. Strong biases are typically accompanied by a perception of unity wherein the stimuli appear to originate from a common source. We probe whether and how different task instructions can affect perceptual biases and uncertainty. Subjects are modeled as observers performing inference on a generative model of the stimuli. The visual and auditory cues result under the model either from a cause at a single location or from two independent locations. The Bayes-optimal inference strategy combines predictions from the two discrete causal hypotheses weighted by the belief in each of them. An alternative strategy first selects the most probable hypothesis and bases subsequent inferences on that hypothesis alone. These two strategies make qualitatively different predictions (Natarajan et al., 2009). Recent studies that have utilized the above model to explain perceptual biases (Koerding and Tenenbaum, 2006; Koerding et al., 2007; Stocker and Simoncelli, 2008) have reached different conclusions with regards to the inference strategy that subjects appear to use. We analyze data from 3 new behavioral experiments that have the same ventriloquist paradigm but differ only in task instructions: participants are required to perform auditory localization alone, localization followed by a unity judgment or to report unity perception alone. We perform a comparative analysis of the data from the three experiments, and fit our model of subject behavior to the data. The dominant account of the ventriloquism effect is that target localization is an early sensory process that is largely immune to top-down influences (Vroomen and deGelder, 2004). Our results challenge this notion — we show that subjects' localization bias depends strongly on whether or not they are instructed to make unity judgments. Our computational model suggests that task-specific instructions change both subjects' expectations and their inference strategy. Subjects are described well by Bayesian model averaging when instructed only to localize; when reporting unity percepts alongside localization, subjects use only the model that best fits their unity judgments. The only fitted model parameter value that differed significantly across experiments was the prior belief that stimuli will be unified. When just localizing a target, subjects strongly believed stimuli would be unified. This belief diminished when also asked to report on unity, and subjects instructed to make unity judgments without localization had very low unity probability. Our results suggest the existence of top-down influences on sensory processing; the task-specific unity priors suggest that subjects have strong task-related expectations regarding stimuli. Our model's task-specific inference strategies can be understood as the result of limited computational resources in the brain. Instructing the subjects to make unity judgements forces them to select a model. When computation is limited it seems reasonable to use that model selection for future inferences, maintaining one consistent explanation of the observations. We hope that these results contribute to our understanding of stimulus representations and models of hierarchical sensory processing.

[doi:](#)

T-10. Adaptation and inference

Adrienne L. Fairhall

University of Washington

FAIRHALL@U.WASHINGTON.EDU

Adaptation occurs in many forms in single neurons and neural systems. While some forms of adaptation serve to optimize the flow of information through the system, the role of adaptation of the firing rate to changes in

the stimulus statistics is less clear. In some cases, this adaptation appears to encode the stimulus envelope, with a transfer function that approximates fractional differentiation. From another perspective, adaptation can be interpreted as on-line estimation of the time-varying statistics. We show that a simple model for this on-line estimation can predict experimentally observed timescales of adaptation.

[doi:](#)

T-11. The same neurons form a visual place code and an auditory rate code in the primate SC

Jungah Lee
Jennifer M Groh
Duke University

VISION.JUNGAH.LEE@GMAIL.COM
JMGROH@DUKE.EDU

A computational hurdle in multisensory integration is that visual and auditory signals potentially have different representational formats. In the early visual pathway, neurons have receptive fields that tile the visual scene and produce a "place code" for stimulus location. In contrast, the binaural computation performed in the auditory pathway has been suggested to produce a "rate code" for sound location. In this latter format, neurons in the IC and auditory cortex respond to a broad range of locations with activity levels that scale proportionate with sound position, exhibiting maximum responses to sounds at extreme contralateral positions (Groh JM et al., 2003; Werner-Reiss U and Groh JM, 2008). The superior colliculus (SC) has been thought to employ a place code for sensory and saccade-related activity, with the same neurons controlling auditory as well as visual saccades. However, there is little quantitative information addressing the coding of sound location in monkey SC. Here, we examined individual auditory neuron in the primate SC to determine whether they have receptive field (place code) or monotonic spatial pattern (rate code). A monotonic code for sound location would mean a discrepancy between visual and auditory processing in the SC. We recorded sensory and saccade-related activity from 180 neurons in the intermediate and deep SC of two rhesus monkeys. Noise bursts or LED came from one of 9 locations between ± 24 degrees in the horizontal dimension. The monkeys made saccades from different initial eye positions to these visual or auditory targets in an overlap saccade task. To quantify the representational format, we fit each neuron's response functions to gaussian and sigmoid curves. The idea was that gaussian functions would be substantially better than sigmoids at fitting the tuned response patterns characteristic of a place code, but that both sigmoid and broad half-gaussians would be equally successful at fitting the monotonic tuning patterns characteristic of a rate code. We found that most neurons had monotonic response patterns for auditory stimuli along the axis of the contralateral ear, even though the same neurons had non-monotonic response patterns to visual stimuli. For the auditory trials, the sigmoid functions were as good as gaussians at capturing the response patterns. For visual trials, the gaussian functions showed significantly better performance than sigmoids in fitting the tuning curves. This pattern was true for both sensory and saccade-related activity. Our findings imply that a read-out algorithm is required to reconcile the discrepancy. The algorithm should be able to convert the visual and auditory signals into a motor command to work on either place or rate code to produce the same accurate saccade for both types of signals. In line with this, we have developed several models that involve transformation of signals from a place code to a rate or vice-versa (Groh and Sparks, 1992; Groh, 2001; Porter and Groh, 2006). Acknowledgements This work was supported by CRCNS grants [R01 NS0942]. JAL was also supported by the Korea Research Foundation Grant funded by the Korean Government [KRF-2008-356-H00003].

[doi:](#)

T-12. Visual influences on information representations in auditory cortex

Christoph Kayser¹
Nikos Logothetis¹
Stefano Panzeri²

KAYSER@TUEBINGEN.MPG.DE
NIKOS.LOGOTHETIS@TUEBINGEN.MPG.DE
STEFANO.PANZERI@IIT.IT

¹Max-Planck-Institute for Biological Cybernetics

²Italian Institute for Technology

Combining information across different sensory modalities can greatly facilitate our ability to detect or recognize sensory stimuli. Recent work demonstrates that sensory integration is a distributed process, commencing in lower sensory areas and continuing in higher association cortices. Here we investigate the impact that visual stimuli have on the representation of sounds in auditory cortex. To this end we analyze neural responses to naturalistic audio-visual stimuli recorded while monkeys perform a visual fixation task. To characterize multisensory influences we quantify stimulus information provided by different putative neural codes. In particular, we have previously shown (Kayser et al. Neuron 09) that neural activity in auditory cortex is stimulus related on multiple temporal scales, including slow modulation of firing rates, millisecond precise temporal spike patterns and the relative timing of spiking activity to slow (<10Hz) ongoing network activity (phase-of-firing). In the context of multisensory stimuli, we find (Kayser et al. Curr Biol. In Press) that visual stimulation renders spiking responses more reliable across trials (repeats of the same stimulus), and more reliable in time (temporal precision of spikes). This increased reliability enhances the stimulus information provided by neural activity on slow (firing rate) and fast (spike patterns) time scales. As shown by feature extraction, this information gain pertains mostly to temporal sound properties, such as sound envelope, and is much reduced when incongruent visual and auditory stimuli are presented. Overall, these results demonstrate that multisensory influences enhance sensory representations already at early stages in cortex, and do so by enhancing the reliability of stimulus representations on multiple time scales.

[doi:](#)

T-13. Neuroethology of social attention

Michael L. Platt
Duke University Medical Center

PLATT@NEURO.DUKE.EDU

Humans and other animals pay attention to other members of their groups to acquire valuable social information about them, including information about their identity, dominance, fertility, emotions, and likely intent. In primates, attention to other group members and the objects of their attention is mediated by neural circuits that transduce sensory information about others, translate that information into value signals, and motivationally scale motor control signals to bias orienting behavior. This process unfolds via a subcortical route mediating fast, reflexive orienting to animate objects and faces and a more derived route involving cortical orienting circuits mediating nuanced and context-dependent social attention. Ongoing studies probe individual and species differences in the neural mechanisms that mediate social attention, the genetic origins of these differences, and their implications for differences in social behavior and social structure using naturalistic, ecologically valid social contexts.

[doi:](#)

T-14. Implications of correlated neuronal noise in decision making circuits for physiology and behavior

Ralf M. Haefner
Sebastian Gerwinn
Jakob H. Macke
Matthias Bethge

RALF.HAEFNER@TUEBINGEN.MPG.DE

Max-Planck-Institute for Biological Cybernetics

Understanding how the activity of sensory neurons contribute to perceptual decision making is one of the major questions in neuroscience. In the current standard model, the output of opposing pools of noisy, correlated sensory neurons is integrated by downstream neurons whose activity elicits a decision-dependent behavior [1][2]. The predictions of the standard model for empirical measurements like choice probability (CP), psychophysical kernel (PK) and reaction time distribution crucially depend on the spatial and temporal correlations within the pools of sensory neurons. This dependency has so far only been investigated numerically and for time-invariant correlations and variances. However, it has recently been shown that the noise variance undergoes significant changes over the course of the stimulus presentation [3]. The same is true for inter-neuronal correlations that have been shown to change with task and attentional state [4][5]. In the first part of our work we compute analytically the time course of CPs and PKs in the presence of arbitrary noise correlations and variances for the case of non-leaky integration and Gaussian noise. This allows general insights and is especially needed in the light of the experimental transition from single-cell to multi-cell recordings. Then we simulate the implications of realistic noise in several variants of the standard model (leaky and non-leaky integration, integration over the entire stimulus presentation or until a bound, with and without urgency signal) and compare them to physiological data. We find that in the case of non-leaky integration over the entire stimulus duration, the PK only depends on the overall level of noise variance, not its time course. That means that the PK remains constant regardless of the temporal changes in the noise. This finding supports an earlier conclusion that an observed decreasing PK suggests that the brain is not integrating over the entire stimulus duration but only until it has accumulated sufficient evidence, even in the case of no urgency [6]. The time course of the CP, on the other hand, strongly depends on the time course of the noise variances and on the temporal and interneuronal correlations. If noise variance or interneuronal correlation increases, CPs increase as well. This dissociation of PK and CP allows an alternative solution to the puzzle recently posed by [7] in a bottom-up framework by combining integration to a bound with an increase in noise variance/correlation. In addition, we derive how the distribution of reaction times depends on noise variance and correlation, further constraining the model using empirical observations. [1] Shadlen, MN, Britten, KH, Newsome, WT, Movshon, JA: J Neurosci 1996, 16:1486-1510 [2] Gold, JI, Shadlen, MN: Ann Rev Neurosci 2007, 30:535-574 [3] Churchland et al. CoSyNe Abstract 2009 [4] Cohen, MR, Newsome, WT: J Neurosci 2009, 29:6635-6648 [5] Cohen, MR, Newsome, WT: Neuron 2008, 60(1):162-173 [6] Kiani, R, Hanks, TD, Shadlen, MN: 28:3017-3029 [7] Nienborg, H, Cumming, BG: Nature 2009, 459:89-92

[doi:](#)

T-15. Efficiency, redundancy and sparse coding: Just a portion of Barlow's legacy

David J. Field
Cornell University

DJF3@CORNELL.EDU

The last 50 years has seen a variety of theories proposed to account for the behavior of neurons in the visual pathway. Many of these theories have focused on techniques found to be mathematically elegant (Fourier analysis, textons, wavelet coding, etc) but not on the nature of the signal to be coded. Although Horace Barlow has authored a wide number of major papers, his focus on efficiency and the redundancy of the natural signal is possibly the most influential. In this talk, I will look back at some of the proposals that Horace Barlow has made and

look at where we are today. In particular, I will describe the influence of his work on my own research related to sparse coding and the statistics of natural scenes.

[doi:](#)

T-16. Predictions of visual performance from the statistical properties of natural scenes

Wilson S. Geisler

GEISLER@PSY.UTEXAS.EDU

University of Texas, Austin

Five decades ago Horace Barlow argued that sensory scientists should explore the relationship between the design of an organism's sensory circuits, the organism's natural tasks, and the stimulus properties relevant to those tasks. In the past two decades advances in physical measurement technology, computational power and statistical modeling have made it possible to begin exploring this relationship in detail. In this talk I will briefly summarize our recent efforts to determine what stimulus features are optimal for performance in specific visual tasks, how those features should be combined to optimally perform those tasks, and how human performance compares with optimal performance. Our methods for determining optimal stimulus features and optimal performance are both based on the concepts of Bayesian statistical decision theory. Our results suggest that a quantitative analysis of the natural scene statistics that support natural tasks can often provide novel quantitative predictions of visual performance and deep insight into the design of the visual system.

[doi:](#)

T-17. Eyes 3, 4 and 5 would be most mystifying structures if one did not know that flies flew

Simon B. Laughlin

S.LAUGHLIN@ZOO.CAM.AC.UK

University of Cambridge

Blowfly flight is stabilized against body rotation by fusing information from several sensors [1,2]. We demonstrated how an ensemble of identified VS neurons in the lobula plate combines information on head rotations from two optical sensors, the ocelli and the compound eyes. The ocelli are wide field optical integrators that respond to movements of the horizon across their visual fields. We made intracellular recordings from VS neurons and identified them by dye injection. We measured the tuning of VS neurons' responses to head rotations sensed by the ocelli using a novel "virtual reality" stimulus. This stimulus uses an optical model of the three ocelli viewing the ground and sky to calculate the optical signals generated in all three ocelli as the head rotates around a given horizontal axis. Sets of these rotation signals are then delivered to the ocelli via three fibre-optic light guides to measure a VS neuron's response to head rotations, as sensed exclusively by the ocelli. Each VS neuron we recorded received an ocellar signal that is cosine tuned to one of 3 horizontal axes of rotation. However, the VS neurons are known to extract 9 horizontal axes of rotation from compound eyes [3]. VS neurons combine these two sets of axes simply - each cell takes the ocellar axis closest to its compound eye axis. This combination fuses information from a faster low accuracy input and a slower high accuracy input. Because the ocellar tuning is broader and ocellar signals are faster (one third the latency and high pass filtered) we suggest that the fusion of both inputs will not unduly degrade either. We refer to this strategy of fusing mutually supportive data streams as complementary coding. The ocellar axes align with the axes of flight instability and this matching of sensory and motor representations should improve efficiency [2]. References 1. Hengstenberg, R. (1993). Multisensory control in insect oculomotor systems. *Rev. Oculomot. Res.* 5, 285-298. 2. Taylor, G.K., and Krapp, H.G. (2007). Sensory systems and flight stability: what do insects measure and why? In *Insect Mechanics and Control*, Volume 34, J. Casas and S.J. Simpson, eds. (London: Academic Press), pp. 231-316. 3. Krapp, H.G., Hengstenberg,

B., and Hengstenberg, R. (1998). Dendritic structure and receptive-field organization of optic flow processing interneurons in the fly. *J. Neurophysiol.* 79, 1902-1917.

[doi:](#)

T-18. A generative model of the covariance structure of images

Geoffrey E. Hinton

HINTON@CS.TORONTO.EDU

University of Toronto

It is possible to model images using a Markov Random Field in which pairwise interactions between nearby pixels are used to define an energy function that favors smooth patches. A more powerful model, first suggested by Geman and Geman, is to use a gated MRF in which binary hidden variables determine the weights on the lateral connections between pixels. The states of these binary variables are easy to infer from an image and they provide a description of the image structure that is very useful for recognizing objects. I will describe a way to greatly reduce the number of parameters in such a model so that it can be learned fairly easily. The model can be made to learn nice topographic maps by using local connectivity between a layer of "simple" cells and a layer of "complex" cells. When applied to RGB patches of natural images, this new generative model learned a map that consisted mainly of fairly high-frequency gray-level filters that learned to be exactly balanced in the RGB channels. The map also contained a color blob consisting of low-frequency red-green and yellow-blue filters. The division of filters into spatially separate gray-level and color-opponent types does not seem to require any innate specification. It can be explained by the the 3-way interaction of a general unsupervised learning algorithm, local connectivity, and the statistical structure of natural images. This is joint work with Marc'Aurelio Ranzato.

[doi:](#)

T-19. Evidence for a neural model to evaluate symmetry in V1

Horace B. Barlow

HBB10@CAM.AC.UK

University of Cambridge

50 years ago Hubel and Wiesel discovered simple and complex cells in V1, but there is still no consensus on their functional roles. It is agreed that complex cells are more often selective for direction of motion than simple cells, that there are differences in the way they combine information within their receptive fields, and that complex cells probably receive most of their input from simple cells, but what this serial hierarchy achieves is not understood. There is another puzzling dichotomy that we think is related, namely that of cross-correlation, which is widely accepted as the operation performed on the input image by simple cells, and auto-correlation, which some think underlies the perception of Glass patterns, and possibly motion. We propose the hypothesis that complex cells signal auto-correlations in the visual image, but to evaluate them they require the preliminary analysis done by simple cells, and also pinwheels - structures intervening between simple cells and complex cells that were quite unknown to Hubel and Wiesel. We shall first present psychophysical evidence, using a new kind of random dot display, which suggests that both cross-correlation and auto-correlation are performed in early vision. We then point to recent evidence on the micro-circuitry of pinwheels, and mappings of their intrinsic activity, which shows how pinwheels might enable complex cells to respond selectively to autocorrelations in the input image that activates the simple cells. Auto-correlation is a powerful tool for detecting symmetry, and many may be surprised by evidence that such an abstract property is detected so early in visual perception.

[doi:](#)

T-20. Timing in the auditory cortex

Anthony M. Zador

ZADOR@CSHL.EDU

Cold Spring Harbor Laboratory

Attending to moments in time is a powerful cognitive mechanism for exploiting temporal structure in behaviors such as hunting moving prey or playing music in an ensemble. We have developed a novel auditory two-alternative choice paradigm in which we could study the cortical mechanisms of temporal expectation in rats. I will present experiments designed to test the hypotheses that temporal expectation influences early auditory cortical processing, and that its influence on the neuronal activity of auditory cortex is causally related to the observed improvements in performance.

[doi:](#)

T-21. Differential sensitivity of different sensory cortices to behaviorally relevant timing differences

Yang Yang

Anthony M. Zador

YANGYANG@CSHL.EDU

ZADOR@CSHL.EDU

Cold Spring Harbor Laboratory

Animals can detect the fine timing of some stimuli. For example, in subcortical structures, submillisecond interaural time differences are computed to determine the spatial localization of sound. Although cortical neurons have not been shown to achieve comparable submillisecond precision, neurons in auditory, visual and barrel cortex can lock with millisecond precision to the fine timing of some stimuli. However, the ability of these cortical neurons to fire precisely does not demonstrate such fine timing is behaviorally relevant. To bridge the gap between physiology and behavior, we have previously used electrical microstimulation to determine the temporal precision with which fine differences in cortical spike timing could be used to drive decisions. We found that in rat auditory cortex, animals could be trained to use timing differences as short as 3 msec to drive decisions (Yang et al, Nat Neurosci 11, 1262-3). Is the auditory cortex unique in its ability to utilize such fine timing differences to drive behavior? Because audition is often considered to be a "fast" modality—one in which subtle differences in temporal structure can be behaviorally relevant—it would not be unreasonable to speculate that the auditory cortex had evolved special mechanisms for rapid processing. On the other hand, it is appealing to hypothesize that the cortex operates according to general principles shared across different regions; in this view, the ability to make use of millisecond-scale differences in neuronal activity would not be unique to the auditory cortex. To distinguish these hypotheses, we compared the ability of different sensory areas to resolve subtle differences in neural timing. In the visual cortex, we found that although animals could be trained to resolve differences as short as 15 msec in neuronal activity, they could not resolve differences as short as 5 msec. This lower limit of 5-15 msec was significantly higher than the limit of 3 msec we observed in auditory cortex, and is consistent with the view that visual cortex is "slower" than auditory cortex. Surprisingly, we found that the barrel cortex was even "faster" than auditory cortex, with a lower limit below 1 msec. Our results suggest that different cortical areas are differentially able to derive behaviorally relevant information from the fine timing of neural activity.

[doi:](#)

T-22. The Poisson clicks task: long time constant of neural integration of discrete packets of evidence

Bingni W. Brunton^{1,2}
Carlos D. Brody³

BWEN@PRINCETON.EDU
BRODY@PRINCETON.EDU

¹Princeton Neuroscience Institute

²Molecular Biology Dept, Princeton University

³HHMI and Princeton University

With the goal of studying neural mechanisms of integration, we have developed a discrimination task that explicitly focuses on promoting integration of evidence over time. During each trial, subjects hear a series of randomly timed clicks from two well-separated speakers. The subjects must report which speaker played the greater total number of clicks. This requires keeping a running counter of clicks (i.e., integrating clicks) over the stimulus period. Task difficulty is controlled by the magnitude of the difference in click rate for the two speakers. Each click is a discrete quantum of evidence, occurring at a single, well-defined and experimenter-controlled timepoint. In this design, then, each trial's inputs to a putative neural integrator of evidence are precisely known. We call the task the "Poisson clicks" task. We used the well-known drift-diffusion framework to model the neural integrator as a noisy, possibly biased, finite time constant (τ) integrator. A positive τ indicates a tendency to forget evidence that arrived more than τ seconds ago. A negative τ indicates a tendency to make a decision based only on the initial τ seconds of the stimulus. The ideal observer and integrator would have infinite τ . Current proposals of neural architectures for integration have variously suggested positive, negative, and infinite τ . Data from monkeys have suggested a magnitude of τ of several hundred milliseconds. We trained rats to perform our discrimination task. Knowledge of the specific trial-by-trial inputs to the integrator allowed obtaining accurate estimates of its parameters. Complete likelihood landscapes for the model parameters were computed, using both numerical integration of the drift-diffusion model's Fokker-Planck equations and closed-form solutions. We found that (1) best-fitting τ are overwhelmingly positive, consistent with neural architectures approximated by leaky integrator models and inconsistent with unstable integrator architectures. (2) Well-trained subjects can achieve remarkably long time constants τ of up to 1000 ms, suggesting that the task was indeed successful at promoting integration. Analysis of stimulus-behavior correlations further confirmed the long integration time constant. (3) During learning, the signal-to-noise ratio of the integrator quickly stabilizes to its final value, but τ grows only slowly, achieving its final value after ~4 months of training. This indicates that in this task, associational learning is completed quickly, and the brunt of perceptual learning consists of fine-tuning of the integrator, once again consistent with the task being focused on integration. These findings establish the Poisson clicks task as particularly appropriate for studying neural integration. Our findings further show that rats are capable of integrating evidence over long time constants, and thus establish rats as a viable model system for studying neural integration.

[doi:](#)

T-23. Action video games as exemplary learning tools

Daphne Bavelier
University of Rochester

DAPHNE@BCS.ROCHESTER.EDU

Although the adult brain is far from being fixed, the types of experience that promote learning and brain plasticity in adulthood are still poorly understood. Surprisingly, the very act of playing action video games appear to lead to widespread enhancements in visual skills in young adults. Action video game players have been shown to outperform their non-action-game playing peers on a variety of sensory and attentional tasks. They search for a target in a cluttered environment more efficiently, are able to track more objects at once and process rapidly fleeting images more accurately. This performance difference has also been noted in choice reaction time tasks with video game players manifesting a large decrease in reaction time as compared to their non-action-game playing peers. A common mechanism may be at the source of this wide range of skill improvement. In particular,

improvement in performance following action video game play can be captured by more efficient integration of sensory information, or in other words, a more faithful Bayesian inference step, suggesting that action gamers may have learned to learn.

[doi:](#)

T-24. Pupillometric evidence for a role of locus coeruleus in dynamic belief updating

Matthew Nassar¹
Robert C. Wilson²
Rishi Kalwani¹
Benjamin Heasley¹
Joshua Gold¹

MATTNASSAR@GMAIL.COM
RCW2@PRINCETON.EDU
RISHIKI@MAIL.MED.UPENN.EDU
BENJAMIN.HEASLY@GMAIL.COM
JIGOLD@MAIL.MED.UPENN.EDU

¹University of Pennsylvania

²Princeton University

Many decisions are based on subjective beliefs about the probability and utility of potential outcomes. These beliefs are typically adjusted upon observation of a new outcome. One biologically supported rule for this belief-updating process is the delta rule, according to which a belief is updated by some fraction of the error made in predicting an outcome. The fractional term, the learning rate, regulates the influence of incoming information on a stored belief and is often set as a constant. However, theoretical and experimental work suggests that the brain does not use a constant learning rate but instead adaptively regulates the influence of new outcomes [A.J.Yu and P.Dayan NIPS 157-164,2002; T.E.J.Behrens et.al. Nature Neuroscience 10:1214-1221,2007]. Little is known about the underlying neural mechanisms. Here we identify a physiological correlate of this adaptive regulation process involving changes in pupil diameter, which are thought to reflect the activity of the norepinephrine-locus coeruleus (NE-LC) system. Human subjects performed a novel inference task that required them to make predictions about a random quantity whose value fluctuated from trial to trial because of either noise or changepoints in the statistics of the generative process. A computationally tractable reduction of the ideal-observer model for this task is a form of delta rule in which the learning rate depends on the number of outcomes that occurred since the most recent changepoint (run length) and the probability that a changepoint occurred on the current trial (changepoint probability). The model can predict trial-to-trial changes in subject learning rates. Such rational adaptation of learning rate has been proposed to involve the NE-LC and other neuromodulatory systems [A.J.Yu and P.Dayan Neuron 46:681-692,2005]. To test this idea, we measured pupil diameter in subjects performing a version of the task with an isoluminant display, a condition under which pupil diameter is correlated with the activity of LC neurons R.Kalwani, unpublished data. We found that pupil diameter changed relative to task events on multiple timescales. Transient changes in pupil diameter occurred on each trial after the presentation of a new outcome. The magnitude of these phasic responses was correlated with model estimates of changepoint probability and subject learning rates. More gradual changes in pupil diameter occurred on the timescale of many trials. This baseline pupil diameter was greatest 1-2 trials after a changepoint and correlated inversely with run-length estimates produced by the model. To examine whether LC plays a causal role in setting the learning rate we used a salient, but task irrelevant, auditory stimulus to evoke phasic pupil (and thus, presumably, LC) responses. Despite knowing that these stimuli were irrelevant to the task, subjects increased their learning rates after stimuli that caused large phasic pupil responses. Together, these data suggest that the NE-LC system helps to dynamically regulate the influence of new outcomes based on information about run length and changepoint probability.

[doi:](#)

T-25. Detection and estimation of defocus in natural images

Johannes Burge¹
Wilson S. Geisler²

JBURGE@MAIL.CPS.UTEXAS.EDU
GEISLER@PSY.UTEXAS.EDU

¹Center for Perceptual Systems, University of Texas, Austin

²University of Texas, Austin

For most mammals, and many other animals, vision begins with an optical system that focuses light near the plane of the photoreceptors. These optical systems generally have limited depth of focus. Thus, the images of objects lying outside the current focus distance are defocused (blurred) by an amount that increases with the distance of the object from the focus distance. The defocus signals at the retina play an important role in many aspects of vision including accommodation, the estimation of scale, distance, and depth, and the control of eye growth. However, little is known about the neural computations visual systems use to detect and estimate the magnitude of defocus under natural conditions. We investigated how to optimally estimate defocus blur in images of natural scenes, given the optical systems of primates. First, we selected a large set of well-focused luminance-calibrated natural image patches. Next, we filtered each image patch with point-spread functions derived from a realistic wave-optics model of the primate (human) eye at different levels of defocus. Finally, we used a statistical learning method, based on Bayesian ideal observer theory, to determine the spatial-frequency filters, or spatial receptive fields, that are optimal for estimating image defocus for the natural patches. We found that near the center of the visual field, the spatial-frequency filters that are optimal for estimating moderate defocus in natural image patches form a systematic set that is concentrated in the range of 5-15 cyc/deg, the range known to drive human accommodation. Furthermore, we found that the optimal filters can be closely approximated by a linear combination of a small number of difference-of-Gaussian filters. Cells with such center-surround receptive field structure are commonplace in the early visual system. These results therefore predict that retinal neurons sensitive to this frequency range should contribute strongly to the retinal and/or post-retinal mechanisms that detect and estimate defocus in and near the fovea. The optimal filters were also used to perform the task of detecting, discriminating, and identifying defocus levels for natural image patches of 1 deg diameter. Consistent with human psychophysical data, detection thresholds were higher than discrimination thresholds. Also, once defocus exceeds 0.25 diopters, we found that 0.25 diopter steps in defocus can be identified with better than 86% accuracy. In conclusion, we found that a systematic set of biologically plausible spatial-frequency filters is optimal for estimating defocus, given the primate optical system. Similar analyses could be carried for the optical systems of other animals or machine vision systems. The estimated optimal filters, and associated optimal decoding rules, provide a rigorous starting point for developing principled hypotheses for the neural mechanisms that encode and exploit optical defocus signals.

doi:

T-26. Invariant contrast coding in photoreceptors

Uwe Friederich
Daniel Coca
S. A. Billings
Mikko Juusola
University of Sheffield

U.FRIEDERICH@SHEFFIELD.AC.UK
D.COCA@SHEFFIELD.AC.UK
S.BILLINGS@SHEFFIELD.AC.UK
M.JUUSOLA@SHEFFIELD.AC.UK

In natural environments, visual stimuli continuously change in time and space. Light intensities of just a single scene can vary 10,000-fold, whereas the dynamic signalling range of photoreceptors is less than 100-fold. Adaptation mechanisms tune a photoreceptor's input-output relationship to extract relevant information in the output. However, it remains an open question, how adaptation shapes neural responses when the statistics of naturalistic scenes change. By combining electrophysiological experiments with an empirical modeling methodology, we show that *Drosophila* photoreceptors adapt to preserve light contrast information at different luminance levels. We derived a deterministic model that can simulate accurately contrast encoding at each luminance level by a simple

gain adjustment. This approach separates adaptation from noise in the neural responses and estimates unbiased transfer characteristics. We show that the gain depends on the statistical properties of the stimulus and present an adaptation model that continuously readjusts its output to changes in the stimulus history. Intracellular voltage responses (output) of photoreceptors to changing light contrast patterns (input) were measured using sharp microelectrodes. The same contrast patterns, delivered from a point source, were repeated at fixed luminances that changed instantaneously up to 10,000-fold. Experimental input-output data was used to estimate and validate the structure and parameters of deterministic NARMAX models at each light level. Based on these empirical models and their analytically computed multidimensional frequency response functions, we studied the neuron's signal transfer at each luminance level. A major advantage of this approach is the iterative estimation of a noise model that allows unbiased estimates even if noise is not additive or white. Here, classical approaches, which assume linear superposition between the average response (signal) and noise, can misestimate the real signal transfer, especially when the input SNR is low. We found that after initial adaptive trends, frequency response functions throughout all luminance levels have consistent shapes and are merely shifted by a constant gain. Thus, a contrast pattern encountered at different luminances evokes a structurally very similar neural response pattern. This general coding property allowed us to design a unified deterministic NARMAX model, which accurately predicts photoreceptor responses at each tested light level by adjusting its input gain. We further show that the dynamics of the voltage responses to continuously changing naturalistic contrasts can be replicated by a dual model structure in which a separate "adaptation model" tunes the gain of the NARMAX model based on the stimulus history over multiple time scales. *Drosophila* photoreceptors are part of the lamina network, in which local and global synaptic feedbacks modulate visual processing. We repeated the same experiments with mutant photoreceptors that are synaptically isolated from the network because they cannot produce the neurotransmitter, histamine. By comparing mutant and wildtype models, we show that the network mostly influences the photoreceptors' contrast coding to bright inputs. The lack of synaptic connectivity seems to reduce the range of environmental intensities to which photoreceptors can adapt. However within this limited range, contrast coding appears to be unaffected, suggesting that adaptation takes place mainly in the phototransduction.

[doi:](#)

T-27. Spike-triggered covariance and synthetic image replay reveal nonlinearities in V1 color processing

Gregory Horwitz^{1,2}

GHWITZ@U.WASHINGTON.EDU

¹University of Washington

²Washington National Primate Research Center

The processing of color information in cortical area V1 remains poorly understood. The vastness of the stimulus space is a fundamental problem: not every stimulus pattern can be displayed during a standard neurophysiology experiment. In a previous study, we addressed this problem by stimulating V1 neurons in awake, fixating monkeys with random, colorful patterns and used dimensionality-reduction techniques (averaging and principal components analysis) to isolate features of the stimulus patterns to which the neurons responded¹. The stimulus tuning of some neurons was well characterized by a single dimension in the stimulus space. For these neurons, we observed structure in the spike-triggered average stimulus and little or no structure in the principal components of the spike-triggered stimuli. Other neurons responded to multiple spike-triggered stimulus features. These features were manifest in the spike-triggered average and in one or more of the principal components. A subpopulation of neurons had a spike-triggered average that indicated sensitivity to a preferred color throughout the receptive field and a first principal component that indicated sensitivity to luminance edges. These results are inconsistent with a linear model of cone signal integration. We investigated the relationship between chromatic and luminance signals in these neurons by fitting a 2-dimensional linear-nonlinear cascade model to the data. The fits suggested that these neurons carry color-opponent signals, the gain of which is enhanced by luminance contrast. In the current study, we tested this prediction directly by performing the spike-triggered averaging and principal components analyses on the data stream as it was being collected. Spike-triggering stimulus features were extracted from these analyses and linearly combined to create a battery of synthetic images. Synthetic images were displayed

at the neuron's receptive field for 200 ms each, and spikes following each presentation were counted. Results of this experiment confirmed the model predictions: a subset of V1 neurons responded modestly to their preferred color in isolation but strongly to their preferred color superimposed on a luminance edge. We conclude that color processing across a subpopulation of V1 neurons is enhanced at luminance edges. This physiological result may be related to interactions between color and luminance observed psychophysically. Acknowledgments We thank J. Gold and C. Hass for assistance with the computer communication code. This work was supported by the McKnight Foundation and NIH grant RR000166. References [1] Blue-yellow signals are enhanced by spatiotemporal luminance contrast in macaque V1. G. D. Horwitz, E. J. Chichilnisky, and T. D. Albright, *Journal of Neurophysiology* 93:2263-2278, 2005.

doi:

T-28. Metamers of the ventral stream

Jeremy Freeman¹
Eero P. Simoncelli²

FREEMAN@CNS.NYU.EDU
EERO.SIMONCELLI@NYU.EDU

¹Center for Neural Science, NYU

²New York University

How is image structure encoded in the extrastriate ventral visual pathway? Direct characterization of the stimulus selectivity of individual extrastriate cells has proven difficult. However, one robust population-level property of all visual areas is that receptive field sizes grow with eccentricity. It has also been reported (Gattass et al., 1988) that the rate of growth increases along the ventral stream. We hypothesize that this successive increase in pooling region size causes information loss. A well known example occurs in the retina, where spatial pooling in the periphery means that high spatial frequency information is lost. In general, stimuli that differ only in terms of information discarded by the visual system will be indistinguishable to a human observer. Such stimuli are called metamers. Here, we probe the population-level computations of the ventral stream using novel metameric stimuli. Starting from any prototype image, we generate stimuli that match in terms of the responses of a simple model for extrastriate ventral computation. The model is based on measurements previously used to characterize visual texture (Portilla & Simoncelli, 2000). The model decomposes an image using a bank of V1-like filters tuned for local orientation and spatial frequency, computing both simple and complex-cell responses. Extrastriate responses are then computed by taking pairwise products amongst these V1 responses, and averaging within overlapping spatial regions that grow with eccentricity. Stimuli are generated by using gradient descent to adjust a random (white noise) image to match the model responses of the original prototype. Previous work showed that the same statistics, averaged over an entire image, allow for the analysis and synthesis of homogenous visual textures. If this model accurately reflects representations in early extrastriate areas, then images synthesized to produce identical model responses should be metameric to a human observer. For each of several natural images and pooling region sizes, we generate multiple samples that are statistically-matched but otherwise as random as possible. We use a standard psychophysical task to measure observers' ability to discriminate between image samples, as a function of the rate at which the statistical pooling regions grow with eccentricity. When image samples are statistically matched within small pooling regions, observers perform at chance (50%), failing to notice substantial differences in the periphery. When images are matched within larger pooling regions, discriminability approaches 100%. We fit the psychometric function to estimate the pooling region over which the observer estimates statistics. The result is consistent with receptive field sizes in macaque mid-ventral areas (particularly V2). Our model also fully instantiates a recently proposed explanation (Balas et al., 2009) of the phenomenon of "visual crowding", in which humans fail to recognize a peripheral target object surrounded by background clutter. In our model, crowding occurs because multiple objects fall within the same pooling region and the model responses cannot uniquely identify the target object. We synthesize images that are metameric to classic crowding stimuli (e.g. groups of letters), and find that stimulus configurations that produce crowding yield synthesized images with jumbled, unidentifiable objects.

doi:

T-29. The control of visual information by prefrontal dopamine

Tirin Moore

TIRIN@STANFORD.EDU

Stanford University

A principal function of the prefrontal cortex (PFC) is executive control, and this control includes the modulation of sensory signals during goal-directed behavior. Dopamine (DA)-mediated activity within the PFC is thought to play an important role in executive control, yet whether it contributes to sensory filtering is not known. I will discuss experiments that demonstrate an involvement of dopamine-mediated PFC activity in the modulation of visual representations. Local blockade of D1 receptors within the frontal eye field (FEF) increases behavioral target selection within the corresponding part of visual space and enhances the magnitude, stimulus discriminability and response reliability of visual responses in area V4. In contrast, local inactivation of FEF activity produces opposite effects. These results demonstrate that prefrontal D1 receptors are involved in top-down modulation of visual cortical representations.

[doi:](#)

T-30. High frequency entrainment of thalamic neurons by basal ganglia output in the singing bird

Jesse H. Goldberg

JESSEG1@MIT.EDU

Michale S. Fee

FEE@MIT.EDU

Massachusetts Institute of Technology

The basal ganglia (BG) are implicated in motor control and learning, and are associated with neuropsychiatric disorders such as Parkinson's, Huntington's, and dystonia. The major output of the BG circuit-pallidal neurons that are GABAergic and tonically active-are thought to control movements by inhibiting their targets, thalamocortical neurons in the ventral 'motor' thalamus. In the dominant gating model of BG output, thalamic spiking is suppressed by the tonic high firing rate of pallidal neurons (gate closed). Thalamic neurons are disinhibited (gate open) by decreases, or pauses, in pallidal activity, permitting movement. This view has long informed our understanding of BG-related disease. Hyperkinetic disorders like dystonia are thought to result from insufficient pallidal inhibition of thalamus, while hypokinetic diseases like Parkinson's result from excess inhibition. However, the gating model fails to explain why both pallidus and thalamus are activated during movements and why high frequency deep brain stimulation in either the pallidus or the thalamus can relieve both hypo- and hyperkinetic motor pathology. Recently, the songbird has emerged as a model system to study pallidothalamic transmission. Songbirds have a specialized BG-thalamocortical circuit, called the anterior forebrain pathway (AFP), devoted to song learning. The input structure of the AFP, Area X, contains striatal elements as well as a pallidal projection to the medial portion of the dorsolateral thalamus (DLM). Because the pallidal axon terminal forms a specialized calyx that can be recorded extracellularly in DLM, it is possible to record from connected pallidothalamic pairs and directly test models of BG-thalamic interaction. Here we examine how BG output controls thalamic activity in the behaving animal by recording from connected pallidothalamic pairs in DLM of singing juvenile birds. In contrast to the predictions of the gating model, we find that increases in pallidal firing rates do not suppress thalamic spiking. Instead, connected pallidal and thalamic neurons fire in concert at hundreds of spikes per second and are paradoxically coactivated during singing. In addition, thalamic spikes are time-locked to the most recent pallidal spike with submillisecond precision, and show no dependence on prior history of pallidal or thalamic activity. These results suggest a novel model of basal ganglia output-that of high frequency entrainment-where basal ganglia outputs control thalamic spike timing on a spike-by-spike basis.

[doi:](#)

T-31. Beside the point: Motor adaptation without feedback error correction in task-irrelevant conditions

Sydney Y. Schaefer¹
 Iris L. Shelly²
 Kurt A. Thoroughman²

SYSCHAEFER@GMAIL.COM
 ILS1@CEC.WUSTL.EDU
 THOROUGHMAN@WUSTL.EDU

¹Department of Physical Therapy, Washington U.

²Department of Biomed Eng, Washington U.

There are several theories on how movement errors and task goals drive motor adaptation. Feedback error learning theory posits that the sensory feedback with respect to a specified task goal (or target) comprise the error signal that drives adaptation (Kawato et al. 1987). Distal supervised learning theory, however, considers error to be the difference between expected and actual performance (Jordan and Rumelhart 1992); therefore, the performance could be relevant or irrelevant to the task. Here we ask: Can motor adaptation occur even when errors are irrelevant to the task goal? Many classic and recent experiments ask subjects to perform planar reaches to a point target defined by two dimensions: a direction and an extent. Here, we introduce experiments in which the goal of the task (i.e. the target) was defined by a single dimension. Subjects (n=8) reached on a digitizing tablet from a fixed start location to one of three targets: a point (located 90° and 10 cm from start), an arc, or a ray. For the arc, subjects were allowed to move in any direction, but to a specific extent. For the ray, subjects could extend their arm to any distance, but to a targeted direction. Subjects were provided visual feedback about the location of their hand via a cursor during the entire movement. Subjects first moved to the point as baseline performance. Next they adapted to one of two visuomotor perturbations: cursor rotation that perturbed the direction of movement (30° clockwise) or cursor gain that perturbed the extent of movement (minimized by factor of 0.8). Therefore, visuomotor perturbation was either relevant or irrelevant to the task goal. Subjects responded to task-relevant perturbations with mid-movement corrective responses, evidenced by mid-movement curvature in direction perturbation ($p < 0.001$) and increased movement time in extent perturbation ($p < 0.01$) early in training. Critically, subjects produced no feedback-mediated corrective responses in task-irrelevant direction ($p > 0.2$) or extent ($p > 0.4$) perturbations. Our perturbation conditions therefore induced clear differences in movement goal and in engagement of feedback control. Despite these task and feedback differences, relevant and irrelevant perturbations induced similar motor memories. To ascertain adaptation in task-relevant and irrelevant conditions, we interspersed catch trials in which we removed the perturbation and switched back to the point target. Even though the goal of the baseline and catch trials was the same, subjects demonstrated after-effects in catch trials that indicated behavioral adaptation in response to the perturbation. Adaptation therefore generalized from the one-dimensional tasks (arc, ray) to the two-dimensional task (point). Moreover, such after-effects were present regardless of whether the perturbation was relevant to the task. Across subjects, there was no significant difference in after-effects when the perturbation was relevant or irrelevant to the task goal, during both rotation ($p = .23$) and gain adaptation ($p = .32$). Adaptation to task-irrelevant perturbations supports the hypothesis that the nervous system can rely on sensory prediction, absent task-dependent motor error or error-driven feedback control, to drive motor learning.

[doi:](#)

T-32. Conscious or not? How neuroscience is building a bridge to understanding recovery following severe brain injury

Nicholas D. Schiff
 Weill Cornell Medical College

NDS2001@MED.CORNELL.EDU

This presentation will review new studies of recovery of consciousness following severe brain injuries. Potential circuit-level mechanisms underlying several seemingly unrelated observations will be emphasized.

[doi:](#)

T-33. Hippocampal processes underlying episodic memory

John Lisman

LISMAN@BRANDEIS.EDU

Brandeis University

A great deal is now known about the properties of hippocampal cells and their role in episodic memory. The properties of the input to the hippocampus (the grid cells of the entorhinal cortex) are known, as are the place field properties of hippocampal cells. The detailed wiring diagram of the excitatory connections is known. Finally, the hippocampus shows prominent oscillations in the gamma and theta frequency range. These organize a dramatic example of temporal coding, a phenomenon called the phase precession. We have been attempting to integrate these findings into a model that accounts for the data and explains how the system stores and recalls episodic memories. Our work points to four principles. 1) Gamma oscillations have two functions: they determine which cells fire by a network-mediated winner-take-all (WTA) process and they control the timing of that firing, thereby producing synchronized cell assemblies. 2) There are about seven gamma cycles nested in each theta cycle. These dual oscillations form a discrete theta phase code. The phase precession can be understood as a cued recall of sequential positions along a path, organized by a theta-gamma discrete phase code. There are now reasons to suspect that the theta-gamma code is used widely in the brain. 3) The input/output transformation of granule cells in the dentate gyrus turns grid cell representations to place fields. This transformation can occur without learning; it can be quantitatively accounted for by the random summation of grid cells inputs and the gamma-mediated WTA process. 4) Samplinsky postulated that accurate memory sequence recall requires iterative interaction between two kinds of synaptic weights. These can now be mapped onto the reciprocal connections between dentate and CA3. The two kinds of weights are: a) Heteroassociative weights (which produce chaining between sequential positions in the sequence and which are postulated to be at the feedback connections from CA3/mossy cells to dentate granule cells) and 2) Autoassociative weights (which clean up errors produced after each step in the chaining process) and which are postulated to be at CA3 recurrent synapses. Taken together, this framework provides a data-constrained explanation of how sequential positions along a path are stored in the hippocampal episodic memory store.

[doi:](#)

T-34. Coordinated hippocampal firing across related spatial locations develops with experience

Anabelle C. Singer¹

ANNABELLE.SINGER@UCSF.EDU

Mattias P. Karlsson²

KARLSSONM@JANELIA.HHMI.ORG

Ana R. Nathe¹

ANA@PHY.UCSF.EDU

Margaret F. Carr¹

MARGARET.CARR@UCSF.EDU

Loren M. Frank¹

LOREN@PHY.UCSF.EDU

¹Keck Center and Department of Physiology UCSF

²HHMI, Janelia Farm Research Campus

The world is full of repeating elements, like city blocks, rolling hills, or trees evenly spread through a forest. However, we do not fully understand how neural representations organize these spatial elements. On the one hand, neurons may encode the similarities among these elements to extract general principles about the environment and to facilitate the application of learned information to new experiences. On the other hand, neurons might encode each element very differently to easily distinguish between them and form unique associations with each element. The hippocampus is required for spatial learning and separating between and generalizing across similar experiences. Hippocampal place cells can show similar firing patterns across locations, but the functional significance of this activity and the role of experience and learning in generating it are not understood. We hypothesized that these similar coding patterns reflect learned generalizations across different places and episodes. We therefore examined hippocampal place cell activity in the context of spatial tasks with multiple similar spatial trajectories. We found that some hippocampal neurons fire in multiple similar locations in environments

with multiple similar spatial trajectories even when animals must behaviorally distinguish among the trajectories. The prevalence of this path equivalent coding increased as animals learned specific rewarded sequences and the relationships between trajectories. Furthermore, path equivalent firing is not simply due to single cells acting independently. Rather pairs of path equivalent cells that repeated together in multiple segments had correlated moment to moment activity, suggesting they were part of functional ensembles. While the similarity of cells' receptive fields in a single location was not predictive of these moment to moment correlations, similarities over multiple locations were. To our knowledge, this is the first demonstration that 1) the firing properties of two place cells in one location can predict their patterns of spatial activity in another location and 2) that correlations in moment to moment activity is related to cells firing similarly over multiple locations but not a single location. These correlations also increased with experience, supporting the hypothesis that learning drives the development of path equivalent ensembles. Thus, our data indicate that, in environments with many repeating elements, ensembles of cells are recruited together to represent general features of the environment. While path equivalence is common in the two environments with repeating elements that we examined, cells that fire in multiple similar locations in an environment fire at different peak rates in different locations. Furthermore, only about half the cells had path equivalent activity. Therefore, path equivalence could provide important information about similarities across different locations while, at the population level, the system could distinguish among those locations using differences in firing rate across different locations in path equivalent cells and cells that fire differently across different trajectories. We propose that this path equivalent activity could be a mechanism to generalize efficiently across related experiences in the hippocampus. General information about related experiences could then be used to encode common behavioral associations across experiences. This work was supported by NIH grant MH077970.

doi:

T-35. Temporal transformations in olfactory encoding promote rapid detection of natural odor fluctuations

Katherine Nagel

Rachel Wilson

Harvard Medical School

KATHERINE_NAGEL@HMS.HARVARD.EDU

RACHEL_WILSON@HMS.HARVARD.EDU

Natural olfactory stimuli form turbulent plumes that fluctuate rapidly in time and may provide information about the location of an odor source. However, olfactory neurons respond to odor with complex and prolonged spike dynamics that would seem unsuitable for encoding such stimuli. We used *Drosophila* as a model system for understanding how complex response dynamics arise in the first layer of an olfactory system and what function they might serve. Using pharmacological and genetic tools, we showed that we could obtain separate measures of transduction currents and spikes from olfactory receptor neurons (ORNs) in vivo. This revealed that ORN response dynamics arise from two non-trivial temporal transformations: one imposed by odor transduction, and one by spike generation. Odor transduction dynamics were fast, and were well-described by a model of ligand binding and receptor activation. This model was sufficient to account for differences in response dynamics across odors and receptors, asymmetries between onset and offset dynamics, and receptor-specific adaptation and recovery. Spiking dynamics were captured by a differentiating linear filter that was similar across odors and ORNs. Spike dynamics could be altered by genetic knockdown of sodium channel expression levels, suggesting that they are tightly regulated. This is consistent with Hodgkin Huxley models showing that the temporal selectivity of a neuron depends on its balance of sodium and potassium conductances and suggests that the differentiating transformation arises from known models of spike generation. In the context of rapid and intermittent natural stimuli, the transduction and spiking transformations lead to fast and highly sensitive reporting of plume encounters, while in response to longer synthetic pulses they produce more complex dynamics. Our analysis suggests that complex spike patterns in olfaction are an unavoidable consequence of the biophysics of odor transduction, coupled to adaptations for rapid encoding of transient plume encounters.

doi:

T-36. Experimental evolution to probe gene networks underlying cognition in *Drosophila*

Josh Dubnau

DUBNAU@CSHL.EDU

Cold Spring Harbor Laboratory

One of the great challenges to understanding genetic impact on human cognitive disorders is that clinical outcomes often are influenced by interactions among groups of genes. A good example is schizophrenia, where no clear genetic mechanism has emerged. In some individuals, complex disorders such as schizophrenia likely emerge from co-inheritance of multiple common gene variants each of which would have little clinical impact on their own. Despite its widespread relevance, mechanisms by which multi-gene interactions modulate phenotype are ill understood because almost all mechanistic studies of gene interaction are limited to pair-wise combinations. To investigate this question, we have developed and implemented a novel approach in *Drosophila*, using the biologically important and clinically relevant cAMP pathway as a model. Our approach uses the power of selective breeding to evolve combinations of gene variants capable of ameliorating the learning defect of a mutation in the rutabaga adenylyl cyclase gene. This method is novel in several ways. First, unlike a classical suppressor screen, our use of experimental evolution has allowed us to explore the impact of higher order combinations of gene interactions. Also, unlike a classical selective breeding experiment, we have constrained the genetic variability to a set of 23 identified and molecularly characterized loci with known involvement in memory. Our strategy models the multi-gene interactions that influence naturally occurring variation in complex phenotypes such as learning, but also makes it feasible to fully genotype the causative loci across multiple animals. This has given us unprecedented access to the underlying molecular genetic mechanisms.

[doi:](#)

I-1. Intrinsic dendritic plasticity maximally increases the computational power of CA1 pyramidal neurons.

Romain Cazé^{1,2}

Mark D. Humphries³

Boris Gutkin³

ROMAIN.CAZE@GMAIL.COM

M.D.HUMPHRIES@SHEFFIELD.AC.UK

BORIS.GUTKIN@ENS.FR

¹INSERM U960

²ENS DEC

³Group for Neural Theory, LNC, DEC, ENS Paris

What additional computing power do dendrites add to a neuron? Previous work using artificial neural networks has suggested that active dendrites improve the computing power of CA1 pyramidal neurons by increasing the number of possible input/output relationships (Poirazi et al 2001). However, several key questions remain open: what characterizes these new input/output behaviors? Is there a dendritic morphology which maximally increases the computational power of such « dendritic » neurons? And which physiological parameters of the neuron should change to reach this maximal computational power? In order to address these issues we start out with the approach of Poirazi et al (2003) and consider the dendritic tree of the CA1 pyramidal neuron as a two layer neural network with excitatory connections. We begin by showing that we can exhaustively characterize the entire space of possible two-layer networks. Extending results on single layer networks by Minsky and Papert (1988) we prove that any two layer neural network, with real, positive synaptic weights and thresholds, has a discrete, weightless equivalent with the same input/output mapping. Using this discretization, we are able to enumerate all possible input/output functions of a two layer neural network (for up to 6 distinct inputs). Analytically we identify the input/output combinations that could only occur in two-layer but not in one-layer networks. Such input/output combinations, should they be identified in a CA1 pyramidal neuron, would indicate that the cell functionally implements a two-layer neural network. For instance, given four independent Schaeffer collaterals (A, B, C and D) impinging on a CA1 pyramidal neuron, our analysis predicts that if both A+B and C+D elicit a response, but A+C, A+D, B+C

or B+D do not, then the dendritic tree is equivalent to a two-layer network, and cannot be reduced to a single layer. Our analysis also shows that a surprisingly simple dendritic tree morphology is required to maximize the computational power of CA1 neurons: the number of dendrites should match and not exceed the dimensionality of space spanned by the inputs. Our model thus predicts that the ratio of oblique dendrites of the CA1 pyramidal cell to the Schaeffer collaterals should be around one in order to efficiently maximize the range of input/output functions. Finally, we find that modification of the somato-dendritic coupling strengths does not change the input/output mapping of the neuron. Conversely, we show that modifying dendritic branch excitability is necessary to access the full range of input/output functions. Thus, our model suggests that intrinsic dendritic plasticity is key to maximizing the computational power of a CA1 pyramidal neuron, consistent with recent experimental demonstrations of dendritic plasticity in CA1 pyramidal cells (Losonczy et al. 2008).

doi:

I-2. Analytical study of history dependent timescales in a generic model of ion channels

Daniel Soudry
Yariv Meir
Ron

DANIELSO@TX.TECHNION.AC.IL
RMEIR@EE.TECHNION.AC.IL

Electrical Engineering, Technion

Recent experiments have demonstrated that the timescale of adaptation of a single neuron in response to periodic stimuli slows down as the period of stimulation increases. At a sub-neuronal level, experiments on sodium and calcium ion channel populations have shown that the timescale of the recovery from inactivation following a long duration of membrane depolarization increased with the length of the depolarization period. We refer to this type of behavior as history-dependence. The origin of this history dependence is generally thought to result from the large inactivation state space hinted at by single channel patch clamp experiments. Previous modeling approaches, based on this idea, have already been suggested in the literature, but fall short in accurately reproducing this behavior. We model the slow inactivation of a channel as a continuous-time semi-Markov process consisting only of two states, the Markovian "available" state, and the non-Markovian "inactivated" state. The residence time probability density function (RTPDF) of the Markovian available state is exponential, while the inactivated state is non-Markovian, with a power-law RTPDF. Both RTPDFs are voltage dependent. These channel RTPDFs are measurable quantities. We reproduce for the first time, to our knowledge, the main experimental finding observed in the channel populations experiments, namely an exponential recovery process with a history-dependent timescale. Using these results we narrow down the options for the model parameters at different voltages, and explicitly address the issue of long memory phenomena. The model introduced here also provides many predictions. Qualitatively, we predict that spiking stimuli change the timescale of channel recovery from inactivation only negligibly, and that the rate of this recovery must be voltage dependent. Quantitatively, we derive an exact dynamic equation that fully defines an input-output relation between the membrane voltage and the channel availability, and solve it exactly in many important cases. Additionally, we develop expressions that describe all joint moments in the single channel and population. The potential contribution of this model goes beyond the specific system addressed in this work. Current models of channels and receptors tend to suffer from an embarrassment of riches. In order to explain behaviors over an ever expanding range of timescales, these complex models often include multiple inactivation states. Since the number of states and their parameters are not directly observable, these models tend to be highly specific and are likely to suffer from over-fitting. Furthermore, such models always have an upper bound on their timescale. In this work, we introduce and thoroughly analyze a type of model that does not suffer from these limitations. Despite its simplicity, it provides a generalization of previous models, is based only on measurable quantities, does not possess an upper bound on its timescale and exhibits considerable analytical tractability. As such, it stands as an appealing alternative to previous approaches. In particular, given the direct impact of channel dynamics on neuronal behavior, we expect that this experimentally well motivated, yet mathematically tractable model will form a sound foundation for realistic neuronal models, spanning multiple time scales and exhibiting history-dependence.

doi:

I-3. Fast Kalman filtering on quasilinear dendritic trees

Liam Paninski

Columbia University

LIAM@GATSBY.UCL.AC.UK

The problem of understanding dendritic computation remains a key open challenge in cellular and computational neuroscience. The major difficulty is in recording physiological signals (especially voltage) with sufficient spatiotemporal resolution on dendritic trees: multiple-electrode recordings from dendrites are quite technically challenging, and provide spatially-incomplete observations, while high-resolution imaging techniques provide more spatially-complete observations, but with significantly lower signal-to-noise. One avenue for extending the reach of these currently available methods is to develop statistical techniques for optimally combining, filtering, and deconvolving these noisy signals. State-space filtering methods are attractive here, since these methods allow us to quite transparently incorporate 1) realistic, spatially-complex multicompartmental models of dendritic dynamics and 2) time-varying, heterogeneous observations (e.g., spatially-scanned multiphoton imaging data) into our filtering equations. The problem is that the time-varying state vector in this problem — which includes, at least, the vector of voltages at every compartment — is very high-dimensional: realistic multicompartmental models often have on the order of $N \sim 10^4$ compartments. Standard implementations of state-space filter methods (e.g., the Kalman filter) require $O(N^3)$ time, and are therefore impractical for applications to large dendritic trees. However, we may take advantage of three special features of the dendritic filtering problem to construct efficient filtering methods. First, dendritic dynamics are governed by a cable equation on a tree, which may be solved using symmetric sparse matrix methods in $O(N)$ time. Second, current methods for imaging dendritic voltage provide low SNR observations, as discussed above. Finally, in typical experiments we record only a few image observations ($n < 100$ or so coarse pixels) at a time. Taken together, these special features allow us to approximate the Kalman equations in terms of a low-rank perturbation of the steady-state (zero-SNR) solution, which in turn may be obtained in $O(N)$ time using efficient matrix solving methods that exploit the sparse tree structure of the dynamics. The resulting methods provide a very good approximation to the exact Kalman solution, but only require $O(N)$ time and space. In addition, a number of extensions of the basic method are possible: for example, we can incorporate spatially blurred or scanned observations; temporally filtered observations and inhomogeneous noise sources on the tree; “quasi-active” resonant membrane dynamics; and even in some cases nonlinear observations of the membrane state. Simulation results using the resulting filter allow us to quantify exactly how much information we can expect to extract about dendritic dynamics from recordings at a given SNR.

doi:

I-4. Dendritic spine plasticity can stabilize synaptic weights

Cian O'Donnell

Matthew F. Nolan

Mark C. W. van Rossum

University of Edinburgh

CIAN.ODONNELL@ED.AC.UK

MATTNOLAN@ED.AC.UK

MVANROSS@INF.ED.AC.UK

Stabilization of synaptic weights is important for long-term memory. Most existing attempts to explain synaptic weight stability assume the existence of elaborate molecular signalling mechanisms. Here we propose a simpler alternative model in which changes in dendritic spines size following plasticity results in stabilization of synaptic strength by modifying local calcium dynamics. Recent experimental studies demonstrate that the size of dendritic spines is increased or decreased following induction of synaptic potentiation or depression respectively (Matsuzaki et al, 2004; Harvey et al, 2008). However, the consequences of altered spine size for signaling events within the spine are not clear. Using a biophysical computer model of a dendritic spine and a common calcium-dependent

plasticity rule, we find that different NMDAR conductance to spine-size relationships can result in stable, unstable or even bistable synaptic weight dynamics. When we use parameter estimates from the experimental literature, the model predicts that real spines fall into the 'stable' category. Our model is sufficient to explain the experimental observations that weak synapses are most susceptible to plasticity protocols and that large spines are the most persistent in vivo. We built reduced versions of our stable and unstable synapses and compared their behavior on a model integrate-and-fire neuron subject to physiological input patterns. We found that the stable synapse model, but not the unstable model, can lead to unimodal synaptic weight distributions similar to those found experimentally. To compare memory storage under the two conditions, we selectively potentiated a subset of synapses, subjected the neuron to ongoing activity and allowed the synapses to follow either the stable or unstable plasticity rules. The stable synapses always retained the memory for a longer period than the unstable synapses. In summary, we propose a biophysical model of how synaptic weights can be stabilized using only known properties of dendritic spine geometry and synaptic receptor distribution. We also investigated the implications of these learning rules for synaptic weight distributions and memory storage. This link can act both as a framework for interpreting experimental data and as a base for future theoretical studies of memory. - Matsuzaki M, Honkura N, Ellis-Davies GCR, and Kasai H. Structural basis of long-term potentiation in single dendritic spines. *Nature*, 429:761-6 (2004). - Harvey CD, Yasuda R, Zhong H and Svoboda K. The spread of Ras activity triggered by activation of a single dendritic spine. *Science*, 321: 136-140 (2008).

[doi:](#)

I-5. Model of synaptic plasticity based on self-organization of PSD-95 molecules in spiny dendrites.

Dmitry Tsigankov
Stephan Eule

DMITRY@NLD.DS.MPG.DE
EULE@NLD.DS.MPG.DE

Max-Planck Institute for Dynamics and Self-Organization

We present a model of stochastic molecular transport in spiny dendrites. In this model the molecules perform a random walk between the spines that trap the walkers. If the molecules interact with each other inside the spines the trapping time in each spine depends on the number of molecules in the respective trap. The corresponding mathematical problem has non-trivial solutions even in the absence of external disorder due to self-organization phenomenon. We obtain the stationary distributions of the number of walkers in the traps for different kinds of on-site interactions between the walkers. We analyze how birth and death processes of the random walkers affect these distributions. We apply this model to describe the dynamics of the PSD-95 proteins in spiny dendrites. PSD-95 is the most abundant molecule in the post-synaptic density (PSD) located in the spines. It is observed that these molecules have high turnover rates and that neighboring spines are constantly exchanging individual molecules. We propose that the geometry of individual PSD-95 clusters determines the dependence of trapping times on the number of molecules inside the trap and thus can vary from spine to spine. Furthermore, we suggest that activity-dependent reorganization of the PSD changes the geometry of PSD-95 cluster and thus can lead to synaptic plasticity in a form of long-term potentiation (LTP). In the model this is achieved by spine specific activity-dependent ubiquitination of PSD-95 molecules, which transiently reduces the amount of PSD-95 in the spine but change the geometry of the PSD-95 cluster in such a way that self-organization process results in the overall increase of the number of PSD-95 molecules associated with LTP. We also show that such a dynamics of the PSD-95 molecules can set up the conditions for anomalous diffusive transport inside spiny dendrites and predict the distribution of the PSD sizes which has features of both exponential and Poisson distributions.

[doi:](#)

I-6. Trajectory prediction combining forward models and historical knowledge

Jill O'Reilly
Tim Behrens

JOIRELLY@FMRIB.OX.AC.UK
BEHRENS@FMRIB.OX.AC.UK

FMRIB Centre, University of Oxford

To interact effectively with the environment, brains need to extract predictive information from their environment. Prediction can involve different computational processes. In motor control, forward models are important for extrapolating the results of motor commands. In other situations, (e.g. foraging) the goal may be to extract the statistical properties of the environment, a process well modelled by reinforcement learning algorithms. How are these computationally different types of prediction, which may be associated with distinct brain structures, combined to optimise behaviour? We designed a paradigm in which participants must combine historical, probabilistic knowledge with online forward modelling to resolve uncertainty in prediction. In this paradigm, participants must predict the endpoint of a visually observed trajectory (the flight-path of a 'space invader') - a forward modelling/extrapolation task. However, probabilistic historical information is also informative - the trajectory-endpoints follow a Gaussian distribution so that the 'space invaders' are most likely to land in one part of the screen. The contributions of forward modelling (extrapolation of the current trajectory) and historical/probabilistic knowledge depend on how informative each is. We manipulated the informativeness of trajectory data (by adding white noise to the trajectory seen by the participant) and the historical distribution (by changing its variance) over the course of the experiment. When the observed trajectory was noisy, participants chose endpoints closer to the historical mean - i.e. they used historical knowledge to resolve uncertain predictions from a forward model. In Bayesian terms, on each trial the participant has both a PRIOR expectation of the endpoint (from historical knowledge), and observed DATA (from forward modelling of the current trajectory). These are combined to give a POSTERIOR distribution from which we hypothesise the prediction about the endpoint is made. Using data from 19 human participants, we found that responses were indeed distributed about the posterior mean. We modelled the endpoints selected by the humans as having a Gaussian distribution about either the posterior mean, prior mean or trajectory endpoints. Using maximum-likelihood estimates for the standard deviation of each distribution, the log likelihood of the posterior model was significantly higher than for the trajectory-only model (Bayes factor 35) or the prior-only model (BF:130). This indicated that human participants really do combine forward-modelling and historical knowledge in a perceptual prediction task. The historical prior must be learnt over many trials. We modelled this process using a Bayesian 'computer participant' which estimated the mean and variance of the historical distribution at each trial t plus the parameters of a transitional Beta distribution, through which the prior for $t+1$ was generated. Estimates of the historical distribution taken from the Bayesian learner predicted human responses better than a baseline model which assumed knowledge of the historical distribution without accounting for learning. We are collecting fMRI data using this paradigm. By modelling the contributions of prior knowledge and forward modelling to endpoint prediction on each trial, we hope to find out which brain structures are involved in each computational mode of prediction, and where the two types of information are combined.

[doi:](#)

I-7. Dynamics of fronto-parietal synchrony during working memory.

Nicholas M. Dotson
Rodrigo F. Salazar
Charles M. Gray
Montana State University

NDOTSON@CNS.MONTANA.EDU
RSALAZAR@CNS.MONTANA.EDU
CMGRAY@NERVANA.MONTANA.EDU

The fronto-parietal cortical network plays a key role in controlling working memory. These processes are thought to involve the temporal coordination of activity within and between areas of the prefrontal (PFC) and posterior parietal (PPC) cortical regions. Previously, we reported task-dependent fronto-parietal coherence, in the beta (12-25 Hz) and gamma range (26-40 Hz) frequency range of the local field potential (LFP). Here we characterize

the time-dependent dynamics of these interactions across single trials. We recorded the LFP from 2-6 sites simultaneously in both the PFC and PPC of a macaque monkey performing a rule-based, delayed match-to-sample task in which the monkey was required to remember either the location or the identity of the sample object. The LFP was bandpass filtered (10-50 Hz) and a moving window correlation analysis (200 ms window, ± 50 ms time lag) was performed on all pairs of signals on each trial. The peak correlation coefficients were compared to shuffled surrogate distributions from the same data ($p < .05$). Significant correlation coefficients, and their time lags, were evaluated with respect to their probability and variance across trials. During stable performance ($>80\%$ correct responses), the probability and magnitude of correlation within and between the PFC and PPC was transiently suppressed during the sample stimulus, increased monotonically during the delay period (slope of the linear regression different from zero at $p < .05$), and dropped again following the match stimulus. These changes were accompanied by a decrease in the variance of the correlation phase lag during the delay that reached a minimum at the time of the earliest match onset. These effects occurred in 31.8/30.6% ($n=85$) of PFC pairs, 64.7/60.8% ($n=102$) of PPC pairs, and 48.6/45.7% ($n=243$) of PFC-PPC pairs during the location/identity rule. In a subset of the pairs, correlation probability was observed to differ across the set of 9 sample stimuli during the delay period. These results demonstrate a time-dependent increase in the magnitude and precision of LFP synchrony within and between the PFC and PPC that is task and stimulus specific and peaks just prior to the match stimulus. In conclusion, working memory processes modulate the probability of neuronal synchrony.

[doi:](#)

I-8. Bayesian optimal use of visual feature cues in visual attention

Benjamin Vincent
University of Dundee

B.T.VINCENT@DUNDEE.AC.UK

How do we utilise visual cues in order to guide attention? While popular accounts (feature integration theory, guided search, and low-level salience) vary in a number of ways, they all assume a max-of-sensory-outputs selection mechanism. Such a max-observer can be expected to be optimal in contrived situations where targets and distracters have equal variability along some sensory dimension such as luminance, contrast or orientation. However it may be overoptimistic to assume that this equal variability is always the case in natural visual environments. So the question arises, is the attention system suboptimal by making the assumption of equal variance about target and distracter features? If so, then current popular models are adequate and notions of optimality in attentional phenomena would be in doubt. If not, and human performance is close to Bayesian optimal, then existing models would be challenged. To test this, a simple psychophysical present/absent detection task was used, with Gabor target and distracter stimuli. By adding additional orientation uncertainty to distracters, the equal variance assumption is violated. Will subjects perform optimally and exceed the performance of a max-observer? Experiments resulted in receiver-operating-characteristic (ROC) curves for individuals and were compared to predictions of the max-observer and a Bayesian optimal observer. Predictions were obtained through Monte Carlo simulations. A side-by-side comparison of models is fair, as both have a single parameter (which is estimated from the data) corresponding to degree of internal orientation uncertainty. The max-observer has a decision criteria based upon a sensory dimension, while the Bayesian optimal observer has a decision criteria based on the posterior probability of target presence. Under conditions which violate the strict equal variance assumptions, it was found that human performance was not in fact suboptimal. Performance at the task greatly exceeded the highest possible performance obtainable by the max-observer, providing strong evidence against the use of a max-of-sensory-outputs selection mechanism. The fact that the decision criteria is based upon a sensory dimension greatly limits the possible performance of a max-observer. In contrast, the Bayesian optimal observer provided very good fits to human performance (and ROC curve) data. Human performance at this task is not explicable by a max-of-sensory-outputs selection mechanism, which is problematic for feature integration theory, guided search, and low-level salience. These results provide a strong proof-of-concept, in a controlled psychophysical setting, that visual features are evaluated for the posterior probability of target presence. Thus this work adds to the small but growing research that suggests attentional phenomena are by-products of near optimal inference processes.

doi:

I-9. Neural correlates of spatial short-term memory in the rodent frontal orienting field

Jeffrey C. Erlich
Max Bialek
Carlos D. Brody
HHMI & Princeton University

JERLICH@PRINCETON.EDU
MBIALEK@GMAIL.COM
BRODY@PRINCETON.EDU

We have trained rats on a memory-guided delayed orienting task, inspired by delayed saccade tasks performed by human and non-human primates. The delayed orienting task requires rats to 'fixate' by holding their nose in a central nose port while a light is on in the port. During the fixation, a brief sound (the "target stimulus") indicates to the rat whether it should orient to the right or to the left after the end of fixation in order to obtain a reward. The end of fixation is indicated by the offset of the center light. We randomly interleave two types of trials: memory and non-memory. In memory trials, the target stimulus is presented shortly after the beginning of fixation and 700 ms before its end. In these trials, the rat must inhibit the impulse to respond immediately and must also remember the reward location (either as a retrospective memory of the stimulus or a prospective motor plan/prediction of reward location). These two elements, inhibition and planning, make the delayed-orienting task suitable for understanding the neural basis of cognitive control. In non-memory trials, the target stimulus is presented at the end of the fixation period, and no memory is required for correct performance. The two types of trials allow us to distinguish activity associated with spatial memory from activity associated with fixation only. The electrophysiological properties of neurons in the rodent orientation system remain largely unexplored. One cortical area of potential interest is the Frontal Orienting Field (FOF, +2 AP, ± 1.3 ML [mm from Bregma]). Previous anatomical data has suggested that the FOF is homologous to primate frontal eye fields: in particular, the FOF connects strongly to both pre-frontal cortex and the superior colliculus. Nevertheless, almost no electrophysiological recordings from the FOF exist. To our knowledge, no previous studies have recorded from FOF using repeated, identically prepared trials requiring an orientation movement, whether memory or non-memory-based. We have now recorded from over 500 single-units from the FOF while rats performed the delayed orienting task. We found that a significant portion of neurons in this region (approx 30%) show spatial working memory: after the offset of the stimulus, and before the onset of the movement, firing rates are selective for the rat's upcoming choice. We have preliminary lesion and pharmacological evidence indicating that unilateral disruption of activity in the FOF creates a contralateral impairment. For example, rats with left FOF lesions are impaired at making rightward responses. Thus, we have both correlational physiological evidence and causal evidence that the FOF is an important element of the rat's neural circuit for spatial orientation and spatial short-term memory.

doi:

I-10. When to recall a memory? Epoch dependent memory trace with a power law of timescales in ACC neurons.

Alberto Bernacchia
Hyojung Seo
Daeyeol Lee
Xiao-Jing Wang
Department of Neurobiology, Yale University

ALBERTO.BERNACCHIA@YALE.EDU
HYOJUNG.SEO@YALE.EDU
DAEYEO.LEE@YALE.EDU
XJWANG@YALE.EDU

To control behavior, it is useful to recall appropriate past events at specific moments (epochs), and the memory of a past event may temporarily fade out and revive later. In reward-seeking tasks, choice behavior depends

on the memory of past rewards combined with current sensory information. The amount of reward obtained previously from a particular action or state is often a good predictor of future reward. Accordingly, the value (or reward expectation) can be defined as a weighted sum of past rewards, where the contribution of each reward to the value decays exponentially in time. While the neuronal activation in the cortex and basal ganglia has been suggested to encode the value, electrophysiological studies have shown that the reward-related signal of single neurons (memory trace) does not display simple exponential decay; instead it fades away and reactivates at different epochs in successive trials. In addition, the same cortical areas encode the trial epoch in a variety of tasks, which we call the epoch code. We investigated the dynamics of memory traces by asking how the codes for memory and epoch are combined. We recorded the activity of single neurons in the anterior cingulate cortex (ACC) of monkeys performing a matching pennies task: in each trial the monkey was required to shift its gaze towards one of two targets, where the correct (rewarded) one is selected by a computer programmed to simulate a rational player. We found that the apparently complex dynamics of the memory trace can be reduced to an exponential decay (or the sum of two exponentials) rescaled by the epoch code, and the majority of recorded neurons implements this factorization. Furthermore, we found that the distribution of memory timescales across neurons follows a power law, while the distribution of response amplitudes is exponential. The theory of tensor products predicts that if a population of neurons encode two variables by the product of two complete sets of tuning curves, a downstream neuron can in principle represent any function of those two variables. We propose to extend this notion to the temporal domain: the two variables are the current trial epoch and the sequence of past rewards, and a downstream neuron would in principle be able to recall the appropriate memory at a specific epoch, or any arbitrary memory-epoch association. Recent evidence in support of our hypothesis is that the epoch code seems to be characterized by a complete set of tuning curves in prefrontal neurons. The power law distribution of memory timescales implies that neurons track the value on a wide range of timescales, making it possible for ACC to flexibly select the appropriate timescale depending on the task at hand. Furthermore, we show that the power law distribution of timescales is consistent with a network model operating near a critical regime, suggesting that ACC performs computations near the edge of chaos. Finally, the exponential distribution of response amplitudes is consistent with a normal matrix governing the interactions in the network model, providing hints about the underlying circuit architecture.

doi:

I-11. Neural encoding of decision uncertainty in prefrontal cortex

R. James Cotton
Allison Laudano
Andreas S. Tolias

Baylor College of Medicine

RCOTTON@CNS.BCM.EDU
ALAUDANO@CNS.BCM.EDU
ATOLIAS@CNS.BCM.EDU

Uncertainty is ubiquitous in perception and decision making and makes inference particularly difficult (e.g. object categorization). Therefore, understanding how the brain represents and computes with uncertain information is a fundamental quest in neuroscience. Sensory data come with bottom-up uncertainty since they are corrupted by noise. Decisions are typically made by combining bottom-up sensory information with prior or top-down knowledge learned from the past. To achieve statistically optimal (or Bayes-optimal) behavior it is necessary for the brain to represent all these sources of uncertainty and combine them appropriately. Behavioral studies have demonstrated that humans perform close to Bayes-optimally when combining multiple sensory cues about an underlying stimulus, supporting the hypothesis that the brain represents and computes with probability distributions. However, the neural mechanisms used by the brain to compute under uncertainty remain elusive. We studied if top-down uncertainty was encoded in the prefrontal cortex. We have trained a monkey in a probabilistic classification task. The orientation of a stimulus (drifting grating) was drawn from one of two overlapping probability distributions. The animal had to infer from which class the stimulus came, although this could often not be determined with certainty. We recorded single-unit activity from the PFC using tetrodes while the monkey performed this task. We found that neurons not only encoded the class decision of the animal but PFC activity was also correlated with the posterior probability of class (i.e. class certainty). Specifically, neurons that preferred class A increased their firing when the probability of class A given the stimulus orientation was higher. PFC neurons also

exhibited this effect during the delay period when the animal was required to remember its decision in the absence of the stimulus. Thus, we demonstrate that the PFC, a brain area thought to be at the peak of the decision-making hierarchy, encodes relevant probabilities associated with the inference problem the brain is solving.

doi:

I-12. Clocking perceptual processing speed: from chance to 75% correct in 30 milliseconds

Terrence R. Stanford
Swetha Shankar
Dino Massoglia
Gabriela Costello
Emilio Salinas

TSTANFORD@WFUBMC.EDU
SSHANKAR@WFUBMC.EDU
DMASSOGL@WFUBMC.EDU
GCOSTELL@WFUBMC.EDU
ESALINAS@WFUBMC.EDU

Wake Forest University School of Medicine

The neurobiology of choice behavior has been intensely studied with tasks in which a subject makes a perceptual judgement and indicates the result with a motor action. Combined psychophysical and neurophysiological experiments have thus characterized perceptual decision-making capacity as a function of signal quality, strength, and subjective value. However, perceptual choices can be difficult for two fundamentally different reasons: (1) because the relevant sensory signal is weak or uninformative (try driving under heavy fog), or (2) because even though the signal is strong, it cannot be processed fast enough (try returning a 130 mph tennis serve). This issue (timing) has been harder to tackle: how long does it take to make a perceptual judgement, versus executing a motor action to report it? When does a subject commit to a particular choice, and what neural mechanisms determine that time point? A major limitation is that reaction times (RTs) are affected by numerous sensory and motor factors, such as motivation, task difficulty and speed-accuracy trade-offs. To minimize such confounds, we designed a two-alternative forced-choice task in which the relevant sensory information is given after the signal to initiate a saccade. In this compelled-saccade task, performance varies between chance and 100% correct, but motor execution and mean RT change little because, in each trial, the relevant sensory information influences a saccadic choice that is already ongoing. This design allows us to construct a new curve - the tachometric curve - that reveals a subject's perceptual discrimination capacity with little contamination from other, non-perceptual processes. In particular, the slope of the tachometric curve is a direct psychophysical manifestation of a subject's perceptual processing speed. With this technique we find that monkeys can make accurate color discriminations in less than 30 ms. This result, and the tachometric curve itself, should depend only on the perceptual difficulty of the task and the perceptual capacity of the subject, but not on other task contingencies. And indeed, additional experiments indicate that the tachometric curve is highly insensitive to variations in motor behavior, even very large ones. All these psychophysical results are accurately replicated by a race-to-threshold model with two variables, which represent the two possible oculomotor plans that may develop in each trial. The model correctly predicts a wide variety of shapes for the RT distributions observed during the task for correct and error trials. Our approach also provides a novel tool for elucidating how neuronal activity relates to sensory versus motor processing, as demonstrated with data from neurons in the Frontal Eye Field. In these cells, perceptual information acts by accelerating and decelerating the ongoing motor plans associated with correct and incorrect choices, as predicted by the race-to-threshold model, and the time course of these neural events parallels the time course of the subject's choice accuracy. In conclusion, the compelled-response design directly reveals how a subject's perceptual performance unfolds in time, and by providing a new tool for correlating the time courses of psychophysical and neuronal responses, it opens up a new avenue for investigating choice behavior.

doi:

I-13. The effect of time pressure on decision making

Shinichiro Kira¹
Michael N. Shadlen²

SHINKIRA@U.WASHINGTON.EDU
SHADLEN@U.WASHINGTON.EDU

¹Dept. Physiology & Biophysics, NPRC; U. WA

²HHMI, Dept. Physiol & Biophys, NPRC, U. WA

The accuracy of many decisions can be improved by acquiring additional information, albeit at the cost of time. When a human or monkey is asked to decide the direction of motion in a noisy display, accuracy improves with viewing duration in a manner that is consistent with perfect integration of independent samples of evidence - at least up to a point. One idea is that decisions end when the accumulated evidence reaches a threshold level, or bound. This idea explains the tradeoff between speed and accuracy of many decisions. However, it does not explain how a decision maker would incorporate a time pressure (or cost) that changes dynamically. We hypothesized that, under time pressure, subjects adjust their criterion for terminating decisions in a time-dependent manner. To test this we manipulated the cost of elapsed time during a simple perceptual decision. Human subjects discriminated the direction of dynamic random dot motion in a choice-reaction time paradigm. They indicated their decisions by making an eye movement whenever ready. Two directions and 6 motion strengths were randomly interleaved. Subjects gained and lost points for correct and incorrect choices, respectively. After each trial they received feedback about their total score and their average rate (points per minute). After 16-20 sessions, we introduced a manipulation that was unknown to the subject: a fraction of trials were terminated prematurely by the experimenter, as if the computer had sensed a fixation error. The manipulation caused subjects to make faster, less accurate decisions. A drift-diffusion model with stationary bounds explained the choices and mean reaction time (RT) for correct choices. To explain the RT on incorrect choices and the shape of the RT distributions, we incorporated a time-dependent stopping rule - that is, collapsing bounds. The rate of this collapse was significantly faster when subjects experienced the premature terminations of trials. These results underscore a flexible termination strategy for decisions. This flexibility is influenced by the subjective cost of decision time. Acknowledgements: We thank R. Kiani for software and advice. This work was supported by HHMI, EY011378, DA022780, and RR000166. S. Kira was supported by a Nakajima Foundation Predoctoral Fellowship.

[doi:](#)

I-14. Decision-related activity in area V2 for a fine disparity discrimination task.

Hendrikje Nienborg¹
Bruce G. Cumming²

NIENBORG@SALK.EDU
BGC@LSR.NEI.NIH.GOV

¹Salk Institute

²National Eye Institute, NIH

In humans and monkeys, judgments of small differences in stereoscopic depth rely on relative binocular disparity, i.e. the comparison of differences in nearby disparities in the stimulus. The first stage in the primate cortex that contains a subgroup of neurons selective for relative disparity is visual area V2. We have previously shown that the activity of disparity selective neurons in V2 is correlated with a macaque monkey's perceptual decision in a 'coarse' disparity discrimination task that relies on judgments absolute, not relative disparity. This suggests a link between the activity of disparity selective neurons and the perception of absolute disparity. Here, we ask whether disparity selective neurons are also correlated with a monkey's judgment in a 'fine' disparity discrimination task relying on relative disparity. Two monkeys were trained to perform a 'fine' disparity discrimination task. They were presented with a circular random-dot-pattern consisting of a central circle and a surrounding annulus. The animals' task was to determine whether the center was protruding or receding relative to the surround. We presented the dots of both center and surround at 100% inter-ocular correlation. The 'fine' disparity discrimination task relies on relative disparity judgments, i.e. comparing the disparity of the center relative to the surround. Upon mastering this task, the monkeys were also trained on the 'coarse' disparity discrimination task. In each experiment, we

kept the surround at a fixed disparity, but between experiments we adjusted the surround disparity depending on the disparity tuning of each neuron. We recorded the activity of 80 disparity selective neurons in V2 of two macaques while they made fine disparity judgments. We quantified the correlation between the neuronal activity and the perceptual judgment as 'choice-probability' (Britten et al. 1996). The mean choice-probability was 0.55, which was significantly larger than 0.5 ($p < 0.001$), indicating that V2 neurons carry decision-related activity in this task. This effect was significant in each monkey (mean 0.55 in both monkeys, $p < 0.05$ and $p < 0.001$, respectively). Our result contrasts with reports from area MT, for which choice-probabilities were found for the 'coarse' (Uka, DeAngelis, 2004), but not for the 'fine' disparity discrimination task (Uka, DeAngelis, 2003). Interestingly, decision related activity was unrelated to the degree to which a neuron was selective for relative disparity ($r = 0.08$, $p = 0.81$, $n = 74$). We next examined the relationship of decision-related activity between the two disparity-based tasks, the 'coarse' task (data from Nienborg, Cumming, 2006, 2007, 2009), and the 'fine' task. We observed similar distributions of decision-related activity for the neurons measured in the two tasks and a significant correlation on a cell-by-cell basis of the size of decision-related activity between the two tasks ($r = 0.41$, $p < 0.01$, $n = 41$). Previous theoretical and empirical studies (Shadlen et al. 1996; Cohen, Newsome, 2009) found that decision-related activity depends critically on the structure of the inter-neuronal correlation matrix of the sensory neurons used in a perceptual task. It has recently been shown that inter-neuronal correlation changes dynamically between different tasks (Cohen, Newsome, 2008). Our results suggest that inter-neuronal correlation across these two different disparity-based tasks was surprisingly stable.

[doi:](#)

I-15. Changes in functional connectivity in LIP during a free choice task

Annegret Falkner¹

Michael E. Goldberg²

ALF2111@COLUMBIA.EDU

MEG2008@COLUMBIA.EDU

¹Columbia University

²Dept. of Neuroscience, Columbia University

Activity in the monkey lateral intraparietal area (LIP) encodes the relative salience of locations in visual space. When multiple stimuli compete for attentional priority, a single winner must emerge on this map to be used as the upcoming saccade target, though it is unclear how the "winning" process is functionally implemented. Peaks of activity in LIP could rise and fall independently and a winner could emerge as one peak hits a designated threshold. Alternatively, peaks of activity could actively compete via mutually inhibitory interactions, such that a winning peak actively suppresses the response to a competitor. We tested these 2 hypotheses by recording from 2 LIP cells simultaneously during a free choice task where the monkey was required to choose between one of two saccade targets presented under varying reward probability. We isolated cells on separate electrodes and insured that a saccade target presented in the receptive field (RF) of one cell did not excite the other. During the task, we recorded from the locations of both targets simultaneously and looked at trial to trial correlated noise during the task at separate epochs during the saccadic decision, keeping both the reward contingency and saccade direction constant. In the epoch prior to the onset of the targets, noise in LIP was positively correlated across trials in many pairs of neurons. In the epoch prior to the saccade, the noise was negatively correlated in the population. Moreover, stronger negative noise correlations were associated with longer saccade latencies without concurrent changes in firing rates, suggesting that suppressive mechanisms may be more strongly engaged during periods of indecision. Since correlated noise can indicate either a common input or mutual connectivity between cells, negative noise correlations in LIP suggest a competitive inhibitory mechanism in LIP between response peaks.

[doi:](#)

I-16. Role of secondary motor cortex in withholding impulsive action: Inhibition or competition?

Ana Fonseca
Masayoshi Murakami
Maria Vicente
Gil Costa
Zachary Mainen

ARSECCA@GMAIL.COM
MASAYOSHI.CSHL@GMAIL.COM
MINESVICENTE@GMAIL.COM
COSTA.GIL@GMAIL.COM
ZACH.IGC@GMAIL.COM

Champalimaud Neuroscience Program at IGC

In a now-famous set of experiments, Mischel et al. (1968), tested self-control in 4-year-old children by giving each child one marshmallow. The child was told she could either eat the marshmallow now or, if she waited while he stepped out for a few minutes, she could have two when he returned. Surprisingly, the ability to wait for the delayed reward in this task turned out to be a good predictor for academic and social success years later. This and many other studies helped to lead to the current idea that impulse control, as measured in delayed gratification tasks, represents a fundamental cognitive function. According with one view, impulse control results from a general inhibitory mechanism, thought to be localized in the frontal cortex, but whose mechanisms are not well understood. Curiously, however, Mischel et al. observed that children with successful impulse control did not wait passively, but distracted themselves by singing or playing with their hands. Could it be that impulse control arises not from a general inhibitory mechanism but from the ability to generate alternative, competing actions? To investigate these issues, we are studying analogous impulse control behavior in rats. In the first task variant, rats were required to wait at a nose poke. A first tone was presented at a fixed short delay (0.4 s) and a second one at a longer exponentially-distributed delay (~2 s mean). Responses between the two tones received a small reward, while responses after the second tone received a larger reward. Video analysis showed that during waiting subjects engaged in various simple motor actions such as grasping or chewing the waiting port. To study the underlying neural bases, we made single-unit recordings from medial prefrontal cortex (mPFC), an area associated with inhibitory control, and secondary motor cortex (M2), an area involved in motor planning. We tested for the ability of single neurons to predict the amount of time the subject would wait before responding. We found that 20%(109/548) of M2 neurons and only 7%(8/122) of mPFC neurons showed predictive activity, a significantly smaller fraction ($P < 0.001$, χ2 test). Next, we tested whether waiting-predictive neurons were general to waiting or specific to the action involved in waiting. To do this, we required subjects to alternate in blocks of trials between waiting at a nose poke and pushing a lever. We found that 43/171 M2 neurons were predictive of nose poke waiting time, but only 6 of these were also predictive of lever press waiting time. Neurons predictive in both tasks typically showed distinct temporal activation profiles. These results suggest that the ability of firing of motor cortex neurons to predict waiting time appears to result not only from an involvement in general inhibition, but due to linkage with the planning or initiation of specific actions that occur during waiting. These findings suggest that the ability to generate "self-distractions" - alternative actions that compete with the pursuit or consumption of a tempting reward - may be a general feature of impulse control across species.

[doi:](#)

I-17. Dorsomedial prefrontal cortex encodes value information during a sequential choice task

Chung-Hay Luk
Jonathan D. Wallis
Helen Wills Neuroscience Institute, UCB

CH_LUK@BERKELEY.EDU
WALLIS@BERKELEY.EDU

Many choices require evaluating possible options one after another as exemplified in wine tasting. The underlying neuronal mechanisms in such a sequential choice paradigm are largely unknown though. Hence we have recorded neuronal and local field potential activity from dorsomedial and dorsolateral prefrontal cortex (PFdm and

PFdl, respectively) as subjects performed a sequential choice task. We expected PFdm to encode the valuation of choice options, owing to its strong anatomical inputs from areas processing reward. In our task, two monkeys (*Macaca mulatta*) chose between two different juices on a trial-by-trial basis. During the sampling phase, the subject made two sample responses separated by delays, each of which resulted in the delivery of a small drop of one of three juices (apple, orange or quinine). During the choice phase, the subject then chose to repeat one of the responses, and received a larger amount of the juice that had been associated with that response earlier in the trial. Thus, in order to receive juices that are more preferable at the choice phase of the task, the subject had to maintain information about the first sampled reward and which response produced it to compare that reward to the subsequent reward. We recorded the activity of 112 PFdm neurons and 172 PFdl neurons from 180 recording sites as the subjects performed the task. Following the sampling of the first juice, a similar proportion of neurons encoded the action producing the reward in PFdm (46%) and PFdl (48%), whereas encoding of the juice reward was prominent in PFdm (60%) but not PFdl (28%). The neuronal activity correlated with high power around the 40 Hz gamma range. Moreover, reward-selective neurons showed a monotonic relationship between their firing rate and the subject's preference for the juice, suggesting that PFdm neurons encoded the juice as a value signal. PFdm neurons encoded the value of the second juice relative to the first, typically showing a higher firing rate when the second juice was less preferred than the first. These findings suggest that options in a sequential choice are evaluated with respect to previous options. By maintaining the value of the first juice and then encoding the value of the second juice relative to the first, PFdm neurons provide the appropriate information to enable the subjects to make their choice.

doi:

I-18. Striatal activity consistent with model-based, rather than model-free prediction errors

C. Shawn Green¹
 Peng Zhang¹
 Nathaniel Daw²
 Daniel Kersten¹
 Sheng He¹
 Paul Schrater¹

CSGREEN@UMN.EDU
 ZHANG870@UMN.EDU
 DAW@CNS.NYU.EDU
 KERSTEN@UMN.EDU
 SHENG@UMN.EDU
 SCHRATER@UMN.EDU

¹University of Minnesota

²New York University

Over the past several decades there has been considerable interest in uncovering the neural underpinnings of human decision-making. One particularly fruitful vein of research has focused on similarities between processing done in certain brain areas and computations that must be performed by model-free reinforcement learning algorithms. Of specific note is the repeatedly replicated correlation seen between activity in the ventral striatum and the size of trial-by-trial model-free reinforcement learning prediction errors. Striatal activity is high when the difference between the expected value of reward and the actual reward is large, while striatal activity is at a minimum when the amount of reward received is very close to the predicted value. Other brain areas such as the ventromedial prefrontal cortex, anterior cingulate, and amygdala have been similarly linked to model-free reinforcement learning parameters. In contrast, we have recently shown behaviorally that human choices are more consistent with a model-based, rather than a model-free learning algorithm. To demonstrate this, several versions of the standard sequential binary choice task were employed. Each of the versions contained contextual cues suggestive of a certain generative process for outcomes (i.e. D whether the outcomes were temporally independent, coupled, etc); importantly however, the actual generative process (and thus the observed outcome statistics) was identical in all versions. While a model-free system would be insensitive to such contextual cues (as such algorithms simply compute the expected values of states), human behavior was instead altered by these cues in a manner that was well predicted given the set of beliefs about the generative process promoted by the contextual cues. Our question is whether the activity in those brain regions associated with prediction error would be modulated by contextual manipulations. If these areas are truly calculating model-free values, they should be

insensitive to irrelevant contextual cues as the actual observed outcome statistics are identical across conditions. If, on the other hand, activity depends on beliefs about the process generating those statistics, predictable differences should be observed. The data favored the latter hypothesis. In particular, when contextual cues were provided that were inconsistent with the true generative process, activity in the ventral striatum was correlated with model-free prediction error values. However, when contextual cues were provided that were consistent with the true generative process, the correlation between prediction errors and ventral striatal activity disappeared. We interpret these findings as being consistent with the notion that reward computations in the ventral striatum reflect the predictions of an internal generative model of the task.

[doi:](#)

I-19. From integrate-and-fire neurons to Generalized Linear models

Srdjan Ostojic¹
Nicolas Brunel^{2,3}

SO2310@COLUMBIA.EDU
NICOLAS.BRUNEL@UNIV-PARIS5.FR

¹Center for Theoretical Neuroscience, Columbia University

²CNRS UMR 8119

³Universite Paris Descartes

In the recent years generalized linear models (GLMs) have become a popular way of describing neural activity elicited by a time-varying input. In these models, the output spike train is obtained by processing the input through a so-called Linear-Nonlinear-Poisson (LNP) cascade, which consists of three consecutive stages: (i) a linear temporal filter is applied to the input; (ii) the outcome is transformed non-linearly to obtain a firing rate; (iii) a Poisson process generates spikes randomly with the prescribed instantaneous rate. Such a decomposition in three sequential processing stages is mathematically appealing, however from a biophysical perspective it seems difficult to identify three distinct mechanism that would correspond to the three stages. An alternative model is therefore to model neural data using integrate-and-fire (IF) models which incorporate some essential biophysical mechanisms. In this poster, we examine the relationship between IF and GLM models by representing a pool of integrate-and-fire neurons as a LNP cascade. To this end, we exploit known analytic results for IF models, that we complement with numerical simulations. More specifically, we consider a pool of uncoupled IF neurons receiving a time-varying input that is identical across all neurons, and which we call the signal. In addition, every neuron receives an independent white-noise input which corresponds to background activity of the surrounding network. We represent the relationship between the signal and the time-varying output firing rate by an LNP cascade. We first compute analytically the linear temporal filter by linearizing the firing dynamics around the baseline activity set by background noise. We show that this filter can be approximated by a sum of exponentials. We then use this filter to determine the static non-linear transformation. This transformation can be determined analytically in two limits: (i) for input variations of small amplitude (ii) for inputs varying slowly. We use numerical simulations to examine the non-linear transformation for the general case. Finally we show analytically that spike generation in an IF neuron can be described either as a refractory Poisson process or as bursting Poisson process, depending on the amplitude of background noise. The corresponding ISI distribution can be accurately approximated by a sum of two exponentials. We conclude that the output of a pool of IF neurons can be described to a reasonable level of accuracy by an LNP cascade. However the precise elements of this cascade depend on the background noise received by the neurons. Our analysis thus points out that a GLM model provides a local approximation of full non-linear dynamics, but different GLM models are necessary to describe different "working points" of a fully non-linear neural network.

[doi:](#)

I-20. Excitatory-inhibitory correlations result from opposing correlating and anticorrelating forces

Cyrus Omar¹
 Jason W. Middleton²
 Daniel J. Simons²
 Brent Doiron³

CYRUS@CMU.EDU
 JMIDDLET@PITT.EDU
 CORTEX@PITT.EDU
 BDOIRON@PITT.EDU

¹Center for the Neural Basis of Cognition, Carnegie Mellon University

²Neurobiology, University of Pittsburgh

³Mathematics, University of Pittsburgh

Correlated variability is thought to facilitate numerous cognitive and computational functions in the brain. However, how sensory stimuli and neural circuitry shapes trial-to-trial correlated variability remains poorly understood. Here we study, using the whisker-barrel system in rodents, pairs of Layer 2/3 regular spike (RS) and fast spike (FS) cells, putative excitatory (E) and inhibitory (I) cells, respectively. In the barrel cortex, the balance between E and I activity has been shown to play an important role in the trial-averaged response of feedforward inhibitory circuits. We discuss how these same circuits shape trial-to-trial covariability of E-I activity. We observe, in lightly sedated rats, that E and I activity is positively correlated during the spontaneous network state and that this correlation extends over long timescales. Using a linear response framework we show that the correlation between E and I cells can be predicted from the simultaneously recorded LFP. This suggests that excitatory and inhibitory elements of layer 2/3 populations are slowly correlated by a global synchronizing field under spontaneous conditions. In response to a single whisker deflection, we observe a decrease in trial-to-trial correlation of E-I activity, while the firing rates of the two cells increase significantly. This is inconsistent with previous results showing an increase in output correlation accompanies an increase in firing rate for a pair of uncoupled neurons receiving correlated input [1]. In order to understand this stimulus-induced decorrelation, we incorporate coupling in a simple stochastic mean-field model of E and I populations. We find the following features to be necessary to explain our data: 1. Strongly correlated background activity, acting as a source of positive spontaneous correlations 2. Sufficiently strong feedforward inhibition, which acts to anti-correlate E and I activity and competes with the correlating background field 3. A non-linear input-to-firing rate transfer function, which diminishes the anti-correlating effect of feedforward inhibition during spontaneous conditions. A whisker deflection transiently shifts network dynamics towards the linear region of the inhibitory population's transfer function, unlocking the full anti-correlating effects of the feedforward inhibitory circuit and leading to the observed sensory-evoked decorrelation between E and I activity. This model provides a simple mechanism by which E-I activity can become de-correlated in a stimulus-dependent manner despite forming a strongly coupled network subject to highly correlated background input. Our model makes predictions about the source of heterogeneity in Layer 2/3 circuitry that would lead to the observed broad distribution of pairwise correlations. In addition, extensions of our model explain the absence of a stimulus-induced increase in correlation between the activities of excitatory cells. The role of inhibitory interneurons in maintaining relatively decorrelated firing activity in response to sensory stimuli has significant implications for population based coding schemes. de la Rocha, J., Doiron, B., Shea-Brown, E., Kreizer, J., Reyes, A. Correlation between neural spike trains increases with firing rate. *Nature* (2007)

[doi:](#)

I-21. Saliency and surround interactions via natural scene statistics: A unifying model.

Ruben Coen-Cagli¹
 Peter Dayan²
 Odelia Schwartz¹

RUBEN.COENCAGLI@EINSTEIN.YU.EDU
 DAYAN@GATSBY.UCL.AC.UK
 ODELIA.SCHWARTZ@EINSTEIN.YU.EDU

¹Department of Neuroscience, Albert Einstein College of Medicine

²Gatsby Computational Neuroscience Unit, UCL

Spatial context in images leads to striking effects at neural and perceptual levels, including surround modulation in neurons, and perceptual illusions. Context also plays a critical role in determining the salience of points in visual space, for instance controlling popout, which involves spatial segmentation; and contour integration, which is associated with grouping. Indeed, it has been proposed [1] that one major function of the primary visual cortex (V1) is to build a visual saliency map; a theory that has been realized in a model of the dynamical, recurrent, interactions among nearby cortical neurons. Here, we consider such interactions from the point of view of a well-found account of natural scene statistics, and so provide a computational theory of contextual effects. When trained on an ensemble of natural images, our model reproduces some perceptual saliency data involving grouping and segmentation; we show that it also unifies a wider range of neural and perceptual contextual effects, including surround modulation and the contrast dependency of receptive field size in V1, and the perceptual tilt illusion. We focus on the Gaussian Scale Mixture generative model (GSM) of natural scene statistics. The model characterizes the probabilistic process by which the statistical dependencies amongst V1-like linear simple cell activations might arise: the multiplication of a multidimensional Gaussian (each dimension corresponding to local filter structure) by a common scalar variable, the mixer. Extracting the structure underlying the image involves inverting the model; i.e., dividing by the common mixer to remove the dependencies. The GSM is thus closely related to neural models of divisive gain control, and to theories of efficient coding. Here we extend the GSM to account for issues of grouping and segmentation in larger swathes of images. First, we focus on the question of inferring for a given scene whether the activations of filters in center and surround locations are statistically coordinated. Filters assigned to the same gain control pool duly share a common mixer and are jointly divisively normalized. Second, we infer the covariance of the underlying Gaussian variables from natural scene statistics. Such covariances can themselves lead to suppressive and facilitatory influences in the model. We show that the model qualitatively reproduces neurophysiology data in V1, namely: the orientation tuning of surround suppression and the contrast dependency of area-summation curves observed in V1. In addition, because the model generalizes divisive gain control, it has the potential to capture a broader range of surround modulations of firing rate, including the facilitation observed in some cases. We then consider a collection of model neurons with different orientations, and show that a standard population model reproduces perceptual saliency effects such as orientation popout, collinear facilitation, and contour integration in simple displays of oriented bars. We also show that the model, with a larger number of oriented units, captures attraction and repulsion effects in the tilt illusion [2]. [1] Z. Li.. Trends Cogn Sci, 6(1):9-16, 2002. [2] O. Schwartz, T.J. Sejnowski, P. Dayan.. Journal of Vision, 9(4)(19):1-20, 2009.

doi:

I-22. A three-layer model of natural image statistics

Aapo Hyvarinen
Michael Gutmann
University of Helsinki

AAPO.HYVARINEN@HELSINKI.FI
MICHAEL.GUTMANN@HELSINKI.FI

Background: Statistical modelling of natural images is an established approach for modelling receptive fields in the early visual system. Most work has considered hierarchical models of one or two layers, corresponding to simple and complex cells. Extending hierarchical models of natural image statistics to more than two layers could provide predictions of visual processing beyond complex cells, possibly in extrastriate cortex. However, the specification of such models is very difficult due to the lack of suitable models of nonlinear visual processing, and estimation of such models is technically difficult because the probabilistic models are usually intractable (e.g. unnormalized). Methods: We propose a three-layer model of natural images which includes as the first two layers the well-established energy model of complex cells. The third layer is constructed as a linear pooling of the logarithms of complex cell outputs. The motivation for using logarithms is that it is a concave function (increases slower than linearly) and thus a linear combination of logarithms is related to an AND operation of complex cell outputs. Based on this nonlinearity, we propose an energy-based probabilistic model. The model is unnormalized and thus difficult to estimate with conventional methods; score matching estimation is also difficult due to the three-layer structure. Therefore, we have developed a new principle for estimating such energy-based unnormalized probabilistic models. The basic idea is to learn a classifier to discriminate between natural images and some artificially generated noise. Theoretical analysis of the estimation principle shows that it provides a

statistically consistent estimator with a relatively small asymptotic variance. Results: We estimated the new three-layer model with natural image input, using our new estimation principle. The first two layers learned receptive fields which were very similar to existing models of simple and complex cells. The third layer learned features which were related to end-stopping and detection of second-order contours (i.e. contours defined by difference in Fourier spectrum instead of luminance). Discussion: The model combined with the new estimation principle is a promising approach for providing predictions of receptive fields of cortical cells, going beyond the classical models of simple and complex cells. In future work, it will be extended to more than three layers, which seems to be straightforward if computationally demanding.

[doi:](#)

I-23. Learning L_p spherical potentials for Markov Random Field models of natural images

Urs Koster^{1,2}

Michael Gutmann¹

Aapo Hyvärinen¹

KOSTER@CS.HELSEINKI.FI

MICHAEL.GUTMANN@CS.HELSEINKI.FI

AHYVARIN@HELSEINKI.FI

¹University of Helsinki

²Redwood Center for Computational Neuroscience

Markov Random Fields (MRF) with linear filters estimated from the data have recently received a resurgence of interest for modeling the statistics of natural images. While showing good performance in image processing (e.g. denoising tasks), the filters obtained in previous work such as the Field of Experts (FoE, Roth and Black) seem to be at odds with the receptive field structure of cells in the primary visual cortex (V1): rather than localized, oriented Gabor-like receptive fields, the FoE models learns discontinuous high-frequency features, which show no clear relation to simple cell receptive fields. Here we revisit these findings, using a generalization of the existing FoE model. We consider potentials which are computed as the L_p norm of a set of filter outputs raised to a power q , which is in contrast to the previous linear combination of weighted, rectified filter outputs. This model is inspired by Independent Subspace Analysis (Hyvärinen et al.) and captures energy dependencies between linear filter outputs which are common in natural images. Two special cases of the model are of particular interest: For the L_1 norm and $q=1$, the model reduces to the FoE, albeit with a Laplacian rather than student-t like nonlinearity. For the L_2 norm, the clique-potentials are spherically symmetric. In addition to considering these special cases, we estimate the optimal nonlinearity (i.e. norm and power) from the data and analyze how the model fit is affected by this choice. The estimation of the potential functions in a MRF is generally a difficult problem, because the probability density cannot be normalized in closed form, rendering maximum likelihood estimation impossible. In previous work, Contrastive Divergence and Score Matching have been used for the estimation, but they have drawbacks in computational efficiency and in handling densities lacking smoothness, respectively. Therefore we choose to estimate the model with a novel estimation method, Noise Contrastive Estimation (Gutmann and Hyvärinen). This method trains a classifier to distinguish data samples from noise, and has been shown to lead to a consistent estimator with a small variance. We estimated the model on training images of size 15x15 and 45x45 pixels with linear filters of size 5x5 and 15x15 pixels respectively. The learned filters are localized, oriented and band-pass like those reported in previous patch-based models, corresponding well to simple cell receptive fields found in V1. Furthermore, we analyze the performance of the different models in a denoising and filling-in task to elucidate how the fit of the model to the data relates to the performance in real-world tasks.

[doi:](#)

I-24. Grasping image statistics

Omid Aladini
Constantin A. Rothkopf
Jochen Triesch

ALADINI@FIAS.UNI-FRANKFURT.DE
ROTHKOPF@FIAS.UNI-FRANKFURT.DE
TRIESCH@FIAS.UNI-FRANKFURT.DE

Frankfurt Institute for Advanced Studies

The human ability to parse images and recognize objects therein is still unmatched by artificial system. A wealth of research has demonstrated that many visual tasks can be analyzed by comparing human and animal performance to normative models understanding perception as Bayesian inference of latent causes generating these images. Previously, many visual tasks have been formulated in terms of the statistics of image features [e.g. 1,2,3] conditional on some scene property, including whether an object contour is present or not. This work has shown that indeed such image statistics obtained from natural images can be used to invert the generative model and infer whether a contour is present at an image location. Almost all of these models have used large collections of labeled images as a means to extract joint statistics of features and image labels such as object category or object contour. Recently, generative models have been proposed, which learn object categories and segmentations in an unsupervised way [e.g. 4,5] from collections of images without labels. Contrary to artificial system, humans do have the distinct advantage of being able to obtain tactile information from manual interaction with the environment in addition to visual input. Here we explore only, how such tactile interaction can be used together with visual input in order to learn image statistics without requiring labeled image databases. A humanoid agent was simulated in a virtual 3d graphics environment. It moves its hand repeatedly on stereotyped trajectories that can be obtained by servoing. The arm moves until it is fully extended or the hand hits a surface. During this process images are captured from the point of view of the agent fixating at a point slightly ahead of its hand. Using this data image features are calculated at the point of fixation including oriented bandpass filter responses at eight orientations and four spatial scales, oriented energy responses with the same set of filters, local luminance gradient direction, and contrast polarity. Furthermore, given that the 3d layout of the scene is fully known, we additionally store whether an object edge was present with each such data item. The joint statistics of image features and the binary variable coding whether the agent was touching a surface or not can be estimated from the described data set. It is now possible to compute the posterior probability that the hand will intersect a surface at the next time step, given the current observation. We show that this posterior indeed reflects the presence of object edges in the scene, without doing, object segmentation or contour grouping. Thus, the presented work demonstrates that image statistics that can be useful for a large variety of visual tasks may not only be learned from images but benefit from learning in conjunction with tactile interaction within the scene. [1] Konishi, Yuille, Coughlan (2003) Image & Vision Computing (21) [2] Geisler, Perry (2009) VisNeuroSci (26) [3] Ren, Fowlkes, Malik (2008) IntJComputVis (77) [4] Bart,Porteous,Perona,Welling. (2008) CVPR [4] Sivic, Russell, Zisserman, Freeman, Efros (2008) CVPR

[doi:](#)

I-25. Beyond magical numbers: towards a noise-based account of visual-short term memory limitations

Wei Ji Ma
Wen-Chuang Chou
Baylor College of Medicine

WJMA@BCM.EDU
PHILIPPEDECHOU@GMAIL.COM

Visual short-term memory (VSTM) is a short-term buffer of visual information that allows animals to detect changes between subsequent scenes. VSTM performance is widely believed to be limited by a fixed number of items that can be memorized (the "magical number 4"), but this notion has recently become subject of intense debate. In change detection experiments (Wilken and Ma, 2004), we found that this limited-capacity model fails to describe empirical receiver-operating characteristics. Instead, we proposed that observers' judgments are based on combining noisy stimulus observations. Our model explained the data well if we assumed that stimulus noise

increases with set size, N . We introduced a delayed estimation paradigm to measure the resulting decrease in precision directly. Recent work (Zhang and Luck, 2008) has questioned our interpretation of this result and attributed the decrease to a higher proportion of random guessing. Here, we present new delayed estimation experiments and a neural model to settle this controversy. Observers viewed between 1 and 8 widely spaced colored discs at fixed eccentricity for 100 ms. Colors were drawn independently from a color wheel. After 1 second, one location was marked and the observer reported the color of the disc that had been at that location (the target) in one of two ways: a) clicking on a color wheel; b) scrolling through all colors using arrow keys. A limited-capacity model predicts: 1) an observer's capacity, K , is independent of response modality; 2) when $N \leq K$, the target color is always reported; 3) any instance of not reporting the target color is due to random guessing; 4) when reporting the target color, response variance is independent of N (in the Zhang & Luck slot model: when $N > K$). Instead we find that: 1) observers' capacity is 36% higher in the scrolling than in the color wheel paradigm; 2) when $N \leq K$, subjects do not always report the target color; 3) when subjects do not report the target color, they often reports the color of another item, consistent with findings of Bays and Husain (2009); 4) response variance increases continuously with N . We next conducted a two-alternative forced-choice experiment in which subjects indicated for a given test color, which of two marked locations contained that color. The results of this experiment were consistent with that of the first one. These findings can be explained by a population coding model in which approximately the same number of spikes is used to encode any number of items. Fisher information per item is then roughly inversely proportional to N , a trend we see in the data. We implemented this quasi-neural argument in a simple neural network characterized by spatial averaging (to account for confusions of items' positions) and divisive normalization. The latter operation is common in models of selective attention. This network can explain all behavioral data without assuming an item or slot limit of any sort. We argue that visual short-term memory must be reconceptualized in terms of noise and inference, and that its limitations are likely tied to attentional ones.

doi:

I-26. Identifiability of nonlinear receptive field models from sensory neurophysiology data

Christopher DiMattina¹
Kechen Zhang²

CHRIS_DIMATTINA@YAHOO.COM
KECHEN@BME.JHU.EDU

¹Electrical Engineering & Computer Science, Case Western Reserve University

²Johns Hopkins University

One approach to systems-level sensory neurophysiology which has seen a recent resurgence of interest is system identification (SI), where the goal of neurophysiology experiments is to estimate the parameters of a generative sensory processing model. Perhaps the most familiar example of this approach is the estimation of the linear receptive field (RF) from complex stimulus ensembles like random noise or natural stimuli. Extending SI methods to characterize nonlinear sensory neurons, it is necessary to fit complicated nonlinear models like hierarchical neural networks or radial basis function models, thereby presenting new computational challenges because different sets of model parameters may define virtually identical input-output relationships. In this work, we consider the identifiability of nonlinear parametric models from sensory neurophysiology data. We derive a partial differential equation as a simple mathematical criterion for a model to have a continuum of parameter values that yield identical input-output functions. Applying this criterion to standard three-layer neural networks, we show that in theory continuous parameter confounding can occur only when the hidden unit gains are given by power, exponential or logarithmic functions. Since the standard hyperbolic tangent or and logistic gain functions used in neural modeling are not of these mathematical forms, one might think that continuous parameter confounding would not be a problem when fitting standard models to neural data. However, these standard gain functions may be well approximated over limited ranges of their inputs by one or more of the forms permitting parameter confounding, and therefore with poorly chosen stimulus sets it is possible to observe continuous equivalence classes of functionally equivalent neural networks, thereby making it impossible to recover the parameters of the true functional model which gives rise to the data. In order to identify a nonlinear model, our simulations suggest that one should avoid the independently distributed random stimuli typically used in SI experiments. We demonstrate that choos-

ing training data in an adaptive manner using optimal experimental design in on-line experiments may be able to reduce or eliminate continuous parameter confounding, suggesting that the use of active on-line data collection not only speeds up convergence but may in fact be necessary to accurately estimate the parameters of nonlinear models like hierarchical neural networks.

doi:

I-27. Detecting a change by a single neuron

Hideaki Kim¹
Barry J. Richmond²
Shigeru Shinomoto¹

KIM@TON.SCPHYS.KYOTO-U.AC.JP
BJR@LN.NIMH.NIH.GOV
SHINOMOTO@SCPHYS.KYOTO-U.AC.JP

¹Kyoto University

²Laboratory of Neuropsychology NIMH/NIH/DHHS

Detecting a change in state of a dynamical process from observing noisy signals is a well-known and difficult statistical problem that has practical applications in places such as manufacturing to detect substandard products [1,2]. Late detection increases inferior goods, whereas false alarms lead to unnecessary production stoppages. There is a parallel in animal behavior. It would be advantageous for individuals to detect subtle changes in the environment as early as possible to avoid predators or to locate prey without incurring a high false alarm rate. In practice, animals seem to detect changes even in weak sensory signals reliably and efficiently [3]. These signal changes occur trains of irregularly timed neuronal spikes. Although the physiological mechanisms underlying this apparently efficient statistical computation have not been identified yet, recent studies show that even single neurons could detect changes in the rate of an incoming spike train [4,5,6]. To investigate the potential for a single neuron to detect a change point, we carry out a systematic analysis using a single leaky integrate-and-fire model. By combining a Bayesian criterion [7] for measuring the quality for a change-point detector and the first passage time of Ornstein-Uhlenbeck process, we develop a method for optimizing the neuronal parameters, including the membrane time-constant and the threshold, given initial and final rates of an incoming spike train. We find that a single leaky integrate-and-fire neuron can achieve performance close to that of Bayes-optimal detection algorithm [8]. Given a reasonable number of synaptic connections and the rate of the input spike train, the values of the membrane time-constant and the threshold that maximize change point detectability are close to those seen in biological neurons, suggesting that biological neurons could be acting as change point detectors. References 1. Zacks, S. (1982). *Stat. Anal. Donnees* 7, 48-81. 2. Bhattacharya, P. (1994). In *Change-point Problems*, E. Carlstein, H. Müller, and D. Siegmund, ed. (Hayward: Institute of Mathematical Statistics). pp. 28-56. 3. Nelson, M. E., and MacIver, M. A. (1999). *J. Exp. Biol.* 202, 1195-1203. 4. Ratnam, R., Goense, J. B. M., and Nelson, M. E. (2003). *Neurocomputing* 52-54, 849-855. 5. Goense, J. B. M., and Ratnam, R. (2003). *J. Comp. Physiol. A* 189, 741-759. 6. Yu, A. J. (2007). In *Advances in Neural Information Processing Systems*. (Cambridge: MIT Press), 19, pp. 1545-52. 7. Shiryaev, A. N. (1963). *Theor. Probab. Appl.* 8, 22-46. 8. Peskir, G., and Shiryaev, A. (2002). In *Advances in Finance and Stochastics*, K. Sandmann, and P. Schönbacher, ed. (Berlin: Springer), pp. 295-312.

doi:

I-28. Complexity and performance in simple neuron models

Skander Mensi¹
 Richard Naud¹
 Michael Avermann²
 Carl Petersen²
 Wulfram Gerstner¹

SKANDER.MENSI@EPFL.CH
 RICHARD.NAUD@EPFL.CH
 MICHAEL.AVERMANN@EPFL.CH
 CARL.PETERSEN@EPFL.CH
 WULFRAM.GERSTNER@EPFL.CH

¹EPFL - LCN

²EPFL, Brain-Mind Institute

The ability of simple mathematical models to predict the activity of single neurons is important for computational neuroscience. For neurons, stimulated by a time-dependent current or conductance, we want to predict precisely the timing of spikes and the sub-threshold voltage. During the last years several models have been tested on this type of data. One of the major outcome is that, from a certain degree of complexity, all are very efficient. However, models have never been systematically compared with the same protocol. We study a class of integrate-and-fire models (IF), with each member of the class implementing a selection of possible improvements: exponential voltage non-linearity [1], spike-triggered adaptation current [2], spike-triggered change in conductance, moving threshold [3], sub-threshold voltage-dependent currents [4]. Each refinement adds a new term to the equations of the IF model. This IF family is extendable and adaptable to different neuron types and is able to deal with complex neural activities (i.e. adaptation, facilitation, bursting, relative refractoriness, ...). To systematically explore the effects of a given term of the model a new fitting procedure based on linear regression of voltage change [5] is used in combination with a novel method to extract dynamic threshold and spike-triggered adaptation. This method is fast, robust and allows the extraction of all the models parameters from a few seconds of patch-clamp recordings during injection of a fluctuating current. To investigate the effect of our approach, we applied it to artificial data from Hodgkin-Huxley-like models, as well as experimental data from fast spiking and pyramidal cells. We observe that it is possible to tune the model so that it can reproduce the activity of neurons with high reliability (i.e. almost 100 % of the spike time and less than 1 mV of sub-threshold voltage difference) on new data that was not used for parameter optimization. Using this framework one can classify IF models in terms of complexity and performance and evaluate the importance of each term for different stimulation paradigms. 1. Fourcaud-Trocme N., Hansel D., van Vreeswijk C., and Brunel N. (2003), How spike generation mechanisms determine the neuronal response to fluctuating inputs, *J. Neuroscience* 23:11628-11640. 2. Brette R. and Gerstner W. (2005), Adaptive Exponential Integrate-and-Fire Model as an Effective Description of Neuronal Activity, *J. Neurophysiol.* 94: 3637 - 3642. 3. Badel L., Lefort S., Brette R., Petersen C., Gerstner W. and Richardson M.J.E. (2008), Dynamic I-V Curves Are Reliable Predictors of Naturalistic Pyramidal-Neuron Voltage Traces, *J Neurophysiol* 99: 656 - 666. 4. Richardson M.J.E., Brunel N., and Hakim V (2003), From Subthreshold to Firing-Rate Resonance. *Journal of Neurophysiology* 89:2538-2554. 5. Paninski, L., Pillow, J. & Simoncelli, E. (2004). Comparing integrate-and-fire-like models estimated using intra- and extra-cellular data. *Neurocomputing* 65: 379-385.

[doi:](#)

I-29. Short-term synaptic plasticity and sensory adaptation as Bayesian inference

Ian H. Stevenson¹
 Beau Cronin²
 Mriganka Sur²
 Konrad Kording¹

I-STEVENSON@NORTHWESTERN.EDU
 BEAU.CRONIN@GMAIL.COM
 MSUR@MIT.EDU
 KK@NORTHWESTERN.EDU

¹Northwestern University

²Massachusetts Institute of Technology

Neurons in the sensory system exhibit changes in excitability that unfold over many time scales. These fluctuations produce noise and could potentially lead to perceptual errors. However, to reduce these errors, postsynaptic

neurons can adapt and counteract changes in the excitability of presynaptic neurons. Here we introduce a model of how neurons could optimally adapt to minimize the influence of changing presynaptic neural properties on their outputs. The resulting excitability estimation model, based on Bayesian inference, reproduces many of the properties of previous models of short-term synaptic plasticity. Additionally, this model explains a range of physiological data from experiments which have measured the overall properties and detailed time-course of sensory adaptation in the early visual cortex. In this framework, short-term plasticity and adaptation are the result of a strategy to compute reliably with a nervous system that changes on many timescales. The central problem in estimating fluctuations in presynaptic excitability is that firing rate information can be ambiguous. High firing rates may occur because of strong sensory drive or, alternatively, because presynaptic neurons are highly excitable. In order to adapt in a way that preserves sensory information, the nervous system needs to resolve this ambiguity. Specifically, the nervous system can use information about the way excitability typically changes over time and information about the way sensory drive typically changes over time. Here we assume that excitability drifts on multiple timescales around a steady state point and that sensory drive is sparse. For each timescale, the optimal adaptation model estimates the current excitability using an extended Kalman filter. The response of the postsynaptic neuron is then given by the synaptic input divided by the total estimate of the presynaptic excitability. Similar to previous work on gain control, the effect of this optimal adaptation rule is that inputs from presynaptic neurons with high excitability will tend to have low gain. Under this rule short term increases in firing rate are typically attributed to high drive while prolonged increases in firing rates are attributed to high excitability. Here we show that the optimal adaptation rule reduces response variability in the presence of fluctuating presynaptic excitability. Moreover, the predicted time-course of the excitability estimation matches the time-course of a prominent synaptic depression model (Tsodyks and Markram, 1997). This similarity raises the possibility that the timescales of synaptic depletion and recovery may be predicted by the statistics of excitability changes in presynaptic neurons. At the systems level, by simulating simple cells receiving tuned inputs, this model can also reproduce the structure and time-course of repulsive tilt-adaptation observed in primary visual cortex (Dragoi et al., 2000). Under the excitability estimation, heterogeneity in the adapted tuning curves appears to be well explained by initial differences in the pre-adaptation tuning curves.

doi:

I-30. The speed of time

Misha Ahrens^{1,2}
Maneesh Sahani³

MBA22@CAM.AC.UK
MANEESH@GATSBY.UCL.AC.UK

¹Cambridge University

²Janelia Farm Research Campus

³Gatsby Unit, UCL

The content of a visual stimulus affects its perceived duration. In particular, stimuli that change rapidly generally appear to last longer than equal-duration stimuli that change more slowly. Exactly why this is, and how stimulus-driven biases interact with internal estimates of elapsed time, is not known. One possible account is that observers expect a certain rate of discrete "events" in the world (Frische, 1964). The rapidly changing stimuli contain more of these events per unit time, and so observers believe that more time has elapsed. We have proposed a different view as part of a general framework for psychological time estimation (Cosyne, 2008). In our account, observers have an expectation regarding the temporal correlations in the environment, which they use to translate the amount by which the stimulus changes during presentation into probabilistic information about its duration. Here, we contrast these two accounts, showing that they make very different predictions regarding the variability of timing estimates, and testing these predictions experimentally. Internal duration estimates are known to follow a Scalar law, whereby the distribution of estimates scales in proportion to its mean (this, in turn, suggests that the variability in estimates follows a Weber law). We measured the variability of observers' estimates of the duration of colored (i.e. space-time smoothed) Gaussian noise stimuli for a range of different true durations, and found (1) that estimates of the duration of these dynamic stimuli were less variable than estimates of the durations of static stimuli; and (2) that the Scalar law continued to hold in the presence of these random dynamic stimuli. Both models are consistent with observation (1), as the added information in the stimulus should help to reduce

variability. We have previously shown that the change-based approach also yields scale-invariant estimates, and it is thus consistent with the second observation too. By contrast, we show here that the event-counting model leads to sub-Weberian growth in variability for these stimuli, and that even when combined with an internal estimate that retains the Scalar property it cannot account for observers' behavior. Our results thus (1) provide further evidence that the reason the psychological speed of time is affected systematically by stimuli is that observers exploit these stimuli to derive additional information about the passage of time; and (2) argue against simple counting-based sources of this additional information, lending support to the idea that the information derives from more extensive knowledge of the statistics.

[doi:](#)

I-31. Reconstruction of sparse circuits using multi-neuronal excitation (RES-CUME)

Tao Hu
Dmitri Chklovskii

HUT@JANELIA.HHMI.ORG
CHKLOVSKIID@JANELIA.HHMI.ORG

HHMI, Janelia Farm Research Campus

One of the central problems in neuroscience is reconstructing synaptic connectivity in neural circuits. Synapses onto a neuron can be probed by sequentially stimulating potentially pre-synaptic neurons while monitoring the membrane voltage of the post-synaptic neuron. Reconstructing a large neural circuit using such a "brute force" approach is rather time-consuming and inefficient because the connectivity in neural circuits is sparse. Instead, we propose to measure a post-synaptic neuron's voltage while stimulating sequentially random subsets of multiple potentially pre-synaptic neurons. To reconstruct these synaptic connections from the recorded voltage we apply a decoding algorithm recently developed for compressive sensing. Compared to the brute force approach, our method promises significant time savings that grow with the size of the circuit. We use computer simulations to find optimal stimulation parameters and explore the feasibility of our reconstruction method under realistic experimental conditions including noise and non-linear synaptic integration. Multi-neuronal stimulation allows reconstructing synaptic connectivity just from the spiking activity of post-synaptic neurons, even when sub-threshold voltage is unavailable. By using calcium indicators, voltage-sensitive dyes, or multi-electrode arrays one could monitor activity of multiple post-synaptic neurons simultaneously, thus mapping their synaptic inputs in parallel, potentially reconstructing a complete neural circuit.

[doi:](#)

I-32. On the connections between SIFT and biological vision

Kritika Muralidharan
Nuno Vasconcelos

KRMURALI@UCSD.EDU
NUNO@ECE.UCSD.EDU

Statistical Visual Computing Laboratory, UCSD

In the past decade, research in object recognition has firmly established the efficacy of image representations based on histograms of dominant gradient orientation. The SIFT descriptor, in particular, could be considered today's default (low-level) representation for object recognition, adopted by hundreds of computer vision papers. It is heavily inspired by known computations of the early visual cortex, but has no formal detailed connection to computational neuroscience. Simultaneously, a seminal development in computational neuroscience research has been to explain the ability of individual cells to adapt their dynamic range to the strength of the visual stimulus, by the implementation of gain control through divisive normalization. In this work, we propose a novel representation of local image orientation that will show that these two, apparently disjoint, developments are, in fact, tightly coupled. We start by formulating the central motivating question for descriptors such as SIFT or HOG - "how

to represent locally dominant image orientation", as a decision-theoretic problem. An orientation is defined as dominant, at a location of the visual field, if its Gabor response at that location is both 1) distinct from that of other orientations and 2) large. An optimal statistical test is then derived to determine if an orientation response is distinct. The core of this test is the posterior probability of each orientation at a location, given its Gabor response. The dominance of an orientation within a neighborhood R is then defined as the expected strength of responses in R , which are distinct. Exploiting known properties of natural image statistics, we then show that this measure of orientation dominance, denoted as bioSIFT, can be computed with the sequence of operations of the standard neuro-physiological model: simple cells composed of a linear filter, divisive normalization, and a saturating non-linearity, and complex cells that implement spatial pooling. This connection between computer vision and neuroscience provides additional justification to both the success of SIFT in computer vision, and the importance of divisive normalization in the brain. It also points to the importance of contrast normalization in vision. To illustrate this, we show that the simple replacement of non-normalized Gabor filter responses, with the normalized orientation descriptors of bioSIFT, produces very significant gains in the recognition accuracy of the HMAX network - a biologically-inspired object recognition architecture. The enhanced network outperforms the previous best HMAX results in the literature, and has performance competitive with that of comparable state-of-the-art non-biological recognition architectures. The proposed descriptor is also shown to exhibit the trademark properties of V1 neurons such as independence, sparseness, cross-orientation suppression and a contrast response that fits the Naka-Rushton equation. The independence properties are not exploited by current SIFT-based recognition architectures, which rely on a computationally expensive probabilistic representation (visual words) of feature dependence. We illustrate the potential of bioSIFT for computationally efficient classification by designing a gist classifier that exploits feature independence. This is shown to have good performance on a gist-based image classification task.

doi:

I-33. The pairwise phase consistency: A bias-free measure of rhythmic neuronal synchronization

Martin Vinck¹

Marijn van Wingerden¹

Thilo Womelsdorf²

Pascal Fries³

Cyriel Pennartz¹

M.A.VINCK@UVA.NL

E.J.M.VANWINGERDEN@UVA.NL

THIWOM@IMAGING.ROBERTS.CA

PASCAL.FRIES@ESI-FRANKFURT.DE

C.M.A.PENNARTZ@UVA.NL

¹University of Amsterdam

²University of Western Ontario, Canada

³Ernst Strüngmann Institute, Frankfurt

Oscillatory activity is a widespread phenomenon in nervous systems and has been implicated in numerous functions. Signals that are generated by two separate neuronal sources often demonstrate a consistent phase-relationship in a particular frequency-band, i.e. they demonstrate rhythmic neuronal synchronization. This consistency is conventionally measured by the PLV (Phase-Locking Value) or the spectral coherence measure. Both statistical measures suffer from significant bias, in that their sample estimates overestimate the population statistics for finite sample sizes. This is a significant problem in the neurosciences where statistical comparisons are often made between conditions with a different number of trials, or between neurons with a different number of spikes. We introduce a new circular statistic, the PPC (pairwise phase consistency). We demonstrate that the sample estimate of the PPC is a bias-free and consistent estimator of its corresponding population parameter. Our numerical simulations show that the population PPC is linearly related to the population PLV for a large range of PLVs. The variance and mean squared error of the PPC and PLV are compared. A procedure is proposed to weigh phases by the signals' amplitudes to obtain a more robust measure of phase consistency, while avoiding the influence of amplitude co-variations, which is a known problem for the coherence measure. Finally, we demonstrate the practical relevance of the method in actual neuronal data recorded from the orbitofrontal cortex of rats that engage in a 2-odour discrimination task. We find a strong increase in rhythmic synchronization of spikes

relative to the Local Field Potential (as measured by the PPC) for a wide range of low frequencies (including the theta-band) during the anticipation of sucrose delivery in comparison to the anticipation of quinine delivery.

doi:

I-34. Physiology in *Drosophila* motion-sensitive neurons during walking and flight behavior

Johannes D. Seelig¹

M. Eugenia Chiappe¹

Gus K. Lott¹

Michael B. Reiser²

Vivek Jayaraman¹

SEELIGJ@JANELIA.HHMI.ORG

CHIAPPEE@JANELIA.HHMI.ORG

LOTTG@JANELIA.HHMI.ORG

REISERM@JANELIA.HHMI.ORG

VIVEK@JANELIA.HHMI.ORG

¹Janelia Farm Research Campus, HHMI

²Janelia Farm Research Campus

Drosophila melanogaster is a genetic model organism with many experimental advantages, including the ability to genetically manipulate specific sub-populations of neurons. Selected neurons in the fruit fly central brain can also be targeted for in vivo electrophysiology and two-photon calcium imaging. This powerful combination of physiology and genetic tools-such as cell-type-specific labeling, activity sensing, light-activation and silencing-is increasingly being applied to questions in systems neuroscience [1]. Our goal is to understand circuit computations underlying sensory-motor transformation in the fly brain. This requires recording not just neural activity but also the fly's behavior. Towards this end, we will present two novel experimental setups that we have developed: (i) Two-photon calcium imaging while the fly is walking on an air-supported 'Buchner' ball [2, 3] in a virtual arena [4]. (ii) Two-photon calcium imaging when the tethered fly is flying in a virtual arena. In order to assess the quality of behavior in our imaging preparation, we focused on a well-known behavior, the optomotor response. This turning response to visual motion can be induced by presenting a large field vertical grating moving horizontally in front of the fly. We show that tethered walking and flying flies reliably perform optomotor behavior during two-photon imaging. To assess the quality of the calcium imaging, we investigated calcium responses in lobula plate tangential cells (LPTCs). This small group of neurons has been exhaustively studied in fixed preparations in larger flies, and their responses to large-field stimuli (such as those that induce the optomotor response) have been carefully characterized [5]. In blowflies, LPTCs respond to large field motion stimuli with strong calcium responses, making these cells ideal candidates to test the capabilities of the behaving *Drosophila* preparation. Recording from Horizontal System (HS) LPTCs, we find reliable calcium responses in behaving fruit flies using a new genetically encoded calcium indicator, GCaMP3.0 [6]. HS neurons show strong (~150%DF/F) and stable responses to motion in their preferred direction. These calcium responses correlate with the fly's own walking optomotor response in that direction. Do responses of HS neurons change depending on the behavioral state of the animal? We will present results that compare responses in walking flies to those in flying flies. In our poster, we will describe the details of setups that allow us to perform stable recordings from identified motion-sensitive interneurons in the visual system of walking and flying *Drosophila*. These recordings represent the first physiological recordings in behaving *Drosophila* and provide a platform for future explorations of decision-making and sensory-motor transformations in this powerful genetic model organism. 1. Olsen and Wilson (Trends Neurosci., 2008) 2. Gotz and Wenking (J Comp Physiol, 1973) 3. Bohm, Schildberger and Huber (JEB, 1991) 4. Reiser and Dickinson (J Neurosci Meth, 2008) 5. Borst and Haag (J Comp Physiol A, 2002) 6. Tian et al. (Nat Meth, 2009)

doi:

I-35. The frequency of hippocampal theta oscillations and unit firing can be manipulated by changing the temperature

Eva Pastalkova¹
György Buzsáki²

PASTAK@JANELIA.HHMI.ORG
BUZSAKI@RUTGERS.EDU

¹Janelia Farm Research Campus, HHMI

²CMBN, Rutgers University

In order to study how theta oscillations organize the temporal patterns of neurons within the hippocampal-entorhinal system, we developed a method which allows us to manipulate the frequency of theta oscillations for short periods of time. We increased or decreased the local temperature within the medial septum using an insulated golden wire, which was connected to a Peltier device above a skull of an animal. We show that local cooling of the medial septum was followed by the decrease of the theta frequency and local heating by the increase of the theta frequency in both, hippocampus and entorhinal cortex. Correspondingly, firing of interneurons as well as pyramidal neurons and their interaction was faster during MS heating and slower during MS cooling. The change of the theta and unit firing frequency did not depend on whether an animal was running on a running wheel or in a maze suggesting that the firing of neurons is controlled by the theta oscillation rather than by the external sensory cues. Thus, heating and cooling of the medial septum can be used to manipulate the frequency of the theta oscillations in the hippocampus and entorhinal cortex, facilitating the study of the relationship between LFP oscillations, neuronal firing and sensory stimuli.

doi:

I-36. Compressed sensing in the brain: role of sparseness in short term and long term memory

Surya Ganguli¹
Haim Sompolinsky²

SURYA@PHY.UCSF.EDU
HAIM@FIZ.HUJI.AC.IL

¹UCSF

²ICNC, Hebrew Univ. and CBS, Harvard Univ.

One of the most exciting advances in signal processing is the field of compressed sensing (CS) [1]. In CS, sparse high-dimensional stimuli are represented by lower dimensional dense measurements, which are linear mixtures of the stimuli. CS shows that by using the computationally tractable L1 norm as a sparse prior, the high-dimensional stimuli (for example, human speech, natural movies, fMRI data) can be fully reconstructed from the compressed measurements. In this work, we have extended CS theory and applied it to reveal new fundamental properties of neural learning and memory in the biologically relevant scenario of sparse coding. Our first extension of CS addresses the capabilities of neuronal circuits to buffer temporal signals, subserving working memory. Working memory systems can store complex temporal sequences, such as speech, over many seconds even though the individual neurons involved, when isolated, forget their inputs rapidly over milliseconds. Previous work [2,3], which assumed normally distributed signals, has shown that a circuit of N neurons can store temporal signals of duration at most N (in units of the neuronal time constant). Furthermore, it was shown [4], that in the presence of noise, recurrent networks cannot outperform an equivalent feedforward network which functions as a delay line. Here we show that sparse signals can be faithfully reconstructed from instantaneous network activity even for a duration of time that exceeds the number of neurons in the network. This enhanced capacity for storing sparse signals is realized by 'orthogonal' recurrent networks and not by feedforward networks. We compute analytically the memory capacity and distribution of errors in signal reconstruction as a function of network size, signal sparsity, and the distribution of nonzero components in the signal. We also analyze the ability of neuronal networks to learn rules when the rule to be learned can be realized by a network with sparse connectivity. We show that in contrast to classical results concerning learning and generalization in neural networks, such networks can actually learn rules, and generalize correctly, using a remarkably small number of examples, smaller than the

dimensionality of the input space. Finally, we explore the properties of synaptic learning in sensory motor tasks which incorporate a sparsity prior on synaptic connectivity patterns in the form of the L1 norm. We analyze the dynamics of such learning models and yield predictions for the performance of the system which can be readily tested in experiments. [1] E.J. Candes, M.B. Wakin, IEEE Sig.Proc.Mag. 2009. [2] H. Jaeger GMD Rep. 148 2001. [3] O.L.White, D.D. Lee, H. Sompolinsky, PRL 2004. [4] S.Ganguli, D.Huh, H.Sompolinsky, PNAS 2008.

[doi:](#)

I-37. A spike timing computational model of hippocampal-frontal dynamics underlying navigation and memory

Laurence Jayet¹
Philip H. Goodman²
Mathias Quoy³

LJAYET@MEDICINE.NEVADA.EDU
GOODMAN@MEDICINE.NEVADA.EDU
MATHIAS.QUOY@U-CERGY.FR

¹Brain Computation Lab, University of Nevada Reno

²University of Nevada Reno

³University of Cergy

Understanding and intervening in central nervous system disorders such as dementia, epilepsy, and stroke requires a deeper understanding of mechanisms between hippocampal and neocortical regions that produce meaningful action. A basic behavior shared by all mammals is the task of navigating in a novel environment, which requires reliable short-term episodic memory. The most well-established electrophysiological findings in the hippocampal-neocortical system are the phenomenon of place fields and hippocampal "place cells", modulated by 6-10 Hz "theta" inhibition, first described by O'Keefe and Dostrovsky in 1971 (Brain Res 1971; 34:171). More recently, Hafting et. al reported entorhinal "grid cells" (Nature 2005; 436:801), which further studies showed are likely responsible for stabilizing place fields (Eur J Neurosci 2008; 27:1933). In 2009, Harvey et al. (Nature 2009; 461:941) reported the first in vivo patch recordings of hippocampal CA1 cells in an awake behaving mouse navigating a virtual maze. Despite this series of discoveries, we are aware only of high level spatial-temporal theories that attempt to explain the role of grid cells (Moser et al. Annu Rev Neurosci 2008; 31:69). Here, we present what we believe is the first comprehensive spike-timing, conduction-based synaptic model of the hippocampal formation (HF)-neocortical system that includes a role of grid cells in stabilizing rather than establishing place cell activity. Our model attempts to explain the mechanisms of both place field formation and stabilization during a computer-simulated rodent maze navigation, exhibiting subthreshold dynamics consistent with the recent in vivo recordings by Harvey et al. This model utilizes recent theoretical microcircuitry dynamics being developed at UNR, called "Recurrent Asynchronous Irregular Nonlinear" (RAIN) networks, that are self-sustaining once activated, and silenced under certain perturbations. This RAIN-HF model has been implemented in segments, that to date successfully reproduce (1) spontaneously activating and de-activating RAIN networks corresponding to place cell activity, (2) interacting RAIN networks that incorporate Karp channels resulting in intracellular and field potential theta inhibitory oscillation with biological irregularity. This model also includes (3) a feedback loop between the hippocampus (place cells) and the entorhinal cortex (grid cells) to stabilize place field formation, and (4) bidirectional monosynaptic connections from prefrontal cortex to represent a role for executive functions and planning. The RAIN-HF model is framed so that predictions can be biologically represented and tested experimentally in vitro and in vivo. Further implications for physiology and pathophysiology of memory are addressed.

[doi:](#)

I-38. Interaction of hippocampo-neocortical neuronal assemblies during learning and sleep

Adrien Peyrache¹
 Francesco P. Battaglia²

ADRIEN.PEYRACHE@GMAIL.COM
 F.P.BATTAGLIA@UVA.NL

¹UNIC, CNRS

²Universiteit van Amsterdam

Sleep is determinant for memory consolidation. Neuronal assemblies formed during previous waking experiences reactivate spontaneously in the hippocampus and the neocortex during sleep episodes, and in particular Slow Wave Sleep (SWS), possibly important for memory stabilization during offline periods. Hippocampal replay events are associated with high frequency (~200 Hz) oscillations in the LFP called ripples. Hippocampal reactivation may lead reactivation of neocortical assemblies and thus contribute to the stabilization of long term memory trace in the neocortex. Learning of complex behavioral strategies is thought to recruit the prefrontal cortex (PFC) which receives a unilateral and monosynaptic projections from the hippocampus. Simultaneous recordings of hippocampal Local Field Potentials (LFP) and large ensemble of isolated units in the PFC while the animal was learning different behavioral contingencies, showed that neurons tend to self-organize in pools of correlated units. Those subgroups of cells were identified by mean of principal component analysis on awake binned spike trains. The projection of instantaneous population vectors from a sleep period onto a given principal component from waking quantifies the timecourse of reactivation of cells during sleep. Reactivation, averaged over whole SWS episodes, was higher in the sleep following the task (sleep POST) than during a control sleep epoch prior to task (sleep PRE). The reactivation took the form of short lasting events (~100 ms) and strong co-firing among the subgroups of cells corresponding to waking assemblies. Amplitude and inter-event intervals for these transients followed power-law distributions. Replay events in the PFC tend to occur at times of hippocampal ripples, hence pointing out an enhanced coordination between the two structures. Furthermore, this relationship was modulated by learning as prefrontal assemblies formed upon rule acquisition were the ones which were most likely replayed with hippocampal reactivations in subsequent sleep episodes. In another experiment, simultaneous recordings of up to ~120 isolated neurons over a large portion of the neocortex were carried out while the animal was performing repetitive and over-trained behavior. Neuronal assemblies, extracted with the same method as above, showed also prominent reinstatement during sleep but both in PRE and POST epochs. However, these reactivations occurred more often at times of hippocampal ripples in sleep POST than in sleep PRE, suggesting an enhanced hippocampo-neocortical dialogue in sleep following immediately task performance, even for pre-existing neuronal patterns. The time resolution of the measure can not determine if these reactivations are lead by one or the other structure, or if they are just coordinated. These assemblies were distributed over the neocortex, often recruiting neurons in both hemispheres. Hence, the reactivation process is largely distributed and could contribute to plasticity over long-range populations as suggested by previous findings on cell pair recordings in different neocortical structures. In new conditions, performance optimal neuronal patterns are reactivated with the hippocampus, most certainly to be consolidated. Conversely, re-experience in familiar situations would transiently make the already consolidated neocortical memory traces coordinated to hippocampal assemblies, thus allowing them to be transformed or updated in case of environmental changes.

doi:

I-39. Spontaneous activity in a self-organizing recurrent network reflects prior learning

Andreea Lazar¹
 Gordon Pipa²
 Jochen Triesch³

LAZAR@FIAS.UNI-FRANKFURT.DE
 MAIL@G-PIPA.DE
 TRIESCH@FIAS.UNI-FRANKFURT.DE

¹Frankfurt Institute for Advanced Studies

²Max Planck Institute for Brain Research

³Frankfurt Institute of Advanced Studies

In the neocortex, spontaneous activity in the absence of sensory input exhibits nonrandom spatiotemporal patterns [Tsodyks et al., 1999, Fiser et al., 2004]. Following the repetitive presentation of a given visual stimulus, the spontaneous activity patterns show similarities with the sensory-evoked responses [Han et al., 2008]. It has been hypothesized that spontaneous activity reflects prior information, learned via plasticity based on past experience, which when integrated with sensory-evoked activity enables Bayesian inference [Berkus et al., 2009]. We explore the characteristics of spontaneous activity following unsupervised learning of spatio-temporal stimuli in a self-organizing recurrent network (SORN) shaped by synaptic and neuronal plasticity. The SORN model [Lazar et al., 2009] consists of a population of excitatory cells and a smaller population of inhibitory cells. The connectivity among excitatory units is sparse and subject to a simple spike-timing-dependent plasticity rule. Additionally, synaptic normalization keeps the sum of an excitatory neuron's afferent weights constant, while intrinsic plasticity regulates a neuron's firing threshold to maintain a low average activity level. The network receives input sequences composed of different letters and learns the structure embedded in these sequences in an unsupervised manner. Following a learning interval, we omit the input and analyse the characteristics of spontaneous activity. We find that the network revisits states similar to those embedded during input stimulation and that it follows similar trajectories through its high-dimensional state space. Furthermore we show that the spontaneous activity reflects the statistical properties of the data: during spontaneous activity the network preferentially visits states that are similar to evoked activity patterns for inputs with a higher prior probability. Our results establish a novel link between STDP-based unsupervised learning in recurrent networks and concepts of statistical inference. Acknowledgments: This work was supported by the Hertie Foundation, grant PLICON (EC MEXT-CT-2006-042484), and GABA Project (EU-04330). We thank Sophie Deneve for literature suggestions.

doi:

I-40. Spike timing-dependent plasticity interacts with neural dynamics to enhance information transmission

Guillaume Hennequin¹

Jean-Pascal Pfister²

Wulfram Gerstner³

GUILLAUME.HENNEQUIN@EPFL.CH

JEAN-PASCAL.PFISTER@ENG.CAM.AC.UK

WULFRAM.GERSTNER@EPFL.CH

¹EPFL, SV and Brain-Mind Institute

²University of Cambridge, Dpt. of Engineering

³EPFL, LCN

Spike Timing-Dependent Plasticity (STDP) - the fact that synapses vary their strengths in a way that depends on the pre- and postsynaptic spike times at a millisecond timescale - has been computationally studied in various ways during the last two decades. A particularly fruitful approach has been the use of minimal phenomenological models [1, 2]. These can be fitted to experimental data [3, 4], resulting in data-grounded and compact models from which the functional role of STDP can be inferred only indirectly. A complementary approach has been to tackle the reverse problem : a hypothesis is made about the functional role of plasticity, and an optimal learning rule to achieve this goal is derived [5, 6, 7, 8]. For example, maximizing the mutual information between input and output spike trains yields a learning rule with STDP-like features [9]. Here we extend the framework of information maximization to include spike-frequency adaptation (SFA) of the postsynaptic neuron. In the resulting learning rule, potentiation that occurs after a pre-before-post pairing event decreases with the distance to the previous postsynaptic spike. This is in agreement with minimal models where triplets of spikes (pre-post-post or post-pre-post) are essential building blocks [3]. Intuitively, we can understand the triplet effect in the limit of highly reliable neurons. Optimal information transfer (at a fixed average firing rate) would be achieved by Poisson distributed output spikes. Therefore, plasticity has to work against refractoriness and SFA, so that events with short intervals post-post need to be enhanced. We compare our optimal learning rule with the minimal triplet rule [3] and a standard pair-based STDP rule [1, 2] on a task where a single postsynaptic neuron receives input with a given spatio-temporal statistics. We show that, with the infomax and the triplet rule, the neuron specialises on spatial and temporal aspects of the stimulus whereas the standard pair-based rule picks up only spatial aspects. [1]

Gerstner et al. Nature, 1996. [2] Song and Abbott. Neuron, 2001. [3] Pfister and Gerstner. J. Neuroscience, 2006. [4] Clopath et al. Nature Precedings, 2009. [5] Linsker. Neural Computation, 1989. [6] Intrator and Cooper. Neural Networks, 1992. [7] Bell and Sejnowski. Neural Computation, 1995. [8] Toyozumi et al. PNAS, 2005. [9] Toyozumi et al. Neural computation, 2007.

[doi:](#)

I-41. Integration of new and old auditory memories in the European starling.

Derek Zaraza
Daniel Margoliash
University of Chicago

DGZARAZA@UCHICAGO.EDU
DAN@BIGBIRD.UCHICAGO.EDU

A central question in neurobiology is how new memories are integrated with old ones. We addressed this question by tracking the representation of newly learned and previously learned songs in secondary auditory cortex (the caudomedial mesopallium, CMM) of European starlings (*Sturnus vulgaris*) during training on a Go-Nogo operant task. Starlings produce complex songs made up of distinct motifs (basic units of song recognition ~1 s in duration), many of which are unique to the individual. In a first study we conducted daily multisite recordings from CMM as birds ($n=7$) overtrained (for over 1 month) on one set of Go-Nogo stimuli were then switched to new stimuli. Prior to the switch there was no difference in the average neural selectivity (measured across 85 well-isolated neurons) for training stimuli vs. untrained stimuli. Following the switch we observed a rapid increase in selective neuronal response for both new training stimuli and old training stimuli, followed by a return to baseline levels of selectivity as starlings reached asymptotic levels of performance (total of 234 neurons recorded across 3-8 days). The rate of increase and decrease in selectivity was correlated with the birds' behavioral performance. These results suggest that CMM establishes dynamic representations during memory formation but that these representations are not the substrate of long-term memory storage. To explore this hypothesis, in a second study we are examining the role of task learning in the increased selectivity for previously learned stimuli. Naive birds trained on a Go-Nogo task for only 3 days, with behavior remaining near chance levels, were then switched to new stimuli which they then slowly acquired over circa 4-11 days of training. In preliminary data, neither of two birds showed any increase in neuronal selectivity or significant learning on the task prior to the switch (90 neurons). We observed, however, an increase in selective neuronal responses (189 neurons) for both groups of training stimuli following the switch. The time course of neural plasticity corresponded closely with the time course of learning, which was much longer than that of birds in the first study. These preliminary results suggest that plasticity in CMM plays a role in associating learned song stimuli with other related stimuli and responses. In this way CMM may be involved in the integration of memories rather than the storage of one aspect of a memory. Reactivation of prior memories in CMM may be task specific, specific to the most recent prior memories, and/or be regulated by other features of behavior that have yet to be determined.

[doi:](#)

I-42. An avian basal ganglia circuit contributes to fast and slow components of songbird vocal learning

Timothy Warren
Evren Tumer
Michael Brainard
Keck Center, UCSF

TWARREN@PHY.UCSF.EDU
EVREN@PHY.UCSF.EDU
MSB@PHY.UCSF.EDU

In many forms of motor learning, the identity of the neural circuits that produce a learned behavior changes over time, even as the behavior itself remains stable. The mechanisms governing this consolidation of motor

learning are not well understood. Previous studies (refs. 1 and 2) suggest vocal learning in songbirds occurs initially in an avian basal ganglia circuit and then gradually consolidates in a downstream motor pathway. Here we test whether and how this avian basal ganglia circuit contributes to consolidation of adult pitch learning in the downstream motor pathway. We drove pitch learning in specific syllables of adult Bengalese finch song using a pitch-contingent auditory feedback paradigm. We then probed the basal ganglia circuit's contribution to this learning by interfering with activity in LMAN, the basal ganglia circuit's output nucleus which projects to RA, a premotor nucleus in the motor pathway. We drove an initial, large (3-4 sd) pitch shift in a specific song syllable over 2-3 days and then maintained this shift for the following week by keeping feedback stable. Pharmacological inactivation of LMAN during the initial days of the shift caused a partial (~40 percent) reversion of pitch from its learned level toward its original baseline level. Over the following week, the magnitude of pitch reversion caused by inactivating LMAN gradually declined to zero. The initial pitch reversion we observed suggests LMAN contributes to the initial expression of vocal learning. The diminishing effect of inactivating LMAN over the week suggests learning gradually consolidates in the motor pathway to become LMAN independent. To confirm these effects resulted from inactivating LMAN rather than inadvertently inactivating neighboring brain areas, we interfered with LMAN-RA synaptic transmission in a parallel set of experiments. This caused similar effects on pitch learning to inactivating LMAN, allowing us to localized the inactivation effects we observed to LMAN. One possible mechanism underlying this consolidation in the motor pathway is a serial transfer of learning from LMAN to the motor pathway. Alternatively, the motor pathway might gradually learn to produce adaptive behavior on its own through a slow, LMAN-independent process. To distinguish between these possibilities, we tested whether slow, LMAN-independent pitch learning occurred in birds with bilateral LMAN lesions. These LMAN lesioned birds exhibited no significant adaptive learning after multiple days of exposure to the pitch-contingent feedback paradigm. This failure to learn suggests LMAN is required for both an initial, fast learning process as well as the slow, later learning which occurs in the motor pathway. Our results support a model in which LMAN initially instructs vocal change by acutely and adaptively patterning RA activity to bias vocal output. This acute biasing of RA activity then enables a serial transfer of learning to the motor pathway via Hebbian unsupervised learning within the motor pathway. We suggest this general framework, in which initial supervised learning in one circuit drives unsupervised learning in a downstream circuit, may be a general feature of basal ganglia-dependent motor learning. References 1. Kao, Doupe, and Brainard. *Nature*, 2005. 2. Andalman and Fee. *PNAS*, 2009.

[doi:](#)

I-43. The BOLD response in the nucleus accumbens quantitatively represents the reward prediction error

Eric E. J. DeWitt
Paul Glimcher
New York University

EDEWITT@NYU.EDU
GLIMCHER@CNS.NYU.EDU

We show that the BOLD response in the nucleus accumbens (NAcc) quantitatively represents the reward prediction error (RPE) defined in standard models of reinforcement learning. Behavioral measurements of individual subjects engaged in a reinforcement learning task and matched neural measurements in the NAcc allowed us to identify a psychometric-neurometric match between the behavioral and neural learning functions. Subjects made a series of choices between two lotteries in each of a series of behavioral and fMRI imaging sessions. At the beginning of each session, subjects were given a cash endowment (\$25 behavioral sessions; \$45 fMRI). Each binomial lottery in the offered pair had the same prizes (+/- \$2) but the probabilities of winning and losing were different, and randomly drawn from a fixed set. On each trial, the subject was cued to make a choice between the two lotteries on offer and then, after a variable delay, the lottery was played, the dollar outcome revealed and the gain or loss added to the subject's cumulative earnings. To limit the influence of temporal uncertainty on the expected reinforcement, a visual 'clock' counted down to the precise time at which the outcome of the selected lottery was revealed. After each trial, there was a constant hazard that the probabilities for the chosen lottery were redrawn from the fixed set. At the end of each session their accumulated earnings were paid to the subject in cash or subtracted from their endowment. We performed a logistic regression to characterized the subject's

choice probability as a function of the history of reinforcement. Adapting the analysis used in Bayer and Glimcher (2004), we also performed a linear regression that predicted the BOLD response in the anatomically defined NAcc for each subject as a function of the history of reinforcement. By comparing these two reinforcement history weighting functions we could perform a psychometric-neurometric match that independently related the history of reward to neural activity and choice. We observed a psychometric-neurometric match between these neural and behavioral functions at two points in the trial: immediately prior to choice where the neural and behavioral functions reflect the expected reward and at the time of reinforcement where the function reflects the difference between expected and obtained reward—both signals of the reward prediction error employed in temporal difference RL models. We conclude that the human nucleus accumbens quantitatively reflects the calculation of a RPE signal of the kind proposed in temporal difference RL models.

doi:

I-44. Internal time temporal difference model of neural valuation

Sivaramakrishnan Kaveri^{1,2}
Hiroyuki Nakahara^{1,2}

SIVA@BRAIN.RIKEN.JP
HN@BRAIN.RIKEN.JP

¹Lab for Int Theor Neurosci, RIKEN BSI

²Dpt Comp Intell & Sys Sci, Tokyo Inst of Tech

The temporal difference (TD) learning framework has become a major paradigm for the understanding of value-based decision making and related neural activity (e.g., dopamine activity). Most of the current TD models use assumptions of tapped-delay-line formalism to represent experimental time. This formalism essentially avoids examining how time is processed in the neural valuation process that is modeled by those TD models, and thus time representation in the neural valuation process remains poorly understood. We propose a TD formulation that separates the time of the observer (experiment) and the operator (neural valuation process or TD model), which we call conventional and internal time, respectively. We describe the formulation and theoretical characteristics of a TD model using internal time for valuation, called internal-time TD. The internal-time TD framework allows us to better understand the time representation of the neural valuation process, separately from the timing of external events, and its possible consequences in neural value-based decision making. The TD model's computations are characterized at both short and long time scales. At long time scales, the discounted value of rewards is inversely proportional to their delays although the exact form of this relationship is debatable. An internal-time TD value function, with non-linear mapping between internal time and conventional time, exhibits the co-appearance of exponential and hyperbolic discounting at different delays and preference reversals observed in inter-temporal choice tasks. Further, we demonstrate that an internal-time TD, composed of multiple parallel neural systems, can produce the behavioral choices as suggested by studies of inter-temporal choice (McClure et al 2003, 2007; Tanaka 2007). At short time scales, TD computations such as TD error are expressed differently, depending on the time frame and time unit. Internal time TD allows us to examine this operator-observer problem in relation to the time representation used by previous TD models. We examine the dynamic construction of internal time and its effect on TD computations. We find that the temporal uncertainty of a fixed delay reward increases with the delay as reported in Fiorillo & Schultz [2008] and Kobayashi & Schultz [2008], due to the effect of noise in internal time. Internal time modulation can lead to decreased or increased valuation of rewards, thereby we suggest an internal time hypothesis for serotonin functions. Finally, internal time is shown to be formulated as a function of the probability distribution of rewards over time. On the basis of this view, internal time modulation may account for the instrumental response distributions of multiple rewards observed in interval-timing studies.

doi:

I-45. Does one simulate the other's value-based decision making by using the neural systems for his own?

Shinsuke Suzuki¹
 Norihiro Harasawa¹
 Kenichi Ueno²
 Sivaramakrishnan Kaveri^{1,3}
 Justin Gardner⁴
 Noritaka Ichinohe⁵
 Masahiko Haruno⁶
 Kang Cheng^{2,7}
 Hiroyuki Nakahara^{1,3}

SHINSUKE@BRAIN.RIKEN.JP
 HARASAWA@BRAIN.RIKEN.JP
 UKEN@RIKEN.JP
 SHINSUKE@ADDS.CO.JP
 JUSTIN@BRAIN.RIKEN.JP
 NORITAKA@CC.HIROSAKI-U.AC.JP
 MHARUNO@LAB.TAMAGAWA.AC.JP
 KCHENG@POSTMAN.RIKEN.JP
 HIRO@BRAIN.RIKEN.JP

¹Lab for Int Theor Neurosci, Riken BSI

²fMRI Support Unit, RIKEN BSI

³Dpt Comp Intell & Sys Sci, Tokyo Inst of Tech

⁴Gardner Research Unit, RIKEN BSI

⁵Dept Neuroanatomy, Hirosaki Univ

⁶Brain Science Institute, Tamagawa Univ

⁷Lab for Cognitive Brain Mapping, RIKEN BSI

In social contexts, another person's behavior often affects the outcome of one's value-based decision making which involves valuation and learning, i.e., reinforcement learning. So predicting the other's decision making is indispensable. In a broader context, it is often said in the field of social cognitive neuroscience that prediction about the other is made by 'simulating the other'; this simulation is often posited as using some systems like the so-called "mirror neuron system" that has been implicated in both one's own actions/sensations/emotions and one's perceptions of those by the other. For social value-based decision making, this could be translated to the question of whether one simulates the other's value-based decision making by using the same systems for one's own decision making or not. Thus, we asked (a) whether one actually simulates both valuation and learning of the other; and (b) whether this simulation uses the same neural mechanisms used for one's own decision making. To address these issues, we conducted an fMRI experiment with model-based analysis, using two tasks: an instrumental learning task and a "predict-other" task in which subjects predict choices of another person who plays the first task. These tasks allow us to directly compare corresponding neural regions between the subject's decision making in the first task and the subject's simulation of the other's decision making in the second task. In the instrumental task, subjects have to learn the reward probability for each option (learning) and compute an action-value (valuation) to make optimal choices. We confirmed that the subjects' behavior was fitted well with a computational, reinforcement learning model. In the predict-other task, subjects should simulate the valuation and learning of the other to make correct predictions. To address if the subjects actually did so, we fitted the subject's behavior with a computational model ("simulation model") that simulates the other's learning and valuation. We compared the goodness of fit of the simulation model with two other models which involved only simulation of the other's valuation or no simulation at all. We found that the simulation model accounted well for the subjects' behavior and provided a better fit than the other two models. Using model-based fMRI, we investigated the relationship (e.g., overlap or separation) between neural correlates of the subject's own decision making and those of the subject's simulation of the other's decision making. Our preliminary results indicate that several brain regions are correlated with the reward prediction error (PE) for both the subject and for the other, as might be expected for a PE-mirror neuron systems. We also found other regions correlated only with the other's PE, but not the subject's PE. These results suggest that one actually simulates both valuation and learning of the other and although preliminary, further that the same neural systems for one's reward prediction (i.e. PE) are used for simulating the other, while additional systems are also involved.

doi:

I-46. Prefrontal neurons solve the temporal credit assignment problem during reinforcement learning

Wael F. Asaad
Emad N. Eskandar

Massachusetts General Hospital

WFASAAD@ALUM.MIT.EDU
EESKANDAR@PARTNERS.ORG

Animals can learn new behaviors by forming associations between stimuli and actions under the guidance of appropriate reinforcement. The association of stimuli, actions and reinforcement is relatively straightforward when they overlap temporally. However, a dilemma is posed by reinforcers that are delayed with respect to their antecedent causes. This problem arises because the assignment of credit for a particular reward to a preceding event is less certain; many stimuli and several actions may have led up to that reward, so which is responsible? How does a learning system attribute these rewards to the causal event? This is the temporal credit assignment problem. To solve this problem, mechanisms must be in-place to represent information about the relevant preceding events at the time of reinforcement. Therefore, we designed a task that created a temporal credit assignment problem, and sought to determine if, during reinforcement, neurons in the lateral prefrontal cortex (PFC) could selectively represent an earlier, reward-predicting stimulus. Monkeys learned, by trial and error, which of four simultaneously-presented cues were associated with later reward. They indicated a choice by executing a saccade, after a blank delay, to the former location of one of those cues. If the correct cue had indeed appeared there, a generic reinforcer (a green circle) signaled a correct selection, followed by reward. Critically, the green circle did not reveal which cue was correct. Thus, for learning to take place, the occurrence of this feedback must be linked in some fashion to the one particular cue, out of the four shown earlier, that predicted it. Because cue arrangement varied randomly on every trial, the location of the response could not predict reward; rather, this was an "object learning" task. As a control, we interleaved a "spatial learning" task in which reward was determined by spatial location, irrespective of which cue had appeared there. This task was identical in all sensory and motor respects to the object task, but differed only in the rule. Here, there was no temporal credit assignment problem because the reinforcement overlapped in time with the predictive feature, the selected location. We found that 1) individual PFC neurons selectively represented the correct object at the time of reinforcement, and over the entire, unscreened population of PFC neurons there was more feedback-period object selectivity in the object-learning than in the spatial-learning task; 2) some neurons had feedback-period object-selectivity even if they had no visually-evoked response at the time of actual cue presentation, earlier in the trial; 3) in neurons that had both cue-period object selectivity and feedback-period object-selectivity, the rank-ordering of this selectivity was rarely (only 11% of the time) identical. Therefore, PFC neurons did indeed actively generate a representation of the relevant stimulus during reinforcement, specifically when there was a temporal credit assignment problem. However, the neuronal re-representation of this information during feedback differed from that observed during the actual presentation of the rewarded stimulus, arguing that population activity states - not individual neurons as "feature detectors" - are the objects of reinforcement.

[doi:](#)

I-47. Changes in the response rate and response variability of area V4 neurons during saccade preparation

Nicholas Steinmetz
Tirin Moore
Stanford University

NICK.STEINMETZ@GMAIL.COM
TIRIN@STANFORD.EDU

The visually driven responses of macaque area V4 neurons are modulated during the preparation of saccadic eye movements, but the relationship between presaccadic modulation in area V4 and saccade preparation is poorly understood. Recent neurophysiological studies suggest that the variability across trials of spiking responses provides a more reliable signature of motor preparation than mean firing rate across trials. We compared the

dynamics of the response rate and the variability in the rate across trials for area V4 neurons during the preparation of visually guided saccades. As in previous reports, we found that the mean firing rate of V4 neurons was enhanced when saccades were prepared to stimuli within a neuron's receptive field (RF) in comparison with saccades to a non-RF location. Further, we found robust decreases in response variability prior to saccades and found that these decreases predicted saccadic reaction times for saccades both to RF and non-RF stimuli. Importantly, response variability predicted reaction time whether or not there were any accompanying changes in mean firing rate. In addition to predicting saccade direction, the mean firing rate could also predict reaction time, but only for saccades directed to the RF stimuli. These results demonstrate that response variability of area V4 neurons, like mean response rate, provides a signature of saccade preparation. However, the two signatures reflect complementary aspects of that preparation.

doi:

I-48. Neurophysiological evidence for basal ganglia involvement in speed-accuracy tradeoff in monkeys

Masayuki Watanabe¹
Thomas Trappenberg²
Douglas Munoz¹

MASAYUKI@BIOMED.QUEENSU.CA
TRAPPENBERG@GMAIL.COM
DOUG_MUNOZ@BIOMED.QUEENSU.CA

¹Queen's University
²Dalhousie University

Speed accuracy tradeoff (SAT) is one of the fundamental behavioral phenomena observed in a wide variety of creatures ranging from insects to humans; time to perform a behavioral response can be shortened with the cost of its accuracy (Chittka et al. 2009, Trends in Ecology and Evolution 24: 400). Several cognitive models suggest intuitive mechanisms to control SAT (Bogacz et al. 2006, Psychol. Rev. 113: 700). However, despite the rich knowledge from behavioral analyses, neural mechanisms underlying SAT still remain unclear. Based on recent neuroimaging and computational studies in humans, it has been hypothesized that the basal ganglia (BG) are the key structures that control SAT (Forstmann et al. 2008, PNAS 105: 17538). In this study, we addressed this hypothesis more directly by neurophysiological experiments in behaving monkeys. We trained two monkeys to perform the antisaccade task; in response to the appearance of a peripheral visual stimulus, the monkeys were required to generate a saccade to the opposite direction of the stimulus. This simple requirement dissociates the following two theoretical saccade commands: an automatic saccade toward the stimulus and a volitional saccade to the opposite direction of the stimulus. This dissociation results in a conflict between automatic and volitional saccade commands, and monkeys as well as humans generate erroneous saccades to the stimulus occasionally (Munoz and Everling 2004, Nature Rev. Neurosci. 5: 218). We manipulated SAT by controlling visual fixation before peripheral stimulus appearance. It has been shown that saccade reaction times are shorter when a central fixation point, on which subjects maintain eyes, disappears before stimulus appearance (gap condition) than when the fixation point remains visible (overlap condition). In the antisaccade task, this manipulation also influences error rates; higher error rates in the gap condition than the overlap condition (Fisher and Weber 1993, Behavioral and Brain Science 16: 553). We carried out single neuron recordings in the caudate nucleus, the main input stage of the oculomotor BG (Hikosaka et al. 2000, Physiol. Rev. 80: 953). We have shown previously that the antisaccade task can dissociate caudate neurons encoding automatic saccades predominantly (automatic neurons) and those encoding volitional saccades (volitional neurons) (Watanabe and Munoz 2009, Eur. J. Neurosci. in press). Here, we found that both automatic and volitional neurons increased baseline activity before stimulus appearance in the gap condition more than the overlap condition, consistent with shorter reaction times in the gap condition. However, signals issued by automatic and volitional neurons before saccade initiation were less than ideal for antisaccade control in the gap condition compared to the overlap condition; automatic neurons issued stronger erroneous saccade commands while volitional neurons issued weaker correct antisaccade commands in the gap condition than the overlap condition. We suggest that signals issued by the caudate nucleus constrain behavioral outcome that obeys SAT. We are currently developing a computational model that integrates winner-take-all circuits in the superior colliculus, the key structure for saccadic decision (Trappenberg et al. 2001, J. Cog. Neurosci.

13: 256), and dynamic BG output estimated based on our experimental findings.

[doi:](#)

I-49. Modelling basal ganglia and superior colliculus in the antisaccade task

Thomas Trappenberg¹

Masayuki Watanabe²

Douglas Munoz²

TRAPPENBERG@GMAIL.COM

MASAYUKI@BIOMED.QUEENSU.CA

DOUG_MUNOZ@BIOMED.QUEENSU.CA

¹Dalhousie University

²Queen's University

A major challenge for agents in a natural environment is responding to a multitude of stimuli with often conflicting demands on a response system. Many recent studies have addressed such decision processes with competitive integrator models that accumulate evidence to aid appropriate responses [M. Usher and J.L. McClelland. *Psychological Review*, 108, 550-592, 2001]. A good example is visual orienting. Since we can only look at a specific direction at any moment in time, a decision system must be able to make sensible choices. An important component of this decision system is the superior colliculus (SC). We have previously modelled this area with a dynamic version of a winner-takes-all (WTA) mechanism based on physiological evidence, and we have shown that this model can explain common decision patterns in a variety of behavioral paradigms [Trappenberg et al, 2001]. This included the antisaccade task that requires subjects to orient away from a visual stimulus. This simple requirement dissociates the following two theoretical saccade commands: an automatic saccade command toward the stimulus and a volitional saccade command away from the stimulus. The automatic and volitional saccade commands stimulate SC neurons encoding opposite saccade directions, which causes a response conflict. Here we analyze the consequences of such decision processes on the frequency of trials with erroneous saccades toward the stimulus. We propose a new model of the SC based on our previous model. While this model simplifies the dynamic of the SC, it captures the competition between automatic and volitional saccades in a saccade likelihood map. We investigate how the percentage of error trials varies when modulating the relative strength of inputs encoding automatic and volitional saccade commands. We then incorporate a recent model of the basal ganglia (BG) based on our single neuron recordings in monkey caudate nucleus, the input stage of the oculomotor BG [M. Watanabe and D.P. Munoz, *European Journal of Neuroscience*, in press, 2009]. We have identified the following three types of caudate neurons; (1) automatic neurons (ANs) encoding contralateral automatic saccade commands, (2) volitional neurons (VNs) encoding contralateral voluntary saccade commands, and (3) ipsi neurons (INs) encoding ipsilateral volitional saccade commands, which could be ideal to suppress erroneous automatic saccade commands. We hypothesized that ANs and VNs give rise to the direct pathway facilitating saccade initiation while INs give rise to the indirect pathway suppressing saccade initiation. Here we implement this BG model and integrate it with the basic SC model. We found that the BG model can facilitate the necessary modulations of the SC to produce less error saccades, in agreement with our behavioral findings. We also found that the aid of the BG can produce behavioural ranges of error saccades when the volitional signals are weaker than automatic signals, in agreement with physiological evidence [S. Everling and D.P. Munoz, *The Journal of Neuroscience*, January 1, 2000, 20(1):387-400]. We are currently studying how the necessary pathways in the BG can be learned through reinforcement learning.

[doi:](#)

I-50. Integration of visual and proprioceptive information for reaching in multiple parietal areas.

Leah McGuire
Phillip Sabes

LMCGUIRE@PHY.UCSF.EDU
SABES@PHY.UCSF.EDU

University of California San Francisco

When reaching to swat a mosquito on one's arm, a person may rely exclusively on somatosensory information to target the movement, or may also look at the arm to gather additional sensory information and improve his chances of squashing the mosquito. Many psychophysical studies show that humans integrate information from multiple sensory modalities in a statistically optimal fashion, thereby minimizing variability. However little is known about how this integration is carried out in the brain. The purpose of this study was to elucidate the neural mechanisms of integration by studying cortical responses in two rhesus macaques during reaches to an array of visual (VIS), proprioceptive (ipsilateral hand, PROP) or visual and proprioceptive (VIS+PROP) targets. Both monkeys showed reduced reach endpoint variability in the VIS+PROP task compared to the unimodal VIS or PROP tasks. Reduced variability is a behavioral hallmark of sensory integration, and neural models of integration suggest that it is achieved through enhanced neural responses to bimodal stimuli. Parietal Area 5 shows enhancement of responses to static proprioceptive position signals with vision of a realistic monkey arm (Graziano et. al. 2000), supporting the idea that the posterior parietal cortex plays an important role in sensory integration. We set out to characterize the integration of visual and proprioceptive signals in this region during reach planning. We recorded from Area 5 as well as the nearby areas MIP and Area 7, which are involved in reach planning. We found that many cells responded during multiple tasks with similar spatial tuning across tasks, though the degree of modulation was typically task-dependent. This finding contrasts with observations in other cortical areas (e.g., MST, VIP) and allows us to assess how response magnitude changes during integration with minimal complications due to differences in spatial tuning across tasks. Unimodal task preferences were heterogeneous in Area 5 and MIP. We first compared plan-related (delay period) activity of each neuron in the bimodal task to its best unimodal response. In Area 5 and MIP, preferred target responses displayed both enhancement and suppression (similar to other cortical studies of integration), with few of the additive or super-additive effects seen in superior colliculus (Stein and Meredith 1983). In contrast, the least-preferred target responses showed suppression of activity. These findings may be explained by response normalization (Ma and Pouget 2008) or by sharpening of tuning curves in the bimodal condition, but are inconsistent with simple additive or enhancement models of integration. Further, neurons in Area 7 showed suppression across all target locations during the delay period and neurons in Area 5 and MIP showed similar suppression during the movement period. These results challenge a simple additive or enhancement model of sensory integration. An understanding of the mechanism of sensory integration must account for the suppressive effects that we and others have observed, as well as the heterogeneity of response patterns across areas and across time. In particular, these results suggest that the process of integration may be better understood at level of the larger cortical circuit.

[doi:](#)

I-51. Sensory integration in PMd: position-dependent dynamic reweighting of vision and proprioception

Matthew Fellows
Philip Sabes

MRFELLOWS@PHY.UCSF.EDU
SABES@PHY.UCSF.EDU

UCSF, Dept. of Physiology and Keck Center

It has been observed behaviorally that human and non-human primates combine information from multiple sensory modalities. In many situations this combination has been found to be statistically near-optimal, meaning that the multi-modal (i.e., integrated) percept is formed by weighting each modality in proportion to its precision. However, sensory modalities are not represented homogeneously across cortex - for a given modality, both the

strength of sensory representation and its functional role vary from area to area. How, then, do disparate local cortical representations manifest as globally optimal behavior? To begin addressing this question, we quantified – both behaviorally and in neurons – the integration of visual and proprioceptive information about hand location across a range of arm positions. We recorded simultaneous activity from large ensembles of dorsal premotor cortex (PMd) neurons in a monkey performing “out-to-center” reaching movements from multiple start locations to a central target, which he was required to fixate. We provided artificial visual feedback (“FB”, a cursor) about the animal’s otherwise unseen hand location. This allowed us to impose a discrepancy (“FB shift”) between the visual and proprioceptive feedback. We used FB shifts in different directions to measure how the animal weighted each modality when estimating hand location for reach planning. Optimality arguments predict that the integrated estimate of hand location should depend on the relative uncertainty of vision and proprioception. In general, these uncertainties are anisotropic and vary over space. That is, for hand localization, proprioceptive uncertainty changes with hand position, and visual uncertainty depends on hand position relative to gaze location. Therefore, due to the anisotropies, FB shifts in different directions should induce localization errors that differ in both magnitude and direction. In addition, due to the position-dependence, this pattern of shift-induced errors should change with start location. We quantified these error patterns both from the animal’s behavior and from population activity in PMd, and compared them. Behavioral analysis showed that the animal flexibly integrates vision and proprioception, with the weighting of each modality varying as a function of arm configuration: the observed error patterns showed systematic, position-dependent anisotropies. We observed flexible sensory reweighting in the PMd population activity as well, with error patterns that were similar to behavior in both orientation and position-dependence. However, the dependence on vision was greater for PMd than for behavior: for a given FB shift, the integrated hand location read-out from PMd is closer to the cursor than the location read-out from behavior. The spatial dependence we observed supports the idea that multi-modal integration, as measured with PMd activity and with behavior depends, at least in part, on the (anisotropic) variability of the sensory modalities. However, the greater dependence on vision in PMd suggests that sensory integration varies across the cortical reach circuit. For example, we expect that integration in primary motor cortex would look more like behavior than does PMd. This highlights the question of how diverse local cortical circuits are used in concert to drive overall behavior.

[doi:](#)

I-52. A new notion of criticality: Studies in the pheromone system of the moth

Christopher L. Buckley

Thomas Nowotny

CCNR, Informatics, University of Sussex

C.L.BUCKLEY@SUSSEX.AC.UK

T.NOWOTNY@SUSSEX.AC.UK

The concept that the nervous system operates close to a critical point in its dynamics has been receiving a growing amount of attention in the last few years [1,2,3,4]. There are a plethora of computational models that posit critical dynamics as a partial but parsimonious explanation of information transmission [2], storage [5] and computational [6] properties of the nervous system. To date, the majority of computational models of critical brain dynamics have focused on excitatory networks of excitable spiking neurons with sparse activity [3]. This is, perhaps, because such models are thought to be adequate models of the cerebral cortex, [2,3] but also because it has allowed researcher to draw formal comparisons with second order phase transition in ‘stick slip’ and Ising models. It is well-known that in many brain areas activity is heavily modulated, if not dominated by inhibitory interneurons [7]. For example, we have been investigating a subsystem of the antennal lobe of the moth which is characterised by dominantly inhibitory synapses and high baseline spike rates. It is not clear what relevance the established notion of critical brain dynamics has for such neural subsystems that do not fit into the framework of excitatory networks with sparse activity. In this work we develop, and formally describe, an alternate notion of criticality in terms of the rate dynamics of a network that can be applied to this different regime. In particular, we focus on the macro glomerular complex (MGC) of the moth which plays a key role in pheromone processing. The MGC comprises a set of recurrently connected GABA_B inhibitory cells which have baseline firing rates of around 20Hz. We model this system as a set of Hodgkin Huxley neurons with spike frequency adaptation connected via first order (“ $\alpha\beta$ ”) synapses [8]. Leveraging the fact that GABA_B synapses act at a much slower timescale than the membrane dynamics we are able to reduce this conductance based model to a formally equivalent rate

model. This allows us to analyse the nonlinear dynamics of the system in detail. Specifically, we describe the rate response of the MGC to pheromone as transient excursions from a globally asymptotically stable fixed point attractor. We show how this dynamics can account for several of the phenomena observed in the MGC and discuss the relationship of this model to reservoir computing paradigms [9]. In particular we formally show how critical rate dynamics allows sensitive response to inputs while maximizing the dynamic range of the MGC model to inputs of different amplitude. We argue that critical rate dynamics is more general and more widely applicable than current notions of criticality in neuroscience. [1] Chialvo, Balenzuela and Fraiman 2008 [2] Beggs and Plenz 2004 [3] Levina, Herrmann and Geisel 2007 [4] Kinouchi and Copelli 2006 [5] Halderman and Beggs 2005 [6] Bertschinger and Natschlager 2004 [7] Buzsáki, Kaila and Raichle 2007 [8] Destexhe, Mainen, Sejnowski 1994 [9] Jaeger 2001

[doi:](#)

I-53. Cellular imaging in behaving mice reveals learning-related specificity in motor cortex circuits

Takaki Komiyama¹
Takashi R. Sato¹
Daniel H. O'Connor¹
Ying-Xin Huber¹
Daniel Hooks¹
Bryan M. Gabitto²
Mariano Svoboda¹
Karel

KOMIYAMAT@JANELIA.HHMI.ORG
SATOT@JANELIA.HHMI.ORG
OCONNORD@JANELIA.HHMI.ORG
HUBERD@JANELIA.HHMI.ORG
HOOKSM@JANELIA.HHMI.ORG
MARIANO.GABITTO@GMAIL.COM
SVOBODAK@JANELIA.HHMI.ORG

¹Janelia Farm Research Campus, HHMI

²Columbia University

Cortical neurons are connected into highly specific neural networks, but the relationship between network dynamics and behavior is poorly understood. Two-photon calcium imaging can monitor activity of multiple, spatially defined cells in the mammalian cortex. Here we applied this technique to image activity in the motor cortex of mice performing a learned choice behavior. We developed an odor discrimination task in which head-fixed mice learned to lick in response to one of two odors and withhold licking for the other odor. Mice routinely learned this task within a single behavioral session. Microstimulation and transsynaptic tracing with pseudorabies virus identified two non-overlapping candidate tongue motor cortical areas. Imaging in layer 2/3 revealed neurons with diverse response types in both areas. Activity in approximately half of the imaged neurons distinguished trial types associated with different actions. Many neurons showed modulation coinciding with or preceding the action, consistent with their involvement in motor control; these neurons were more prevalent in the area identified by transsynaptic tracing. Neurons with different response types were spatially intermingled. However, nearby neurons (within ~150 micrometers) showed pronounced temporally coincident activity. These temporal correlations, which were apparent during task performance and inter-trial intervals, were particularly high for pairs of neurons with similar response types, and increased with learning specifically for similar response type pairs. We propose that correlated activity in specific ensembles of neurons is a signature of learning-related circuit plasticity underlying motor behavior. Our findings reveal a fine-scale and dynamic organization of the frontal cortex which likely underlies flexible behavior.

[doi:](#)

I-54. Neural mechanisms underlying the reduction in behavioral variability during trial-and-error learning

Alexis Dubreuil^{1,2}

Yoram Burak³

Timothy Otchy²

Bence Ölveczky²

ALEXIS.DUBREUIL@ENS-CACHAN.FR

YBURAK@FAS.HARVARD.EDU

TIM.OTCHY@GMAIL.COM

OLVECZKY@FAS.HARVARD.EDU

¹Ecole Normale Supérieure Cachan

²OEB and CBS, Harvard University

³Center for Brain Science, Harvard University

Motor exploration is essential for trial-and-error learning, yet as learning progresses motor variability is often reduced to yield a stereotyped performance. Here we explore the neural mechanisms underlying this reduction in behavioral variability and how it relates to the learning-driven maturation of motor circuits in the zebra finch, a songbird that shows decreased vocal variability with song learning. Song variability is driven by a basal ganglia circuit that projects to the motor cortex analogue brain region RA through nucleus LMAN [1]. In very young songbirds variable LMAN activity dominates the motor program [2], resulting in variable vocalizations, but as learning progresses, HVC, a premotor area providing the other main input to RA, takes over and, in adult birds, drives a precise and robust song through its stereotyped input to RA. The song learning process is thought to be driven by synaptic reorganization in the HVC-RA network, yet the variable LMAN input to RA remains intact also in adult birds: If HVC is lesioned, LMAN is capable of driving highly variable song, similar to what is observed in young birds [2] [3]. In this study, we examine the extent to which LMAN induced variability is reduced as a consequence of the strengthening and pruning of HVC-RA synapses. In our model of the HVC-RA-LMAN circuit, RA neurons are driven by temporally precise inputs from HVC and by random Poisson spike trains from LMAN. Maturation of HVC-RA synapses is modeled as a gradual shift from broadly distributed synaptic strengths to a more bimodal distribution, while synaptic inputs from LMAN to RA are kept fixed. Comparing our model with recordings from RA projection neurons made in juvenile zebra finches during singing, we can account for two distinct trends observed in the recordings: (1) Firing patterns gradually become more sparse and bursty, with stereotyped, higher firing rate events emerging. (2) Trial-to-trial variability of RA firing patterns is gradually reduced with age. Furthermore, our model predicts that at any given age LMAN's effect on variability in RA should depend on the instantaneous firing rate of RA bursts. Analysis of song-aligned spike trains from RA neurons recorded in zebra finches of different ages confirmed this prediction, showing a strong statistical relationship between firing rate and variability. In agreement with the model, this relationship depends only weakly on age, whereas the distribution of firing rates significantly evolves as learning progresses. Our results suggest a direct mechanistic link between the shaping and maturation of a learned motor program and the reduction in behavioral variability. References: [1] Ölveczky BP, Andalman AS, Fee MS (2005) Vocal Experimentation in the Juvenile Songbird Requires a Basal Ganglia Circuit. *PLoS Biol* 3(5): e153. [2] Aronov D, Andalman AS, Fee MS . 2008. A specialized forebrain circuit for vocal babbling in the juvenile songbird *Science* 320 630-634. [3] Thompson, J.A., Wu, W., Bertram, R. & Johnson, F. Auditory-dependent vocal recovery in adult male zebra finches is facilitated by lesion of a forebrain pathway that includes the basal ganglia. *J. Neurosci.* 27, 12308-12320 (2007).

[doi:](#)

I-55. Evidence for a central pattern generator built on a heteroclinic channel instead of a limit cycle

Kendrick M. Shaw
Hui Lu
Jeffery M. McManus
Miranda J. Cullins
Hillel J. Chiel
Peter J. Thomas

KMS15@CASE.EDU
HXL75@CASE.EDU
JMM45@CASE.EDU
MJC35@CASE.EDU
HJC@CASE.EDU
PJT9@CASE.EDU

Case Western Reserve University

To survive, an animal must generate patterns of behavior reproducibly enough to achieve its intended goal but flexibly enough to adapt to changing conditions. In many cases, this balance is achieved by coupling a central pattern generator (CPG) to feedback from the biomechanical system in which it is embedded. How this integration works in detail remains an open question. Limit cycles (isolated closed trajectories) are often used to model the dynamics of this type of neuromechanical system. Recently it has been suggested that another framework, stochastic stable heteroclinic channels, may provide an alternative class of models (Rabinovich et. al. 2006). In such a system an underlying deterministic dynamics is perturbed by superimposed noise. The deterministic dynamics contain a chain of saddle equilibria, that is points that repel the system along one direction and attract the system along others, connected by heteroclinic paths, which are trajectories connecting the equilibria. The noise may represent the variable effects of intrinsic influences (network or ion channel fluctuations) or external influences (environmental variability) or both. In addition to the noise we also consider deterministic perturbations such as proprioceptive feedback. We are interested in the utility of these two classes of models for understanding the behavior of a central pattern generator coupled to the periphery. To explore this interaction, we have constructed a highly simplified neuromechanical model of the feeding apparatus of the marine mollusk *Aplysia californica*. The model consists of a CPG whose behavior in isolation is described by a homoclinic orbit. This pattern generator can then be coupled to a simple biomechanical model via motor activation and proprioceptive feedback. In the absence of noise and proprioceptive input, the cycle times of the central pattern generator grow quickly with time. As has been previously described, small amounts of additive noise are able to rescue the cycling so that the average cycle time approaches a constant proportional to the log of the noise. In addition, we found that adding a small amount of proprioceptive input was also able to rescue the pattern generator. We observed empirically that the cycle time is again proportional to the log of the strength of proprioceptive input. With further increases in proprioception, the homoclinic cycle changes to a relatively smooth limit cycle. We have compared these results with EMG and ENG activity from intact behaving animals, reduced preparations, and isolated ganglia. Plots of the average activity during biting in the intact animal show limit-cycle like behavior. In reduced preparations the cycle begins to dwell longer in particular portions of the cycle, and these dwell times become even larger in the isolated ganglion. These changes are consistent with a model in which reducing the influence of proprioceptive feedback exposes an underlying central pattern generator built around a stable heteroclinic channel rather than a limit cycle.

[doi:](#)

I-56. A neural microcircuit using spike timing for novelty detection

Christopher R. Nolan
Gordon Wyeth
Michael Milford
Janet Wiles

CNOLAN@ITEE.UQ.EDU.AU
WYETH@ITEE.UQ.EDU.AU
MILFORD@ITEE.UQ.EDU.AU
JANETW@ITEE.UQ.EDU.AU

School of ITEE & Queensland Brain Institute, The University of Queensland

Place cells in the rodent hippocampus acquire their characteristic spatial selectivity during initial exploration of an unknown environment. Within a single environment, these cells develop fields that individually select only a

fraction of the space, but together cover the whole space, despite potentially sharing a large number of sensory cues. Once learned however, place fields can remain stable even during significant cue removal. That is, with respect to allothetic cues, the cells can be both highly discriminatory during learning yet robust to variation in a known environment. The aforementioned characteristics of place cells present a problem of competing requirements: how does an animal distinguish between a sensory stimulus that is 'close enough' to something it has experienced before, and one that is novel? Intuitively, it makes sense that the threshold for this distinction should vary based on the task at hand. Although the hippocampal CA3 network is often proposed as an autoassociative network performing pattern-completion on known inputs, these proposals have assumed that some innate ideal threshold is an intrinsic property of the network. We have developed a spiking neural microcircuit to explore methods of explicitly differentiating between novel and known inputs. In its canonical form, this circuit consists of two simple spiking cells connected via two routes: one monosynaptic route subject to spike-timing dependent plasticity and one disynaptic route through an auxiliary cell with fixed-weight synapses. We demonstrate that in this network, a race exists between activation of the target cell via the monosynaptic pathway and the activation of the same cell via the disynaptic pathway. The result of this race - the relative spike timing between the target cell and the auxiliary cell - signals the familiarity of the input. The structure of this microcircuit is similar to that of the entorhinal cortex (EC), the dentate gyrus (DG) and subregion CA3 of the hippocampus. We demonstrate that if each of the three cells in the microcircuit is instead a cell group, then (with the addition of some inhibitory interneurons) input patterns can be recognized as novel or known by virtue of relative spike timing between CA3 and DG on a pattern-by-pattern basis. We are currently testing this model with real-world data derived from video footage, which is processed into spiking units mimicking grid cells and cells tuned to respond preferentially to certain visual cues. Our goal is to test whether traversing through a previously learned environment will reactivate previously encoded representations and avoid acquiring new responses, mimicking the robustness and selectivity of real place cells.

doi:

I-57. Using natural stimuli to estimate receptive fields in neurons that employ sparse coding

Guy Isely¹
Christopher Hillar²
Fredrich Sommer²

GISELY@GMAIL.COM
CHILLAR@MSRI.ORG
FSOMMER@BERKELEY.EDU

¹Redwood Center for Theoretical Neuroscience, UC Berkeley

²University of California, Berkeley

Estimating receptive fields with natural stimuli has become an important approach for analyzing response properties of sensory neurons. It has been noted that natural stimuli contain correlations that have to be accounted for in analyzing receptive field properties. The appropriate correction has been worked out for the linear coding model [4] and there are methods to estimate receptive fields if the assumed coding model contains a pointwise nonlinearity [3]. Here we ask the question how one can estimate receptive fields if the underlying coding process is sparse coding, a model for sensory neurons that has been very successful in explaining the emergence of experimentally observed response properties [1, 2]. To explore this question we perform recording experiments on the sparse coding model, probing its activity with stimulus samples and determining the receptive fields of its neurons from stimulus response pairs. Naive reverse correlation with natural stimulus probes finds receptive fields that are similar in structure to the feed forward connections in the model. However, correcting for the autocorrelations in these probes corrupts that structure with high frequency noise. The linear generative mathematics inherent in the sparse coding model can explain why naive reverse correlation without the correction for stimulus autocorrelations provides a better estimate of the receptive field structure of sparse coding neurons. In particular, a mathematical analysis of sparse coding networks suggests a correction for the autocorrelations in the responses of their neurons rather than a correction for stimulus autocorrelations. Since these response autocorrelations are relatively small for a sparse coding network sampled with a sufficiently large natural stimulus ensemble, uncorrected naive reverse correlation with natural stimulus probes is a good approximation. More generally, an implication of our

analysis is that the best linear predictor of a neuron's firing rate is potentially not the best estimate of the spatial structure of connections that drive its stimulus response. Thus, surprisingly, under the assumption that the neurons one records from are performing sparse coding or similar computations involving later inhibition, naive reverse correlation may provide a better estimate of the structure their receptive fields than seemingly more rigorous techniques that correct for stimulus autocorrelations. [1] Olhausen BA and Field DJ (1996) Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381: 607-609. [2] Rehn M, Sommer FT (2007) A network that uses few active neurones to code visual input predicts the diverse shapes of cortical receptive fields. *J. Computational Neuroscience* 22:135-146. [3] Sharpee T et al. (2008) On the importance of static nonlinearity in estimating spatiotemporal neural filters with natural stimuli. *J. Neurophysiology* 99: 2496-2509. [4] Theunissen FE et al. (2001) Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli. *Network* 12: 289-316.

doi:

I-58. Comparison of V1 receptive fields mapped with spikes and local field potentials

Felix Biessmann^{1,2}

Frank Meinecke¹

Anwesha Bhattacharyya³

Julia Veit³

Robert Kretz⁴

Klaus-Robert Müller⁵

Gregor Rainer³

FBIESSMA@TUEBINGEN.MPG.DE

MEINECKE@CS.TU-BERLIN.DE

ANWESHA.BHATTACHARYYA@TUEBINGEN.MPG.DE

JULIA.VEIT@UNIFR.CH

ROBERT.KRETZ@UNIFR.CH

KLAUS-ROBERT.MUELLER@TU-BERLIN.DE

GREGOR.RAINER@UNIFR.CH

¹TU Berlin, Dept. Machine Learning

²MPI Biological Cybernetics, Tuebingen

³Visual Cognition Lab, University of Fribourg

⁴Anatomy Unit, University of Fribourg

⁵TU Berlin

Extracellular neurophysiological recordings are typically separated in two frequency bands. Low frequency content, also called local field potentials (LFPs), reflect subthreshold integrative processes of a population of neurons. High frequency content, or multi-unit activity (MUA), contains the information conveyed by action potentials, or spikes. Spikes reflect neuronal output and are generally considered as the main currency of information in the brain. For a long time receptive field mapping methods have focused exclusively on spiking information, although some recent studies have begun to address spatial characteristics of LFP responses (Xing/Yeh/Shapley, 2009, *J Neurosci*). In order to compare the information about visual stimuli carried by the LFP signal and spiking activity we mapped receptive fields in primary visual cortex of the tree shrew using spike count and LFP time-series recorded at different cortical depths. We presented white noise checkerboard patterns and sparse noise patterns and computed the standard spike triggered average (STA) receptive fields. Moreover we extracted the LFP timeseries, in different frequency bands, and the spike histograms following each stimulus and computed receptive fields for each signal employing standard canonical correlation analysis (CCA) between stimulus and LFP and spike response, respectively. Receptive fields as estimated from LFP data have two main advantages over traditional STA estimates. For one, LFP receptive fields do not suffer from binning artefacts, in contrast to STA receptive fields. Besides, CCA allows for computing a temporal filter for the respective neural signal. Receptive fields estimated using spikes were very similar to those computed from LFP signals, also for LFP bands below 20Hz. In particular the spatial extent of receptive fields computed from LFPs was comparable to that of spikes, in line with previous studies reporting a small spatial focus of LFP selectivity (Katzner et al. 2009, *Neuron*, Xing/Yeh/Shapley, 2009, *J Neurosci*). The receptive field size of both LFP and spikes varied with cortical depth. In summary our results confirm that in early stages of the visual processing hierarchy LFP signals contain to a large extent the same information about the visual stimulus as the spiking activity. In line with the above mentioned studies on non-human primates our findings suggest that the spatial selectivity of LFP signals with respect to the

visual stimulus is comparable to that of spikes.

[doi:](#)

I-59. A novel method to estimate information transfer between two continuous signals of finite duration

Jouni Takalo
Irina Ignatova
Matti Weckström
Mikko Vähäsöyrinki
University of Oulu

JOUNI.TAKALO@OULU.FI
IRINA.IGNATOVA@OULU.FI
MATTI.WECKSTROM@OULU.FI
MIKKO.VAHASOYRINKI@OULU.FI

The rate of information transfer provides an objective and rigorous measure on how much information a system can convey to its output from its time-varying input. Instead of the general measure of information transfer, the information capacity estimates as originally formulated by the Shannon in 1948¹ are still the standard choice for neuroscience research. In fact, the Shannon's method has for a long time been the only practical choice to analyze information processing of continuous signaling, such that takes place in many sensory and inter-neurons as well as in the dendrites of the spiking neurons. However, estimation of the information capacity assumes that the input has Gaussian statistics, the system is linear and time-invariant and that the noise is Gaussian and additive. These conditions are rarely met in neurons under naturalistic stimulus conditions, which necessitate the development of new methods to obtain reliable estimates. Here, we introduce a novel method to estimate the information rate between two continuous signals of finite duration. This method only requires assumptions of stationarity and ergodicity, i.e. that the statistical properties of the signals do not change over the course of time. Based on a recursively formulated estimator, it uses two signal processing methods to reduce computational cost. First, as the processing in the neural systems introduces lags or latencies, the estimation is prioritized to the time delays between the input and output signals that carry most of the mutual dependencies. Secondly, it uses a linear dimensionality reduction to concentrate the analysis on the most significant features of the signals. The method is validated with simulated input-output data from noisy linear filtering process, where Shannon's method gives accurate estimates for benchmarking. We then apply it to a highly non-linear data set from insect photoreceptors recorded under naturalistic stimulus conditions. Reliable estimates are obtained from single input-output realizations of 45 s in duration. The information rates are shown to be ca. 50 % higher than those estimated by the Shannon's information capacity, which is not surprising considering that none of the necessary assumptions of the Shannon's method are met. The presented method provides a practical tool for estimating information rate between two continuous signals. Unlike the popular Shannon's method, it does not pose restrictions to the signal and noise statistics or to the linearity of the system under study. Although a method was recently introduced² that, in principle, shares the same generality, it has not been taken into common use possibly due to the strict requirements it poses for the experimental design. We conclude that, in addition to being widely usable in neuroscience applications, our method could be introduced to analysis of other biological systems, such as reliability of biochemical networks and regulatory power of a transcription factor in gene expression.

[doi:](#)

I-60. The firing irregularity as the firing characteristic orthogonal to the firing rate

Takeaki Shimokawa
Shigeru Shinomoto
Kyoto University

SHIMOKAWA@TON.SCPHYS.KYOTO-U.AC.JP
SHINOMOTO@SCPHYS.KYOTO-U.AC.JP

A single train of neuronal spikes may be characterized with not only the rate of the spike occurrence but also the irregularity of the inter-spike intervals. Recently, we reported that the degree of irregularity in neuronal firing is closely related to the function of the cortical area: neuronal firing is regular in motor areas, random in the visual areas, and bursty in the prefrontal area [1]. We also developed a Bayesian estimation method for estimating both the firing rate and the irregularity, moment by moment for a given sequence of spikes [2]. In spite of the importance of analyzing firing irregularity, there has been yet no principled consideration for the proper method of gauging the irregularity. Here we search for a metric of irregularity under an information-theoretical principle, requiring it to be orthogonal to the mean interspike interval (ISI). It is found that only the mean log ISI satisfies the orthogonality condition [3]. Under the given values of the firing rate and the irregularity, respectively defined by the inverse mean ISI and the mean log ISI, the distribution function that maximizes the entropy is the gamma distribution. We set the gamma distribution in a Bayesian method for estimating the rate and the irregularity. By applying the method to spike sequences derived from different ISI distributions, such as the log-normal and inverse-Gaussian distribution, we confirmed that the estimation method can extract the modulations of the firing rate and the irregularity reasonably well. [1] Shinomoto S., Kim H., Shimokawa T. et al (2009) Relating neuronal firing patterns to functional differentiation of cerebral cortex. *PLoS Comput Biol* 5, e1000433. [2] Shimokawa T. and Shinomoto S. (2009) Estimating instantaneous irregularity of neuronal firing. *Neural Comput* 21, 1931. [3] Shimokawa T., Koyama S., and Shinomoto S. (2009) A characterization of the time-rescaled gamma process as a model for spike trains. *J Comput Neurosci*, doi: 10.1007/s10827-009-0194-y.

doi:

I-61. Optimal information transfer in the cortex through synchronisation

Andres Buehlmann¹

Gustavo Deco²

ANDRES.BUHLMANN@UPF.EDU

GUSTAVO.DECO@UPF.EDU

¹Universitat Pompeu Fabra, Barcelona

²Universitat Pompeu Fabra and ICREA

Gamma band synchronisation has been found in many cortical areas and in a variety of tasks. Some authors have proposed that neuronal synchronisation drives the interactions among neuronal groups, a hypothesis that has been referred to as communication through coherence (CTC). Although several experimental studies have presented results supporting the CTC hypothesis, some questions remain. Is CTC restricted to the gamma band? What is the influence of the total gamma power in the signal? Do phase and power only correlate or is there a causal dependence between the two? Here, to address these questions we use a biophysical, conductance based model network with realistic spiking properties. Using a model has the advantage that we can generate more data than in an experiment, making it possible to use a better statistical measure for the mutual interaction than just rank correlation. Instead, we use transfer entropy (TE). The TE is an information theoretical measure with the advantage that it does not only measure the coherence between two signals, but is able to distinguish between driving and responding elements. The model we use in this study consists of integrate-and-fire neurons. One of two pools of excitatory neurons receives input (Poisson spike train) which it passes to a neighbouring pool, connected by feedforward and feedback connections. We show that the coherence as measured by the Spearman rank correlation coefficient depends on the phase relation in the gamma band, confirming the experimental finding of Womelsdorf et al. (2007). Secondly, after applying TE to measure the information exchange between two pools, we find that TE very similarly depends on the phase shift, i.e., that there is an optimal phase relation where the TE is maximal. Thirdly, we reveal such dependence also in the beta band. Fourthly, we demonstrate that the TE increases as a function of the power in the gamma band. Lastly, we show that the information exchange gets faster if the gamma band synchronisation increases. In sum, we provide support for the CTC hypothesis and make the prediction that CTC is a general mechanism, not restricted to the gamma band. References: Womelsdorf T, Schoffelen JM, Oostenveld R, Singer W, Desimone R, Engel AK, Fries P (2007) Modulation of neuronal interactions through neuronal synchronization. *Science* 316:1609-1612.

doi:

I-62. Two dimensions for the price of one: the efficient encoding of vertical disparity

Jenny Read

J.C.A.READ@NCL.AC.UK

Institute of Neuroscience, Newcastle University

Usually, in order to encode information about a stimulus property X , one needs a population of neurons tuned to a range of values of X . For example, in order to encode information about photons' wavelength, the retina has to contain cones tuned to long, medium and short wavelengths; people with only a single cone type are color-blind. Here, I present an interesting example of a situation where this does not quite hold. Binocular disparity is a two-dimensional quantity, with components H and V representing the horizontal and vertical differences between where an object projects to in the two eyes. Humans are sensitive to both components, using them to deduce information about 3D scene structure and object distance. Here, I consider a neuronal population of disparity sensors based on the highly successful stereo energy model. I show that a population tuned to a range of different H_{pref} , but all tuned to the same V_{pref} , can nevertheless encode values of V away from V_{pref} , including both their magnitude and sign. This is of interest because in natural images, the range of vertical disparities encountered at any point on the retina is typically very narrow, much less diverse than the range of horizontal disparities. One would therefore expect the brain to contain a much narrower range of vertical disparity tuning than of horizontal disparity tuning: $SD(V_{pref}) \ll SD(H_{pref})$. My results show that, in fact, $SD(V_{pref})$ can be reduced to right down to zero while still retaining information about V . Potentially, this would enable the brain to represent 2D disparity very efficiently, encoding 2D disparity with a purely 1D population (in the sense that the distribution of preferred disparities, (H_{pref}, V_{pref}) , lies along a line). This apparent paradox arises because for cells tuned to oblique orientations, the response surface $F(H, V)$ is inseparable. Thus, while the cell responds best to $V = V_{pref}$ when probed at its optimal horizontal disparity $H = H_{pref}$, at non-optimal horizontal disparities it responds best to vertical disparities on either side of V_{pref} . It has been suggested that these inseparable response surfaces are later converted to separable ones during the cortical processing of disparity, but the present work suggests that the initial, inseparable response may have a role to play in enabling an efficient encoding of 2D disparity.

doi:

I-63. Analysis of subsets of higher-order correlated neurons based on marginal correlation coordinates

Hideaki Shimazaki

Sonja Gruen

Shun-ichi Amari

RIKEN Brain Science Institute

SHIMAZAKI@BRAIN.RIKEN.JP

GRUEN@BRAIN.RIKEN.JP

AMARI@BRAIN.RIKEN.JP

Recent studies on multiple parallel neural spike data raised the question whether it is critical to include higher-order interactions to describe relevant aspects of experimentally found synchronous spike patterns [Schneidmann et al. 2006, Shlens et al. 2006, Montani et al. 2009, Roudi et al. 2009]. Up to now it is only partially disclosed what aspects of synchronous spiking activities are ignored by analyses concentrating on lower-order correlations only [Amari et al. 2003; Staude et al. in press]. To address this issue we describe a new coordinate system, the marginal correlation coordinates, consisting of the information-geometric measure of spike correlation introduced formerly [Amari 1985, Amari & Nagaoka 2000, Nakahara & Amari 2002, Amari 2009]. The coordinates are constructed by considering sequentially increasing subsets of r neurons of a full N neuron system ($r=1, \dots, N$). The set of the r -th order marginal correlations is obtained as the highest-order parameters of the log-linear models applied to subsets of size r . Due to this hierarchically marginal structure, already determined correlations are not altered by adding neurons to the analysis although the considered order of correlations is increased. We prove that the marginal correlation coordinates are hierarchically orthogonal as was also shown for the mixed coordinates. To illustrate the relation of the marginal correlations to the occurrence probabilities of spike patterns of various

orders we first study a 3 neuron system. In particular we examine the iso-clines of occurrence probabilities of singlet, doublet, and triplet spike patterns in the marginal correlation coordinates. We find (i) an increase of marginal pairwise correlations first increases the occurrences of doublets and triplets, but decreases singlets. A further increase of pairwise correlations, however, decreases doublets as they are replaced by triplets to realize strong pair correlations. (ii) Increasing triplewise correlation increases triplet and singlet patterns, but decreases doublets so that pair correlations induced by triplets are canceled. (iii) Changes in the firing rates lead to a strong change in the occurrence probabilities of the spike patterns, although the overall structure of the iso-clines is not significantly altered. Next, we extend the analysis to a larger system of 8 neurons. Presence of only the 8th-order parameter as an interaction term in a log-linear model introduces marginal correlations of lower orders (<8) with descending contributions the lower the order. If the contributions of lowest orders (2-3) are negligibly small, the occurrence of the 8-th order synchrony is necessarily sparse because its rate is bounded by the chance level synchrony of subsets of lower-order, in particular for a system of realistically low firing rates. These results indicate the necessity of the analysis of higher-order correlation for detecting assemblies in particular if they exhibit sparse excess synchronous spikes of higher order.

doi:

I-64. Spike latency code for orientation discrimination and estimation by primary visual cortical cells

Oren Shriki¹
Adam Kohn²
Maoz Shamir¹

OREN70@GMAIL.COM
ADAM.KOHN@EINSTEIN.YU.EDU
SHMAOZ@BGU.AC.IL

¹Ben-Gurion University

²Albert Einstein College of Medicine

Accumulating evidence from behavioral experiments and imaging studies shows we are able to discriminate between different visual objects at a remarkable speed. The high computation speed yields a strong constraint on any possible answer to the question: how is information from one brain region communicated to another? Conventional neural coding research has ignored the temporal structure of the neural response and focused, in many cases, on the neural mean response over long timescales that are often beyond the relevant behavioral timescale. It has been suggested that the temporal structure of the neural initial response and in particular response latency are used by the central nervous system for fast communication of information between different brain regions. However, the accuracy of such a scheme has not been analyzed rigorously. Here we addressed this question in the framework of orientation coding by primary visual (V1) cortical cells of the monkey. To this end, simultaneous recordings of multiple V1 neurons over many repetitions per each orientation were performed. The spike data from these recordings were used to investigate the utility of first spike latency for encoding information about the orientation of visual stimuli. Cells in V1 are known to code for the orientation of a grating stimulus by their rate of firing. Typically, V1 cells show a maximum firing rate in response to a 'preferred orientation'. We find that many cells in the monkey V1 also show tuning of their first spike time latency to the orientation of the stimulus. Most cells have the shortest latency at the preferred orientation of their rate tuning curve. Moreover, by transforming the latency tuning curve to units of rate we find that the two tuning curves have a very similar tuning width. Using various statistical measures, we quantified the performance of a highly nonlinear readout mechanism, which estimates stimulus orientation by the preferred orientation of the cell with the shortest first spike latency, the temporal-Winner-Take-All (tWTA). In the context of a two-alternative forced-choice paradigm, we find that the tWTA discrimination accuracy is comparable to that of a conventional rate-code readout, which takes into account the total number of spikes fired by the cell in response to the visual stimulus. The accuracy of the tWTA readout can be further increased by considering a generalized n-tWTA readout, which estimates the orientation as the preferred orientation of the cell which fired the first n spikes. We find that the stimulus orientation can be estimated by n-tWTA with a relatively small bias for $n \geq 2$. The study demonstrates that a readout based on spike latency, or the first few spikes fired, may significantly improve response speed at a small cost to the accuracy of the decision.

doi:

I-65. Dissecting the action potential onset rapidness on the response speed of neuronal populations

Wei Wei^{1,2}Fred Wolf¹

WEI@NLD.DS.MPG.DE

FRED-WL@NLD.DS.MPG.DE

¹MPI for Dynamics and Self-Organization²BCCN Goettingen

Neuronal populations can track varying signals through population coding. In cortical neurons, there are two fundamentally different effective input channels: the mean synaptic current and the amplitude of synaptic noise. The mean input increase maximally if the excitatory input increases while inhibition decreases. The amplitude of noise increases when excitatory and inhibitory input exhibit correlated changes as is a key property of the balanced state [1][2]. Experiments indicate that ensembles of cortical neurons behave like a low-passed filter with a high cut-off frequency and that the response speed for noise coded signals is faster than that for mean current coded signal [3][4][5]. It has been shown numerically that details of action potential (AP) generation mechanism of single neurons play an important role in the dynamical response of neuronal populations [6][7]. To clarify this relation, we constructed a new dynamic model of AP generation in which the onset rapidness r of AP initiation is a freely variable parameter and which is analytical solvable. This r -tau model reduces to the leaky integrate-and-fire model (LIF) for infinite r and to the perfect integrator model for zero r . For finite r the impact of dynamic AP generation for linear response becomes accessible to rigorous analysis. We found that the linear response decomposes into two parts: one part approaches zero when the absorbing boundary is moved to infinity, indicating an artifact of the model; the other part possesses only a weak dependence on the boundary and reproduces the results of LIF neurons for r approaching infinity. This part represents the dynamics of AP generation. When the onset rapidness is large, the cut-off frequency in population response for noise coded signal will be proportional to the onset rapidness, while for mean current coded signal it is constrained by the membrane time constant. Since the onset rapidness of APs was found experimentally to be very large [8] (see, however, [9] [10]), our model suggests an explanation why the response speed can be much faster for noise coded signal than for mean current coded signal. References: 1. van Vreeswijk CA and Sompolinsky H, Science 1996, 274:1724-1726. 2. Renart A et al, Frontiers in Systems Neuroscience 2009, Conference Abstract: Computational and systems neuroscience. Doi: 10.3389. 3. Silberberg G, Bethge M, Markram H, Pawelzik K, Tsodyks M, J Neurophysiol 2004, 91:704 -709. 4. Köndgen H, Geisler C, Fusi S, Wang XJ, Lüscher HR, Giugliano M, Cereb Cortex 2008, 18:2086 -2097. 5. Boucsein C, Tetzlaff T, Meier R, Aertsen A, Naundorf B, J Neurosci 2009, 29:1006 -1010. 6. Fourcaud-Trocmé N, Hansel D, van Vreeswijk C, Brunel N, J Neurosci 2003, 23:11628 -11640. 7. Naundorf B, Geisel T, Wolf F, J Comput Neurosci 2005, 18:297-309. 8. Naundorf B, Wolf F, Volgushev M, Nature 2006, 440:1060-1063. 9. McCormick DA, Shu Y, Yu Y, Nature 2007, 445, doi:10.1038/nature05523. 10. Naundorf B, Geisel T, Wolf F, Nature 2007, 445:E2-E3. :

doi:

I-66. Diversity of efficient coding solutions for a population of noisy linear neurons

Eizaburo Doi¹Liam Paninski²Eero P. Simoncelli¹

EDO@NYU.EDU

LIAM@STAT.COLUMBIA.EDU

EERO.SIMONCELLI@NYU.EDU

¹New York University²Columbia University

Efficient coding is a well-known principle for explaining early sensory transformations (Barlow, 1961). But even in the "classical" case of a linear neural population with Gaussian input and output noise, the optimal solution depends heavily on the choice of constraints that are imposed on the problem. These can include constraints on output capacity (which is necessary to prevent the solution from diverging) and number of neurons in the population. With the exception of (Campa et al., 1995), previous literature assumes the number of neurons is equal to the input dimensionality, and furthermore, that the receptive fields are identical (i.e., the population performs a convolution). In addition, previously published examples are based on a single capacity constraint, such as the variance (loosely analogous to average spike rate) of the outputs (Atick & Redlich, 1990, Atick, Li, & Redlich, 1990, van Hateren, 1992), or the norm (analogous to the sum of squared synaptic weights) of filters (Campa et al., 1995). Each of these potential constraints has some approximate mapping onto biologically relevant (i.e., metabolic) costs, implying that a complete formulation should include a generalized cost function that combines power, weights, and population size. Toward this end, we examine a more general formulation of the efficient coding problem. We assume a discrete and finite input signal that is Gaussian with known covariance structure (as in natural signals), corrupted by additive white Gaussian noise. We assume a neural population of arbitrary size with linear receptive fields (RFs), each of whose outputs is corrupted by additive white Gaussian noise. And finally, we assume a cost function that is a linear combination of the number of cells, the L2 norm of the RF weights, and the output power. Given these constraints, we solve for a population of RFs that maximizes the information transmitted about the stimulus. The problem is convex, and thus can be solved with standard optimization methods. We note several important attributes of this formulation. First, it can achieve both over- and under-complete solutions, as is required to explain biological systems. In the retina, for example, previous efforts have assumed a convolutional solution (homogenous population of RFs, one per input cone) (Atick & Redlich, 1990, Atick, Li, & Redlich, 1990, van Hateren, 1992), but the ratio of cones to ganglion cells varies dramatically with eccentricity. Second, the joint cost function allows the theory to automatically select the optimal population size. And finally, even in the case of a convolutional population, the spectral properties of the optimal solution depend critically on the choice of cost function, and can assume shapes ranging from low-pass, to band-pass, to high-pass. For example, although previous literature (Atick & Redlich 1990) obtained low-pass solutions in the case when the input SNR was low (i.e., low contrast stimuli), the inclusion of the penalty for RF weights allows the possibility of low-pass (and highly redundant) solutions even when the input SNR is high. We conclude that identifying the relative significance of different costs in biological systems is critical to testing the efficient coding hypothesis.

doi:

I-67. Orientation and direction selectivity in the population code of the visual thalamus

Garrett B. Stanley¹

Jianzhong Jin²

Yushi Wang²

Gaelle Desbordes³

Michael J. Black⁴

Jose-Manuel Alonso²

GARRETT.STANLEY@BME.GATECH.EDU

JJIN@MAIL.SUNYOPT.EDU

YUSHIWANG@SUNYOPT.EDU

DESBORDES@GMAIL.COM

BLACK@CS.BROWN.EDU

JALONSO@SUNYOPT.EDU

¹Biomedical Engineering, Georgia Tech/Emory University

²Suny College of Optometry

³Georgia Tech/Emory

⁴Brown University

Neurons in the visual thalamus respond to natural scenes by generating synchronous trains of spikes on the timescale of 10-20 ms (Butts et al., 2007; Desbordes et al., 2008) that are very effective at driving cortical targets (Alonso et al., 1996; Usrey et al., 2000; Roy & Alloway, 2001; Bruno & Sakmann, 2006; Kumbhani et al., 2007). Here we demonstrate that this synchronous activity contains unexpected rich information about fundamental properties of visual stimuli. We report that the synchronous activity of thalamic cells with overlapping receptive

fields can be sharply tuned for the orientation and the direction of motion of the visual stimulus. We show that this stimulus selectivity is robust, remains relatively unchanged under different contrasts, stimulus velocities and temporal integration windows and cannot be predicted from linear models. Finally, we demonstrate that the direction of motion of a visual scene can be decoded from very short observations of synchrony (on a single trial basis) within small groups of thalamic cells with highly overlapped receptive fields that are likely to converge at the same cortical target. Taken together, these findings suggest a novel population code in the synchronous firing of neurons in the early visual pathway that could be the building blocks for higher level representations of motion within the visual scene. This work was supported by NSF CRCNS Grant IIS-0904630 (GBS, MJB, JMA), NSF IIS-0534858 (MJB), the National Eye Institute (JMA), and the Research Foundation of the State University of New York (JMA).

doi:

I-68. Visual hyperacuity despite fixational eye movements: a network model

Ofer Mazor¹
Yoram Burak²
Markus Meister²

OMAZOR@FAS.HARVARD.EDU
YBURAK@FAS.HARVARD.EDU
MEISTER@FAS.HARVARD.EDU

¹Harvard University

²Center for Brain Science, Harvard University

The retina transmits a representation of the visual environment to the brain. Recent studies have focused on the many forms of information processing that occur within the retina. Yet much less is known about how the brain interprets the raw retinal output. Here we address this question by focusing on a visual acuity task with well-characterized behavioral performance. Human observers can resolve the separation of two parallel lines with a precision many times finer than the sampling resolution of the retina (Vernier hyperacuity). This occurs in the presence of constant involuntary body and eye movements that scan the visual image over the retina with a random trajectory. The image drifts faster than the output cells of the retina (retinal ganglion cells, RGCs) can respond, yet the visual system can extract high resolution stimulus information, in this case the line separation, from the RGC population response. What must the brain do to extract this information? Using a multi-electrode array, we measured the activity of a population of mouse RGCs in response to the presentation of two parallel lines. The effects of fixational movements were simulated by drifting the pair of lines, in unison, across the retina. We also simulated ganglion cell spike trains from the human eye, based on published response models from the primate retina. We then explored strategies for estimating the line separation from the RGC population response. An optimal decoder should treat the drift trajectory as a hidden variable to be estimated along with the stimulus. This approach results in a complex decoding algorithm with large memory requirements - an unlikely computation to attribute to the brain. As an alternative, we introduce a simpler decoding algorithm that does not track the drift trajectory but can still estimate the stimulus with high accuracy. This simplified decoder uses successive time windows of the RGC response to make independent estimates of stimulus likelihood that are accumulated to form a final estimate. We found that this scheme performs the Vernier task almost as well as a comparable decoder that does track the drift trajectory. Furthermore, its precision matches human performance: The algorithm can estimate line separation to less than one quarter of one RGC receptive field for presentations lasting only 100-200 ms. Finally, we propose a simple two-layer neural network, composed of coincidence detectors and temporal integrators, that implements this computation. We show that the network performs well when the implementation details - number of cells, connectivity, timescale of coincidence detection - are consistent with the known properties of visual cortex. Thus, by closely examining the output of the retina under natural conditions, we can make quantitative proposals about cortical processing that underlies human visual performance.

doi:

I-69. Sparse coding in modular networks

Eva Dyer
Don Johnson
Richard Baraniuk
Rice University

E.DYER@RICE.EDU
DHJ@RICE.EDU
RICHB@RICE.EDU

Simple cells spanning the input layer of the primary visual cortex (V1) are believed to form the basis from which the visual cortex constructs a representation of its visual environment. The sparse coding hypothesis suggests that the receptive fields of these cells have adapted over time to produce sparse population codes in response to natural scenes and other ecologically relevant visual stimuli. Recently, a neurally plausible mechanism for sparse coding was proposed that models the behavior of a population of simple cells and interneurons via recurrent inhibition and stimulus-driven excitation (1). In order to produce sparse codes that tradeoff the number of active neurons with the error incurred in the approximation, the strength of the connection between any two neurons within the network is determined by the coherence between their receptive fields. To find a sparse code that is globally optimal, the network must be nearly fully connected. We posit that the organization of cells into densely connected microcircuits or modules, suggests that the visual cortex may employ a coding strategy that produces locally optimal sparse representations instead of requiring denser connectivity to achieve the sparsest global solution. To extend the sparse coding framework in (1) to incorporate these constraints, we study a number of cost functions capable of promoting sparsity at the level of single cells and at the level of sub-populations. The networks that result from these objectives exhibit modular structure and when sufficiently overcomplete, exhibit small-world topologies. To show the viability of this new approach, we present a hierarchical model for sparse coding that describes the dynamics of collection of orientation minicolumns all encoding a small region in retinotopic space. In addition to preserving the fine-scale details contained in the activity of individual neurons, we show that our model is also capable of producing a coarse-scale representation that may be used in contour integration and in other tasks where a reduced representation is all that is required. References 1. C. Rozell, D. Johnson, R. Baraniuk, and B. Olshausen. Sparse coding via thresholding and local competition in neural circuits. *Neural Computation* 20:2526-2563, 2008.

[doi:](#)

I-70. Suppression of intrinsic cortical response variability is state- and stimulus-dependent

Benjamin White¹
Larry Abbott²
Jozsef Fiser³

BLWHITE@BRANDEIS.EDU
LFABBOTT@COLUMBIA.EDU
FISER@BRANDEIS.EDU

¹Program in Neuroscience, Brandeis University
²Department of Neuroscience, Columbia Univ.
³Department of Psychology, Brandeis University

Neural responses to identical sensory stimuli can be highly variable across trials, even in primary sensory areas of the cortex. This raises the question of how the brain reliably transmits sensory-evoked responses and guides appropriate behavior. Internally-generated, spontaneous activity is ubiquitous in the cortex, and is a leading candidate for causing much of the observed response variability. However, the interaction between spontaneous and evoked activity is still poorly understood. Recent theoretical analyses have suggested that the dynamic recurrent network of the cortex might operate in a regime bordering the chaotic domain where the fidelity of stimulus transmission is enhanced. It has also been proposed that a hallmark of this operation is a strong temporal frequency-dependent noise-suppression in response to sensory stimulation, which is determined by the properties of the network rather than the stimulus. To test these predictions, we investigated spontaneous and visually-evoked extracellular neural activity from 57 mostly multi-units (MUs) in the primary visual cortex (V1) of 6 rats. We recorded from the rats under 5 conditions: while fully awake and while under 4 different levels of isoflurane

anesthesia. The anesthetized conditions were included to investigate the responses of the neural circuitry as its dynamic behavior was gradually eliminated. Anesthesia ranged from very light to deep, and stable levels were verified by various physiological parameters such as breathing rate, reflex response, and local field potential structure. Rats were head-fixed in a sound- and light-attenuating box while passively viewing flashing stimuli on a monitor 6 inches away. Five different stimulus conditions were presented to all rats in all states. Full-field flashing visual stimuli were presented at four frequencies, from 1 Hz to 7.5 Hz, and spontaneous neural activity was also recorded during periods of complete darkness. Stimulus appearance order was interleaved and randomized. We found that variability, as quantified by the Fano-factor, in spontaneous neural firing is actively and selectively suppressed by visual stimulation both in the awake and anesthetized states. However, the pattern of suppression was different across states: in the awake case, it followed the theoretical prediction, showing a significant dip in the Fano-factor across the different temporal stimulus frequencies. This frequency-dependency vanished with increased anesthesia. In addition, we found that the lowest level of noise and the largest amount of suppression across all evoked conditions compared with the spontaneous condition occurred in the awake state. Importantly, power spectrum analysis showed that this pattern of frequency-dependent noise suppression could not be explained by differences in intrinsic neural oscillations. These results suggest that there exist an active noise suppression mechanism in the primary visual cortex of the awake animal which is tuned to optimally support the propagation and coding of signals across frequencies, and that this mechanism essentially disappears as the dynamics of the network are eliminated. Acknowledgement This work has been supported by the Swartz Foundation.

doi:

I-71. Background synaptic activity modulates spike train correlation

Ashok L. Kumar¹
Maurice J. Chacron²
Brent Doiron³

ALK@CMU.EDU
MAURICE.CHACRON@MCGILL.CA
BDOIRON@PITT.EDU

¹Center for the Neural Basis of Cognition, Carnegie Mellon

²Dept. of Physiology, McGill Univ.

³Dept. of Mathematics, Univ. of Pittsburgh

Sensory neurons have many synaptic inputs, not all of which directly transmit sensory information. Instead, some inputs serve to modulate the neuronal response to relevant stimuli. One example of neural modulation is that the intensity of background synaptic input determines the postsynaptic neuron's leak conductance and membrane potential variability. This in turn controls the gain of its firing rate response to an excitatory stimulus [1]. State dependent gain control in sensory neurons is necessary for proper stimulus processing in low and high contrast environments, and allows attentional state to modulate signal transfer. Most work that characterizes the modulatory influence of background activity has focused on single cell responses, but it is reasonable to expect that such activity may also modulate population responses. We study how background activity shapes correlations between pairs of neurons using simulations, a theory from non-equilibrium statistical mechanics, and data recorded from electrosensory pyramidal neurons in weakly electric fish. We find that increasing background activity can decrease the output spike train correlation of two neurons over long time scales but increase pairwise synchrony on short time scales. We study these effects in simulated integrate-and-fire neurons. Using a linear response theory [2,3], we find that changes in correlation are related to single cell response properties including gain and integration timescale. We first demonstrate these changes in correlation using pairs of model neurons receiving partially correlated conductance-based synaptic activity. The neuron pairs received low or high levels of background activity, and spike train statistics were compared for these two states. We next demonstrate similar state-dependent changes in a recurrent network of neurons in which correlations arise through coupling. Finally, we test our predictions on simultaneously recorded extracellular spike data from primary sensory nucleus in the electrosensory system of weakly electric fish. We induce a state change in the neuronal response properties by driving the animal with stimuli mimicking a prey or a communication signal [4]. Consistent with our theoretical findings, the timescale and correlation of recorded neuron pairs vary depending on these processing states. The findings connect state-dependent changes in single cell response properties, such as firing rate gain, to changes

in correlated activity in a neuronal population. [1] Chance, F. S., Abbott, L. F. & Reyes, A. D. Gain modulation from background synaptic input. *Neuron* 35, 773-782 (2002). [2] de la Rocha, J., Doiron, B., Shea-Brown, E., Josic, K. & Reyes, A. Correlation between neural spike trains increases with firing rate. *Nature* 448, 802-806 (2007). [3] Richardson, M.J.E. Spike-train spectra and network response functions for non-linear integrate-and-fire neurons. *Biological Cybernetics* 99, 381-392 (2008). [4] Chacron, M.J., & Bastian, J.. Population coding by electrosensory neurons. *J. Neurophysiol.* 99, 1825-1835 (2008).

[doi:](#)

I-72. Dynamic population coding with recurrent networks of integrate and fire neurons

Martin Boerlin
Sophie Deneve

MARTIN.BOERLIN@ENS.FR
SOPHIE.DENEVE@ENS.FR

Group for Neural Theory, LNC, DEC, ENS Paris

Our nervous system is capable of representing and computing with sensory and motor variables that are ambiguous and changing over time. These variables are encoded by populations of spiking neurons whose activity is highly variable and weakly correlated. Several studies have used rate-based models to describe probabilistic computation with static stimuli (e.g. Ma et al. 2006). However, this approach faces several shortcomings. First, it neglects a crucial dimension of perception: time. Population codes need to be constructed, integrated and combined on the time scale of interspike intervals while the underlying stimuli vary dynamically. Moreover, information contained in population codes can be kept in memory for a significant period of time, even in the absence of sensory input. Finally, rate models rely on stochastic spike generation rules that add noise to the code. Counteracting this noise requires averaging over large numbers of neurons which makes these codes unsuitable to be implemented in small neural populations. Here, we propose a new interpretation of population coding in the context of temporal sensory integration. We consider spikes, rather than rates, as the basic unit of probabilistic computation. Spike generation in our model is deterministic and results from a competition between an integration of evidence from feed-forward inputs and a prediction from recurrent connections. A neuron therefore acts as a "predictive encoder", only spiking if its input cannot be predicted by its own or its neighbor's past activity. We show that this can be performed by simple recurrent networks of integrate and fire neurons. Decoding in our model reduces to a leaky integration of spikes weighted by a kernel that can be learned from the neurons' tuning curves and spike count covariance matrix. We demonstrate that such networks can integrate and combine sensory and motor inputs optimally, i.e. without losing information, and retain this information over time. In particular, small networks of only tens of neurons can encode stable memories reflected by sustained, asynchronous spiking activity at low rates, in the absence of sensory inputs. Such persistent activity has previously been difficult to achieve in small networks. Our model sets a guideline on how to choose the network structure without the burden of laborious fine tuning. Despite being deterministic, these networks generate weakly correlated, Poisson-like distributed spike trains. The trial to trial variability of the neural responses, however, is a direct reflection of sensory noise alone and is not due to other intrinsic noise sources. Nonetheless, optimal coding is robust to the addition of synaptic noise. Our model suggests that our brain might not be as noisy as it first appears. In particular, the spike times of a neuron should become more predictable if they are conditioned on simultaneously recorded spike trains of other neurons in the population. This would provide an experimental test for deterministic spiking population codes, as opposed to stochastic rate codes. References: Ma et al., *Nature Neuroscience* 9:1432-1438, 2006.

[doi:](#)

I-73. Noise correlations in area MSTd are weakened in animals trained to perform a discrimination task

Yong Gu¹
 Sam Fok²
 Adhira Sunkara²
 Sheng Liu²
 Gregory DeAngelis³
 Dora Angelaki²

GUYONG2004@GMAIL.COM
 SAM.B.FOK@GMAIL.COM
 AS29@CEC.WUSTL.EDU
 SLIU@PCG.WUSTL.EDU
 GDEANGELIS@CVS.ROCHESTER.EDU
 ANGELAKI@CABERNET.WUSTL.EDU

¹Washington University School of Medicine

²Washington University in Saint Louis

³University of Rochester

Introduction: Behavioral performance is based on population activity, and the accuracy of population coding is constrained by correlated noise among neurons. Weak but significant inter-neuronal noise correlations have been described in many areas (e.g. V1, MT), suggesting that correlated noise may be prevalent among sensory cortical neurons. Here we examine noise correlations in the dorsal medial superior temporal area (MSTd). Neurons in this area are tuned to direction of self-motion based on both visual (optic flow) and vestibular inputs, and have been implicated in multisensory integration for heading perception. **Methods:** Pairs of single-units were recorded while animals fixated a head-fixed target and experienced a variety of self-motion stimuli. Two groups of animals were examined: those trained to perform a fine heading discrimination task ('trained', n=4) and those that were not trained to perform any task other than fixation ('naïve', n=3). Three stimulus conditions were presented: (1) a 'vestibular' condition in which animals were translated along different heading directions by a motion platform; (2) a 'visual' condition in which the same headings were simulated using optic flow, and (3) a 'combined' condition in which synchronized inertial motion and optic flow signaled the same heading direction. Each movement followed a 2s Gaussian velocity profile and noise correlations (r_{noise}) were computed from spike counts measured during the middle 1s of the stimulus period. **Results:** The average value r_{noise} was not significantly different between the three stimulus conditions ($p > 0.4$), indicating that noise correlations in MSTd are not dependent on the sensory modality of stimulation. However, noise correlations depended on the presence and strength of the heading stimuli: r_{noise} was weakest during the middle of each trial when the motion stimulus was strongest. Noise correlations also depended on the distance between neurons. The average r_{noise} was significantly > 0 when the distance between neurons in a pair was < 1 mm, and not significantly different from zero for more distant neurons. Interestingly, noise correlations in 'trained' monkeys were significantly smaller than those in 'naïve' monkeys ($p = 0.004$); in fact, noise correlations during stimulus presentation in trained animals were effectively zero. This difference was not a confound of the relationship between noise correlation and signal correlation, where signal correlation quantifies the similarity of tuning between neurons in a pair (ANCOVA). Tuning properties and neuronal variability, as assessed by mean firing rate, fano factor, direction discrimination index (DDI) and tuning width, were not different between the two groups of animals. Thus, the weaker noise correlations in 'trained' animals appear to be a result of training itself or perhaps attention. **Conclusions:** These results suggest that correlated noise among nearby and similarly-tuned MSTd neurons is weakened in animals trained to perform a fine heading discrimination task. The training-induced reduction of correlated noise may serve to increase the fidelity of population codes for heading in MSTd.

doi:

I-74. Sound texture perception via synthesis

Josh H. McDermott¹
 Andrew J. Oxenham²
 Eero P. Simoncelli³

JHM@CNS.NYU.EDU
 OXENHAM@UMN.EDU
 EERO.SIMONCELLI@NYU.EDU

¹Center for Neural Science, New York University

²Dept. of Psychology, University of Minnesota

³HHMI / Center for Neural Science, NYU

Many natural sounds, such as those produced by rainstorms, fires, and swarms of insects, result from the superposition of many rapidly occurring acoustic events. We refer to these sounds as "auditory textures", and their temporal homogeneity suggests that their defining characteristics are statistical. To explore the statistics that might underlie the perception of natural sound textures, we designed an algorithm to synthesize sounds from statistics extracted from real sounds. The algorithm was inspired by those used to synthesize visual textures, in which statistical measurements from a photographic texture image are imposed on a sample of white noise (Heeger & Bergen, 1995; Portilla & Simoncelli, 2000). Because we are interested in biologically plausible representations, we studied statistics of responses of a standard auditory filterbank that approximates the information available in the auditory nerve. Statistics were first measured from the subbands of a natural sound texture; the subbands of a noise sample were then adjusted using a gradient descent method until their statistics matched those measured in the original. If the imposed statistics capture the perceptually important properties of the texture in question, the synthesized result ought to sound like the original sound. We found that simply matching the marginal statistics (variance, skew, kurtosis) of individual filter responses and their envelopes was generally insufficient to yield perceptually satisfactory results, producing compelling synthetic examples only for certain water sounds. We observed that many sound textures contained structure in frequency and time, evident in pair-wise envelope correlations (between different subbands, and between different time points within each band). Imposing these envelope correlations greatly improved the results, frequently producing synthetic textures that sounded natural and that listeners could reliably recognize. Sounds signals that were successfully synthesized in this way included bubbling water, thunder, insect, frog, and bird choruses, applause, running animals, and frying eggs, among many others. Despite these successes, there were cases for which synthesized sounds sounded notably different from the corresponding original sound, despite having the same marginal statistics and envelope correlations. Examples of failures included sounds with abrupt broadband onsets, pitch-varying harmonic structure, or strong reverberation. These failures indicate that the statistics we imposed are insufficient to capture these sound qualities, and that the auditory system must be utilizing additional measurements. Our current efforts are directed towards identifying new statistics to account for these sound properties. Our results suggest that statistical representations could underlie sound texture perception, and that in many cases the auditory system may rely on fairly simple statistics. Although we lack definitive evidence that the precise set of statistics used in our model are instantiated in the auditory system, we note that they are of a form that could plausibly be computed with simple neural circuitry. Our method provides a means of testing the perceptual importance of such statistics, and of generating new forms of experimental stimuli that are precisely characterized, yet share important properties with real-world sounds.

doi:

I-75. Manipulation of sound-driven decisions by microstimulation of auditory cortex

Petr Znamenskiy^{1,2}
Anthony M. Zador²

ZNAMENSK@CSHL.EDU
ZADOR@CSHL.EDU

¹Watson School of Biological Sciences

²Cold Spring Harbor Laboratory

What are the cortical mechanisms underlying auditory perception? Electrophysiological recordings have revealed that the activity of single neurons in the auditory cortex is correlated with both sensory stimuli and with behavioral responses. However, such studies are correlational only; establishing a causal relationship between neural activity and behavior requires a perturbation. Lesions and reversible inactivation of the auditory cortex have been shown to impair performance on certain auditory tasks. However, by themselves such loss-of-function experiments can be difficult to interpret. For example, because of the plastic and redundant nature of the brain, performance can often recover after lesions. Here we have adopted a complementary gain-of-function microstimulation approach to study the role of auditory cortex in decisions driven by sounds. We trained rats to discriminate low-

and high-frequency acoustic stimuli in a two-alternative choice task. Each stimulus consisted of a sequence of short (30 ms) overlapping pure tones, distributed over a 3 octave range (5-40 kHz). In low-frequency stimuli, tones were drawn predominantly from the lowest octave (5-10 kHz), whereas in high-frequency stimuli they were drawn predominantly from the highest octave (20-40 kHz). Rats learned this task quickly (2-4 weeks). Discrimination performance varied as a function of the fraction of low and high octave tones in a stimulus, yielding smooth psychometric curves. Choice triggered average analysis revealed that rats integrated perceptual evidence for hundreds of milliseconds. To assess the causal role of auditory cortex in performing this task, we microstimulated primary auditory cortex during behavior and assessed the effects on choice bias. Owing to the tonotopic organization of primary auditory cortex, we could preferentially activate neurons with particular frequency preferences. We found that microstimulation of auditory cortex could have dramatic effects on rats' performance on the task, biasing the fraction of rats' responses ($p < 0.001$) to one side by as much as 35%. Our results support the hypothesis that neurons within the auditory cortex mediate this frequency discrimination task.

doi:

I-76. A generalized linear model for estimating receptive fields from midbrain responses to natural sounds

Ana Calabrese¹
Joseph Schumacher²
David Schneider¹
Sarah Woolley³
Liam Paninski⁴

AMC2257@COLUMBIA.EDU
JS2880@COLUMBIA.EDU
DMS2159@COLUMBIA.EDU
SW2277@COLUMBIA.EDU
LIAM@STAT.COLUMBIA.EDU

¹Columbia University

²Program in Neurobiology and Behavior, Columbia University

³Department of Psychology, Columbia University

⁴Department of Statistics, Columbia University

Understanding neural responses to natural stimuli has become an essential part of characterizing neural coding. For this reason, it is becoming increasingly important to develop unified models for neural responses to stimuli with a wide range of statistical properties, from white noise to the fully natural case. In the auditory system, the stimulus-response properties of single neurons are often described in terms of the spectrotemporal receptive field (STRF), a linear kernel relating the spectrogram of the sound stimulus to the instantaneous firing rate of the neuron. STRFs provide useful information about neurons' tuning properties, such as the spectral and temporal acoustic patterns to which neurons are maximally responsive. Traditionally, STRFs have been estimated as the spike-triggered average multiplied by the inverse of the stimulus covariance matrix to account for pairwise correlations in the stimulus [1]. However, when nonlinear neurons are probed with natural stimuli, which contain strong higher-order correlations, normalized reverse correlation methods (NRC) produce systematic biases (or deviations) in the estimates of the underlying filter. Furthermore, due to nonlinear effects, the calculation of the STRF depends on the corpus of the stimuli. Recently, several methods have been proposed to characterize the tuning properties of auditory neurons from responses to natural stimuli that reduce the impact of stimulus-correlation biases on the estimated STRFs [2, 3]. These algorithms differ in their functional models, cost functions, and regularization methods. Here we describe the stimulus-response relationship with a generalized linear model (GLM). In this model, each cell's input is described by: 1) a stimulus filter (STRF); and 2) a post-spike filter, which captures dependencies on spike history. As opposed to most previous models, this model provides precise spike timing information that allows accurate predictions of neural spike trains to novel stimuli. Using maximum likelihood techniques, we fit the model to the spiking activity of zebra finch auditory midbrain neurons in response to conspecific vocalizations (songs) and modulation-limited (ML) noise for which the maximum spectral and temporal modulation boundaries matched zebra finch song. In order to obtain accurate fits, we add the L1 regularizer to the likelihood function, yielding a sparse solution. Comparison of GLM and NRC methods shows that a GLM has better predictive power for songbird auditory midbrain neurons. We compare GLM and NRC STRFs in terms of their basic tuning properties and show that GLM STRFs are more consistent between stimulus ensembles.

With ML noise stimuli, STRFs computed with the NRC and GLM methods were similar, as would be predicted theoretically. With song stimuli, STRFs from the two models can differ profoundly. These results suggest that the L1-penalized GLM method provides a more accurate characterization of spectrotemporal tuning than does the NRC method when responses to natural sounds are studied in these neurons. [1] F.E. Theunissen et al. (2001) *Network: Comput. Neural Sys.* 12: 289-316. [2] T.O. Sharpee et al. (2004) *Neural Comput.* 16(2): 223-250. [3] S.V. David et al. (2007) *Network: Comput. Neural Sys.* 18(3): 191-212.

doi:

I-77. Identification of excitation and inhibition in the auditory cortex using nonlinear modeling

Nadja Schinkel-Bielefeld¹

Stephen V. David²

Shihab A. Shamma²

Daniel A. Butts²

SCHINKEL@UMD.EDU

SVD@UMD.EDU

SAS@UMD.EDU

DAB@UMD.EDU

¹Department of Biology, University of Maryland

²University of Maryland, College Park

Understanding the function of neurons in the auditory cortex relies on characterizing non-linear neuronal processing in the context of complex stimuli. In particular, it is well known that A1 neurons receive both excitatory and inhibitory inputs, often representing the same or overlapping frequencies. The resulting neuronal processing often depends on the precise nature of this balance between excitation and inhibition and its dynamical interplay in complex stimulus contexts. For example, balanced excitation and inhibition is assumed to sharpen neuronal responses in time and frequency, flanking inhibition can enhance sensitivity to a center frequency, and completely unbalanced excitation and inhibition can lead to intensity tuning. Here, we address these issues using non-linear modeling of extracellular recordings from neurons in primary auditory cortex of a passively listening ferret in the context of speech stimuli. The standard approach to mapping stimulus selectivity using extracellular data is through measurements of the spectro-temporal receptive field (STRF), which offers a first-order characterization of the features that the neuron responds to. However, STRF-based linear models have difficulties identifying inhibition, especially if excitation and inhibition are balanced or the spontaneous firing rate is low. For the separate characterization of excitation and inhibition, usually intracellular recordings are necessary. We use a newly developed Generalized Non-Linear Modeling (GNM) approach to characterize A1 neurons. This approach is based on efficient maximum likelihood estimation techniques developed for Generalized Linear Models, but incorporates additional static non-linearities operating on the output of each linear element of the model. Because these non-linearities can be determined from the data itself, depending on the processing present they can alternately represent a simple linear transform, rectification, saturation, or any combination. Importantly, this framework also allows for the identification of multiple spectrotemporal features that influence the response and their associated non-linearities, such as distinct excitatory and inhibitory contributions. The GNM approach can be readily applied to highly correlated stimuli such as speech, and has similar data requirements to standard STRF estimation experiments. We find that the GNM performs as well or better on cross-validated speech data than models built using a standard STRF-based Linear-Nonlinear (NL) model. The GNM typically identifies the spectrotemporal tuning observed in the spike-triggered average, and in most cases it also finds a second suppressive or "inhibitory" kernel, based on spectrotemporal tuning that is usually not apparent in the STRF. Time kernels for excitation and inhibition generally look similar with excitation leading inhibition, while the relative frequency tuning is much more variable from neuron to neuron. While in many cases a sharpening of responses results from inhibition that balances or is slightly wider than excitation, there are also cases with additional inhibitory peaks that do not correspond to excitation, and can be distributed over more than three octaves. We thus offer a detailed analysis of how putative excitatory and inhibitory inputs are tuned in complex stimulus contexts in primary auditory cortex. More generally, we demonstrate a modeling framework to identify multiple elements that contribute to the neuronal responses in complex stimulus contexts using extracellular data.

doi:

I-78. Top-down influences on intensity coding in primary auditory cortex

Liberty S. Hamilton
Shaowen Bao

University of California, Berkeley

LHAMILTON@BERKELEY.EDU
SBAO@BERKELEY.EDU

Both bottom-up (stimulus driven) and top-down (task-relevant) demands have been shown to influence cortical plasticity in adult rats. Previous research has shown that plasticity in A1 differs in rats trained to recognize sound loudness versus sound frequency, but it is not clear whether task demands may change neural coding strategies within one of these stimulus dimensions. We trained adult female Sprague Dawley rats on two different behavioral tasks using the same sound stimuli: (1) an intensity recognition task and (2) an intensity discrimination task. For (1), an absolute intensity recognition task, rats performed a two alternative forced choice task in which they listened to pure tone pip trains played at either 45 dB SPL or 60 dB SPL, and had to make a left nose poke when they heard a 45 dB at any frequency and the right for 60 dB at any frequency in order to receive a food reward. For (2), a relative intensity discrimination task, rats listened to trains of pure tone pips at either 45 dB or 60 dB, which, after a hold period, would then alternate between quiet and loud. The rat would receive a food reward by initiating a poke while the sounds alternated in level, indicating its ability to discriminate between the two sound levels. In both conditions, rats were trained using tone pips between 10.8 kHz and 21.2 kHz at either of the decibel levels, and training lasted approximately one month. Multiunit activity in response to pure tone pips from 1 kHz to 32 kHz and -5 dB SPL to 70 dB SPL was recorded from primary auditory cortex (A1) of trained rats and naïve controls under sodium pentobarbital anesthesia. From these data, we constructed classical receptive fields to determine the best frequency, frequency bandwidth, and best level of each recording site. We then calculated rate level functions (RLFs) for each recording site by calculating the average number of spikes at each intensity level, collapsed across frequencies. While we found no difference in frequency representation across groups, we found an increase in the number of nonmonotonic RLFs in the intensity recognition group compared to intensity discrimination and control rats. Both the fraction of nonmonotonic sites and the degree of nonmonotonicity for neurons increased in the recognition group compared to discrimination and controls, suggesting distinct mechanisms for encoding sound intensity in response to top-down behavioral demands. Such differences may be related to changes in the balance between excitatory and inhibitory synaptic activity in auditory cortex.

[doi:](#)

I-79. Neural encoding of global statistical features of natural sounds.

Maria Neimark Geffen
Thibaud Taillefumier
Marcelo Magnasco

Rockefeller University

MGEFFEN@ROCKEFELLER.EDU
TTAILLEFUM@MAIL.ROCKEFELLER.EDU
MGNSCLB@MAIL.ROCKEFELLER.EDU

We investigate the statistical structure of water sounds, and show that water sound can be generated as a fractal controlled by two global parameters. Psychophysical studies show that human perception of this signal as a natural object differs as the parameters are changed, suggesting that these parameters are sufficient to represent the family of water sounds. We propose that the early auditory system may extract and encode the statistics of these stimuli by performing a scale-invariant analysis of the sound signal. We next use these sounds in electrophysiological experiments to map the neural circuit that might encode the global statistical parameters of these signals. Neural responses in the primary auditory cortex of awake rats (A1) correlated with the changes in the global parameters of the synthetic sounds, both at the level of the firing rate, and the specific timing and timing precision of individual firing events. Changes in global parameters of the stimuli further resulted in temporal changes in the receptive field of the A1 neurons. The temporal receptive field was slower for stimuli with longer auto-correlation structure, and faster for stimuli with faster auto-correlation structure. Adding a time-dependent non-linear feedback term improved the fit of the linear-non-linear model. We propose the hypothesis that encoding of natural environmental sounds requires comparison of global statistics of these sounds at a range of timescales.

doi:

I-80. Contour representation of sound signals

Yoonseob Lim¹Barbara G. Shinn-Cunningham¹Timothy Gardner²

YSLIM@BU.EDU

SHINN@CNS.BU.EDU

TIMOTHYG@BU.EDU

¹Department of Cognitive and Neural Systems, Boston University²Department of Biology, Boston University

Continuous edges, or contours, are powerful features for object recognition, both in neural and machine vision. Similarly, auditory signals are characterized by sharp edges in some classes of time-frequency analysis. Linking these edges to form contours could be relevant for auditory signal processing. However, the mathematical foundations of a general contour representation of sound have not been established. Sinusoidal representations of voiced speech and music have been explored, but these approaches do not represent broadband signals efficiently. Here we construct a two-dimensional contour representation that is generally applicable to any time series, including sound. Starting with the Short Time Fourier Transform (STFT), the method defines edges by coherent phase structure at local points in the time-frequency plane (zero crossings of a complex reassignment matrix). Continuous edges are grouped to form contours that follow the ridges and valleys of the traditional STFT. Local amplitudes are assigned by calculation of fixed points in an iterated reassignment mapping. The representation is additive; the complex amplitudes of the contours can be directly summed to reproduce the original signal. This re-synthesis matches the original signal with a signal-to-noise ratio of 15 dB or higher, even in the challenging case of white noise. In practice, this level of precision provides perceptually equivalent representations of speech and music. For many sounds of interest, a subset of the full contour collection can provide an accurate representation. To find this compact subset, an over-complete set of contours are calculated using multiple filter bandwidths. Contours are then ranked by power, length and curvature, and subject to lateral inhibition from neighboring contours. The top-ranking contours in this distribution provide a sparse representation that emerges without any prior suppositions about the nature of the original signal. By combining contours from multiple bandwidths, the representation achieves high precision in both time and frequency. As such, the method is relevant to a wide range of time-frequency tasks such as constructing receptive fields of auditory neurons, characterizing animal vocalizations, pattern recognition and signal de-noising. We speculate that neural auditory processing involves a similar contour representation. Each stage in the analysis is a plausible operation for neurons: parallel and redundant primary processing in multiple bandwidths, grouping by phase coherence, linking by continuity and lateral inhibition.

doi:

I-81. Neural activity as samples from a probabilistic representation: evidence from the auditory cortex

Pietro Berkes¹Stephen V. David²Jonathan B. Fritz²Mate Lengyel³Shihab A. Shamma²Jozsef Fiser¹

BERKES@BRANDEIS.EDU

SVD@UMD.EDU

RIPPLE@ISR.UMD.EDU

M.LENGYEL@ENG.CAM.AC.UK

SAS@ISR.UMD.EDU

FISER@BRANDEIS.EDU

¹Brandeis University²University of Maryland³University of Cambridge

In the past years, there has been a paradigm shift in the field of cognitive neuroscience as a number of behavioral studies demonstrated that animals and humans can take into account statistical uncertainties of task, reward, and their own behavior, in order to achieve optimal task performance. These results have been interpreted in terms of statistical inference in probabilistic models. However, such an interpretation raises the question of how cortical networks represent and make use of the probability distributions necessary to carry out such computations. Recently, we have proposed that neural activity patterns correspond to samples from the posterior distribution over interpretations of the sensory input, a hypothesis that is consistent with several experimental observations (e.g. trial-to-trial variability). Last year, using this framework, we verified experimentally that the distribution of spontaneous activity in such probabilistic representations adapts over development to match that of evoked activity averaged over stimuli, based on recordings from V1 of awake ferrets. In the present study, we define and test two novel predictions of this framework. First, we predict that the match between evoked and spontaneous activity should be specific to the distribution of neural activity evoked by natural stimuli, and not to that evoked by artificial stimulus ensembles. We expect this match to hold for instantaneous neural activity, and for temporal transitions between activity patterns. Second, if this hypothesis captures the general computational strategy in the sensory cortex, it should be valid across sensory modalities. To test these predictions, we analyzed single unit data (N=32 over 6 recordings) recorded simultaneously from multiple electrodes in the primary auditory cortex (A1) of awake ferrets in three stimulus conditions: a natural condition consisting in a stream of continuous speech, a white noise (0-20 kHz) condition, and a spontaneous activity condition where the animal was listening in silence. Speech was chosen since its spectrotemporal characteristics are similar to those of natural sounds. We analyzed the neural data, which was discretized in 25 ms bins, binarized, and the distribution of instantaneous, joint activity, and the transition probability from one activity pattern to the next was estimated in the three conditions. We measured dissimilarity between the silence and stimulus condition distributions using Kullback-Leibler divergence. The robustness of our results was estimated using a bootstrapping technique. In agreement with our predictions, we found that the distribution of speech-evoked activity is consistently more similar to spontaneous activity than the distribution of noise-evoked activity, for both the instantaneous distribution of activity and for transition probability. These results provide new evidence for stimulus specific adaptation in the cortex that leads to preference for natural stimuli, and also provide additional support for the sampling hypothesis. Our findings in A1 complement our earlier data from V1, suggesting that the match between spontaneous and evoked activity might be a universal hallmark of representation and computation in sensory cortex. Acknowledgements: This work has been supported by the Swartz Foundation, by the Swiss National Science Foundation, and by NIH (R01DC005779, K99DC010439)

doi:

I-82. Two photon imaging of tactile responses during frequency discrimination in awake head fixed rats

Florent Haiss¹
 Johannes Mayrhofer¹
 David Margolis¹
 Mazahir T. Hasan²
 Fritjof Helmchen¹
 Bruno Weber¹

FHAISS@PHARMA.UZH.CH
 JOHANNES.M.MAYRHOEFER@GMAIL.COM
 MARGOLIS@HIFO.UZH.CH
 MAZAHIR.HASAN@MPIF-HEIDELBERG.MPG.DE
 HELMCHEN@HIFO.UZH.CH
 BWEBER@PHARMA.UZH.CH

¹University of Zurich

²Max Planck Institute for Medical Research

A mainstay in systems neuroscience is the use of two alternative forced choice paradigms in the head-fixed primate. Although favorable in many respects, a rodent model of this type of behavioral task has not been established up to now. Head restriction is a prerequisite for a precise presentation of tactile stimuli to the vibrissae and enables two-photon imaging of the neural activity in the behaving animal. Rats were trained to discriminate sinusoidal whisker vibration stimuli. Using a dedicated motorized mechanical system, the head fixation is released following stimulus presentation. In this way, the animal is allowed to perform a left or right rotation of the head to

receive a reward from a water-spout located on either side. By choosing one of the spouts the animal reported which stimulus it perceived. Stimulation with a target frequency (190 Hz) of the left/right whisker would lead to a reward on the respective water spout. After reward retrieval, the head is slowly brought back to centre position. The rats were able to discriminate frequencies in the range of 60 to 190 Hz and associate them to the location of the reward. The animals were able to perform this task with 90% of correct responses. We furthermore combined the described head fixation mechanics with two photon imaging of barrel cortex neurons. A recombinant adeno-associated virus carrying constructs for the genetically-encoded calcium indicator Yellow Chameleon 3.6 (YC 3.6) was injected in the somatosensory cortex. This procedure allows imaging of action potential evoked calcium signals in multiple neurons during the behavioral task. Continuous imaging of identical neuronal populations across trials and days was possible due to a chronic cranial window and the stable expression and function of YC 3.6. Imaging was performed up to two months after infection of the neurons. In conclusion, the two alternative forced choice paradigm presented here opens new possibilities for imaging neuronal activity related to complex discriminative behavior and learning in rats.

[doi:](#)

I-83. From form to function: deriving preferred stimuli from neuronal morphology

Jonas Mulder-Rosi
Graham Cummins
John Miller

MULDERROSI@CNS.MONTANA.EDU
GIC@LORAX.ORG
JPMILLERUSA@GMAIL.COM

Montana State University

Sensory neuroscience experiments are often used to determine a neuron's response tuning to experimentally provided stimuli. The morphology and connectivity of the neuron are then understood in terms of this 'preferred stimulus'. This procedure is inevitably limited to some extent by the stimuli originally provided. Varying, or even considering, all possible stimulus dimensions to which an organism may be sensitive is an unavoidable difficulty of this approach. This leaves open the possibility that measured tuning curves may indicate only local maxima within a larger stimulus space. Arguments for the importance of the studied stimulus dimension are often provided in the form of biologic relevance or complementary tuning of other neurons in the population. The difficulty of selecting stimuli may be sidestepped by turning the problem on its head: moving from morphology to function. If enough is known about a neuron and its upstream partners, we should be able to compute the stimulus to which the neuron is globally most sensitive. We have taken just such an approach using interneurons in the well-studied cricket cercal system. The cricket cercal system detects near-field air motion and allows the animal to respond to predators and conspecifics. The targets of the system's entire primary afferent population have been mapped, and models of the conserved projecting interneuron morphology have been imbedded in this sensory map. We focused on two pairs of bilaterally symmetric interneurons which compose a classic model system of population coding. These four cells have evenly spaced tuning curves sensitive to the direction of air currents laterally around the animal. Recent work in our lab has shown that the interneurons in this system are sensitive to variation in air current dynamics at the centimeter scale. Stimulus variation at this scale has previously been avoided experimentally. We created highly accurate three-dimensional models of these neuron's passive dendritic trees and their probabilistic presynaptic afferent partners. These models were used to obtain tuning curves for these neurons in response to stimuli varying in direction not just globally, but with centimeter scale spatiotemporal variation. These modeling experiments predicted tuning curves along previously unmeasured stimulus dimensions. We were then able to test these predictions in vivo using a novel recording method which allowed us to simultaneously obtain spike times from both classes of previously studied interneurons (as well as the rest of the ascending interneuron population). The agreement between our predicted tuning curves and our recorded response was good, especially considering the relative simplicity of the neuronal models used. Interestingly, while these four neurons equitably cover the dimension of bulk airflow direction, their tuning for smaller scale variation appears to be mutually orthogonal. We have shown the validity of using models to derive functional tuning curves for neurons. These models are naïve to the limitations of stimulus generation in a laboratory setting and thus suggest global "best stimuli" rather than

simply local maxima. This has allowed us to uncover additional tuning dimensions for a classic model system.

[doi:](#)

I-84. Receptive field mapping of local populations in mouse visual cortex using two-photon calcium imaging

Vincent Bonin
Mark H. Histed
R. Clay Reid

Harvard Medical School

VINCENT_BONIN@HMS.HARVARD.EDU
MARK_HISTED@HMS.HARVARD.EDU
CLAY_REID@HMS.HARVARD.EDU

Sensory stimuli are processed through individual neurons' receptive fields. These receptive fields have historically been measured from electrical recordings, which provide little information about the location of the recorded neurons, and only sparsely sample local populations. As a result, little is known about how receptive fields are organized at the microscopic scale, within a cortical column. Here we examine receptive fields in the visual cortex of the mouse, for which it is not known how different aspects of the stimulus are mapped in local populations. To address these issues, we have constructed an optimized, video-rate two-photon microscope which can produce data at a resolution approaching that of electrode recordings. Two-photon calcium imaging provides anatomical information about neurons and also can sample nearly all neurons in a small volume of cortex. However, because of slow imaging rates and limited signal-to-noise, it has been difficult to use this technique to measure receptive field properties. The new microscope can resolve the time course of responses in single trials at a resolution of ~30 ms. To reconstruct receptive fields, we used a visual stimulus optimized to produce responses easily detectable using calcium imaging. Receptive fields are often measured by stimulating with white noise and calculating the average stimulus preceding a response spike. White noise, however, produces a small number of spikes that are dispersed in time, and these responses are often too small to be detected with in vivo calcium imaging. Instead, we sought a stimulus that would elicit large bursts of spikes. We constructed a colored noise stimulus from a set of wavelet basis functions whose spectrum matches the sensitivity of the neurons (cutoffs ~2 Hz and ~0.15 cycles/deg). Basis function coefficients were sparse (typically 1 / 100 were nonzero), had random signs, and amplitudes inversely proportional to frequency. We measured somatic calcium responses to these stimuli. We bulk-labeled neurons and astrocytes in visual cortex of anesthetized mice with the calcium indicator Oregon Green BAPTA-1 and the glial label SR101. We imaged hundreds of neurons at several depths, 150 - 350 μm below the surface. Stimuli covering 60 x 60 degrees of the visual field were presented monocularly. We measured 3 - 12 trials, each consisting of 16 stimuli lasting 16 seconds, interspersed with 8 seconds of gray screen to reduce the effects of adaptation. We calculated the somatic calcium responses and deconvolved them using an efficient algorithm (Vogelstein et al, in preparation), and thresholded at $> 3\% \Delta F/F$. The neurons responded strongly to these stimuli yielding clear receptive field maps in many neurons. In the best preparations, we obtained significant receptive field maps for a majority of the neurons, consistent with the large number of simple cells this species (Niell and Stryker 2008). These maps were repeatable across trials and their structure was consistent with the responses to single gratings moving in different directions measured in the same neurons. We are currently using this assay to study the organization of receptive fields in mouse visual cortex. This work was supported by NEI and NINDS.

[doi:](#)

I-85. The encoding of fine spatial information in salamander retinal ganglion cells

Frederick Soo¹
 Gregory Schwartz²
 Michael J. Berry II¹

FSOO@PRINCETON.EDU
 GREGWS@U.WASHINGTON.EDU
 GREATSPORTII@GMAIL.COM

¹Princeton University

²University of Washington

Classical models of retinal ganglion cell signaling assume that ganglion cell receptive fields are smoothly varying, approximately Gaussian in profile and arranged in a regular array. This model predicts that the receptive fields of neighboring ganglion cells will be highly overlapping and that the cells will convey largely redundant visual information. In whole-cell voltage clamp recordings from salamander retina, however, the receptive fields of ganglion cells were irregular on a fine scale and non-Gaussian, and the information conveyed by neighboring ganglion cells about a flashed spot stimulus was nearly independent. Large groups of ganglion cells encoded significantly more information than expected from the classical model even when the approximate sizes and positions of the cells' receptive fields were taken into account. The discrepancy was only explained by including the receptive field irregularities in the model. This result suggests that irregularities in spatial receptive field profiles are a positive design feature rather than an unavoidable defect and that subsequent brain circuits can benefit from recognizing such irregularities when they interpret retinal spike trains.

[doi:](#)

I-86. The projective field of single bipolar cells in the retina

Hiroki Asari¹
 Markus Meister²

ASARI@FAS.HARVARD.EDU
 MEISTER@MCB.HARVARD.EDU

¹Department of Molecular and Cellular Biology, Harvard University

²Harvard University

The vertebrate retina contains about 10 types of bipolar cells that convey information from photoreceptors to retinal ganglion cells. Previous studies suggest that they form parallel channels, and that each bipolar cell type contributes a specific visual message to select types of ganglion cells. Here we test this hypothesis by determining the full projective field of a single bipolar cell, specifically the responses it elicits in the population of ganglion cells. In the isolated salamander retina, we controlled bipolar cell activity with an intracellular electrode while recording the firing of ganglion cells with a multi-electrode array and manipulating the intervening circuitry pharmacologically. We found that excitation of a single bipolar cell altered the responses of many ganglion cells. The effect was generally excitatory at short distances and inhibitory at longer distances. Electrical synapses among bipolar cells contributed substantially to the lateral spread of excitation, whereas amacrine cells suppressed the excitatory spread and mediated the inhibitory effects. Within the excitatory region, different ganglion cells showed distinct temporal response patterns. A sustained depolarization of the bipolar cell produced a transient burst of spikes in some ganglion cells but sustained firing in others. This range of response kinetics resulted primarily from the interactions of individual bipolar cell terminals with amacrine cells. Our results highlight the diversity of neuronal circuits that distribute signals from a bipolar cell to various ganglion cells, and suggest considerable cross-talk between bipolar cell channels through gap junctions and via amacrine cells.

[doi:](#)

I-87. Contribution of amacrine transmission to fast adaptation of retinal ganglion cells

Neda Nategh
Mihai Manu
Stephen Baccus
Stanford University

NNATEGH@STANFORD.EDU
MIHAIM@STANFORD.EDU
BACCUS@STANFORD.EDU

Retinal ganglion cells are most sensitive to the visual feature defined by the linear spatio-temporal receptive field. They encode this feature according to a nonlinear sensitivity curve that often has a threshold and saturation. Both the linear receptive field and nonlinearity are adaptive, in that these parameters change depending on the recent statistics of the stimulus. One potentially rich source to generate adaptation is the diverse population of inhibitory amacrine cells, which comprise about thirty types. Amacrine transmission is thought to play a role in retinal adaptation to more complex stimulus statistics (Hosoya et al., 2005), but not for simple statistics such as luminance and contrast. We measured how the signals transmitted through individual amacrine cells contribute to the ganglion cell response by recording intracellularly from single amacrine cells while simultaneously recording spiking activity from the ganglion cell population using a multielectrode array. We presented a randomly flickering visual stimulus drawn from a Gaussian distribution while injecting Gaussian white-noise current into the amacrine cell. By this direct perturbation of the circuit we measured how the interneuron generates adaptation of the ganglion cell visual response. To model the contribution of each amacrine cell to each ganglion cell's visual response, we combined elements of a linear-nonlinear (LN) model, consisting of a linear temporal or spatio-temporal filter followed by a static nonlinearity. The model consisted of the linear receptive field and nonlinearity of the ganglion cell, a modulatory pathway containing the LN model of the amacrine cell, and a transmission filter linking the two pathways. We found that amacrine transmission scales the ganglion cell nonlinear response function by a gain factor. In some cases, we also found that amacrine output modulates the linear receptive field of the ganglion cell, changing it from being more integrating to more differentiating. This modulation is driven by the preferred feature of the amacrine cell, even if this feature is different from that of the ganglion cell. Even at a fixed luminance and contrast, retinal ganglion cells adapt at a fast timescale. For this type of adaptation, an amacrine cell provides contextual information that modulates the ganglion cell visual response. Thus, the space of visual features encoded by the diverse population of amacrine cells defines a multidimensional context that gates and modifies a different space of visual feature encoded by the population of ganglion cells.

[doi:](#)

I-88. Perception of the reverse-phi illusion by *Drosophila melanogaster*

John C. Tuthill
M. Eugenia Chiappe
Vivek Jayaraman
Michael B. Reiser

TUTHILLJ@JANELIA.HHMI.ORG
CHIAPPEE@JANELIA.HHMI.ORG
VIVEK@JANELIA.HHMI.ORG
REISERM@JANELIA.HHMI.ORG

Janelia Farm Research Campus, HHMI

When the contrast polarity of an image is inverted as it moves, human observers report an illusory reversal in the direction of perceived motion. This illusion, called "reverse-phi motion" has been studied extensively with the two stripe apparent-motion paradigm and moving random-dot kinematograms. Humans exhibit nearly equal sensitivity, and comparable spatial and temporal tuning, for standard vs. reverse-phi motion. Neurophysiological correlates of reverse-phi motion have been identified in several vertebrate species, including primates and cats. The salience and predominance of the reverse-phi illusion have made it an important tool for understanding how vertebrates compute visual motion. In this study, we tested whether a genetic model organism-the vinegar fly, *Drosophila melanogaster*-exhibits sensitivity to the reverse-phi illusion. Using an LED-based flight simulator and optical wingbeat analyzer, we compared behavioral responses of tethered flies to full-field rotating and expanding reverse-phi motion across a range of temporal and spatial frequencies. We found that flies exhibit "reverse-optomotor"

responses when presented with panoramic reverse-phi motion. Flies only perceive the reverse-phi illusion if the rates of contrast-inversion and pattern motion are matched and they occur in phase. We will also present results detailing the responses of motion-sensitive neurons in the Lobula Plate to presentations of reverse-phi motion. Flight steering responses to reverse-phi motion are accurately modeled with an array of Hassenstein-Reichardt elementary motion detectors, lending further support to the correlation-type motion detector as a suitable model for motion computation in the fly retina. The spatial and temporal tuning characteristics of reverse-optomotor flight behavior constrain the computational properties of the fly motion detector and may contribute to the ongoing search for the neural basis of the fly motion detector.

[doi:](#)

I-89. The structure of spontaneous and evoked population activity in mouse visual cortex

Sonja Hofer¹

Bruno Pichler¹

Ho Ko¹

Joshua T. Vogelstein²

Nicholas Lesica³

Thomas Mrsic-Flogel¹

S.HOFER@UCL.AC.UK

BRUNO.PICHLER@NIMR.MRC.AC.UK

HO.KO@UCL.AC.UK

JOSHUA@JHU.EDU

LESICA@ZI.BIOLOGIE.UNI-MUENCHEN.DE

T.MRSIC-FLOGEL@UCL.AC.UK

¹University College London

²Johns Hopkins University

³Ludwig-Maximilians-University

How similar are cortical network response patterns under spontaneous and different evoked conditions? In the absence of sensory input, neuronal populations exhibit considerable 'spontaneous' activity which is temporally and spatially correlated across millimetres of cortical surface. Recent work suggests that ongoing network activity, which may reflect underlying circuit architecture, governs the population response to external input. Specifically, patterns of spontaneous activity can resemble those driven by sensory input. In many (but not all) studies, external input seems to do little to alter the similarity of neuronal (spatial) correlations present in spontaneous states. Moreover, similar spatial patterns of population activity can be observed both when a neuron fires action potentials spontaneously and when the same neuron is driven by input. The firing of individual neurons therefore appears to be tightly coupled to the activity of the surrounding neuronal population, suggesting that stereotypical co-activation patterns should be observed during spontaneous and evoked network states. To test this idea at the single cell level, we used in vivo two-photon calcium imaging to assess the similarity of spontaneous and evoked activity patterns within mouse primary visual cortex. Spike-related somatic calcium signals were sampled from complete local populations of 30-60 neurons at 7.5-15 Hz. Visual stimuli included drifting grating sequences of different directions presented either episodically or continuously, and different types of naturalistic movies. 'Spontaneous' activity in darkness was also recorded. We compared different measures of population activity across these conditions, including population sparseness (fraction of responding cells per time bin), pairwise correlations, and pattern correlations. The overall structure of population activity was different across conditions. First, the distribution of percentage of cells active per time bin varied systematically between conditions. Second, as a first approximation of network interactions, we tested how correlated is the activity between pairs of cells during spontaneous and evoked states. The correlation coefficient for each cell pair was frequently either higher or lower between spontaneous and stimulus-driven conditions, but similar across two epochs of the same evoked condition. Pairwise co-activations were also different between different evoked conditions. Third, in order to look at patterns of co-activation in the whole population, we computed population response maps for each cell whenever it fired a spike. Patterns of population activity in these maps were very different between spontaneous and evoked conditions, but similar within the same evoked condition. Moreover, population activity patterns associated with a cell's spiking were also very different between different evoked conditions (i.e. natural movies vs. gratings). In summary, the structure of network activity is strongly dependent on both the presence and type of external stimulus. Similar co-activation patterns are rarely observed across different conditions. Therefore, in contrast

to the idea that population response patterns are largely dependent on intrinsic constraints, our data support the alternative view that sensory input dynamically restructures ongoing cortical activity according to stimulus characteristics.

[doi:](#)

I-90. Modelling molecular mechanisms of light adaptation in *Drosophila* photoreceptor

Zhuoyi Song
Daniel Coca
S.A. Billings
Mikko Juusola

The University of Sheffield

ZHUOYI.SONG@SHEF.AC.UK
D.COCA@SHEFFIELD.AC.UK
S.BILLINGS@SHEFFIELD.AC.UK
M.JUUSOLA@SHEFFIELD.AC.UK

We wish to investigate how light adaptation (LA) happens at the molecular level. What are the molecular dynamics that enable photoreceptors to translate vast environmental light changes into neural responses of limited range? By using a biophysical model of a *Drosophila* photoreceptor, which can replicate many features seen experimentally during adaptation to naturalistic light inputs, we aim to work out how molecular interactions may generate LA in multiple time scales. We chose *Drosophila* photoreceptor as the model system, because much is known about the molecular mechanisms of its phototransduction cascade. LA, is a collective of intra- and inter-cellular processes, which tune a visual neuron's output in respect to recent light input history. A fly photoreceptor adapts to efficiently transfer information about the visual input using its limited range, which is only a small fraction (40-60 mV) of ambient light intensity range that can span >8 log units. Without adaptation, responses to weak stimuli would vanish into neuronal noise, while responses to large stimuli would saturate. But it is an open question how this sophistication comes about. In *Drosophila* photoreceptors, photons are captured by rhodopsin molecules in the photosensitive plasmamembrane, the rhabdomere. Its 30,000 finger-like formations, microvilli, presumably transduce and amplify the energy of each captured photon independently into light-induced-current (LIC) through a chain of biochemical reactions. The photoinensitive membrane then helps to convert LIC into a voltage response. Accordingly, the input to our model is a continuous time series of light intensities in units of photons, which stimulates a model of a phototransduction cascade in a single microvillus. This part of the model intends to replicate known biochemical interactions of major transduction proteins (rhodopsin, metarhodopsin, G-protein, PLC, PIP2, DAG, $\text{Na}^+/\text{Ca}^{2+}$ -exchanger, CAM), several of which are feedback targets for Ca^{2+} , fluxing through light-gated channels [1,3]. It is simulated by Gillespie algorithm to account for the stochastic properties of biochemical reactions [4]. Based on 300 simulated microvilli, the integrated LIC is extrapolated for 30,000 microvilli. Their collective transmembrane current is then used to drive the photo-insensitive cell body, using the Hodgkin-Huxley-formalism to approximate the dynamics of the known voltage-gated ion-channels [2]. The model is validated by performing intracellular measurements from *Drosophila* photoreceptors to continuous light patterns in vivo and by comparing these to the model output for light inputs. Even in this relatively basic form, our model can predict well the waveforms of the voltage responses. The model is used to study how molecular mechanisms of phototransduction enable efficient coding, and how the molecular feedback interactions contribute to fast adaptation. 1. Hardie, R.C., Raghu, P, Nature 413, 186-193 (2001) 2. Vähäsöyrinki, M. Thesis, University of Oulu (2004) 3. Pumir, A. et al. PNAS 10354-10359 (2008). 4. Gillespie D.T., J. Comp. Phys. A. 22, 403-434 (1976)

[doi:](#)

I-91. The role of EAG K⁺ channels in insect photoreceptors

Esa-Ville Immonen
Roman Frolov
Mikko Vahasoyrinki
Matti Weckstrom

ESA-VILLE.IMMONEN@OULU.FI
ROMAN.FROLOV@OULU.FI
MIKKO.VAHASOYRINKI@OULU.FI
MTW@CC.OULU.FI

University of Oulu, Department of Physics

Eag (ether-á-go-go) channels are ubiquitously found across animal kingdom, but their physiological function largely remains obscure. We have found that these channels constitute the principal component of potassium conductance in photoreceptors of several insects, such as American cockroach and field cricket, and describe here specific functional roles for them. Graded voltage signaling in insect photoreceptors is a result of transformation of a huge spatiotemporal range of environmental visual stimuli into a narrow range of amplitude and frequency modulated voltage responses. Graded signals are produced by concerted opening of light-activated channels and processed by the membrane filter. Voltage-activated potassium channels in the light-insensitive part of photoreceptor membrane determine the membrane filter and thereby regulate the speed, range and amplification of voltage responses. Because of the sheer range of visual inputs, co-expression of several types of K⁺ channels with different biophysical properties is necessary to adequately process visual information. For example, in photoreceptors of the *Drosophila* (the fruit-fly) different arrays of K⁺ channels are found, including Shaker-type channels and delayed rectifier channels such as Shab, providing, respectively, amplification of voltage signals at low light levels, and attenuation of depolarization in brighter light. In cockroach photoreceptors, the eag channels could be identified functionally by various specific clues. The channels were inhibited by sub-micromolar concentrations of clofilium and other blockers of eag. They showed a strong Cole-Moore shift (sensitivity of channel activation kinetics to membrane voltage) and deceleration of activation in the presence of external divalent cations, both unique properties of these channels. Eag current was suppressed by light via elevation of cytosolic calcium. The light-dependent inhibition (LDI) was strongest in brightest backgrounds and was eliminated by supplementing the electrode solution with EGTA, but not by removal of external Ca²⁺. The LDI was enhanced in the presence of ryanodine (to stimulate Ca²⁺ release from endoplasmic reticulum) or after substitution of Na⁺ with Li⁺ ions (to inhibit Ca²⁺-Na⁺ exchanger), corroborating the role of intracellular calcium. The eag channels could be shown to be responsible for the resting potential, properties of the membrane filter, and the information capacity of photoreceptors, which were studied using naturalistic sequences of light contrast over the "normal" range of light intensities. Inhibition of eag depolarized the cell and altered properties of membrane transfer function; it increased gain (and range) and latency of photoreceptor responses, decreased membrane corner frequency and the information capacity. Change in gain led to a relative increase in signal transmission in the low frequency range and dampening of higher-frequency signals. A relatively small eag conductance during responses to dim light (which produce little depolarization), in the absence of other conductances, resulted in a small but significant loss of information capacity as estimated by comparison between current and voltage responses. This loss progressively decreased with the increase in light intensity (and membrane depolarization). On the other hand, the LDI did not alter photoreceptor signaling significantly, but was instrumental in reducing metabolic cost of photoreception by approximately 25% in brightest backgrounds by optimizing K⁺ to Na⁺ flux ratio.

doi:

I-92. Memory-related activity in the PFC depends on cell type only in the absence of sensory stimulation

Cory Hussar¹
Tatiana Pasternak²

CORY_HUSSAR@URMC.ROCHESTER.EDU
TANIA@CVS.ROCHESTER.EDU

¹Department of Neurobiology and anatomy, University of Rochester

²University of Rochester

Neurons in the prefrontal cortex (PFC) show direction selective responses to behaviorally relevant visual motion

(Zaksas & Pasternak, 2006, Hussar & Pasternak, 2009). In this study we examined memory-related signals in the PFC during a task where monkeys compared two directions of motion, sample and test, separated by a brief delay. For the analysis of neuronal activity recorded during the sample, the delay and during the comparison test we used spike waveform durations to classify the recorded neurons into narrow-spiking (NS) putative inhibitory interneurons and broad-spiking (BS) putative pyramidal neurons. We found that while responses of both classes of neurons to visual motion used in the task were equally likely to be direction selective, during the memory delay the pattern of activity for the two cell classes was different. BS neurons were significantly more active than NS cells and were more likely to show anticipatory changes in firing rates. Furthermore, BS neurons were also significantly more likely to carry signals reflecting the direction of the preceding sample. These signals were largely transient and appeared in different neurons at different times in the delay, suggesting that the information about the remembered direction is likely to be distributed among PFC neurons. The difference between the two classes of neurons became particularly apparent at the end of the delay when memory-related signals were represented exclusively by BS neurons. In contrast, during the comparison phase of the task, responses to the test of NS and BS cells were similar and on trials when test direction matched that of the preceding sample, both cell types showed lower activity. This match suppression, likely to represent the process of sensory comparison, reflected the difference in the direction between sample and test, decreasing with smaller difference between the two stimuli and disappearing when the monkey was not required to perform direction discrimination. Furthermore, responses during the test of both cell classes reflected the upcoming decision, showing significant choice probability towards the end of the response. These results reveal important differences in the contribution of the putative inhibitory interneurons and of the pyramidal cells to delayed discrimination tasks. Stimulus-driven activity, likely to represent bottom-up signals arriving from sensory cortex, was similar in both classes of cells, suggesting that both cell types participate in the sensory components of the task. However, in the absence of sensory stimulation, delay activity was dominated by putative pyramidal neurons. Since these neurons are a likely source of top-down projections from the PFC to visual and parietal cortical neurons, they may be a source of anticipatory and stimulus-related delay activity frequently observed in these neurons. Supported by NIH grants R01 EY11749, T32 EY07125 & P30 EY01319

doi:

I-93. Compete globally, cooperate locally: Signal integration for behavior depends on cortical separation

Kaushik Ghose
John H. R. Maunsell
Harvard Medical School

KAUSHIK_GHOSE@HMS.HARVARD.EDU
JOHN_MAUNSELL@HMS.HARVARD.EDU

Signals representing external stimuli and internal states are distributed throughout the cerebral cortex. It is commonly assumed that behaviors depend on integrative mechanisms that can select and process arbitrary combinations of relevant signals from these distributed representations. Although some theoretical work has considered these integrative mechanisms, relatively few experimental studies have addressed the neuronal mechanisms and processes that can integrate widely distributed cortical signals. Here we approach the question of signal integration by using a task where the subject was required to simultaneously monitor two groups of neuronal activity and report if either group was active. In this task, by comparing the subject's detection performance in cases where both groups of neurons were jointly active with performance when only one group was active, we can deduce how the subject combined the signals from the two groups of neurons. Because sensory stimuli invariably activate many neurons that are distributed across cortex in a broad and unknown pattern, we used electrical microstimulation to introduce precisely controlled activity at pairs of known cortical sites in a rhesus monkey that was trained to report when it detected stimulation at either site in a two interval forced choice paradigm. We varied the cortical separation between pairs of sites to explore how integration depends on the separation between the activated neurons. One hundred and fifteen pairs of cortical sites in V1 were tested with five electrode spacings (400, 800, 1600 and 2400 μm , and pairs in opposite hemispheres). On every trial the stimulus for each electrode was selected randomly and independently from a set of 6 currents (which always included 0 μA) yielding 36 stimulus

pairs that were each presented 50 times to estimate behavioral detection for that condition. The data were then fitted with a power law summation model with one parameter (k) which captured the nature of the signal summation. For separations of 400 and 800 μm , k was indistinguishable from 1.0, indicating that the signals were summed linearly. In contrast, when the signals were in different hemispheres k was 2.5, indicating a competitive interaction between the signals, approaching a winner-take-all operation. At the intermediate distances, 1600 and 2400 μm , k was intermediate (1.5 and 1.75 respectively), suggesting that the integration mode progresses from linear signal summation toward a winner-take-all operation as signal separation increases over a small cortical distance. In the task executed by the subject, a competitive combination of signals results in poorer detection performance than simple summation, and is a suboptimal mode of signal integration. These results suggest the process that integrates cortical signals to drive behavior is necessarily competitive when it operates on arbitrary groups of widely distributed signals, but can act as a linear summation, if required, when signals are located within about a millimeter on cortex.

doi:

I-94. The role of inhibition in formatting visual information in the retina and LGN

Daniel A. Butts¹
Alexander R. R. Casti²

DAB@UMD.EDU
ACASTI@GMAIL.COM

¹Dept. of Biology and Program in Neuroscience, University of Maryland

²Cooper Union School of Engineering

Despite being well characterized anatomically and physiologically, our understanding of how the visual pathway processes information is relatively impoverished, due in part to our reliance of the "receptive field" as a description of neuronal function. The linear receptive field describes the average visual stimulus that a neuron responds to, and is known to break down in describing cortical neurons (such as "complex cells"), whose response is known to involve nonlinear combinations of more than one visual feature. In fact, even processing in the retina and lateral geniculate nucleus (LGN) involves the combination of multiple stimulus features, due to the presence of inhibitory interneurons at each stage. However, it has not been clear what role such inhibition plays, and how it contributes to neuronal processing, in part because single receptive-field-based descriptions of these neurons cannot separate the effects of excitation and inhibition. Here we apply a new General Nonlinear Modeling (GNM) framework to simultaneously recorded pairs consisting of an LGN neuron and the retinal ganglion cell (RGC) that provides its main input. This modeling framework associates nonlinear processing with each stimulus-processing element, and thus can combine the influences of multiple stimulus-tuned elements to predict the observed spike train. By recording from successive stages of processing simultaneously, we can furthermore distinguish the processing that occurs in the retina from processing that occurs in the LGN, and understand how visual information is successively formatted for the visual cortex. We detect separate putative inhibitory elements that affect processing both in the retina and the LGN. We find that RGCs consistently have a strong inhibitory input that has similar tuning to its excitatory tuning, but delayed in time. This makes RGC responses more precise in time, because the inhibition attenuates the response earlier than the decay of excitation. At the level of the LGN, a second inhibitory input is added, except in this case it is an "opposite-sign", or "pull" inhibition. Additionally, the effects of the "same-sign" inhibition inherited from the retina are much more evident, and combined with the higher threshold and pull inhibition, result in temporally precise, sparse responses. We further probe these mechanisms by looking at their effects as a function of contrast. This reveals that the strength of inhibition changes relative to excitation, and all but disappears at low contrast. This suggests several observed effects of contrast gain control, such as changes in gain, temporal sensitivity, and latency, might be a result of the interplay of excitation with delayed inhibition, and potentially reveals a contrast-independent function of the underlying circuitry. Thus, inhibition likely plays a role at multiple levels in formatting visual information for the visual cortex. By revealing how visual information is "formatted" in the early visual pathway, we provide insight into what is likely relevant to the cortex. Furthermore, we present a general method for probing the role of inhibition in other sensory areas.

doi:

I-95. Towards large-scale, high resolution maps of object selectivity in inferior temporal cortex

Elias B. Issa
 Alex Papanastassiou
 Benjamin B. Andken
 James J. DiCarlo

ISSA@MIT.EDU
 APAPANASTASSIOU@GMAIL.COM
 BEN@ANDKEN.COM
 DICARLO@MIT.EDU

McGovern Inst/Dept of Brain & Cog Sci, MIT

Inferior temporal cortex (IT) has been shown to have large-scale (mm to cm) maps of object category selectivity as well as small-scale (sub-millimeter) organization for object features. These two scales of spatial organization have yet to be linked because they were measured using different techniques (fMRI, optical imaging, and local electrophysiology), each with their own limitations. For example, fMRI has poor spatial resolution while optical imaging has higher resolution at the expense of a narrow field of view of only surface-accessible cortex. Given that much of IT lies inside a major cortical sulcus or at the skull base, what is needed is a method that can access the whole of IT with high resolution. Microelectrode based mapping has such potential: electrodes can reach almost anywhere in IT (high spatial coverage) and record from single cells (high spatial resolution). This potential has not yet been realized because of the difficulty of precisely localizing and co-registering many electrode recordings in vivo. Methods such as histological reconstruction of lesion sites or MRI visualization of electrodes are post-hoc and can introduce additional spatial errors. Here, we have adopted a microfocal stereo x-ray system for localizing electrodes that can be used at an unlimited number of sites and operates virtually in real-time (Cox et al., J. Neurophys. 2008). We have used this system to construct broad scale maps of object category selectivity in IT for comparison to fMRI-based maps. We found a weak but significant correspondence between physiology and fMRI maps collected in the same animal, and this correspondence improved substantially when MUA and LFP signals were smoothed (~3-5 mm) to broader scales suggesting the spatial low pass nature of fMRI. Transformations other than spatial smoothing such as dividing the LFP into power in different frequency bands did not produce noticeable improvement in map correspondence. Currently, we are extending our approach to address fine-scale organization in IT at spatial scales more similar to those obtained in optical imaging studies. Although the ex vivo, skull-based accuracy of our system is 50 microns, in vivo resolution may be limited by tissue movement within the skull. For example, the brain may not be in exactly the same position today as yesterday, and cortex may deform locally during electrode recording. To address these issues, we tracked the movements of implanted internal markers within and across sessions. We also measured the neural 'fingerprint' (selectivity profile across a battery of images) of nearby sites recorded on separate days as an empirical test of how reproducible serial samples are. Finally, we estimated local non-rigid deformations of the brain around the electrode. These measurements will determine the current in vivo, tissue-based accuracy of the serial mapping approach and guide mechanical modeling of brain tissue to compensate position shifts. Going forward, linking broad scale and fine scale maps of neural organization in IT will help reveal structure-function relationships and yield insights into the organization of computation in IT.

[doi:](#)

I-96. Recording a large population of retinal cells with a 252 electrode array and automated spike sorting

Olivier Marre¹
 Dario Amodè¹
 Frederick Soo¹
 Timothy E. Holy²
 Michael Berry¹

OLIVIER.MARRE@GMAIL.COM
 DAMODEI@PRINCETON.EDU
 FSOO@PRINCETON.EDU
 HOLY@WUSTL.EDU
 BERRY@PRINCETON.EDU

¹Princeton University

²Washington University in Saint Louis

Recent theoretical work has suggested that recording the activity of more than 100 neurons in the retina simultaneously might uncover non-trivial collective behavior [1]. Furthermore, understanding the neural code of the retina requires access to the information sent to the brain about a large region of the visual space. For that purpose, we used a dense array of 252 electrodes to record activity in the ganglion cell layer of the salamander retina. The electrode density, which is close to the cell density, has been shown to be high enough to record from nearly all the ganglion cells in a patch of retina for smaller arrays [2]. The large number of electrodes precludes doing spike sorting by hand. We thus designed a highly automated algorithm to extract spikes from these raw data. The algorithm was composed of two main steps: 1) a "template-finding" phase to extract the cell's templates, i.e. the pattern of activity evoked over many electrodes when one ganglion cell fires an action potential; 2) a "fitting" phase where the templates were matched to the raw data to find the location of the spikes. For the template-finding phase, we started by detecting all the possible times in the raw data that could contain a spike. Using the minima and maxima values in the neighborhood of the spike on each electrode, spikes were clustered into groups. We then extracted the template corresponding to each group by a least-square fitting method. In the fitting phase, we matched the templates to the raw data with a method that allowed amplitude variation for each template. For that purpose, we selected the best-fitting template, and decide to include it in the match or not according to a criterion that compared the fitting improvement with a cost function. This latter aimed at forcing the spike amplitudes to be close to 1, and imposing a sparseness constraint corresponding to the fact that the overlapping of many spikes is highly unlikely. This process was then iterated to match additional templates to the raw data. Since a first pass clustering did not capture all the cell's templates, we then repeated these two steps. After the fitting part, we did another clustering by taking the minima and the maxima for each putative spike after having subtracted the surrounding contribution of the other templates. This improved clustering made possible the extraction of new templates, then leading to better fits to the raw data. This alternation of clustering and matching was then run iteratively until no additional templates were found. We tested our algorithm by generating surrogate data where we add an artificial template to the data, and trying to recover these artificial events with the algorithm. The ratio of events recovered could reach 99 %when the template was successfully extracted. [1] G Tkacik, E Schneidman, MJ Berry II & W Bialek, q-bio.NC/0611072 (2006). [2] Segev R, Goodhouse J, Puchalla J & Berry MJ 2nd, Nat. Neur. (2004).

doi:

I-97. Models for the mechanisms of perceptual learning: linking predictions for brain and behavior

Michael Wenger¹
Von Der Heide Rebecca²

MJW19@PSU.EDU
RJV151@PSU.EDU

¹Department of Psychology, The Pennsylvania State University
²The Pennsylvania State University

Standard indicators of the acquisition of visual perceptual expertise include systematic reductions in detection and identification thresholds, along with decreases in mean response times (RTs). Two additional patterns have emerged in recent studies of perceptual learning for gray-scale contrast: systematic increases in false alarm rates in detection (but not identification), and systematic increases in the ability to adapt to variations in perceptual workload (capacity, as measured at the level of the hazard function of the RT distribution). The present effort is an initial step in developing a modeling approach capable of accounting for these behavioral results—with a specific focus on changes in capacity in the present effort—while simultaneously predicting patterns of scalp-level EEG. The approach is intended to allow for the representation of multiple competing hypotheses for the neural mechanisms responsible for these observable variables (i.e., placing the alternative hypotheses on a "level playing field"), and for the ability to systematically relate these hypotheses to formal models for perceptual behavior. The neural modeling approach uses populations of discrete-time integrate-and-fire neurons, connected as networks. The architecture is based on the known circuitry of early visual areas as well as known connectivity into and out of early visual areas. The architecture is shown to be capable of instantiating a set of prominent competing hypotheses for neural mechanisms (Gilbert, Sigman, & Crist, 2001): changes in cortical recruitment, sharpening of

feature-specific tuning curves, changes in synaptic weightings, changes in within-region synchrony, and changes in across-region coherence, in both feed-forward and feed-back relations. In addition, it is shown that under reasonable simplifying assumptions, the models are also capable of making predictions for both observable response behavior and scalp-level EEG. Analysis of the computational models for the set of contrasting hypotheses reveals that (a) although all of the hypotheses are capable of accounting for some of the standard empirical regularities, at least one (the cortical recruitment hypothesis) is shown to be unable to account for all those empirical regularities, ruling it out as a general explanation for the neural mechanisms; and (b) many of the standard measures of EEG (e.g., peak values of early negative and positive components) are unable to distinguish among the competing hypotheses; and (c) measures of synchrony across and within regions do offer potential for testing among competing hypotheses.

doi:

I-98. Disparity tuning of the population responses in the human visual cortex: an EEG source imaging study

Benoit R. Cottureau
Anthony M. Norcia
Suzanne P. McKee

COTTEREAU@SKI.ORG
AMN@SKI.ORG
SUZANNE@SKI.ORG

The Smith-Kettlewell Eye Research Institute

The perception of depth is based on horizontal binocular disparities. In the last 15 years, an increasing number of studies have presented disparity tuning functions of single neurons located in different visual areas. However, it is difficult to determine the overall population response of an area from the tuning functions of individual neurons. In this study, we estimated population-level disparity tuning functions of different visual areas in humans using visual evoked potentials and source localization methods. Using dense dynamic random dot patterns, we modulated a disparity-defined central disk (5° diameter) at 2 Hz back and front across a static annulus (12° diameter) presented in the fixation plane. We recorded the evoked response amplitude to disparities ranging from ± 0.5 arcmin to ± 64 arcmin within fMRI-defined ROI's across the visual cortex. Based on the average signal-to-noise ratios of 12 subjects, we found a tuning function in V1 that increases with disparity up to a maximum between 4 and 16 arcmin and then decreases to the noise level at 64 arcmin. This tuning function was measured at the second harmonic, which was the dominant response component, reflecting an equal response to changes between crossed and uncrossed disparities. The observed tuning function in V1 agrees with the population firing rate calculated by Backus (Backus et al., 2001) from the data on macaque V1 neurons (Prince et al., 2002). In addition to the responses in V1, other visual areas in both the dorsal and ventral pathways exhibit tuned responses to disparity modulation that also reached a peak between 4 and 16 arcmin. Interestingly, all these areas show virtually identical signal-to-noise ratios at all measured disparities, except for visual area V3A. Its signal-to-noise ratio is significantly higher than V1 at its peak, which suggests that this area may have a special role in processing horizontal disparity (Backus et al., 2001, Tsao et al., 2003). The absence of responses to large disparities (at 64 arcmin) in all visual areas is consistent with psychophysical measurements of dMax in random dot displays; dMax is the largest disparity that permits a discrimination between front and back (Glennnerster, 1998). References: 1 - Prince SJ, Pointon AD, Cumming BJ and Parker AJ (2002). Quantitative Analysis of the Responses of V1 Neurons to Horizontal Disparity in Dynamic Random-Dot Stereograms. *J. Neurophysio.* 87: 191-208. 2 - Backus BT, Fleet DJ, Parker AJ and Heeger DJ (2001). Human Cortical Activity Correlates With Stereoscopic Depth Perception. *J. Neurophysio.* 86: 2054-2058. 3 - Tsao et al. (2003), Stereopsis Activates V3A and Caudal Intraparietal Areas in Macaques and Humans. *Neuron*: 39, 555-568. 4 - Glennnerster, A. (1998). dmax for Stereopsis and Motion in Random Dot Displays. *Vision Res.* 38(6): 925-935.

doi:

I-99. Sources of response variability underlying contrast-invariant orientation tuning in visual cortex

Srivatsun Sadagopan¹

Nicholas Priebe²

Ian Finn³

David Ferster³

S-SADAGOPAN@NORTHWESTERN.EDU

NICHOLAS@MAIL.UTEXAS.EDU

IFINN005@NORTHWESTERN.EDU

FERSTER@NORTHWESTERN.EDU

¹Dept. of Neurobiology and Physiology, Northwestern University

²The University of Texas at Austin

³Northwestern University

Simple cells of cat primary visual cortex (V1) exhibit orientation tuning that is invariant to stimulus contrast. Several models have attributed this invariance to feedback or feed-forward inhibition. Recently, however, we have demonstrated that two factors are critical in generating contrast invariant spike outputs from contrast invariant synaptic inputs: 1) that the trial-to-trial variability in membrane potential response smooths the effective relationship between membrane potential and spike rate, and 2) that the trial-to-trial variability decreases with increasing contrast. The source of this contrast-dependent change in membrane potential (Vm) noise is as yet unknown and we have therefore investigated two possible sources: recurrent cortical activity, and feed-forward activity from the lateral geniculate nucleus (LGN). To test for a cortical source, we first recorded intracellularly in-vivo from simple cells and measured Vm responses evoked by briefly flashed sinusoidal gratings. At low contrasts, the distribution of peak Vm was broader or right-skewed, indicating increased trial-to-trial variability compared to high-contrast gratings. We then inactivated the cortical circuit locally using electrical stimulation. In many cases, cortical inactivation had no effect on the contrast-dependent change in variability, suggesting that the local cortical activity, whether excitatory or inhibitory, was not a source of increased trial-to-trial variability. We then asked whether the increase in Vm variability could be explained by feed-forward mechanisms, i.e., the trial-to-trial variability in the responses of LGN neurons. To test this hypothesis, we recorded from pairs of neurons in the LGN with nearby or overlapping receptive fields. Preliminary data show increases in response variability with decreasing contrast in individual LGN neurons. In addition, the correlation between LGN responses in nearby cells increases at low contrasts. Together, these results suggest that response variability in cortical simple cells, and therefore contrast invariance in simple cells, arises primarily from the contrast dependence of spike-count variability and correlations in feed-forward inputs.

[doi:](#)

I-100. Lateral Occipital cortex responsive to local correlation structure of natural images

H. Steven H. Steven¹

Sennay Ghebreab¹

Arnold Smeulders²

Victor Lamme³

H.S.SCHOLTE@UVA.NL

S.GHEBREAB@UVA.NL

A.W.M.SMEULDERS@UVA.NL

V.A.F.LAMME@UVA.NL

¹University of Amsterdam, Dep. of Psychology

²University of Amsterdam, Dep. of Informatics

³University of Amsterdam, Dep of Psychology

It is clear from behavioral experiments that subjects can rapidly access information about visual scenes (Potter, 1976) and that different types of scenes (such as beaches and mountains) differ in terms of low level image statistics (Torralba & Oliva, 2003). Furthermore, the distribution of local contrasts in natural images adheres to the Weibull distribution (Geusebroek & Smeulders, 2005), which is a family of distribution deforming from power-law to normal with two free parameters, beta and gamma. The beta parameter indicates the scale of the distribution, whereas the gamma parameter represents its shape. Spatially coherent scenes with one or a few objects tend to

have a low gamma value, i.e. the distribution of their local contrast values approximates a power-law. In contrast, cluttered scenes with many uncorrelated visual structures typically have a high gamma values corresponding with a Gaussian distribution. We recently showed that the brain is capable of estimating the beta and gamma value of a scene by summarizing the X and Y cell populations of the LGN (Scholte et al., 2009). Here we investigated to what degree the brain is sensitive to differences in the global correlation (gamma) of a scene by presenting subjects with a wide range of natural images while measuring BOLD-MRI. Covariance analysis of the single-trial BOLD-MRI data with the gamma parameter showed that only the lateral occipital cortex (LO), and no other areas, responds stronger to low gamma values (corresponding to images with a power-law distribution) than high gamma values (corresponding to images with a normal distribution). The analysis of the covariance matrix of the voxel-pattern cross-correlated single-trial data further revealed that responses to images containing clear objects are more similar in their spatial structure than images that do not contain objects. This data is consistent with a wide range of literature on object perception and area LO (Grill-Spector et al., 2001), and extends our understanding of object recognition by showing that the global correlation structure of a scene is (part of) the diagnostics that are used by the brain to detect objects.

doi:

II-1. Beyond linear perturbation theory: the instantaneous response of the integrate-and-fire model

Moritz Helias¹
Moritz Deger²
Stefan Rotter²
Markus Diesmann¹

HELIAS@BRAIN.RIKEN.JP
DEGER@BCF.UNI-FREIBURG.DE
STEFAN.ROTTER@BIOLOGIE.UNI-FREIBURG.DE
DIESMANN@BRAIN.RIKEN.JP

¹RIKEN Brain Science Institute

²Bernstein Ctr. f. Comp. Neurosci. Freiburg

The integrate-and-fire neuron model with exponential postsynaptic potentials is widely used in analytical work and in simulation studies of neural networks alike. For Gaussian white noise input currents, the membrane potential distribution is known exactly [1]. The linear response properties of the model have successfully been calculated and applied to the dynamics of recurrent networks in this diffusion limit [2]. However, the diffusion approximation assumes the effect of each synapse on the membrane potential to be infinitesimally small. Here we present a novel hybrid theory that takes finite synaptic weights into account. We show, that this considerably alters the absorbing boundary condition at the threshold: the probability density increases just below threshold. As a result, the response of the neuron to a fast transient input is enhanced much in the same way as found for the case of synaptic filtering [3]. However, in contrast to this earlier work relying on linear perturbation theory [4], we quantify to all orders an instantaneous response that is asymmetric for excitatory and inhibitory transients and exhibits a non-linear dependence on positive perturbation amplitudes. Furthermore we demonstrate that in the pooled response of two neuronal populations to antisymmetric transients the linear components exactly cancel. In this scenario the macroscopic network dynamics is dominated by the instantaneous non-linear components of the response. These results suggest that the linear response approach neglects important features of the rectifying nature of threshold units with finite jumps even for small perturbations. We provide an analytical framework to go beyond [5]. Partially funded by BMBF Grant 01GQ0420 to BCCN Freiburg, EU Grant 15879 (FACETS), DIP F1.2, Helmholtz Alliance on Systems Biology, and Next-Generation Supercomputer Project of MEXT. [1] L. M. Ricciardi and L. Sacerdote. The Ornstein-Uhlenbeck process as a model for neuronal activity. *Biol. Cybern.*, 35:1-9, 1979. [2] N. Brunel and V. Hakim. Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Comput.*, 11(7):1621-1671, 1999. [3] N. Brunel, F. S. Chance, N. Fourcaud, and L.F. Abbott. Effects of synaptic noise and filtering on the frequency response of spiking neurons. *Phys. Rev. Lett.*, 86(10):2186-2189, 2001. [4] B. Lindner and L. Schimansky-Geier. Transmission of Noise Coded versus Additive Signals through a Neuronal Ensemble. *Phys. Rev. Lett.*, 86(14):2934-2937, 2001 [5] M. Helias, M. Deger, S. Rotter, and M. Diesmann. A Fokker-Planck formalism for diffusion with finite increments and absorbing boundaries. *arXiv:q-bio (0908.1960)*, 2009

doi:

II-2. Gamma oscillations in the optic tectum in vitro represent top-down drive by a cholinergic nucleus

C. Alex Goddard¹
 Devarajan Sridharan^{1,2}
 John Huguenard³
 Eric Knudsen¹

CGODDARD@STANFORD.EDU
 DSRIDHAR@STANFORD.EDU
 JOHN.HUGUENARD@STANFORD.EDU
 EKNUDSEN@STANFORD.EDU

¹Dept of Neurobiology, Stanford University

²Dept. of Bioengineering, Stanford Univ

³Dept of Neurology, Stanford University

Gamma-band (25-90 Hz) oscillations are induced in neuronal networks both by attention in vivo and by the attention-related neurotransmitter acetylcholine in vitro. Visual and auditory stimuli drive gamma band oscillations in the avian optic tectum, the homolog to the mammalian superior colliculus, and a structure implicated in attention and orienting movements. The cholinergic nucleus isthmi pars parvocellularis (lpc) has been implicated in these oscillations. In an in vitro slice preparation of the chicken optic tectum, we have discovered oscillations that are nearly identical to those observed in vivo. Here we test the requirement for lpc and acetylcholine (ACh) for the oscillations. We report that while lpc is required, ACh is not; persistent activity in the intermediate (multimodal) layers of the tectum drives lpc, which in turn drives oscillations in the superficial (visual) layers of the tectum. Slices that maintained connectivity with lpc exhibited gamma band oscillations in the visual layers of the tectum. These oscillations, which persist for hundreds of milliseconds, occurred spontaneously or following retinal afferent stimulation. They did not require the addition of exogenous acetylcholine or kainate, as is often the case in other in vitro preparations of the hippocampus and neocortex. The responses consisted of high frequency (> 500 Hz) bursts of spikes that were phase locked to a ~30 Hz oscillation, and were characterized by concurrent activity in all tectal layers and in the lpc. Blockade of synaptic transmission with CNQX in the lpc blocked the oscillations, as did transection of connections between the tectum and the lpc. Surprisingly, blockade of cholinergic transmission with nicotinic and muscarinic antagonists had little effect on the oscillations. Thus, lpc is required for the oscillation, but ACh is not. As ACh is not required for the oscillations, we looked for activity that drives the oscillation. In slices that had no oscillations in the visual layers, persistent activity was observed in the multimodal layers of the tectum. Neurons in the multimodal layers are known to project to the lpc. Three observations suggest that persistent activity in the multimodal layers drives the periodic activity in the lpc: 1) no periodic bursting occurred in the lpc without intact connections to the tectum, 2) multimodal layer neurons fire at ~30 Hz, and 3) EPSPs in lpc were periodic during an oscillation. Thus, the persistent activity of neurons in the multimodal layers periodically drives lpc, which in turn drives gamma band oscillations in the visual layers of the tectum. These findings suggest that the multimodal layers drive the lpc to induce oscillations in the visual layers of the tectum; this pathway represents a mechanism for top-down modulation of visual responses within the tectum. This study also provides a mechanistic understanding of how a cholinergic nucleus is involved in gamma oscillations in the context of an in-vivo-like oscillation. Ultimately, we hope to use this preparation to explore how gamma oscillations and acetylcholine shape information processing in a sensory microcircuit.

doi:

II-3. Parallel channels in the OFF visual pathway emerge at the cone synapse

Charles P. Ratliff
 Steven H. DeVries
 Northwestern University

C-RATLIFF@NORTHWESTERN.EDU
 S-DEVRIES@NORTHWESTERN.EDU

The brain's computational power derives from its massively parallel information pathways. In the visual system, these parallel pathways emerge at the first synapse in the retina, where each cone releases glutamate onto 10 or more types of bipolar cell. Bipolar cells are divided into ON and OFF subtypes, which possess either metabotropic (ON) or ionotropic (OFF) glutamate receptors, and consequently respond with opposite polarity to increments and decrements of light. OFF bipolar cells are further divided based on receptor type (AMPA vs. Kainate), type of synaptic contact (basal vs. invaginating), and number of synaptic contacts. Our previous work has shown that for basic laboratory stimuli, each type of OFF bipolar cell receives a different signal at the cone synapse. Here we show that for more physiologically realistic stimuli, each type of bipolar cell receives a distinct synaptic signal. To study signal transfer at the cone synapse, we performed patch-clamp recordings from pairs of synaptically connected cones and bipolar cells in slices from the cone-dominated ground squirrel retina. Cones were recorded in perforated patch (using amphotericin or beta-escin) and stimulated in voltage clamp to follow a graded voltage signal. Bipolar cells were recorded in voltage clamp at -70 mV to eliminate voltage-gated currents. Picrotoxin and strychnine were used to eliminate inhibitory ligand-gated currents, and thus the current recorded from the bipolar cell represented an isolated synaptic response. The input-output relationship of the synapse was assessed using white noise stimuli (variable mean, standard deviation 2.5 mV or 3 mV, cutoff frequency 100 Hz or 500 Hz), and was modeled using a linear filter followed by a static nonlinearity. Model performance was tested for recordings where the cone also received a natural stimulus. All stimuli were repeated to measure signal-to-noise ratio (SNR) and estimate information rate. We studied three types of OFF bipolar cells (b2, b3 and b7), and found several consistent differences in their characteristic synaptic response. Gaussian stimuli produced skewed responses with longer tails for inward currents, and showed significant kurtosis (b2: 3.8 ± 1.0 ; b3: 16 ± 8.4 ; b7: 5.6 ± 2.9 ; mean $\pm$ SEM). Time-to-peak of the linear filter was different in each type (b2: 2.7 ± 0.2 ms; b3: 6.2 ± 0.4 ms; b7: 4.4 ± 0.4 ms; mean $\pm$ SEM). The linear filter was biased toward the highest frequencies in the b2, and the lowest frequencies in the b3. For each cell type, the linear-nonlinear model provided a reasonable description of responses, but failed to describe (1) the transience of responses, (2) differences in response timing evoked by cone voltage steps of different magnitude, (3) history-dependent differences in response amplitude caused by desensitization of receptors, and (4) the time-course and magnitude of the asymmetry between inward and outward currents. From this data, we conclude that a single cone signal evokes distinct synaptic currents in different OFF bipolar cell types. Each bipolar cell type forms an independent array, and each array forms a different representation of the visual image.

doi:

II-4. The logic of cross-columnar interactions along horizontal circuits

Hillel Adesnik^{1,2}
Massimo Scanziani¹

HILLEL.ADESNIK@GMAIL.COM
MASSIMO@BIOMAIL.UCSD.EDU

¹HHMI
²UCSD

The cerebral cortex constructs a coherent representation of the world by integrating distinct features of the sensory environment. While these features are initially processed in discrete cortical modules, called columns, horizontal connections between columns allow this information to be processed in a context dependent manner, ultimately resulting in a meaningful perception. Despite the wealth of physiological and psychophysical studies addressing the function of horizontal projections, the cellular mechanisms by which they coordinate activity among cortical columns are poorly understood. To address this question we selectively activated horizontal projection neurons in mouse somatosensory cortex, in vivo and in vitro, and determined how the resulting domains of excitation and inhibition impact columnar activity. Surprisingly, we found that horizontal projections suppress the superficial layers while simultaneously activating the deeper output layers of cortical columns. This layer specific modulation of activity does not result from a spatial separation of excitatory and inhibitory domains, but rather from a layer specific difference in the ratio between these two opposing, yet precisely overlapping conductances. Hence, through this novel mechanism of cross-columnar interaction individual columns exploit horizontal projections to compete for cortical space.

doi:

II-5. Bang-bang optimality of energy efficient spikes in single neuron models

Biswa Sengupta¹
 Martin Stemmler²
 Jeremy Niven¹
 Andreas Herz²
 Simon Laughlin¹

BS393@CAM.AC.UK
 STEMMLER@BIOLOGIE.UNI-MUENCHEN.DE
 JEN22@CAM.AC.UK
 HERZ@BIO.LMU.DE
 SL104@CAM.AC.UK

¹University of Cambridge

²BCCN Munich, LMU Munich

About 50-80% of the total energy allocated to the mammalian brain is used for signaling, mainly to drive the Na^+ / K^+ pump (Attwell & Laughlin, 2001). Due to the substantial contribution of action potentials to the energy consumption, the biophysical properties generating an action potential can be matched to make them energy efficient (Alle et al., 2009). By combining different voltage-dependent ionic conductances with varying biophysical properties, different neurons express a myriad of action potential shapes characterized by varying heights and widths (Bean, 2007). How widespread are energy efficient action potentials and what are the major factors defining the energy efficiency of a single action potential? The passive electric properties of a neuron, namely its capacitance and input resistance, set a lower bound for the cost of a given action potential. Inward (Na^+) and outward (K^+) voltage-dependent currents act in a push-pull manner to change the voltage; these currents will invariably overlap, which can cause the energetic cost to rise to more than eleven-fold the baseline cost. We use seven single compartment models to assess the energy cost of vertebrate (cerebellar granule neurons, cortical interneurons, thalamo-cortical relay neurons and hippocampal interneurons) and invertebrate (squid axons, crab axons and bee kenyon cells) action potentials. We show that the energy consumption of a single action potential is directly dependent on the overlap between the Na^+ and K^+ currents that generate them. By optimizing the ionic conductance parameters that affect this overlap, we investigate how close each neuron type is to the minimal energy expenditure. Using sensitivity analysis and perturbation theory for periodic systems we show that an optimal model minimizes the overlap between Na^+ and K^+ currents. Just as in optimal control theory, the ionic conductance parameters allow for a bang-bang control of the underlying Na^+ and K^+ currents eliciting the action potential. Energy efficient action potentials are neuron-dependent but generally have narrower width, shorter height and exhibit an extremely rapid, explosive onset. Our modeling shows that in some neurons, such as the thalamo-cortical relay neurons, the hippocampal GABAergic interneurons or fast-spiking cortical interneurons, the properties of the currents that minimize the energy cost of the action potential are close to the experimentally measured currents, suggesting selection for reducing the energy costs. The experimentally observed properties of the squid giant axon model however, consumes substantial energy when compared to the optimal model that consumes far lower energy. We suggest that the deviation between the parameters generating the theoretical minimum cost of the action potential and those observed in neurons is dependent upon the function of the neuron.

doi:

II-6. Cognitive control recruits theta activity in anterior cingulate cortex for establishing task rules

Thilo Womelsdorf¹

Kevin Johnston²

Martin Vinck³

Stefan Everling²

THIWOM@IMAGING.ROBARTS.CA

KEVINJ@BIOMED.QUEENSU.CA

MARTIN.VINCK@FCDONDERS.RU.NL

SEVERLIN@UWO.CA

¹Department of Physiology & Pharmacology, University of Western Ontario

²Robarts Research Institute

³University of Amsterdam

Accomplishing even simple tasks depends on neuronal circuits to configure how incoming sensory stimuli map onto responses. During task performance these stimulus-response (SR) mapping rules evolve in an attentional control network comprising the anterior cingulate cortex (ACC). However, neither (i) the precise contribution of the ACC to control SR mapping is understood, nor (ii) are the mechanisms underlying the instantiation of SR mapping rules established. Existing studies have shown that neurons within the ACC encode two function of cognitive control: They convey information about currently relevant SR mapping rules, and they provide corrective signals upon incorrect performance, required to optimize efficient SR mapping in future trials. Here, we tested with a task-rule shifting paradigm whether both cognitive control aspects are evident in rhythmic neuronal synchronization, which could serve as the substrate to efficiently link neuronal groups across a larger attentional control network. We recorded from 108 neuronal sites within ACC in two monkeys, performing blocks of trials requiring either pro- or anti-saccades towards/away from a peripheral stimulus. SR mapping rules changed without overt cue upon 30 correct trials on either task. Prior to peripheral stimulus onset, monkeys fixated during a 1 sec. preparatory period, allowing the analysis of task-selective neuronal activity. For time-resolved estimates of oscillatory activity we calculated the power of the local field potentials (LFPs) based on Hanning tapered Fourier transforms during the preparatory period. To test for the influence of LFP theta activity to synchronize the neuronal spiking output we calculated spike-LFP phase consistency with a novel bias-free measure of neuronal coherence. We find that two core aspects of cognitive control are encoded by rhythmic theta activity (5- 11 Hz). First, in 50% of the neuronal sites theta activity was task-selective, predicting which of two SR mapping rules will be established prior to processing visual target information. For task-selective neuronal groups, theta activity was evident during all trials following a task-rule change. Selective theta activity predicted correct task performance with stronger theta activity in correct vs. incorrect trials. Secondly, theta activity was stronger in correct trials that followed error trials, rather than other correct trials, showing that enhanced control demands to update SR mapping rules are reflected in theta modulation. This conclusion is further supported by a subset 15% of sites showing significant theta modulation only in trials immediately following a task change. Importantly, LFP theta activity had consequence for neuronal spiking output with significant spike-LFP phase locking for a large subset of single neurons. Spike-LFP phase locking increased with increasing LFP theta activity, suggesting that cognitive control recruits neuronal synchronization in the theta frequency band. These results show that two aspects of cognitive control are reflected in selective neuronal theta synchronization in the ACC: (i) The instantiation of selective preparatory information signaling the relevant SR mapping rule, and (ii) corrective signals following incorrect SR mapping. We propose that task-selective theta modulation arises in superficial layers to functionally link the mosaic of areas subserving efficient cognitive control.

doi:

II-7. Neural mechanisms of competitive interaction in recurrent maps of a visual pathway

Dihui Lai

Ralf Wessel

Dept. of Physics, Washington University

DLAI@PHYSICS.WUSTL.EDU

RW@PHYSICS.WUSTL.EDU

Visual scenes typically contain multiple objects. Because of the retinotopic organization of the early visual system, objects at different locations are represented by neural populations with partial or complete spatial separation. The competitive interaction between the populations allows the organism to shift activity, and thus attention, towards relevant locations. Although phenomenological models of shifts in spatial attention succeed to reproduce many observations, a mechanistic understanding of how these spatial shifts in activity are achieved at the circuit level is still lacking. To fill this gap, we investigated the dynamics of the avian isthmotectal system. This model preparation was chosen, because oscillatory bursts in response to visual stimuli, shifts in activity to novel visual stimuli, and winner-take-all representation of stronger visual stimuli have been demonstrated and quantified in the isthmotectal system of pigeon and owl. Furthermore, because of its anatomically well-characterized circuitry, the avian nucleus isthmi (parabigeminal nucleus in mammals) and the optic tectum (superior colliculus in mammals) are ideally suited for an investigation of the mechanisms of competitive interaction at the circuit level. A key anatomical feature of the isthmotectal system is that local (cholinergic lpc neurons) and more global (GABAergic lmc neurons) topographic visual information is superimposed. To gain insight into the dynamics of competitive interaction in the avian isthmotectal system, we designed a model network of 1200 leaky integrate-and-fire neurons with spike-rate adaptation. The connectivity of this model network is constrained by anatomical information. The cellular properties of the three model neuron types are constrained by electrophysiological data from chick midbrain slice experiments. Simulations of the experimentally constrained model network reproduce the three in vivo observations: (i) the oscillatory bursts during visual stimulation, (ii) the shift in activity to novel stimuli even when the old stimulus remains present, and (iii) the disproportional high responses to the larger of two stimuli in a winner-take-all manner. Further, changing the model parameters, we found that spike rate adaptation is crucial for the network ability to be sensitive to novel stimuli, even when the novel stimulus is weaker than the existing stimulus. Without spike rate adaptation the circuit reduces to a winner-take-all network, where the competitive interaction is mediated by the recurrent lateral inhibition. In conclusion, this computational analysis reveals how the combination of network architecture, fast synaptic inhibition, and slow cellular spike-rate adaptation mediates the competitive interaction of spatially-separate neural populations. More generally, this investigation provides insight into how the coordinated activity of competing populations of neurons represents spatiotemporal stimuli.

[doi:](#)

II-8. Attention modeled as a two-dimensional neural resource

Dana Ballard

DANA@CS.UTEXAS.EDU

University of Texas at Austin

Although hundreds or even thousands of experiments invoke the concept of 'attention,' there is scant specification of its function at the neural level beyond the increased firing of neurons with respect to different task and stimuli conditions. We attempt a more precise computational definition of neural attention with the aim of accounting for data in different kinds of neural recording experiments and find we require two distinct descriptive axes. One measures the numbers of neurons in a network that are allocated to a task. In a network processing quantitative data, we show that increased numerical precision can be traced to additional neurons, and consequently the increased number of spikes in any individual neuron. A second axis reflects the effects of competition between two or more tasks. Experimental data can be simply explained by positing that individual neurons' spike output is time-shared between tasks. We make these assertions concrete by demonstrating them in a model circuit that learns striate cortex receptive fields from LGN data using gamma-phase coding [1] and a probabilistic version of matching pursuit [2]. gamma-phase coding represents quantitative data with each spike by using a small (0~5ms) lag with respect to a gamma oscillation peak. Probabilistic matching pursuit routes this numerical data through different neurons at different cycles in the gamma signal. Thus while spikes appear random from the vantage point of any particular cell, a deterministic numerical message is sent through the network. A large number of experiments model attention as a gain change wherein the receptive field of a neuron as measured by its peri-stimulus histogram is modified by a scalar multiplier [3]. gamma-phase coding allows numerical data to be sent with different levels of precision. Thus at any instant, a more precise numerical code can be sent by including more neurons with longer lags. However our simulations show that a side effect of this increase in precision is that the receptive field of a neuron, as measured by the peri-stimulus histogram, represents a scalar gain

change. When a neuron has a complex receptive field such as those found in cortical area V4, attending to a subfield can increase its firing rate, but attending outside of its receptive field results in a normalized response to all its subfields' components. For example in the classic experiment of Desimone et al [4] recording from a neuron in V4 with subfields A and B, attending to A produced spikes $s(A)$, but attending outside the receptive field produced spikes $(s(A) + s(B))/2$. A single circuit can explain this phenomenon if its neurons' subfields are coded by separate gamma oscillators that differ in overall phase. Thus in the case of attending to A, the stimulus successfully monopolizes a neurons gamma-phase time slots, whereas in attending outside of the neurons RF, A and B divide the time slots between them. 1. *TRENDS in Neurosciences* 30, 309 (2007) 2. *PLoS Comput Biol* 5(5): e1000373 (2009) 3. *Neuron*, 23, 765, (1999) 4. *Annual Rev. Neuroscience*, 18, 193 (1995)

[doi:](#)

II-9. Evidence for a race model

Dorion Liston¹

Leland Stone²

DLISTON@ARC.NASA.GOV

LSTONE@ARC.NASA.GOV

¹San Jose State University Foundation

²NASA Ames Research Center

In statistical decisions, various decision rules can be equivalent (Green and Swets, 1966). Consider a generic magnitude discrimination between two stimuli (a and b). This discrimination could be implemented in at least two ways: 1) monitor separate magnitude estimates over time for a and b and respond when the first one reaches a threshold (race mechanism), or 2) monitor an estimate of the difference in magnitude between the two and respond when the delta reaches either a positive or negative threshold (difference mechanism). These two decision rules predict identical behavioral performance, but different patterns of response times: the difference model predicts that both stimuli have equal and opposite effects on response time whereas the race model predicts that the selected stimulus has the greatest impact on response time. Methods: For saccade choices in a 2AFC spatial brightness discrimination task, we compared the correlation between reciprocal response time (response rate) and the "absolute" signal strength of the selected and unselected stimulus, and fit the response-time data with a simple model that can scale between a race mechanism and a difference mechanism. Results: For the correlation analysis, we observed that response rate was better correlated with the strength of the selected stimulus (average r value: 0.10) than the strength of the unselected stimulus (average r value: 0.01). For the joint fit analysis, we observed a strong positive relationship between response rate and the selected stimulus (average regression parameter value: 0.09), consistent with both mechanisms, but also observed a weak positive relationship between response rate and the unselected stimulus (average regression parameter value: 0.02). Conclusion: These results are inconsistent with a difference mechanism, but can be explained by a race between decision signals.

[doi:](#)

II-10. Adaptive properties of differential learning rates for positive and negative outcomes

Romain Cazé¹

van der Meer Matthijs²

ROMAIN.CAZE@ENS.FR

MVDM@UMN.EDU

¹Group for Neural Theory, ENS, Paris

²Department of Neuroscience, University of Minnesota

A central concept in theories of Pavlovian and instrumental learning alike is the prediction error, which signals how much better or worse than expected an outcome turned out to be. The impact of prediction errors is controlled

by a "learning rate" parameter, commonly a single value for positive and negative outcomes. However, a single learning rate may not be an accurate description of how humans actually learn. For instance, Frank et al. (2007) found that subjects learned differentially from positive and negative outcomes on a probabilistic two-choice task. Furthermore, genetic polymorphisms associated with striatal D1 and D2-receptor pathways were independently predictive of learning rates associated with positive and negative outcomes, suggesting that differential learning rates may be dissociable at the neural level. While the computational consequences of such differential learning rates have rarely been studied, a common suggestion is that such biases and related asymmetries such as loss aversion are irrational (e.g. Kahneman and Tversky 1972). In this study we sought to identify conditions in which differential learning rates may be adaptive, by comparing the performance of three reinforcement learning agents: one that learns more from positive than negative outcomes (gain learner), its opposite (loss learner), and one that learns equally from both, on a variety of probabilistic choice tasks. We found that when two choices had a high (but different) probability of reward (0.8 and 0.9), the loss learner performed best, while when both choices had a low probability of reward (0.1 and 0.2) the gain learner performed best. In both situations, differential learning rates enabled a better separation of learned reward probabilities, compared to the normal agent's convergence near to the true reward probabilities, which are close together and promote instability in the face of stochastic rewards. We derived analytical expressions for the reward obtained in the steady state as a function of the two learning rates and the distribution of rewards, and show that these results hold independently of the action selection mechanism used (epsilon-greedy or softmax). Thus, from a reinforcement learning perspective, there are situations in which an agent with different learning rates for positive and negative outcomes performs better than an agent with a single, symmetric learning rate. These results suggest that having different neural systems independently support learning from positive and negative outcomes with potentially different learning rates can in fact be adaptive even in simple choice situations. However, real world learning situations often involve more complex operations than the processing of prediction errors alone; for instance, in serial reversal learning, subjects do not "unlearn" previously learned associations but learn to switch between different world-states. Nevertheless, there is a wide range of proposals in psychology and economics that suggest an asymmetric impact of positive and negative outcomes, including not only loss aversion but also variations in happiness set-point (Frederick and Loewenstein, 1999) and optimism bias (Sharot et al. 2007) which can be informed by this reinforcement learning approach.

[doi:](#)

II-11. A Bayesian model of simple value-based decision-making

Debajyoti Ray
Antonio Rangel
CNS, Caltech

DRAY@CALTECH.EDU
RANGEL@HSS.CALTECH.EDU

The ease with which we make simple decisions belies the complexity of the underlying computational problem. Even simple decisions, such as whether or not to consume a potential reward, pose a challenging problem for the brain: it has to aggregate uncertainty information about the stimulus to determine its value. We study the computational problem faced by the brain in a simple decision making paradigm. We propose a normative (Bayesian) account of the decision problem using a Partially Observable Markov Decision Problem framework. In the model the decision-maker trades search costs over gaining more information about the items that might lead to improvements in the quality of the decision. We test the predictions of the model using an eye-tracking experiment in which subjects make choices over pairs of food stimuli. We show that the model captures the key aspects of the choice, reaction time, and fixation search patterns. Furthermore, we show that the model provides better quantitative fits to the data than alternative descriptive models, such as the drift-diffusion model.

[doi:](#)

II-12. The rational control of aspiration in learning

Daniel Acuna¹
C. Shawn Green²
Paul Schrater²

ACUNA002@UMN.EDU
CSGREEN@UMN.EDU
SCHRATER@UMN.EDU

¹Dept. Computer Science and Engineering, University of Minnesota

²University of Minnesota

One of the fundamental questions for any agent entering a novel environment is how to balance exploratory and exploitative actions. This is especially difficult when the total number of states, the number of rewarded states, and the distribution of rewards at rewarded states are initially unknown. How much to explore in an unknown environment is controlled by the agent's beliefs about the value of unexperienced states, which we refer to as the agent's "aspiration." To examine the effect of aspiration on exploratory behavior in humans, we used a well-known challenging test problem for exploration in reinforcement learning - the "chain game." Briefly, this task produces two reasonable policies - a small reward policy that requires little exploration to find and a large reward policy that requires more exploratory (and unrewarded) actions to find. Human strategies fell into two distinct groups; one group performed enough unrewarded exploratory actions to find the larger reward policy, while the second group under-explored and stuck with the low reward policy without ever experiencing the larger reward state. In debriefing, subjects in the former group reported finding the small local maxima quickly, but believed that higher rewards were possible and thus continued exploration. Conversely, subjects in the latter group typically reported an initial exploratory phase, but upon finding only the local maxima and otherwise only unrewarded states, determined that the optimal solution was to exploit the small local maxima. Using model-based Bayesian reinforcement-learning, an agent can be made to mimic either of these groups by manipulating the agent's initial prior belief about the size of the state space and/or the magnitude of potential rewards. In particular, we model these prior beliefs in terms of hyperparameters of a hierarchical Dirichlet process prior on the entries in the state-action transition matrix. Interestingly, these hyperparameters can be learned from data. Given that aspiration is obviously critical in determining exploratory behavior, the question then arises: what factors determine aspiration in humans? One potential source of information is knowledge regarding the reward history of others. To test the effect of this type of knowledge on exploratory choice behavior, subjects were again placed within an environment with an easy-to-find low-reward policy and hard-to-find high-reward policy. After half the total trials (and convergence to the easy to find policy), a "high-score" sheet was shown to the subjects. One group was shown high-reward values and while the other group was shown low-reward values similar to their own score. While low-reward scores produced no change in behavior in the second half of the experiment, the group exposed to high-reward scores showed a complete reinitialization of exploration. Finally, we also tested whether subjects are capable of inferring aspiration directly from the statistics of the environment. Subjects played repeated games with similar reward structures (i.e., similar probabilities of states being rewarded, similar distributions on reward amount, etc). Choice behavior was extremely sensitive to these statistics (e.g., in sparse environments with few rewarded states, subjects typically exploited the first rewarded state they encountered).

[doi:](#)

II-13. Bonsai trees: How the Pavlovian system sculpts sequential decisions

Quentin J. M. Huys^{1,2}
Neir Eshel³
Peter Dayan⁴
Jonathan P. Roiser⁵

QHUY@GATSBY.UCL.AC.UK
NESHEL@GMAIL.COM
DAYAN@GATSBY.UCL.AC.UK
J.ROISER@UCL.AC.UK

¹Gatsby Unit and Neuroimaging Centre, UCL

²Medical School, UCL

³Harvard University

⁴Gatsby Unit, UCL

⁵Institute of Cognitive Neuroscience, UCL

People face decision problems of gargantuan dimensions daily and happily. Efficient pruning of large decision trees is likely a crucial ingredient of this striking ability. Here, we examine the ability of the Pavlovian system to shape human goal-directed decision making in deep sequential choice scenarios. More specifically, we ask to what extent people optimistically inhibit the evaluation of parts of decision trees that lie below large negative reinforcements. We arrange the cost function so that pruning is counterproductive, thus pitting a reflexive Pavlovian tendency (suppression in the face of punishments) against optimal goal-directed search. Three groups of 15 subjects played a novel computerized task in which they used two buttons to navigate between six states. They first learned that each button led to a particular deterministic transition. For instance, from state 3 button 1 led to state 6 whereas button 2 led to state 4. Subjects then learned the costs of particular transitions. The groups differed in terms of the costs associated with the three most costly transitions (-70, -100 and -140 points respectively). Subjects were then repeatedly dropped in a random state, and asked to produce a varying number of sequential button presses such as to maximise the cumulative rewards earned over the entire sequence. They received a proportion of these points in cash at the end of the experiment. We first compared subjects' behaviour to optimal choices. Subjects chose optimally on over 70% of the trials when up to three choices remained. They were at chance when 6 or more choices remained. We then fit a model with two discounting parameters. The specific discounting parameter applied to reinforcements on subparts of the decision tree below the large punishments; the general discounting parameter applied to other reinforcements. We find that this model captures subjects' choices extraordinarily well. Subjects chose the action associated with the higher expectation (after pruning) on close to 90% of the trials, independently of the number of choices remaining. Importantly, subjects pruned after the large negative reinforcements even when it was strongly advantageous not to; that is even when it was profitable to incur large losses because of the large rewards hidden below them. We then compared subjects' specific pruning to various psychometric variables and found that specific pruning was strongly positively correlated with measures of depression (Beck Depression inventory scores; $p=0.0017$; $r=0.4589$) and strongly negatively correlated with a measure of extraverted personality (Revised NEO Personality Inventory, Extraversion factor; $p=0.00247$, $r=-0.4450$). (Both significant at a Bonferroni corrected threshold of 0.005). This work provides fundamental new insight into specific interactions between goal-directed, sequential choices and the aversive side of Pavlovian systems. It shows how Pavlovian influences sculpt decision trees, thereby influencing the possible outcomes of a goal-directed tree search; and how this process varies in a predictable manner with standard measures of psychopathology and personality. We believe that this is the first specific demonstration of how lowly Pavlovian systems may fundamentally affect – and facilitate – the function of higher cognitive systems.

doi:

II-14. The temporal dynamics of human decision under risk

Laurence Hunt^{1,2}

Matthew Rushworth³

Tim E. Behrens³

LHUNT@FMRIB.OX.AC.UK

MATTHEW.RUSHWORTH@PSY.OX.AC.UK

BEHRENS@FMRIB.OX.AC.UK

¹FMRIB Centre, University of Oxford

²Experimental Psychology, University of Oxford

³University of Oxford

The functional neuroanatomy of human value-based decision has been widely investigated using functional MRI (fMRI). A ventromedial portion of prefrontal cortex (VMPFC) has been isolated as especially important, although its precise contribution remains debated. One view holds that VMPFC codes stimulus values (1) that are then compared elsewhere in the brain (2). Another suggests VMPFC codes the value of the option that will become chosen (3), perhaps to enable subsequent computation of a prediction error (4). In yet other studies, VMPFC signals the difference in value between chosen and unchosen options specifically at choice, not outcome, time (5,6) - arguing for a central role for VMPFC in the formation of a decision. This would suggest that an initial representation of stimulus value in VMPFC may evolve into chosen and unchosen value signals as an option is selected. However, the temporal evolution of this signal may occur too quickly to be resolved using fMRI. To test

this hypothesis, we measured simultaneous magnetoencephalography and electroencephalography (M/EEG) in 19 subjects as they performed a simple economic decision task. Subjects weighed the probability of receiving monetary reward on two options with the magnitude of reward that could be attained. Subject choices matched predictions from descriptive economic models of behaviour, and reaction times correlated negatively with the difference in value between the two options. We used trial-by-trial regression to investigate the effects of several variables on M/EEG data at different timepoints through the trial. At the time decision stimuli were presented, and at the time of button-press, we looked for the effect of 'value difference' (chosen-unchosen values) but also of overall 'stimulus value' (chosen+unchosen values). We employed the multiple sparse priors (MSP) source reconstruction algorithm implemented in SPM8 to investigate the neuroanatomical basis of these signals. Aligned to the presentation of the decision, 'value difference' signals localised to VMPFC, peaking approximately 800ms after stimulus onset. Crucially, however, this signal was immediately preceded by a 'stimulus value' signal, also localised to VMPFC, representing the sum of all options available to the subject (500-800ms). Aligning instead to the time of subject response, 'value difference' signals were found preceding the button-press, approximately 300ms before the action was executed. In contrast to the stimulus-aligned signals, these localised to SMA/cingulate motor area and primary motor cortex, but were not preceded by a 'stimulus value' signal, and so might relate to motor preparation rather than the decision process per se. These results suggest that as a decision is made, the signal within VMPFC evolves from representing the stimulus value of all available options to representing the decision that has been made. The role of VMPFC may therefore go beyond the signalling of the chosen option for subsequent computation of a reward prediction error, and may be fundamental to the process of selecting an option during value-based decision. 1. Plassmann, JNeuro, 2007; 2. Kable, Neuron, 2009; 3. Wunderlich, PNAS, 2009; 4. Schoenbaum, NRN, 2009; 5. Boorman, Neuron, 2009; 6. FitzGerald, JNeuro, 2009

doi:

II-15. Approaching avoidance: asymmetries in reward and punishment processing

Quentin J. M. Huys^{1,2}
Roshan Cools³
Martin Goelzer⁴
Eva Friedel⁴
Ray J. Dolan⁵
Andreas Heinz⁴
Peter Dayan⁶

QHUY@GATSBY.UCL.AC.UK
ROSHAN.COOLS@DONDEERS.RU.NL
MARTIN.GOELZER@CHARITE.DE
EVA.FRIEDEL@CHARITE.DE
R.DOLAN@FIL.ION.UCL.AC.UK
ANDREAS.HEINZ@CHARITE.DE
DAYAN@GATSBY.UCL.AC.UK

¹Gatsby Unit and Neuroimaging Centre, UCL

²Medical School, UCL

³Donders Centre

⁴Charité Universitätsmedizin Berlin

⁵Wellcome Trust Centre for Neuroimaging, UCL

⁶Gatsby Computational Neuroscience Unit

In a Pavlovian conditioning setting, an animals responses do not affect the receipt of reinforcements. Yet, such classically conditioned stimuli (CSs) predictive of affective events have a strong tendency to elicit responses which are hard-wired; generally adaptive; and acquired over an evolutionary time-scale. Pavlovian behaviours can be seen as evolutionarily acquired policy generalizations and may explain many significant quirks of human behaviour, including impulsivity, framing effects, and even psychiatric disturbances including addictions and mood disorders. As such, they have been influential in both elucidating and complicating our understanding of the architecture of appetitive and aversive decision-making. A key lacuna in the data is that the various factors that emerge as being central in this evolving architecture have not been systematically tested. We provide such a test using a Pavlovian-Instrumental Transfer design in humans in which we carefully control: go vs nogo; approach vs withdrawal; individual differences in sensitivities to both rewards and punishments; the relationship between instrumental and Pavlovian expectations. Our first key finding is a highly significant ($p=1e-5$) interaction (see sup-

plemental figure) between the affective valence of Pavlovian stimuli and actions. We show that approach actions are promoted by positive Pavlovian stimuli. In the first demonstration of conditioned suppression in humans, we also show that approach actions are inhibited by aversive Pavlovian stimuli. Most importantly, the same actions, with the same associated expectations, but instantiating withdrawal, are promoted by aversive Pavlovian stimuli, but not by appetitive ones. We show that a simple reinforcement learning model can capture these effects; and indeed verifies that they are robust to differences in subjects' acquisition of the instrumental tasks. This indicates that generalization of values occurs not only in terms of an action's associated value (the predicted reinforcement), as has long been reported in the PIT literature; rather, generalization also occurs in terms of the intrinsic affective qualities of actions. Our second key finding concerns the asymmetry between rewards and punishments in the instrumental learning phase. We find that, both when reinforcements are probabilistic and deterministic, subjects initially avoid, but then rapidly come to ignore punishments, coming to rely instead solely on rewards to guide their actions. We interpret this in terms of a fundamental informational asymmetry between rewards and punishments. Finally, we show that subjects' a priori bias against withdrawal actions correlates positively with measures of both anxiety and depression. This is consistent with a previous prediction by us that disturbances in aversive Pavlovian behavioural systems may result in negative affective experiences necessary for the generation of a variety of cognitive constructs resulting in different mood disorders.

doi:

II-16. Neurons in area LIP encode perceptual decisions in a perceptual, not oculomotor, frame of reference

Sharath Bennur
Joshua Gold
University of Pennsylvania

SHARATHS@MAIL.MED.UPENN.EDU
JIGOLD@MAIL.MED.UPENN.EDU

Perceptual decisions require the brain to read out information from sensory cortex to generate a categorical choice. One prominent approach to identify neural substrates of this read-out process has been to use tasks that link the decision to a specific behavioral response and target neurons known to participate in the selection of that response. For example, neurons in the lateral intraparietal cortex (LIP) of monkeys that contribute to oculomotor preparation can encode decisions about the direction of random-dot motion when they are indicated with particular eye movements. However, a weakness of this approach is the inability to identify perceptual processing that is independent of motor planning. Thus, it is unknown if LIP encodes perceptual decisions not linked to specific eye movements, or how decision-related activity in LIP relates to its other sensory and motor properties. To study these issues, we recorded responses from 84 LIP neurons in two monkeys performing a motion-discrimination task that included a flexible association between the direction decision and the eye-movement response. This "colored-target" task required the monkeys to indicate their decision with an eye movement to a target not at a particular location but of a particular color: the monkey was rewarded for looking at the red target following rightward motion, or the green target following leftward motion. One target was always placed in the neuron's response field, the other 180° opposite relative to the motion stimulus. We also controlled the timing of the association by showing the colored targets either before, during, or after motion viewing. This design allowed us to analyze how the encoding process depends on knowledge of the specific sensory-motor mapping. We found that individual LIP neurons encoded three distinct variables at different times during a trial: target color, saccadic choice, and motion direction. Whenever the colored targets appeared, there tended to be a transient response that was selective for target color, regardless of the motion direction or saccadic choice on that trial. In the epoch just prior to the saccadic response, the same LIP neurons encoded the saccadic choice and not just the color of the chosen target or the direction of motion. In addition, during motion viewing many neurons encoded the direction of motion, regardless of when the monkey learned the specific sensory-motor mapping on that trial. Unlike the color- and choice-selective responses, these motion direction-dependent responses were sensitive to the strength of motion in a manner akin to the decision-related signals reported for LIP using tasks with fixed sensory-motor mappings. Moreover, these direction-dependent responses were most affected by errors, further linking them to the perceptual report. The results suggest a general role for LIP in encoding perceptual decisions,

even those not linked to a pre-specified oculomotor response.

[doi:](#)

II-17. Context-dependent gating of sensory signals for decision making

Valerio Mante
William T. Newsome

VALERIO@MONKEYBIZ.STANFORD.EDU
BNEWSOME@STANFORD.EDU

HHMI and Stanford University

Humans and animals can process sensory signals in a remarkably flexible manner. Depending on the context in which they are perceived, identical sensory stimuli can lead to very different motor actions. Such context-dependent sensory-motor associations are thought to rely on the brain's ability to flexibly gate the flow of neural signals between the appropriate sensory and motor areas. The neural processes underlying this flexible gating are largely unknown. To better understand these processes, we trained a macaque monkey to perform two different perceptual discriminations on the same set of visual stimuli. On separate trials, the monkey was instructed to either discriminate the direction of motion or the color of a random-dot display, and to report his choice with a saccade to one of two visual targets. We used a simple model of the behavior to show that the monkey based his choices on the relevant stimulus dimension (motion or color) and largely ignored the irrelevant dimension. While the monkey performed this task, we recorded extracellular responses simultaneously in two cortical areas: area MT, which is thought to represent the sensory evidence relevant for direction discriminations; and the frontal eye fields (FEF), where developing saccade plans have been shown to reflect the integrated sensory evidence in favor of one of the two possible choices. We expect FEF to integrate responses from MT during motion discrimination, but from other areas (possibly V4 or IT) during color discrimination. How is this context-dependent gating reflected in MT and FEF responses? We find that FEF responses reflect the upcoming choice during both motion and color discrimination. Thus, responses from motion- and color-selective areas converge before or at the level of FEF. Responses in MT, on the other hand, reflect the strength of the motion signal during both discriminations—both average MT firing rates as well as neural sensitivities are virtually unchanged across the two contexts. We conclude that gating of sensory responses does not require and is not reflected in the modulation of firing rates in the relevant sensory areas. We hope to better understand how gating occurs by analyzing trial-by-trial correlations between neural responses and behavior, and between neural responses in MT and FEF.

[doi:](#)

II-18. The effect of value normalization and cortical variability on rational choice

Kenway Louie¹
Paul Glimcher²

KLOUIE@CNS.NYU.EDU
GLIMCHER@CNS.NYU.EDU

¹New York University

²Center for Neural Science, NYU

The neural circuits underlying the decision process must represent the values of the available choice options. In the monkey lateral intraparietal area, a visuomotor subregion of posterior parietal cortex responsive to eye movements, neuronal activity is strongly modulated by the value of specific saccades. In recent neurophysiological experiments, we have demonstrated that this value representation is not absolute: neurons code the value of saccades to the response field relative to the values of all available saccade choices. Interestingly, this value normalization is well-described by a divisive normalization model that also characterizes nonlinear phenomena such as gain control and cross-orientation suppression in visual cortex. Does this normalized value representation affect behavior? We explore here the predictions of the divisive normalization model and compare them to observed

choice behavior. Computational simulation of the choice process indicates that a normalized value representation produces specific behavioral irregularities that violate rational choice theory. In particular, preference between two high-valued options appears to be a function of a third, low-valued irrelevant alternative. This effect depends crucially on cortical neuron response variability: as the total value of available options increases, the separation between the distributions of firing rates representing two differently valued options decreases; if variance doesn't decrease appropriately, the options will be increasingly difficult to distinguish. We trained two monkeys to choose between three differently valued stimuli (A, B, and C) in a block design. Stimulus locations and reward associations were fixed within a block, which consisted of 40 single-stimulus trials followed by 40 randomly-presented two- and three-stimulus choice trials. Across blocks, the values of the target options (A and B) were varied to quantify how choice varied as a function of value difference. The value of the distractor option (C) took one of two possible values, but was always lower than any possible target option value. We find that choice preference is context-dependent: the relative preference between the two high-valued options depends on the value of the third option. At the low distractor value, monkeys are more likely to correctly choose the higher valued target option; at the high distractor value, choice behavior becomes more stochastic. Importantly, examination of the data segregated by distractor value shows that this effect is equivalent to a change in the slope of the logistic choice function, an effect predicted by the combination of normalized value and cortical variability. We conclude that the normalized representation of value in choice circuits is observable at the behavioral level, and may play a role in real-world examples of context-dependent choice.

[doi:](#)

II-19. Thalamocortical changes in clinical depression probed by physiology-based modeling

Cliff Kerr^{1,2}

Andrew Kemp³

Chris Rennie^{1,2}

Peter Robinson^{1,4}

CKERR@PHYSICS.USYD.EDU.AU

AKEMP@USYD.EDU.AU

RENNIE@PHYSICS.USYD.EDU.AU

ROBINSON@PHYSICS.USYD.EDU.AU

¹School of Physics, University of Sydney

²Brain Dynamics Center, Westmead Hospital

³School of Psychology, University of Sydney

⁴Brain Dynamics Centre, Westmead Hospital

Clinical depression is a heterogeneous disorder characterized by persistent dysphoria (unpleasant mood) and anhedonia (inability to experience pleasure). Depending on the particular subtype, additional symptoms can include difficulty concentrating, psychomotor slowing, and changes in appetite and sleep. Historically, most studies of clinical depression have focused on neuronal or cognitive changes, with comparatively few examining the systems-level computational neurobiology of the disorder. To help bridge this gap, we use a mean-field model of neuronal dynamics to investigate possible causes of the electrophysiological changes observed in patients with depression. Event-related potentials (ERPs) were elicited from four subject groups (clinically depressed patients with and without the melancholic depression subtype, participants with subclinical depressed mood, and healthy controls) using an auditory oddball paradigm, in which infrequent high-pitched tones requiring a response ("targets") were interspersed with frequent low-pitched tones not requiring a response ("standards"). These ERPs were extracted from subjects' ongoing EEG activity, and analyzed using the physiology-based mean-field model developed by Robinson et al. (Phys. Rev. E 2001 63:021903). This model describes firing rate dynamics in five interconnected populations of neurons: cortical excitatory, cortical inhibitory, thalamic relay, thalamic reticular, and subthalamic. The model also incorporates neuronal properties important for global brain dynamics, including dendritic time constants and axonal propagation velocities. Fitting the model to experimental data allows its parameters to be estimated, which in turn can be related to underlying neurophysiology. Several major differences between healthy controls and other subject groups were found: (i) Dendritic time constants were significantly smaller in subjects with both clinical and subclinical depressed mood, indicative of changes in neurotransmission (for example, a shift from AMPA to NMDA receptors). (ii) Transmission velocities in thalamocortical axons were de-

creased in patients with melancholic depression, potentially explaining the psychomotor slowing associated with this subtype. (iii) Connection strengths between neuronal populations changed dramatically, including decreased cortical excitation, decreased thalamocortical excitation (via the relay nuclei), and increased thalamocortical inhibition (via the reticular nucleus), leading to a substantial change in the balance of excitation and inhibition in patients with depression. These results shed new light on the computational neurobiology of clinical depression, and provide a framework allowing future integration of additional data across a range of spatiotemporal scales.

[doi:](#)

II-20. Multistability as a mechanism for modulation of EEG coherences

Jonathan Drover
Jonathan Victor
Shawniqua T. Williams
Mary Conte
Nicholas Schiff

JOD2017@MED.CORNELL.EDU
JDVICTO@MED.CORNELL.EDU
STW2009@MED.CORNELL.EDU
MMCONTE@MED.CORNELL.EDU
NDS2001@MED.CORNELL.EDU

Weill Medical College of Cornell University

Coordinated activity between multiple cortical areas is necessary for organized behavior and cognitive activity. It is speculated that deep brain structures, specifically the reticular thalamus, play an important role in coordinating this activity. In previous work we developed a mean field model of a thalamocortical network consisting of two thalamocortical modules (each module containing cortical, thalamic relay, and thalamic reticular populations) coupled via a shared population of reticular neurons. We showed that this network is capable of spontaneous transitions, distinguishable by changes in the coherence between the two cortical populations modeled. These transitions can occur when the parameters of the model are such that there are multiple stable attractors. There are two types of attractors that we were interested in: symmetric solutions, where each module maintains a similar activity level; and winner-take-all solutions, where one of the modules suppresses the other. We show that the multistable region has, as boundaries, a winner-take-all generating fold bifurcation and a subcritical pitchfork bifurcation that destabilizes the symmetric solution. We show that this configuration is persistent over a wide range of values of the parameters that determine the strength of the connections within the thalamocortical modules, and a realistic range of time constants. Because this was a striking and consistent feature of model behavior, we sought to determine whether it was present in the human EEG. We analyzed EEG/CCTV recordings from three patients with severe brain injury, characterized by metabolic (resting PET) and anatomical (MRI, DTI) studies and behavioral observations. Via the multitaper method, we calculated time-localized EEG spectra and coherences from 30 segments of artifact-free EEG obtained during eyes-open rest. We then applied principal components analysis to the coherograms obtained from pairs of channels within each hemisphere, revealing bimodal behavior. Thus, time-varying patterns of coherence can be identified in the EEG of human subjects, as well as in the model. In these brain-injured patients, this dynamical feature appeared to correlate with relatively more preserved functional or structural integrity. In sum, a population-based model of thalamocortical interactions robustly demonstrates multistability, and this dynamical feature can be identified in the human EEG, supporting a role for the thalamus in establishing changing patterns of cortical coherence. Moreover, its predominance in the relatively more preserved hemisphere of brain-injured patients suggests an EEG-based approach to assaying the integrity of thalamocortical interactions.

[doi:](#)

II-21. A robust, bilateral line attractor model of the oculomotor system with ordinary neurons

Pedro Goncalves^{1,2}
Christian K. Machens^{1,2}

PEDRO.GONCALVES@ENS.FR
CHRISTIAN.MACHENS@ENS.FR

¹Group for Neural Theory, ENS, Paris

²INSERM Unite 960, ENS, Paris

The oculomotor system controls the position of the eyes during fixations and saccades. A prime candidate for horizontal, fixational control are the so-called "position" neurons, which fire persistently with a frequency that is proportional to the horizontal eye position. Since these neurons can maintain firing over several seconds at any of a continuum of levels, previous modeling studies have suggested that these neurons form a line attractor network [1]. Past research has focused on three properties of this system: (1) Recruitment order: The slope of the position neurons' tuning curves increases as their firing threshold moves towards more eccentric eye positions [2]. (2) Hysteresis and fine-tuning: Line attractor models of the goldfish oculomotor system rely on fine-tuning of synaptic parameters ($<1\%$), yet how a biological system can fine-tune its synapses remains an open issue [1]. Previous modeling work solved this by making neurons or dendrites bistable [3,4], in agreement with the hysteresis found in the tuning curves of position neurons. (3) Bilateral dependency: Silencing of position neurons from one side impairs the functioning of the contralateral neurons in half of the oculomotor range [5]. This led to the suggestion that the two sides of the system work as independent line attractor networks. Modeling work showed that proper coordination of the two networks is possible if individual neurons have high synaptic thresholds. Here we investigate the construction of models that observe all these features, yet rely on standard single neurons - without the need for bistability or high synaptic thresholds. Using a mean-field network approach, we study the class of networks with rank-two weight matrix that obey the recruitment order. Under these constraints both ipsilateral excitation and contralateral inhibition are necessary to maintain eye position. Surprisingly, several of our models naturally reproduce the inactivation experiments. We conclude that the inactivation result does not prove the independence of the two sides, but could be a simple consequence of the recruitment order of position neurons. We solve the robustness problem by assuming that neurons adapt their firing rates, a well-established biophysical process. The adaptation rule, based on [6], leads to a network that is robust to perturbations in its parameters up to 5%, making the overall model more robust than fine-tuned network models. With this adaptation rule, hysteresis emerges as observed in the data, i.e., without resorting to hysteretic units. We therefore suggest that hysteresis is a signature of an active robustness mechanism, but not necessarily of a hidden bistability. We conclude by suggesting further experiments that would allow us to validate or refute existing models. [1] Seung PNAS 1996 [2] Aksay et al. Journal of Neurophysiology 2000 [3] Koulakov et al. Nature Neuroscience 2002 [4] Goldman et al. Cerebral Cortex 2003 [5] Aksay et al. Nature Neuroscience 2007 [6] Moreau, Sontag Physical Review E 2003

doi:

II-22. Cortical activity demystified: a unifying theory that explains state switching in cortex

Alexander Lerchner
Peter E. Latham

LERCHNER@GATSBY.UCL.AC.UK
PEL@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, UCL

An increasing number of experiments show that neocortical activity not only exhibits different states, but can rapidly switch between them. Examples include up and down states [1], synchronized and de-synchronized activity [2], high and low spiking precision [3], and large and small subthreshold correlations [4]. Although there is agreement that these states have important implications for computations, it has not been clear what mechanisms underlie them. Here we provide a unifying explanation by showing that different inputs applied to the same cortical

network can generate all the states, and transitions between the states, that have been observed experimentally. For the model underlying our theory, we need to assume only a small number of well-established properties that are shared by all local networks in the cortex. To analyze the model, we use an extended mean-field theory with temporal fluctuations. As with most mean-field theories of this type, we combine subthreshold activity with spiking activity to derive a self-consistent set of equations. Importantly, the equations contain only a small number of parameters, and even these have values that are constrained by biology, so the network can exhibit only a restricted range of behaviors. The main outcome of the theory is that the activity of the network depends primarily on the external input, and far less so on single-neuron properties or connectivity. For constant input, our theory (and simulations) recovers previous results indicating that networks exhibit weak correlations and irregular spiking activity. Inputs with intermediate structures lead to states in which sub-threshold correlations increase during sudden stimulus increases, accompanied by more precise spike timing during such events. And finally, brief inputs that drive the network only occasionally result in a dynamic state exhibiting membrane-potential "bumps" that are strongly correlated across neurons. These bumps lead to almost perfect synchrony for both excitation and inhibition, but with a characteristic lag of inhibition behind excitation of a few milliseconds. In this regime, the theory predicts that membrane potentials are highly correlated, and the first spikes after stimulus onsets are precisely timed. Our results show that both single-neuron and population activity in cortex are constrained by fundamental dynamics of local cortical networks, which in turn depend on the structure of the network input in a non-trivial manner. The theory predicts that cortical networks can have only a small number of behaviors. It is nontrivial, then, that the behaviors predicted by our theory (and verified with detailed network simulations) are consistent with experimental observations. Beyond providing a mechanistic explanation for a wealth of existing data, our theory can guide design and data analysis for future experiments that aim to probe detailed function and micro-structure of cortical networks. References: [1] Lampl et al. (1999) *Neuron* 22, 361-374. [2] Curto et al. (2009) *J. Neurosci* 29, 10600-10612 [3] Buracas et al. (1998) *Neuron* 20, 959-969 [4] Poulet and Petersen (2008) *Nature* 454, 881-885

doi:

II-23. Pattern separation by adaptive networks: neurogenesis in olfaction

Siu Fai Chow¹
Stuart D. Wick¹
Hermann Riecke²

SIUCHOW2008@U.NORTHWESTERN.EDU
STUART.WICK@GMAIL.COM
H-RIECKE@NORTHWESTERN.EDU

¹Northwestern U.

²Applied Mathematics, Northwestern University

A characteristic aspect of early processing of sensory information by neuronal circuits is a reshaping of activity patterns that may facilitate further processing in the brain. For instance, in the olfactory system the activity patterns that related odors evoke at the input of the olfactory bulb can be highly similar; nevertheless, the corresponding activity patterns of the mitral cells, which represent the output of the olfactory bulb, can differ significantly from each other due to strong inhibition by granule cells and peri-glomerular cells [1]. Due to the high dimensionality of 'odor space' the activation patterns that need to be separated are very complex. This presumably requires bulbar network connectivities that are more complex than those generating the center-surround receptive fields in the retina. We therefore investigate to what extent adaptive inhibitory networks can learn to perform pattern separation, i.e. to enhance the difference between similar patterns. Considering simple firing-rate models, we first present general biophysical considerations that can constrain the ability of networks to learn pattern separation. Then we investigate to what extent adult neurogenesis, as it is observed in the olfactory bulb, can provide a learning mechanism for this task. The stimuli that an animal needs to discriminate are typically sensed at quite different times. Without access to substantial memory it is therefore difficult for the network to learn the connectivity based on the similarity of different stimuli; biologically it is more plausible that learning is driven by simultaneous correlations between the input channels. We investigate the connection between pattern separation and channel decorrelation and demonstrate that networks can achieve effective pattern separation through channel decorrelation if they simultaneously equalize their output levels. In feedforward networks biophysically plausible learning mechanisms fail, however, for even moderately similar input patterns. Recurrent networks do not have that limita-

tion. Even when the connectivity of the recurrent networks is optimized for linear neuronal dynamics they perform very well when the dynamics are nonlinear [2]. Even in adult animals new inhibitory interneurons are persistently incorporated into the bulbar network; less than 50% of them survive, however, in the long term. Since their survival rate depends on the odor exposure of the animal and on behavioral tasks it may perform, this adult neurogenesis may provide an efficient mechanism to restructure the bulbar network to adapt it to the challenges presented to the animal by the olfactory environment. Consistent with this, adult neurogenesis has been found to be correlated with the animal's performance in odor discrimination tasks. In our model new interneurons are integrated persistently into the network and subsequently removed depending on their activity. The networks resulting from this training procedure are able to separate even quite similar stimuli. As observed experimentally, we find that young neurons are more responsive to novel odors. [1] R.W. Friedrich, G. Laurent, "Dynamic Optimization of Odor Representations by Slow Temporal Patterning of Mitral Cell Activity", *Science* 291 (2001) 889. [2] S.D. Wick, M.T. Wiechert, R.W. Friedrich, H. Riecke, "Pattern orthogonalization via channel decorrelation by adaptive networks", *J. Comp. Neurosci.* (2009).

[doi:](#)

II-24. A common-input model of a complete network of ganglion cells in the primate retina.

Michael Vidne¹

Yashar Ahmadian²

Jonathon Shlens³

Jonathan W. Pillow⁴

Jayant Kulkarni⁵

Eero P. Simoncelli⁶

E. J. Chichilnisky⁷

Liam Paninski²

MV333@COLUMBIA.EDU

YA2005@COLUMBIA.EDU

SHLENS@SALK.EDU

PILLOW@MAIL.UTEXAS.EDU

JAYANT.E.KULKARNI@GMAIL.COM

EERO.SIMONCELLI@NYU.EDU

EJ@SALK.EDU

LIAM@STAT.COLUMBIA.EDU

¹Center for Theoretical Neuroscience, Columbia University

²Columbia University

³New York University

⁴University of Texas at Austin

⁵CSHL

⁶HHMI / NYU

⁷Salk Institute

Synchronized firing among retinal ganglion cells (RGCs) has been proposed to indicate either redundancy or multiplexing in the neural code from the eye to the brain. Two major candidate mechanisms of synchronized firing are direct electrical coupling and common synaptic input. Recent modeling efforts (Pillow 2008) suggest that a generalized linear model with coupling between cells is able to accurately capture the synchronized spiking activity in parasol RGCs of the primate retina. But recent experimental work (Khuc-Trong 2008) indicates that electrical coupling between parasol cells is weak, and neighboring parasol cells share significant excitatory synaptic input in the absence of modulated light stimuli. These findings suggest that an accurate model of synchronized firing must include the effects of common noise. Here we develop a new model of synchronized firing that incorporates the effects of common noise, and use it to model the light responses and synchronized firing of a complete network of a few hundred simultaneously recorded parasol cells. We use a generalized linear model augmented with a state-space model to infer common noise, spatio-temporal light response properties, and post-spike feedback which captures dependencies on spike train history. All model parameters are estimated by maximizing the likelihood of the spiking data. Common noise is modeled as an autoregressive process with a correlation time consistent with that observed by Rieke et al. We use fast methods for computing the estimated maximum a posteriori path of the hidden input, by taking advantage of its banded diagonal structure (Paninski 2009). To test the model, we compare average light response properties and two- and three-point correlation functions obtained from the model and the data. The model provides an accurate account of these properties. We also use the model to decode

the visual stimulus, by maximizing the posterior probability of the stimulus given the spiking activity and the model parameters, and compare the results to decoding based on a model with coupling between RGCs but with no common input. We find that the common input architecture is more robust with regard to spike time perturbations than a network with direct coupling between the RGCs, especially when synchronized firing is strong. [1] Pillow, J.W. and Shlens, J. and Paninski, L. and Sher, A. and Litke, A.M. and Chichilnisky, E.J. and Simoncelli, E.P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*. 454:995-999 [2] Trong, P.K. and Rieke, F. (2008). Origin of correlated activity between parasol retinal ganglion cells. *Nature Neuroscience*. 11, 1343 - 1351 [3] Koyama, S. & Paninski, L. (2009) Efficient computation of the maximum a posteriori path and parameter estimation in integrate-and-fire and more general state-space models. *Journal of Computational Neuroscience*. (in print)

doi:

II-25. Closed-form correlation-based identification of recurrent spiking networks

Michael Krumin
Ariel Tankus
Shy Shoham
Technion - IIT

SSSHOHAM@BM.TECHNION.AC.IL
ARIELTA@GMAIL.COM
MIHA@BM.TECHNION.AC.IL

The correlation structure of neural activity is believed to play a major role in the encoding and possibly decoding of information in neural populations. Additionally, some of the most fundamental and widely applied engineering tools for system identification rely on the use of second order statistical properties (correlation or spectral). An increasing arsenal of tools for identifying spike train models from their correlations, rather than from their full observed realizations could form a welcome bridge between 'classical' signal processing ideas and tools and the field of neural spike train analysis. In recent work, we have analyzed the correlation distortion induced by the Linear-Nonlinear-Poisson (LNP) model family and have used the results for controlling the correlation structure of synthetic spike trains, for introducing new methods for 'blindly' identifying neural encoding models, for multivariate autoregressive modeling of spike trains and for performing causality analysis over populations of spiking neurons. However, the range of statistical structures that can be explained by the LNP model class is limited. Hence, the methods based on this model occasionally fail while analyzing correlation structures that are observed in neural activity. Here, we generalize these previously developed methods by combining the LNP model with the multivariate self- and mutually-exciting Hawkes model class. This results in a much more powerful, but still analytically tractable, Linear-Nonlinear-Hawkes (LNH) model family that is capable of capturing the dynamics of spike trains with more complex multi-correlation structure, enabling their analysis. We explore new applications of this framework including highly compact representations of multi-channel spike train data, causal analysis of network information flow, and the identification of cortical networks performing optimal dynamical Bayesian filtering. Acknowledgements This work was supported by Israeli Science Foundation grant #1248/06 and European Research Council starting grant #211055.

doi:

II-26. Neuronal variability and linear neural coding in the vestibular system

Adam Schneider
Maurice Chacron
Kathleen Cullen
McGill University

ADAM.SCHNEIDER@MAIL.MCGILL.CA
MAURICE.CHACRON@MCGILL.CA
KATHLEEN.CULLEN@MCGILL.CA

Neuronal variability is ubiquitous in the central nervous system. It was originally thought of as an obstacle to overcome by the nervous system as stated by Jon von Neumann, "How can a reliable nervous system be made of unreliable parts?" However, recent studies have suggested more positive roles for neural variability such as increasing information transmission through suprathreshold stochastic resonance. The vestibular system benefits from easily characterized sensory stimuli and has several advantages for studying the effects of neural variability on neural coding. Vestibular afferents have been characterized as either regular (low variability) or irregular (high variability) according to their coefficient of variation (CV). Despite spanning a wide range of CV across the afferent population, they are nonetheless known for their faithful linear encoding of head velocity as a firing rate modulation, as shown by Goldberg and Fernandez (1971). These afferents project to neurons within the medial vestibular nuclei that have also been shown to encode sensory stimuli through modulation of firing rate in vivo. This faithful linear encoding is surprising given that in vitro studies have found that these neurons possess several nonlinearly activated ion channels such as calcium-activated potassium as well as hyperpolarization-activated channels such as I-h. How do those nonlinear conductances interact with in vivo conditions in order to make such a nonlinear system behave linearly? We performed numerical simulations of a detailed Hodgkin-Huxley type model based on in vitro data, originally developed by Av-Ron and Vidal (1999) to model in vitro recordings from central vestibular neurons. In vivo conditions were mimicked by the addition of noise to the model and tuning the noise intensity such that the CV generated by the model matches that found in in vivo recordings. We found that this nonlinear model was capable of various exotic phenomena in the absence of noise such as oscillations and burst firing. These nonlinearities strongly interfered with the model's ability to encode sinusoidal and noise stimuli through changes in firing rate. Most surprisingly, addition of noise reduced the effects of these nonlinearities significantly and allowed the model to encode stimuli linearly. These results show that variability introduced by synaptic bombardment can significantly attenuate nonlinearities due to voltage-gated ion channels and allow faithful linear encoding of sensory stimuli.

doi:

II-27. Transition-state theory for integrate-and-fire neurons

Laurent Badel¹

Wulfram Gerstner²

Magnus J. E. Richardson³

LB2686@COLUMBIA.EDU

WULFRAM.GERSTNER@EPFL.CH

MAGNUS.RICHARDSON@WARWICK.AC.UK

¹Department of Statistics, Columbia University

²Ecole Polytechnique Fédérale de Lausanne

³Warwick University

We derive an approximation for the firing-rate response of integrate-and-fire neurons driven by filtered gaussian noise, in the case where the membrane potential fluctuations are small compared with the distance between the mean membrane voltage and the firing threshold. The computation is based on the analogy with the classical problem of noise-activated escape from a trapping potential, and yields results that are valid for values of the synaptic filtering time constant close to the membrane time constant, a situation which may be particularly relevant in-vivo where the membrane time constant can be significantly reduced under the action of background synaptic activity. The approach yields simple, tractable results, and can be extended to include excitatory and inhibitory pathways with distinct filtering time constants, additional voltage-activated subthreshold currents in the formalism of generalized integrate-and-fire models, and allows for the inclusion of spatial correlations between excitation and inhibition, an aspect that is often overlooked in theoretical analyses. The method can also be applied to the time-dependent case to compute the firing-rate response to arbitrary time-varying signals. In this case, the theory captures many interesting features of the dynamical response, such as firing-rate "rebound" in the presence of an Ih-like hyperpolarization-activated subthreshold current, or firing-rate modulations induced by correlations in the synaptic inputs. Furthermore, the method can be applied to pairs of neurons receiving cross-correlated inputs to derive an expression for the cross-correlation of the output spike trains. These results allow for the identification of the role of each parameter in shaping the firing-rate response of integrate-and-fire neurons, and could be useful to analyze models of interacting populations. The approach could also be used to describe linear multi-compartment integrate-and-fire models to study the influence of dendritic filtering.

doi:

II-28. Auditory textures and primitive auditory scene analysis

Richard E. Turner¹
Maneesh Sahani²

TURNER@GATSBY.UCL.AC.UK
MANEESH@GATSBY.UCL.AC.UK

¹Computational and Biological Learning Lab

²Gatsby Computational Neuroscience Unit, UCL

Both the categorical identity of a natural sound and the perceptual analysis of an auditory scene often appear to depend on the statistical properties of the sound and of the acoustic environment. An example of the first case is the sound of "rain"; no two rain sounds are identical because the precise arrangement of falling water droplets is never repeated. Thus, the identity must lie not in the precise waveform but in its statistical properties, such as those relating to the rate of falling rain-drops, the distribution of droplet sizes, and so on. Our perception of other auditory "textures", such as running water, wind or fire, is similar. The statistical properties of a broader class of natural sounds in turn influence how we parse an auditory scene. We have argued previously (Cosyne 2008) that a number of the "Gestalt" phenomena of auditory scene analysis are well captured by the process of inference within an appropriate probabilistic model. Here, we present a new probabilistic generative model for natural sounds, which is both more comprehensive and more numerically tractable than previous such models. The model comprises a set of narrow-band Gaussian carriers that are modulated by a set of positive envelopes. These envelopes are determined by mixing slowly-varying Gaussian processes and passing the result through a positive non-linearity. The parameters of this model – the power, centre-frequencies and bandwidths of the carriers, the time-scale, and depth of the modulation and the patterns of comodulation – were learned from training sounds using approximate maximum-likelihood methods. We show that the model is able to capture many of the important statistics of auditory textures, and can thus be used to synthesise realistic versions of running water, wind, rain and fire. The same generative model was also used to capture many of the basic principles by which listeners appear to understand simple stimuli. In the generative framework, the carriers and the envelopes are latent variables being inferred from the sound waveform. We show that the inferred values of these latent processes correspond to perceptual principles of grouping by proximity, good continuation and common-fate as well as the continuity illusion, comodulation masking release, and the old plus new heuristic (Bregman 1990). Thus, our results suggest that the auditory system is optimised to process sounds with naturalistic statistics, and that a model that captures the statistics of the natural environment can be used to both synthesise perceptually valid "natural" sounds, and to account for a range of perceptual phenomena.

doi:

II-29. Bayesian Pitch

Phillipp Hehrmann
Maneesh Sahani

HEHRMANN@GATSBY.UCL.AC.UK
MANEESH@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, UCL

Pitch is a fundamental perceptual attribute of many sounds. It carries melodic, prosodic and sometimes semantic information, and contributes to auditory scene analysis. However, over a century of study has failed to yield a generally-accepted definition of pitch and model of its perception. No single physical feature of the acoustic waveform correlates perfectly with the reported pitch for all tested sounds, raising the question of whether pitch perception instead relies on a combination of several different features, and in turn, mechanisms. Three implicated features are harmonic regularities in the spectrum, periodicity of the waveform envelope, and the timing of fine structure peaks within the envelope. Here, we propose a computationally unified account of pitch perception. Our hypothesis is that the pitch of a sound is a Bayesian estimate of its periodicity, as inferred from the resulting

time-varying auditory nerve firing rates. Inference proceeds within a generative model in which pitched sounds are first formed by convolving a regular train of impulses with a randomly drawn impulse-response and adding Gaussian noise. These sounds are then transformed to auditory nerve activity by a cascade of bandpass filters and demodulators, corrupted by further noise. This generative process describes responses to many naturally occurring pitched sounds such as voiced speech, vocalisations and musical instrument sounds. We suggest that the aberrant or ambiguous pitch percepts reported for some laboratory constructed sounds arise because the auditory system infers periodicity assuming this same generative model, even when it does not apply. To test our hypothesis, we chose three types of stimuli to exemplify the three physical cues listed above. Evidence for the importance of spectral features comes from harmonic complex sounds with missing fundamental f_0 . The pitch of these sounds corresponds to f_0 , but its strength varies depending on the number of the lowest harmonic present in the spectrum. Evidence for temporal envelope periodicity comes from amplitude modulated noise, which is weakly pitched despite the complete lack of spectral features. Finally, the timing of fine structure peaks must be invoked to explain the ambiguous pitch of amplitude modulated pure tones, which corresponds neither exactly to the envelope modulation rate nor to the spacing of spectral components. In all three cases, our Bayesian periodicity estimate provides a good match with human pitch perception, demonstrating that the model can indeed utilise spectral, temporal envelope and temporal fine structure cues within a single computational framework. Notably, the effect of model mismatch in case of the amplitude modulated tone and noise stimuli are also consistent with human pitch perception, lending further support to our hypothesis. Viewing pitch perception as probabilistic inference relates pitch to an unobserved, yet real physical quantity and provides a normative interpretation of its role in auditory perception, rather than a purely phenomenological match. Whilst this interpretation does not address the issue of neural implementation, it suggests that different neural mechanisms, should they exist, are set up together in a way to achieve a single computational goal.

doi:

II-30. Bayesian belief propagation and border-ownership signals in early visual cortex

Haruo Hosoya

HAHOSoya@IS.S.U-TOKYO.AC.JP

University of Tokyo

Visual cortex is known to have neuronal response properties that strongly depend on stimuli to extra-classical receptive fields. One of such contextual effects is border-ownership, reported by Zhou, Friedman, and von der Heydt (2000), where the response of an edge-selective neuron is modulated by whether the edge belongs to a figure on one side or the other. Although a number of computational models have been proposed to explain such response properties, they assume neural circuits that perform specific image processing. Thus, it is not clear how such processing may arise from a general, learnable network—a presumably important property of visual cortex—and what such border-ownership signals may mean in mathematically crisp terms. We hypothesize that border ownership signals represent a posterior joint probability of a low-level visual property (e.g., edge orientation and phase) and a high-level visual feature (e.g., figure presence), given stimuli to the classical receptive field and to a surrounding context. In order to give a support to this hypothesis, we have used a model based on Bayesian network, for which several authors have recently discovered close relationship with known anatomical and physiological properties of cortical circuitry. As a theoretical result, we have found that posterior joint probabilities mentioned above can readily be found, under certain assumptions, as one of the variables used in the approximate belief propagation algorithm. Further, by using Ichisugi's mapping of this algorithm to cortical six-layer structure, we predict that border-ownership signals can be observed from layer II/III, which is in fact consistent with the report by Zhou et al. on macaque V1. To give a further support, we have conducted a computer simulation with our network. In this, we first performed a conventional training for pattern recognition on our network and then measured response properties of model units to artificial stimuli similar to those used experimentally. The results we obtained were qualitatively similar to physiological data in a rather detailed way. Specifically, we found (1) units of all four combination types in terms of dependencies on the edge polarity and the figure side and (2) units with responses invariant to the size and the shape of the figure, and to a conflicting cue. Further, the

proportions of these units closely resemble those found experimentally. In conclusion, the present study suggests that contextual effects widely found in visual cortex may be interpreted in terms of posterior joint probabilities. In a broader sense, the study shows a usefulness of Bayesian network in the understanding of cortical computation in terms of probability theory, serving as a solid theoretical basis for the prediction of future experimental outcomes.

[doi:](#)

II-31. Orientation maps as Moiré interference of retinal ganglion cell mosaics

Se-Bum Paik¹
Dario Ringach²

SBPAIK@UCLA.EDU
DARIO@UCLA.EDU

¹Department of Neurobiology, UCLA

²University of California Los Angeles

In mammalian primary visual cortex (area V1), neurons respond selectively to the orientation of visual stimuli. In some species, such as monkeys and cats, preferred orientation varies smoothly across the cortical surface, defining quasi-periodic orientation maps. In other species, such as mice, rats and squirrels, there is no detectable orientation map even though many individual neurons have robust orientation tuning. The mechanism of orientation map generation and its variability across species - even within individuals of the same species - remains elusive. Here we propose that orientation maps may result from the interference pattern between the ON and OFF center retinal ganglion cell (RGC) mosaics. We present computer simulations showing that quasi-regular ON and OFF RGC lattices generate periodic Moiré interference patterns, the regularity of which depends on the geometrical parameters such as lattice grid spacing difference and relative angle between the mosaics. We studied the structure of the resulting orientation maps, varying the noisiness of the RGC mosaics, and determined that Moiré interference would still contribute to the initial map structure for the realistic levels of noise observed in experimental RGC lattices. We demonstrate that one critical factor in establishing the regularity of the map is the relative density of ON and OFF RGCs. Equal densities predict a map that is locally disorganized (that is, the lack of orientation maps), while densities differing by 10% are sufficient to induce orientation maps in the cortex. The model also explains the tendency of simple cells to have anti-symmetric receptive fields, the generation of simple and complex cells by statistical wiring, the lack of orientation maps by the blockage of particular cell type activity during development, the consistency of orientation maps before and after monocular deprivation, and the establishment of orientation tuning and maps without structured spontaneous activity in the developing thalamus. We believe the simplicity and the power of the proposed model at explaining published data makes it a serious candidate for further experimental tests.

[doi:](#)

II-32. Estimation and assessment of non-Poisson neural encoding models

Jonathan W. Pillow
University of Texas at Austin

PILLOW@MAIL.UTEXAS.EDU

Recent work on the statistical modeling of neural responses has focused on modulated renewal processes in which the spike rate is a function of the stimulus and recent spiking history. Typically, these models incorporate spike-history dependencies via either: (1) a conditionally-Poisson process with rate dependent on a linear projection of the spike train history (e.g., generalized linear model); or (2) a modulated non-Poisson renewal process (e.g., inhomogeneous gamma process). Here we show that the two approaches can be combined, resulting in a conditional renewal (CR) model for neural spike trains. This model captures both real-time and rescaled-time history effects, and can be fit by maximum likelihood using a simple application of the time-rescaling theorem. We show that for any modulated renewal process model, the log-likelihood is concave in the linear filter parameters only under certain restrictive conditions on the renewal density, which rules out many popular choices (e.g.

gamma with shape $k \approx 8800; 1$). This result suggests that real-time history effects are easier to incorporate than non-Poisson renewal properties. Finally, we show that goodness-of-fit tests based on the time-rescaling theorem quantify relative-time effects, but do not reliably assess accuracy in spike prediction or stimulus-response modeling. We illustrate the CR model with applications to both real and simulated neural data.

[doi:](#)

II-33. Bayesian line orientation perception: Human prior expectations match natural image statistics

Ahna R. Girshick¹
Michael S. Landy¹
Eero P. Simoncelli²

AHNA@CNS.NYU.EDU
LANDY@NYU.EDU
EERO.SIMONCELLI@NYU.EDU

¹Dept of Psychology & Center for Neural Science, New York University

²HHMI / New York University

The visual world is replete with contours, and their location and orientation provide important information about visual scenes. Are visual estimates of local contour orientation determined by sensory measurements alone? Or are percepts biased, in a Bayesian fashion, by prior knowledge of the distribution of contour orientations in the environment? Line orientation provides a good domain for understanding how humans use prior information because orientation statistics of natural images are non-uniform: There is a preponderance of cardinal (vertical and horizontal) orientations, as compared to oblique orientations, in both natural and human-made scenes (Switkes et al., 1978). If observers behave in a Bayesian manner, orientation estimates of noisy stimuli should be biased toward cardinal orientations. We adapted a recently developed technique for estimating priors used by human observers (Stocker & Simoncelli, 2006) to determine human orientation priors, and then compared these to those measured from natural image databases. In the psychophysical experiment, observers performed an orientation discrimination task, comparing either two low-noise stimuli (LvL), two high-noise stimuli (HvH), or a low- and high-noise stimulus (LvH). The first two conditions were used to assess the widths of subjects' likelihood functions, whereas the LvH condition allowed us to infer the shape of observers' prior expectations. A Bayesian observer with a non-uniform prior should exhibit biases in the LvH condition, because the prior will affect the orientation estimate of a high-noise stimulus more than a low-noise stimulus. The stimuli consisted of an array of 38 Gabor patches with orientations either all identical (L) or drawn from a normal distribution with standard deviation approximately 20 deg (H; SD chosen per observer based on a pilot discrimination experiment). The observers' task was to select the stimulus whose mean orientation was more clockwise. On each trial, the mean orientation of the standard stimulus was randomly selected from 12 orientations equally distributed over 180 deg. In the LvH conditions, observers behaved as if the perceived orientation of the high-noise stimulus was systematically biased toward the nearest cardinal orientation. Under the assumption that our observers are acting as Bayesian estimators, we used methods similar to those in (Stocker & Simoncelli, 2006) to extract a prior distribution on orientation that would explain their perceptual biases. We compared these perceptual priors to the distribution of orientation measured in three databases of images which included natural and human-made scenes. We used a Gaussian pyramid (Burt & Adelson, 1983) to represent each image at six different spatial resolutions, computed gradients using pairs of localized rotation-invariant derivative filters (Farid & Simoncelli, 2004), and then locally combined these to compute an estimate of dominant orientation. We found that while histograms of these measurements varied in detail across databases and spatial scale, in all cases the cardinals were significantly more frequent than the obliques. The perceptually derived priors of our observers also varied in detail, but all exhibited substantially higher probability at the cardinals. Thus, human observers exhibit Bayesian behavior consistent with the probabilistic structure of the environment when estimating visual line orientation.

[doi:](#)

II-34. The value of lateral connectivity in visual cortex for interpreting naturalistic images

Xaq Pitkow¹
Yashar Ahmadian²
Ken Miller¹

XAQ@POST.HARVARD.EDU
YASHAR@STAT.COLUMBIA.EDU
KEN@NEUROTHEORY.COLUMBIA.EDU

¹Center for Theoretical Neuroscience, Columbia University

²Dept of Statistics, Columbia University

In many species, orientation-selective neurons in primary visual cortex are preferentially connected to others with similar orientation selectivity. It has been hypothesized that this circuitry is tuned to the statistics of natural scenes, and helps refine the interpretation of feedforward visual inputs: it may provide a mechanism by which neural networks can exploit the natural correlations between nearby edges. A related but more general claim holds for the machine learning algorithm called belief propagation. In this algorithm, model neurons encode the probabilities of local features and send information about their current state to other neurons encoding statistically related features. Following these dynamics, a network of such neurons converges to a consensus state representing the most probable synthesis of the available information. In this study we test how well neurally plausible approximations of belief propagation perform on feature discrimination tasks. Using a simple occlusion model of natural scenes, we are able to compute the probabilities of image features exactly, providing a ground truth against which we compare the performance of these model networks. We calculate how performance varies with the specificity and range of lateral connectivity, and compare the effectiveness of a given connectivity structure for neurons of different response types (simple cells, complex cells, border ownership cells). Finally we discuss the implications of these results for neural coding in primary and secondary visual cortex.

[doi:](#)

II-35. Perturbation of hippocampal cell dynamics by halorhodopsin-assisted silencing of PV interneurons

Sebastien Royer¹
Boris Zemelman¹
Attila Losonczy²
Jeffery Magee¹
Gyorgy Buzsaki³

ROYERS@JANELIA.HHMI.ORG
ZEMELMANB@JANELIA.HHMI.ORG
AL2856@COLUMBIA.EDU
MAGEEJ@JANELIA.HHMI.ORG
BUZSAKI@ANDROMEDA.RUTGERS.EDU

¹Janelia Farm Research Campus, HHMI

²Columbia neuroscience department

³CMBN, Rutgers University

Understanding the input-output relationships of neurons and circuits requires methods with the appropriate spatial selectivity and temporal resolution and duration for mechanistic analysis of neural ensembles. We describe that, in a head-fixed mouse running on a treadmill equipped with a long and cue-rich belt, CA1 hippocampal neurons generate sequences of 'episode fields,' reminiscent of place cells during navigation. We combined the use of multiple-shank silicon probe with integrated etched optical fibers and mice expressing halorhodopsin selectively in PV interneurons to transiently inactivate a restricted population of these CA1 neurons using light. Silencing small local populations of PV neurons transiently changed the firing rate of neighboring pyramidal cells, their temporal relationships to each other and to the theta cycle and the distance vs. phase-offset relationship ('distance-time compression'). PV interneurons are, therefore, critical for setting the temporal delays within theta cycles, a prerequisite for storing sequence representations.

[doi:](#)

II-36. Behavioral state continuously modulates hippocampal information processing

Caleb Kemere¹
Feng Zhang²
Karl Deisseroth³
Loren M. Frank⁴

CKEMERE@PHY.UCSF.EDU
ZHANGF@GMAIL.COM
DEISSERO@STANFORD.EDU
LOREN@PHY.UCSF.EDU

¹UCSF

²Society of Fellows, Harvard University

³Bioengineering, Stanford University

⁴Keck Center and Dept. of Physiology, UCSF

In behaving animals, changes in behavioral state are accompanied by profound changes in the patterns of recorded neural activity. In the rodent hippocampus, exploratory locomotion is associated with the prominent ~8Hz theta oscillation. Intervening periods of relative immobility (e.g., eating, grooming) are associated with high frequency (~200 Hz) sharp wave ripple (SWR) events. Studies involving the electrical microstimulation of hippocampal pathways have contributed to the conclusion that these states are separate both physiologically and functionally. Thus, the dominant view in the field is that theta governs a state where information is encoded in the hippocampal circuit while SWRs underlie a distinct state associated with consolidation. The presence of two clearly distinct states has not yet been proven, however. In addition, the effect of these states on the modulation of information transmission via the dentate gyrus (DG), a key pathway, has never been investigated in the intact animal. We used virally-mediated channelrhodopsin-2 (ChR2) to selectively activate the mossy fibers (MF), the dense projections from the dentate gyrus to hippocampal area CA3. Using an implanted optical fiber, we were able to activate the DG inputs to area CA3 while simultaneously recording downstream activity in behaving animals, including both single neuron spiking and LFPs. Histology revealed ChR2 expression limited to dentate granule cells and their processes, including strong expression in the mossy fibers. In recordings in CA3, repeated optical activation yielded a frequency facilitation effect further consistent with mossy fiber activation. We used low frequency (< 0.2 Hz) pulses at to probe the propagation of excitation along the trisynaptic pathway. Observing the output of this excitation in area CA1, we found that periods of fast movement corresponded to a large (nearly 50%) decrease in transmission. Moreover, we found that the modulation of information processing by behavior is not a switching between theta (moving) and SWR (still) states, but rather a continuum. This modulation manifested as a smooth reduction in fEPSP slope in CA1 as a function of increasing speed, where movement speed could account for up to 60% of the variance in EPSP slope. Intriguingly, the amplitude of SWRs in CA1 shows a similar speed-dependent decline; thus, as SC input is essential for normal CA1 SWRs, our results suggest that changes in the Schaeffer collateral synapse contribute to behaviorally dependent changes in CA1 output. Our results suggest that the simple two state model is not sufficient. Instead, they are consistent with the idea that that hippocampal information processing involves dynamic overlays of the integration of sensory information and more internally generated activity, with the balance shifting towards externally driven activation during locomotion and towards internal processing during stillness.

[doi:](#)

II-37. Hippocampal learning and cognitive maps as products of hierarchical latent variable models

Adam Johnson¹
Zachary Varberg¹
Paul Schrater²

ADAM-JOHNSON@BETHEL.EDU
ZACHARY-VARBERG@BETHEL.EDU
SCHRATER@UMN.EDU

¹Bethel University

²University of Minnesota

The hippocampus supports inference within early learning regimes and on a variety of complex spatial/context memory tasks. The defining characteristic of hippocampus dependent learning is that it provides the basis for a broad class of inferences in domains where sampling is limited relative to the potential dimensionality of inputs. Standard approaches to spatial learning and context learning have used neural network models to describe generalization behavior at the level of place cell maps and at the level of animal behavior. However, these models often overfit training data and compromise inference for novel test data. And in many cases, pattern classification dynamics observed in neural network models can be easily accommodated within probabilistic inference frameworks. We develop a model-based Bayesian approach to learning hierarchical latent variable structure and dynamics for spatial and context learning tasks. This approach expands probabilistic treatments of classical conditioning (Courville et al., 2003) to spatial learning domains for inference related to context dependent learning rules (e.g. "odor/reward position" paired-associate learning). These context dependent learning rules can be found by computing the posterior probabilities over latent variable hyperparameters. As a consequence, resultant behavioral inferences for novel test items reflect a mixture of latent models and are generally more robust. We model the spatial paired-associate task developed by Morris and colleagues (Day, Langston and Morris, 2003; Tse et al., 2007). Tse et al. (2007) showed that following initial training, rats learn new paired associates in a single trial. Single trial learning was dependent on the hippocampus and the consistency of the previously learned paired-associates within a given context. We show that initial training tunes the hyperparameter posterior density and consequently facilitates learning subsequent novel paired-associates - both in the search for newly learned paired-associate food-sites and for the avoidance of known food-sites following presentation of a novel odor. Furthermore, we show that the structure of the task requires the use of a full conjunctive representation for learning new paired-associates but does not require a full conjunctive representation for retrieval of previously learned paired-associates. Together these results show how the hippocampus facilitates new learning but is not always necessary for retrieval of previously learned information. Morris and colleagues suggest the hippocampus supports schema representations that facilitate new learning within a context. While this perspective has provided important insights into animal learning theory, it has lacked a computational counterpart. We suggest that hippocampus dependent schematic representations and cognitive maps are the consequence of a process that probabilistically infers hierarchical latent structure from input data. This statistical approach supports model optimization for inference across test data rather than model optimization across training data. As a result, the latent variable approach provides a much more robust inference engine than standard neural network models. Finally, the extraction of latent structure from input data allows for a simple explanation of hippocampal function and its important role within early learning regimes.

doi:

II-38. Computational role of theta oscillations in delayed-decision tasks

Prashant Joshi

JOSHI@FIAS.UNI-FRANKFURT.DE

Frankfurt Institute for Advanced Studies

Working memory allows temporary storage of task-relevant information needed to perform several complex cognitive functions thereby enabling the brain with the capability to execute delayed decisions so that all relevant sensory information and cues can be integrated to prepare an optimal response [1]. Several experimental studies have reported a sustained elevated discharge of action potentials in cortical neurons during the delay period of a working memory task [2]. This delay-activity is believed to play a key role in the short-term maintenance of information. Recent experimental results indicate that computationally relevant information for working memory tasks is encoded not only in the firing rates of cortical neurons as believed earlier, but also by their phase-of-firing with reference to some baseline periodic signal such as the local field potential (LFP). More precisely, an increase in power in the theta band (4-8 Hz) of the LFP signal was reported during the delay duration, suggesting that the theta band of LFP plays a key role in holding the stimulus "in mind" [3]. These results raise a key question as to how this additional information encoded in the theta signal during the delay period can be used by cortical networks of neurons. We demonstrate here that simple linear readouts connected to generic neural microcircuits composed of spiking neurons can integrate information from phase-of-firing in afferent spike trains into their delayed decisions to fire or not to fire during tasks that require short-term maintenance of memory. The feasibility of

such emergent computations using phase-of-firing coding is presented in a detailed model for experimental data from a delayed-match-to-sample task, where the memory of the stimulus during delay period is encoded in the phase-of-firing of neurons in area V4. Frequency domain analysis of LFP signals demonstrates that our model captures several phenomena reported in experimental studies. Additionally it is shown that this computational paradigm allows several other open-loop sensory processing tasks to be performed in parallel with a high degree of precision. References [1] A. Baddeley. Working memory. *Science*, 255(5044):556-559, 1992. [2] J. M. Fuster and G. E. Alexander. Neuron activity related to short-term memory. *Science*, 173:652-654, 1971. [3] H. Lee, G. V. Simpson, N. K. Logothetis, and G. Rainer. Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate cortex. *Neuron*, 45:147-156, 2005.

[doi:](#)

II-39. Prefrontal and hippocampal coding during long-term memory formation in monkeys

Scott L. Brincat
Earl K. Miller

SBRINCAT@MIT.EDU
EKMILLER@MIT.EDU

Picower Institute for Learning & Memory, MIT

A number of convergent studies have implicated the hippocampus and prefrontal cortex (PFC) as two of the most critical brain areas for formation of new long-term declarative memories. However, their respective roles in memory encoding are still not well understood, particularly at the neural level. To address this question, we recorded spiking and local field potential activity simultaneously from multiple electrodes in the hippocampus and lateral PFC of monkeys learning novel associations between pairs of visual objects, an animal model of declarative memory formation. On each task trial, a cue object instructed recall from long-term memory of its learned associate object; the monkey then decided whether a subsequent test object matched the recalled associate, and was given feedback about the correctness of his response. Each day, four novel object-pair associations were learned through trial and error to a high level of performance, so neural activity could be tracked through the full course of long-term memory acquisition and analytically partitioned into factors reflecting perceptual, mnemonic, and response processes. We found that, in parallel with the monkeys' behavioral learning, prefrontal neurons acquired two types of specific information reflecting the contents of the newly-formed long-term memories-information about the recalled associate object and about the associative match/non-match decision. Hippocampal neurons, in contrast, showed very little information reflecting memory content, but strongly represented trial outcome-whether the monkey's response on each trial was correct or not. Hippocampal outcome information was strongest during the initial stages of learning, when it might be used as a "teaching signal" instructing PFC and other neocortical structures what should be encoded into memory, and subsided as associations became well-learned. Robust trial outcome information was also conveyed in the strength of phase-synchrony between hippocampal spikes and PFC local field potentials, consistent with the idea of a hippocampal teaching signal being communicated to PFC. Interestingly, positive and negative trial outcomes were represented in distinct beta (~12-32 Hz) and theta (~4-8 Hz) frequency bands, respectively, perhaps reflecting their putatively opposing effects on target structures. These results suggest prefrontal cortex plays a central role in long-term memory formation and storage, while the hippocampus-contrary to many accounts-may function primarily in driving memory encoding and/or consolidation in other structures.

[doi:](#)

II-40. How neurogenesis and modulation affect network oscillations in a large-scale dentate gyrus model.

James B. Aimone
Fred H. Gage
Salk Institute

AIMONE@SALK.EDU
GAGE@SALK.EDU

As one of the two locations of adult neurogenesis, the dentate gyrus (DG) region of the hippocampus is unique in the scales of neural dynamics that it experiences. Although at any given instant only a fraction of DG neurons are "new," over longer time scales neurogenesis can provide a significant change in network structure. In a recent study (Aimone et al., Neuron, 2009) we described a computational model that investigates how this temporally extended process can affect the DG's pattern separation function in hippocampal processing, with several clear behavioral predictions. Nevertheless, our model was built to study the long-term effects of neurogenesis and was not designed to investigate new neuron's effects on short-term network dynamics, such as network oscillations. Hippocampal oscillations in the theta and gamma frequencies are considered important during memory encoding and increasingly are also being implicated in memory consolidation. Acetylcholine and serotonin, two neurotransmitters known to affect hippocampal oscillations, have both been shown to interact with DG neurogenesis as well. These relationships introduce the interesting question of whether these new neurons have unique relationships with rhythmic states in the DG. To address this question, we have elaborated our previous model in several key areas. First, we use "Izhikevich" spiking neurons (Izhikevich, IEEE Transactions on Neural Networks, 2003) that are fit to the individual neuron types observed within the DG. Modeling with spiking neurons permits simulation at a millisecond resolution, allowing a better representation of oscillatory dynamics. Second, we have scaled the network to more biologically realistic levels, with simulations of up to 64,000 granule cells (compared to ~300,000 in mouse). This better allows the observation of long distance spatial effects within the network, a potentially necessary feature to observe meaningful rhythmicity. Finally, we have continued to incorporate more specialized details about the network, including an improved representation of interneuron populations and the addition of neuromodulatory systems. The first result that will be discussed is the computational effect of subcortical inputs - such as acetylcholine and GABA from the medial septum and serotonin from the raphe nucleus - on network oscillations in the DG. By targeting distinct interneuron populations, changes in modulatory input shift the population behavior of the entire DG circuit in the model, consistent with theoretical predictions. The second result we will describe is how these oscillations and accompanying neuromodulation affect the behavior of new neurons incorporating into the network, and vice versa. New neurons are unique in their response to GABAergic signaling: initially GABA is depolarizing, and only later in development are young neurons as fully inhibited as mature granule cells. This makes their response to the typical inhibition/disinhibition cycles unique. These results make several predictions about how new neurons may behave differently from mature cells in vivo, which is important since it is currently unclear how to identify the ages of neurons in DG recordings. Furthermore, these results provide a possible network mechanism for the effects of running and REM sleep on neurogenesis, two conditions that have considerable theta rhythms driven by acetylcholine and serotonin.

doi:

II-41. A computational approach to neurogenesis and synaptogenesis using biologically plausible neurons

Lyle N. Long
Ankur Gupta
Guoliang Fang
The Pennsylvania State University

LNL@PSU.EDU
AZG139@PSU.EDU
GZF108@PSU.EDU

Traditional rate-based neural networks and the newer spiking neural networks have been shown to be very effective for some tasks, but they have problems with long term learning and "catastrophic forgetting.". Once a

network is trained to perform some task, it is difficult to adapt it to new applications. To do this properly, one can mimic processes that occur in the human brain: neurogenesis and synaptogenesis, or the birth and death of both neurons and synapses. To be effective, however, this must be accomplished while maintaining the current memories. In this paper we will describe a new computational approach that uses neurogenesis and synaptogenesis to continually learn. Neurons and synapses can be added and removed from the simulation while it runs. The learning is accomplished using a variant of the spike time dependent plasticity method, which we have recently developed [Ankur and Long, IEEE Paper, IJCNN, June, 2009]. This Hebbian learning algorithm uses a combination of homeostasis of synapse weights, spike timing, and stochastic forgetting to achieve stable and efficient learning. The approach is not only adaptable, but it is also scalable to very large systems (billions of neurons). We will demonstrate the approach on a character recognition example, where the system learns several characters and then to learn more it automatically adds more neurons and synapses. Also, it has the capability to remove synapses which have very low strength, thus saving memory. There are several issues when implementing neurogenesis and synaptogenesis in a spiking code. Though they may seem different, they are actually a coupled phenomenon. When several synapses die, it may lead to a neuron which has no synapses and thus require its removal. Conversely, a neurons death may require updating of synaptic information of all neurons it was connected to. These issues and efficient ways to address them will be discussed. The neuron model that we use is the Hodgkin-Huxley model with four coupled, nonlinear ordinary differential equations. We will compare this approach to the simpler leaky integrate and fire approach also. The algorithms have been implemented in C++ and are fast, efficient, and scalable.

doi:

II-42. Why is connectivity in barrel cortex different from that in visual cortex? - A plasticity model.

Claudia Clopath¹
Lars Buesing²
Eleni Vasilaki³
Wulfram Gerstner³

CLAUDIA.CLOPATH@GMAIL.COM
LARS.BUESING@IGI.TUGRAZ.AT
E.VASILAKI@SHEFFIELD.AC.UK
WULFRAM.GERSTNER@EPFL.CH

¹LCN

²Graz University for Technology

³LCN-EPFL

Electrophysiological connectivity patterns in cortex often show a few strong connections in a sea of weak connections. In some brain areas a large fraction of strong connections are bidirectional, in others they are mainly unidirectional. In order to explain these connectivity patterns, we use a model of Spike-Timing-Dependent Plasticity where synaptic changes depend on presynaptic spike arrival and the postsynaptic membrane potential, filtered with two different time constants. The model describes several nonlinear effects in STDP experiments, as well as the voltage dependence of plasticity under voltage clamp and classical paradigms of LTP/LTD induction. We show that in a simulated recurrent network of spiking neurons our plasticity rule leads not only to development of localized receptive fields, but also to connectivity patterns that reflect the neural code: for temporal coding paradigms with spatio-temporal input correlations, strong connections are predominantly unidirectional, whereas they are bidirectional under rate coded input with spatial correlations only. Thus variable connectivity patterns in the brain, mainly unidirectional in barrel cortex versus bidirectional in visual cortex, could reflect different coding principles across brain areas.

doi:

II-43. Rapid feature binding by a learning rule that integrates branch strength potentiation and STDP

Robert Legenstein¹
Wolfgang Maass²

LEGENSTEIN@IGI.TUGRAZ.AT
MAASS@IGI.TUGRAZ.AT

¹Graz University of Technology, IGI

²Inst. for Theoretical Computer Sci. TU Graz

The pyramidal neuron is the principal cell type of the neocortex. Dendritic processing in pyramidal neurons is highly nonlinear, showing NMDA spikes at basal and thin apical tuft dendrites and calcium spikes at the proximal tuft dendrites (Larkum et al. 2009). It has often been hypothesized that nonlinear branch computations are important for the binding of object features such as shapes or color in single neurons (see e. g. Häusser et al., 2003; Migliore et al. 2008). However, it is unknown how such bindings could emerge through biologically realistic learning mechanisms. Recent experimental data (Losonczy et al. 2008) show that not only the strength of synaptic efficacy is plastic, but also the coupling between dendritic branches and the soma (via dendritic spikes). More precisely, the strength of this coupling can be increased both through a coincidence of dendritic branch activations with action potential generation, and through a coincidence of branch activation with Acetylcholine (ACh). This effect has been called Branch Strength Potentiation (BSP). Starting from an error-minimization principle, we derive a learning rule for BSP and synaptic plasticity that turns out to be consistent with experimental data on BSP and STDP. We show that a simple spiking neuron model with nonlinear dendritic branches learns through this rule to bind salient object features. This works both if saliency is modelled by a global modulatory signal (such as ACh), or by strong synaptic input that forces the neuron to fire. The learning rule induces a competitive mechanism between dendritic branches, where the most activated branch experiences the largest amount of plasticity. Learned weight patterns are stabilized in the absence of a global signal via STDP. We show through computer simulations that autonomous rapid binding of features (the presence of a feature is indicated by concurrent firing of a population of presynaptic neurons) is an emergent property of this learning rule. More precisely, our combined learning rule for BSP and STDP induces a weight distribution where combinations of features that are characteristic for salient objects create clusters of strong synapses at single dendritic branches. Hence single neurons learn through this rule to solve "binding problems". For example, a single neuron learns to fire upon the presence of features A AND B, and also upon the presence of features C AND D, but NOT in response to the presence of features A AND C, or B AND D. Furthermore, learning with the new combined learning rule for BSP and STDP turns out to be extremely fast, approximating single trial learning. Altogether this paper initiates the theoretical analysis of BSP and its interaction with STDP, and it exhibits potentially powerful functional properties of combined plasticity of branch strengths and synaptic efficacies. References: Häusser M, Mel B. *Curr Opin Neurobiol.* 2003 Jun;13(3):372-83. Review M. E. Larkum, T. Nevian, M. Sandler, A. Polsky, and J. Schiller. *Science*, 325:756-760, 2009. A. Losonczy, J. K. Makara, and J. C. Magee. *Nature*, 452:436-441, 2008.

[doi:](#)

II-44. Bimodal structural plasticity can explain the spacing effect in long-term memory tasks

Andreas Knoblauch
Honda Research Institute Europe

ANDREAS.KNOBLAUCH@HONDA-RI.DE

The spacing effect means the finding that learning is more efficient if rehearsal is spread over time compared to single block rehearsal. The spacing effect has been reported to be very robust occurring in many explicit and implicit memory tasks in humans and many animals being effective over many time scales from single days to months. For these reasons it has long been speculated about a common underlying mechanism at the cellular level. We propose that structural plasticity in synaptic networks is this common mechanism. According to our model, ongoing structural plasticity can reorganize networks by replacing obsolete synapses and growing new

synapses at locations that are potentially more useful for storing a given set of memories. We have recently shown that such models can increase storage capacity of neural networks with n neurons from less than one bit per synapse found in common Hopfield-type learning models to diverging $\log(n)$ bits per synapse (A.Knoblauch, F.T.Sommer, G.Palm, Neural Computation, in press). Besides this massive performance increase, the model was also able to qualitatively explain several memory phenomena such as Ribot gradients in retrograde amnesia, absence of catastrophic forgetting, and the spacing effect (A.Knoblauch, Connectionist Models of Behavior and Cognition II, pp 79-90, World Scientific, 2009). Here we focus on quantitatively modeling recent psychological findings concerning long-term spacing effects (N.J.Cepeda et al., Psychological Science 19:1095-1102, 2008). There, probands had to memorize a given set of facts in two rehearsal sessions separated by a time gap of variable length. After an additional retention interval (RI) of up to one year the final recall rate was evaluated. The experiments revealed several characteristics of the spacing effect. 1) For any gap duration, recall performance decays according to the well known forgetting curve. 2) For any RI there is an optimal gap maximizing recall rate. 3) The spacing effect is large, e.g., optimal gaps can double recall rate. 4) The spacing effect is asymmetric, i.e., shorter gaps impair performance more seriously than longer gaps. 5) As RI increases the optimal gap increases. 6) As RI increases, the ratio of optimal gap to RI declines. In a first attempt we simulated our original model to reproduce these characteristics. The model easily reproduced characteristics 1-4 but not 5-6 because optimal gaps were independent of RI as confirmed by a theoretical analysis. In an extended model variant we included two synapse populations with two different rates of structural plasticity per time unit. The first synapse population consisted of motile synapses well suited for short-term storage corresponding to small optimal gaps, whereas the other one consisted of more stable synapses better suited for long-term storage corresponding to longer optimal gaps. Simulation experiments of several network variants revealed that such bimodal structural plasticity can account for all six characteristics of the spacing effect. This was true when the two synapse populations were mixed within a single synaptic network, and also when fast and slow plasticity was segregated into two different networks, e.g., corresponding to hippocampus and neocortex.

doi:

II-45. The role of dopamine in long-term plasticity in the rat prefrontal cortex: a computational model

Denis Sheynikhovich¹

Satoru Otani²

Angelo Arleo³

DENIS.SHEYNIKHOVICH@UPMC.FR

SATORU.OTANI@SNV.JUSSIEU.FR

ANGELO.ARLEO@UPMC.FR

¹Lab. of Neurobiology of Adaptive Processes

²INSERM U952, CNRS-UMR7224, Univ Paris 6

³CNRS-UMR7102, Univ Paris 6

The prefrontal cortex (PFC) is thought to mediate executive functions, including strategic organization of behavior (Fuster, 1995). These functions rely on long-term plasticity within the PFC (Touzani et al., 2007). Dopamine (DA) input to the PFC has been shown to modulate the magnitude and direction (i.e. potentiation, LTP; or depression, LTD) of long-term plasticity induced by tetanic stimulation in vitro (Kolomiets et al., 2009; Huang et al., 2004). Moreover, the DA action is different depending on whether tonic (background) or phasic (stimulation-induced) DA levels are manipulated. Whereas a fair amount of theoretical work addresses short-term DA action on the level of single PFC neurons and its importance for working memory, no theoretical models address the role of DA for long-term plasticity in the PFC. The present work attempts to fill this gap by proposing a computational model of induction and maintenance of LTP and LTD in the PFC under the influence of DA. We use a Hodgkin-Huxley-type computational model of a single PFC layer V pyramidal cell and we study neuronal properties that may be responsible for the changes in synaptic efficacy following tetanic stimulation in the presence of DA. We use a variant of Tag-Trigger-Consolidation framework (Clopath et al., 2008) as a model for LTP and LTD induction and maintenance. Distinct properties of our model are a DA-dose-dependent switch from LTD to LTP during induction, and an inverted-U-shape dependence of protein synthesis threshold on the level of background DA. Protein synthesis is responsible for maintenance and late phase of LTP/LTD in the model. The model has been

tested by stimulating the simulated neuron with spike trains of different duration under different doses of DA and has been found to reproduce well the results of the in vitro studies. Our simulations suggest that in order to comply with the in vitro data, prefrontal synapses must contain a protein that is slowly (on the timescale of minutes) activated in the presence of DA in a dose-dependent manner. The activation value at the time of the stimulation, and the internal calcium level determine the direction of plasticity at prefrontal synapses in the model. The calcium level during stimulation is dependent on the strength of the stimulation mainly via the activation of N-methyl-D-aspartate receptors (LTP) or metabotropic glutamate receptors (LTD). We propose several candidate molecules for the putative DA-activated protein. More generally our results support the hypothesis that phasic release of endogenous DA is necessary for the induction of long-term changes in synaptic efficacy, while the concentration of tonic DA determines the direction (i.e. LTP or LTD) of these changes (Kolomiets et al., 2009). References: Clopath C, Ziegler L, Vasilaki E, Busing L, Gerstner W (2008) PLoS Comput Biol. 4(12):e1000248 Fuster JM (1995) Boston: The MIT Press. Huang YY, Simpson E, Kellendonk C, Kandel ER (2004) Proc Natl Acad Sci USA. 101:3236-3241 Kolomiets B, Marzo A, Caboche J, Vanhoute P, Otani S (2009) Cereb Cortex (E-pub ahead of print) Touzani K, Puthanveettil SV, Kandel ER (2007) Proc Natl Acad Sci USA. 104:5632-5637

doi:

II-46. Structural plasticity improves stimulus encoding in a working memory model

Cristina Savin
Jochen Triesch

SAVIN@FIAS.UNI-FRANKFURT.DE
TRIESCH@FIAS.UNI-FRANKFURT.DE

Frankfurt Institute for Advanced Studies

Instead of being fixed, hard-wired structures, cortical networks are capable of significant reorganization. As we learn new skills or adapt to changes in the environment, brain structure changes as well (Yoshida et al 2003, Hihara et al, 2006)- existing synapses are eliminated and new synapses are grown. These structural changes can sometimes be homeostatic, maintaining the stability of the system, while in other cases they may play an important role in shaping the network function (Zito and Svoboda, 2002). Moreover, it was suggested that such cooperative synaptic formation is important for explaining the statistics of synaptic connections observed in rat cortex, which could not emerge by random sparse connections alone (Fares and Stepanyants, 2009). How such activity dependent structural plasticity affects the function of the network remains unclear, however. Here, we investigate this question for a sparsely connected recurrent neural network, trained to perform a delayed match to sample task. As we have done in a previous model (Savin and Triesch, 2009), an output layer reads out the activity in the network providing a behavioral response, which yields a corresponding reward. Then, reward-modulated STDP (Izhikevich, 2003) shapes the synapses within the recurrent network and those connecting to the motor layer. In addition, structural plasticity is implemented in two steps. First, very weak synapses are pruned, as dendritic spines are known to retract in absence of synaptic activity (Lamprecht and LeDoux, 2004). Second, new synapses are grown between neurons which exhibit correlated activity, but are not yet synaptically connected. Moreover, the two processes are balanced, such that the overall connectivity of the network is preserved. When comparing networks implementing structural plasticity to networks with fixed random connectivity, we see that the performance can be significantly improved by network reorganization. The sparseness of the connectivity matrix, with values similar to those observed in the cortex, ensures appropriate dynamics for the network, but makes the performance critically dependent on the particular instance of the fixed weight matrix. In contrast, the activity-dependent synaptic reorganization will correct a 'bad' initial choice of weights, such that the network can encode the input stimuli more reliably. Interestingly, the Fano factor for the distribution of incoming synapses is small, resembling values reported for cortical networks, as seen in (Fares and Stepanyants, 2009). Our results suggest that activity-dependent structural plasticity could play an important role in optimizing the sparse cortical connectivity to best encode information.

doi:

II-47. Vigour in the face of fluctuating rates of reward: An experimental test

Ulrik Beierholm¹
 Marc Guitart Masip²
 Ray Dolan³
 Emrah Duzel²
 Peter Dayan¹

BEIERH@GATSBY.UCL.AC.UK
 M.GUITART@UCL.AC.UK
 RDOLAN@FIL.ION.UCL.AC.UK
 E.DUZEL@UCL.AC.UK
 DAYAN@GATSBY.UCL.AC.UK

¹Gatsby Computational Neuroscience Unit, UCL

²Institute of Cognitive Neuroscience, UCL

³Wellcome Trust Centre for Neuroimaging, UCL

The two key questions underlying behaviour are what to do and how vigorously to do it. The former has been the topic of a near overwhelming wealth of theoretical and empirical work in the fields of reinforcement learning and decision-making. Although the latter concerns motivation, and so is the subject of many empirical studies in diverse fields, it has suffered a dearth of computational models. Recently, Niv et al (2005) suggested that vigour should be controlled by the opportunity cost of time as measured by the average rate of reward. For instance, if a subject is in an environment with a high average rate, then acting languidly implies receiving rewards inefficiently slowly. This coupling of reward rate and vigour can be shown to be optimal under the theory of average-return reinforcement learning. We sought to test the theory by presenting human subjects with slowly, but systematically, fluctuating reward rates for monetary outcomes, and measuring the vigour of their responses in terms of reaction times. 13 healthy volunteers performed multiple rounds of a rewarded odd-ball discrimination task. In each round, subjects were informed about their potential reward (1-100 pence), and were then shown a screen with three visual stimuli, two of which were identical. The task was to press the button corresponding to the odd-one-out within 500 ms. We varied the potential reward over time in order to exercise the purported link with the motivational system. We expected to see subjects' response times being positively affected by the locally-estimated rate of experienced reward. Note that this prediction goes in exactly the opposite direction of an obvious alternative, that the difference between the potential reward on a trial and the local average reward, should control vigour. We performed a linear regression on the subjects' individual response times, using a collection of nuisance predictors (including the immediate reward in each round) and the regressor of interest, namely the time varying average expected reward. We found that after taking into account the nuisance parameters, a significant fraction of the variance in subjects' responses could be explained by this putative 'motivational signal'. Furthermore, this was in the direction predicted by the theory. This result implies that human behaviour can be influenced by experienced reward in a manner that is consistent with predictions from a normative reinforcement learning model. The model goes on to suggest that it is tonic levels of dopamine that code the average rate of rewards, and thus mediate the effects on vigour. We plan to test this via direct manipulations of the dopamine system.

doi:

II-48. Stability and competition in multi-spike models of spike-timing dependent plasticity

Baktash Babadi¹
 L. F. Abbott²

BB2280@COLUMBIA.EDU
 LFA2103@COLUMBIA.EDU

¹Center for Theoretical Neuroscience, Columbia University

²Dept. of Neuroscience, Columbia University

Synaptic competition and stability are desirable but often conflicting features of Hebbian synaptic plasticity paradigms, including the spike-timing dependent plasticity (STDP). In its simple form, STDP involves potentiation of a synapse when the corresponding presynaptic spike precedes a postsynaptic spike, and depression otherwise. As a result, strong synapses that are more likely to cause a postsynaptic spike get stronger, and weak synapses get weaker until they reach the limits of their allowed range. This very same instability makes the STDP rule highly sensitive to the input correlations: correlated spikes have a greater chance to ignite a postsynaptic action potential, so their

corresponding synapses become stronger in competition with uncorrected ones. In light of recent experiments involving more complex spike patterns, the STDP rule has been augmented by taking into account the interactions between multiple pre- and postsynaptic spikes. The effect of these interactions on the stability/competition interplay is a matter of question. Here, we address this question by numerical simulations of a single integrate-and-fire neuron having few thousand plastic excitatory synapses. We evaluate the final distributions of synaptic weights obtained by three different proposed multi-spike STDP rules. In the "suppression model", the effect of each pre- or postsynaptic spike in inducing plasticity is suppressed by the preceding spike in the same neuron. The resultant final distribution of synaptic weights is very stable over a wide range of parameters, leading to a narrow density function. This behavior is in contrast with the simple STDP models where the synaptic distribution is unstable and bimodal. When a subset of the incoming spike trains are correlated, a competition takes place between their corresponding synapses and those of the uncorrelated spike trains. Surprisingly, the uncorrelated synapses are the winners of the competition. This anti-Hebbian behavior is also in contrast with that of the simple STDP models. In the "revised suppression model", not only the previous one but all the preceding pre- and post synaptic spikes exert a suppressive effect on the subsequent spikes in the same neuron. The qualitative features of stability and competition in the revised model is similar to the former model. In the "triplet model", the effect of each presynaptic spike in inducing plasticity is suppressed and the effect of each postsynaptic spike is facilitated by the preceding spikes in the same neuron. Here the final distribution of weights is highly volatile and sensitive to the parameters of the STDP rule: for any given set of parameters, all the synaptic weights either tend to zero or the maximum allowed value. This behavior is an extreme case of instability observed in the simple STDP rule. When a subset of the incoming spike trains are correlated, an intense competition takes place between their corresponding synapses and those of the uncorrelated spike trains. The correlated synapses are the winners, as observed in the simple STDP models. We conclude that multi-spike STDP rules can have radically different consequences at the network level, depending on the exact implementation of the multi-spike interactions.

doi:

II-49. Risk-minimization through Q-learning of the learning rate

Kerstin Preuschoff¹
Peter Bossaerts²

PREUSCHOFF@IEW.UZH.CH
PBS@RIOJA.CALTECH.EDU

¹Social and Neural Systems Lab, University of Zurich

²California Institute of Technology

In reinforcement learning, the learning rate is a fundamental parameter that determines how past prediction errors affect future predictions. Traditionally, the learning rate is kept constant, either tuned to the learning problem at hand, or - in behavioral and imaging experiments - fit to the observed data. In stable environments this approach works well, yet the thus found learning rate will vary depending on the underlying probabilities and number of trials played. However, prediction performance drops in uncertain environments, when keeping the learning rate constant. Adaptable learning rates may help to learn faster and thus improve predictions over time. We have previously proposed a way to adapt learning rates in risky environments with no changes to the underlying stochastic parameters (zero volatility). In such an environment, the learning rate adapts as a function of risk, decreasing with increasing risk. The optimal learning rate depends on how much correlation (covariance) there is between optimal predictions and the immediately preceding prediction error. While this approach works well in theory, a history of optimal predictors is usually unavailable. Here we propose to adapt the learning rate by minimizing the overall prediction risk (i.e., by maximizing the prediction precision). The overall prediction risk can be considered a value function that represents the discounted sum of future prediction errors given a specific learning rate. We can learn the best learning rate by minimizing this value function. This implicitly incorporates additional information about underlying processes and thus accelerates learning. To achieve this, we borrow ideas from Q-learning to translate risk-sensitive reward-learning into learning an action-value function that minimizes prediction risk using past reward prediction errors. The ensuing optimization problem is a function of a decision-makers risk-sensitivity and converges under the same conditions as standard Q-learning algorithms. Using the inverse prediction risk as a value and the reward-learning rate as an action, we show that the resulting policy adjusts the (reward-) learning rate as a function of the decision-makers risk preferences. This learning rate is a function of both the risk and

volatility of the environment. Learning rates decrease with increasing risk and increase with increasing volatility as shown by behavioral data (Behrens et al, 2007) and predicted by previous models (Preuschoff & Bossaerts, 2007; Behrens et al, 2007). Evidence is discussed that suggests that the dopaminergic system, insula and ACC in the (human and nonhuman) primate brain support a risk-minimizing algorithm in addition to risk-sensitive reward learning. Together with the previous model this can be used to incorporate the trade-off between expected reward and risk by adjusting the learning rate in reward-based learning. The model can be generalized to include risk-neutral as well as risk-seeking decision makers. It essentially extracts information about the origin of uncertainty (e.g., risk vs. volatility) to decide on how much weight to put on more recent prediction errors compared to those that occurred many time steps ago.

doi:

II-50. Model averaging as a developmental outcome of reinforcement learning.

Thomas H. Weisswange¹

Constantin A. Rothkopf¹

Tobias Rodemann²

Jochen Triesch¹

WEISSWANGE@FIAS.UNI-FRANKFURT.DE

ROTHKOPF@FIAS.UNI-FRANKFURT.DE

TOBIAS.RODEMANN@HONDA-RI.DE

TRIESCH@FIAS.UNI-FRANKFURT.DE

¹Frankfurt Institute for Advanced Studies

²Honda Research Institute Europe GmbH

To make sense of the world, humans have to rely on the information that they receive from their sensory systems. Due to noise on one side and redundancies on the other side, it is possible to improve estimates of the signal's causes by integrating over multiple sensors. In recent years it has been shown that humans do so in a way that can be matched by optimal Bayesian models (e.g. [1]). Such an integration is only beneficial for signals originating from a common source and there is evidence that human behavior takes into account the probability for a common cause [2]. For the case in which the signals can originate from one or two sources, it is so far unclear, whether human performance is best explained by model selection, model averaging, or probability matching [3]. Furthermore, recent findings show that young children are often not integrating different modalities [4,5], indicating that this has to be learned during development. But which mechanisms are involved and how interaction with the environment could determine this process remains unclear. Here we show that a reinforcement learning algorithm develops behavior that corresponds to cue-integration and model-averaging. The reinforcement learning agent is trained to perform an audio-visual orienting task. Two signals originate from either one or two sources and provide noisy information about the position of these objects. The agent executes orienting actions and receives rewards that are exponentially decaying with the distance from the true position. The value function is represented through a non-linear basis function network. Positions in the two stimulus dimensions are coded through Gaussian tuning curves. The weights used in the computation of the orienting action are adapted during learning using gradient descent. Actions are selected probabilistically based on the current reward predictions using the softmax-function. The agent quickly learns to act in a way that closely approximates the behavior of a Bayesian observer. It inherently learns the reliabilities of the cues and behaves differently depending on the probability for a single cause. The agent obtains more reward than Bayesian observers that always or never integrate cues. When we test with signals for which the behavior of model selection and model averaging differ most, the agent obtains significantly more reward than a Bayesian model selector and matches very closely the reward obtained by the Bayesian model averager. Furthermore, when a single object is the cause for both stimuli, the variance of the distribution of chosen actions is smaller than for actions based on either of the cues alone. Our results show that a caching reinforcement learning agent can learn when and how to do cue integration, without explicitly computing with probability distributions. Moreover the performance of this an agent is matched best by a Bayesian observer that does model averaging. This suggests that reinforcement learning based mechanisms could at least support the development of such behavior. References: [1]Ernst&Banks (2002) Nature 6870 [2]Körding et al. (2007) PloS One 2(9) [3]Shams&Beierholm (2009) in Proc. Cosyne09 [4]Nardini et al. (2008) Curr. Biol. 18(9) [5]Gori et al. (2008) Curr. Biol. 18(9)

[doi:](#)

II-51. Learning to plan: planning as an action in simple reinforcement learning agents

G Elliott Wimmer¹
Matthijs van der Meer²

GEW2105@COLUMBIA.EDU
MVDM@UMN.EDU

¹Columbia University

²Dept. of Neuroscience, Univ. of Minnesota

Current neuroscientific theories of decision making emphasize that behavior can be controlled by different brain systems with different properties. A common distinction is that between a model-free, stimulus-response, "habit" system on the one hand, and a model-based, flexible "planning" system on the other. Planning tends to be prominent early during learning before transitioning to more habitual control, and is often specific to important choice points (e.g. Tolman, 1938), implying that planning processes can be selectively engaged as circumstances demand. Current models of model-based decision making lack a mechanism to account for selective planning; for instance, the influential Daw et al. (2005) model plans at every action, using an external mechanism to arbitrate between planned and model-free control. Thus, there is currently no model of planning that defines its relationship to model-free control while respecting how humans and animals actually behave. To address this, we explored a "T-maze grid world" reinforcement learning model where the agent can choose to plan. The value of planning is learned along with that of other actions (turn left, etc.) and is updated after an N-step fixed policy (the "plan") is executed, offset by a fixed planning cost. The contents of the plan consist of either a random sequences of moves (random plan control) or the sequence of moves that leads to the highest valued state on the agent's internal value function (true plan). Consistent with previous results (Sutton, 1990), we find that planning speeds learning. Furthermore, while agents plan frequently during initial learning, with experience, the non-planning actions gradually increase in value and win out. Interestingly, even in this simple environment, the agent shows a selective increase in planning actions specifically at the choice point under appropriate conditions. We explore a number of variations of the model, including a hierarchical version where action values are learned for the model-free and model-based controller separately. Thus, a simple Q-learning model which includes an added planning action, the value of which is learned alongside that of other actions, can reproduce two salient aspects of planning data: planning is prominent early but gives way to habitual control with experience, and planning occurs specifically at points appropriate to the structure of the task. The fact that these phenomena can be learned in a simple reinforcement learning architecture suggests this as an alternative to models that use a supplemental arbitration mechanism between planning and habitual control. By treating planning as a choice, the model can generate specific predictions about what point in time, where in the environment, and how far ahead or for how long, agents may choose to plan. More generally, the current approach is an example of blurring the boundary between agent and environment, such that actions (like planning) can be inside the agent alone instead of having to affect the environment, and a demonstration that the state space can include internal variables (such as the contents of a plan), similar to the role of working memory (O'Reilly and Frank, 2006; Zilli and Hasselmo, 2007).

[doi:](#)

II-52. Dynamics of frontal eye field and cerebellar activity during smooth pursuit learning

Jennifer Li¹
 Javier Medina²
 Loren Frank¹
 Stephen Lisberger¹

JENNIFER.LI@UCSF.EDU
 JMED@SAS.UPENN.EDU
 LOREN@PHY.UCSF.EDU
 SGL@PHY.UCSF.EDU

¹UCSF

²University of Pennsylvania

Motor learning requires plasticity in neural circuits that lead to alterations in behavior. Many sites undergo changes during learning, but to understand each site's contribution to the learning process it is important to resolve their learning dynamics in relation to one another. Smooth pursuit is a simple behavior with robust learning; motor signals for pursuit are found in two well-characterized brain areas, the frontal eye field's smooth eye movement sub-region (FEFSEM) and the downstream cerebellar flocculus. We compared the activity of these two areas during pursuit learning to gain insight into how neural signals for learning are represented and modified as they flow through the circuit. We recorded from both the FEFSEM and the flocculus during a learning paradigm where the monkey was repeatedly exposed to a change in the direction of target motion 250 ms after the onset of target motion. The monkey quickly learned to produce a smooth eye movement that predicts the direction and timing of the change in target velocity; these behavioral changes were accompanied by changes in the mean activity of FEFSEM neurons and floccular Purkinje neurons (1). Here, we looked for differences in the proportion of neurons that respond to learning, the timing of the learned response within a trial, and the evolution of neural learning across trials. A greater percentage (61%; 22/36) of Purkinje cells than FEFSEM cells (41%; 41/100) exhibited statistically significant changes in mean firing rate as consequence of learning, and were included for further analysis. The timing of the peak of the learned response during a trial was more heterogeneous in the FEFSEM than in the cerebellum, largely due to the presence of a subset of FEFSEM neurons whose firing rate peaked earlier than the range found in the cerebellum. Thus, within a trial, learned changes in the FEFSEM occur on average earlier and are more distributed than those from the flocculus, consistent with the timing of responses from these areas during unlearned, target-driven pursuit. Next, we compared how the FEFSEM and the cerebellum encode learning across trials, based on neural learning curves constructed with the adaptive estimation algorithm (2). We generated corresponding behavioral learning curves by smoothing the trial-by-trial changes in the eye movement with an appropriate causal filter. Estimating the trial latencies between the neural and behavioral learning curves produced broad but similar distributions in the FEFSEM and the flocculus, suggesting that as a whole, both areas learn at the same rate. 1) Medina, J.F. & Lisberger, S.G. Nat. Neurosci. 11, 1185-1192 2) Frank, L.M. et al. J Neurosci. 22, 3817-3130

[doi:](#)

II-53. Idiosyncratic and systematic features of spatial representations in the macaque PRR

Steve W. C. Chang¹
 Lawrence H. Snyder²

STEVE.CHANG@DUKE.EDU
 LARRY@EYE-HAND.WUSTL.EDU

¹Duke Institute for Brain Sciences

²Washington University School of Medicine

Sensorimotor transformations required for reaching towards something we see were originally thought to take place in a series of discrete transitions from one systematic frame of reference to the next, with gaze-centered representations in many occipital and posterior parietal areas, shoulder-centered representations in dorsal pre-motor cortex (PMd), and muscle- or joint-based representations in motor neurons. More recently both empirical and theoretical work has suggested that having many cortical neurons with a range of idiosyncratic represen-

tations provides computational power and flexibility. We now report both systematic and idiosyncratic coding features existing together on neurons in the parietal reach region (PRR). We fit PRR data from single neurons to a model with a Gaussian representation of target location modulated by eye and hand position gain fields. We find that eye and hand gain fields are organized systematically to form a single compound gain field that codes the distance between the eyes and the hand. In contrast to this systematic gain field organization, we find that the frame of reference for target representations is continuous and idiosyncratic for each cell. We applied three previously published classification schemes for distinguishing gaze- and hand-centered frames of reference (singular value decomposition, Euclidean distance, cross-correlation) in order to confirm that PRR shows a broad range of representations, from gaze-centered to hand-centered, with a bias for gaze-centered. Most cells were best fit by a model in which target location was encoded relative to the hand (hand-centered), relative to the eye (gaze-centered), or relative to a location intermediate between the eye and the hand. We refer to these as "in-bound" cells. A minority of cells encoded targets relative to a location that lay outside these bounds. Such "out-of-bound" representations have been observed in the past in both theoretical and empirical studies. We found that our model explains more variance for in-bound compared to out-of-bound cells. Additionally, injecting random noise into in-bound cells can cause them to appear to be out-of-bound cells. These and other facts suggest that the out-of-bound cells in our data arise as an artifact from the corruption of in-bound cells. Our results support that systematic and idiosyncratic signals may be combined together within a single representation. One dimension of coding in PRR (the frame of reference) is continuous and idiosyncratic within each cell, whereas another dimension (the compound eye and hand gain field) is organized systematically. We suggest that an seemingly haphazard organization, such as is seen with frames of reference in many brain areas, occurs when encoding signals that will be used for multiple different and perhaps non-linear computations. In contrast, a systematic organization, such as is seen with eye and hand gain fields in PRR, occurs when encoding signals that will be used for a small number of fairly linear computations. These two coding strategies may occur together within the same population of neurons.

[doi:](#)

II-54. High-performance continuous neural cursor control enabled by a feedback control perspective

Vikash Gilja
Paul Nuyujukian
Cindy Chestek
John Cunningham
Byron Yu
Stephen Ryu
Krishna Shenoy
Stanford University

GILJA@STANFORD.EDU
PAUL@NPL.STANFORD.EDU
CINDYC@STANFORD.EDU
JCUNNIN@STANFORD.EDU
BYRONYU@STANFORD.EDU
SEOULMAN@STANFORD.EDU
SHENOY@STANFORD.EDU

Neural prostheses, or brain-computer interfaces (BCIs), have the potential to substantially increase quality of life for people suffering from motor disorders, including paralysis and amputation. These systems translate recorded neural signals into control signals that guide a paralyzed arm, artificial limb, or computer cursor. Although current laboratory demonstrations provide a compelling proof-of-concept, the field must continue to increase performance to achieve clinical viability. Many BCIs use activity from motor and/or premotor cortex to achieve continuous control. These BCIs can be viewed from a feedback control perspective, as the motor field has done for the native limb: the brain is the controller of a new plant, defined by the BCI. This perspective leads us to two advances that result in significant qualitative and quantitative performance improvements. We tested these advances in closed loop with one rhesus macaque trained in a virtual 3D workspace. On each trial he used a cursor, controlled by the native contralateral limb or a BCI, to acquire a target on a 2D plane within an allotted time period. Neural data were recorded from a 96-electrode array (Blackrock) implanted spanning PMd and M1. Our designs are informed by a feedback model, which assumes the user develops a volitional control signal to achieve a goal given the current state of the world. This signal and task-unconstrained signals (such as sensory feedback,

attention) give rise to neural firing, which we record. Finally, the decoding algorithm estimates desired cursor movements from the neural firing, and updates the workspace. By applying the assumptions of this simple feedback model, we augment a basic position/velocity Kalman filter. We consider the position/velocity Kalman filter to represent "baseline" as it meshes with the performance of and is algorithmically similar to methods common in the literature (e.g., Kim et al., 2008). All experiments used spike counts generated by a threshold detector without spike sorting. Such a system has clinical appeal, particularly for arrays with potentially decreased SNR (these experiments were 22-24 months post implantation). Design iterations were tested within the same experimental session using a blocked "ABA" design. Through this design process, we made two advances that substantially improve performance. First, using a standard Kalman filter, we fit neural data to a guess of the desired volitional control signal, instead of observed or instructed kinematics. Second, we developed a modified velocity-only Kalman filter, whose observation model incorporates cursor position as feedback. The new BCI appears more controllable and produces straighter reaches and crisper stops. Compared to the standard Kalman BCI, mean time to target is reduced by nearly a factor of two. This system can run freely for hundreds to thousands of trials, making point-to-point reaches to targets randomly placed across the workspace. These feedback-perspective based algorithmic innovations, together with initial experimental verification, suggest that approximately a factor of two performance advance is possible, thereby increasing clinical viability. Support: NSF, NDSEG, Stanford Med Scholars, Soros Fndn, HHMI, SGF, JHU APL under DARPA RP2009: N66001-06-C-8005, CDRF, BWF, NIH-NINDS-CRCNS-RO1, NIH Pioneer Award 1DP1OD006409

[doi:](#)

II-55. The emergence of stereotyped behaviors in *C. elegans*

Greg Stephens^{1,2}

William Ryu³

William Bialek⁴

GSTEPHEN@PRINCETON.EDU

WRYU@PHYSICS.TORONTO.CA

WBIALEK@PRINCETON.EDU

¹Lewis-Sigler Institute

²Joseph Henry Laboratories of Physics

³University of Toronto

⁴Princeton University

Many organisms, including humans, engage in stereotyped behaviors and these are often attributed to a deterministic command process within the nervous system. Here we use the locomotor dynamics of the nematode *C. elegans* to suggest an alternative explanation in which stereotyped behavior emerges due to noise within a non-linear dynamical system. In previous work (PLoS Comp Bio 4, e1000028 (2008)) we found that the body shapes of freely-crawling *C. elegans* are well-captured by four 'eigenworms', two of which encode the phase of a locomotory wave that generates forward and backward motion. We also used this representation to infer a non-linear dynamical model for the phase in which forward and backward crawling emerge as attractors of the deterministic dynamics. Here we show that noise induces reversals between forward and backward crawling and that the predicted reversal rate is in good agreement with experiment, with no adjustable parameters. In this model, reversals follow a stereotyped trajectory for the same reason that Brownian escape over a barrier is dominated by a narrowly defined class of trajectories. Stereotypy becomes even clearer in the dynamics with lower noise levels; the real *C. elegans* is just outside the regime where the reversal rate follows an Arrhenius dependence on the noise level. We discuss the implications of our results for *C. elegans* and other organisms.

[doi:](#)

II-56. Preparatory tuning in premotor cortex relates most closely to the population movement-epoch response

Mark M. Churchland
Matthew Kaufman
John P. Cunningham
Krishna Shenoy
Stanford University

CHURCH@STANFORD.EDU
MATT235@STANFORD.EDU
JCUNNIN@STANFORD.EDU
SHENOY@STANFORD.EDU

A common practice in systems neuroscience is to examine neural activity that precedes movement, in the hopes of better understanding the mechanisms that determined that movement. This approach has led to an understanding of the basic mechanisms that trigger saccades, and the more cognitive mechanisms underlying decisions regarding where to saccade. It is generally agreed that preparatory activity encodes the upcoming saccade vector, and that a saccade is triggered when the strength of preparatory activity crosses a threshold. Models of the reach system often assume a similar pattern: rising preparatory activity leading to a similarly tuned burst of movement-related activity. Alternately, we have proposed that preparatory activity does not rise in strength, but rather needs to be brought to a particular state before movement onset. While different, both frameworks suppose that preparatory activity and movement activity are causally linked and thus closely related. Despite this presumed link, most studies of neural tuning account for preparatory and movement-related responses using different, albeit related, parameters (e.g., target location versus reach velocity). Thus, the following open questions remain. First, what is preparatory activity tuned for? Second, how do activity patterns during the two epochs - preparatory and movement - relate to one another? We analyzed four datasets from three monkeys performing delayed-reach tasks. 550 single neurons were recorded from motor and premotor cortex, using single-electrode and array recording techniques. We found that a neuron's tuning during the preparatory epoch typically showed little straightforward relationship to its tuning during the movement epoch (mean tuning correlation = 0.21, 0.19, 0.06, and 0.10 across the datasets). Despite this, the preparatory activity of individual neurons could be predicted rather well by the population-level pattern of movement activity. To assess this, we assumed that each neuron's preparatory activity was determined by a preferred direction in the space of movement-epoch firing-rate patterns. That space was constructed by applying principal component analysis to a matrix where each row contained the movement-epoch responses of all neurons. Preferred directions in this abstract space accounted for preparatory tuning better than did preferred directions in any of the more traditional spaces we tested, including spaces based on target location, hand velocity, and EMG activity. Furthermore, preferred directions were more stable (when assessed across different subsets of conditions) in this abstract space than in any of the other candidate spaces. We conclude that the reach system is very different from the saccadic system. In particular, the motor burst is not a simple burst, but a pattern of activity that bears no superficial relationship to the preceding pattern of preparatory activity. Nevertheless, activity during the two epochs is lawfully related once one takes the entire population into account. This supports the idea that there is a close causal relationship between preparatory and movement activity, despite the lack of congruent tuning at the single-neuron level. It further suggests that preparatory activity exists not to represent external parameters, but to provide an initial state that determines the subsequent pattern of movement activity. Support: BWF, HHW, Stanford/NSF graduate fellowships, NIH-CRCNS-R01

doi:

II-57. Sparse connectivity in short-term memory networks

Dimitry Fisher¹
Emre Aksay²
Mark Goldman¹

DIMITRIFISHER@YAHOO.COM
EMA2004@MED.CORNELL.EDU
MSGOLDMAN@UCDAVIS.EDU

¹Center for Neuroscience, UC Davis

²Weill Medical College of Cornell University

Short-term memory is thought to be stored in patterns of neural activity that persist for several seconds following

a transient stimulus. However, neural mechanisms that underlie this persistent activity are not fully understood. Previous efforts at modeling short-term memory networks have used simplifying assumptions on both the patterns of connections and the nature of responses in the network, such as assuming linear networks in which negative rates were allowed or imposing certain symmetries on the connectivity. Although these studies identified theoretical networks that can generate short-term memory, the fundamental question remains unanswered: what are the short-term memory network architectures in actual biological systems? To address this question, we have developed a general modeling framework applicable to a wide variety of short-term memory settings. Experimental data are directly incorporated into the model while the evaluation of the network connectivity is reduced to a constrained linear regression problem with no a priori assumptions on the form of the connection strengths. The framework uses single-neuron properties known from electrophysiology: neuronal tuning curves and neuronal spiking responses to somatic current injection. Responses to current injection are calibrated by tuning the parameters of a conductance-based model neuron to reproduce the current injection experimental data spike by spike. Best-fitting connection strengths are found for any choice of the (typically unknown) nonlinear synapto-dendritic activation curves, which describe the current flowing into the soma from individual dendrites as a function of presynaptic neuron firing rate. For a network of spiking neurons with stochastic noise based on experimentally measured coefficients of variation, we perform systematic searches of parameter space to determine which sets of synapto-dendritic activations and resulting network connectivities produce good matches to the memory activity observed in the experimental system. This methodology is applied to two networks: (i) a network with monotonic tuning curves - the oculomotor integrator - that calculates and stores the eye position resulting from a sequence of eye velocity commands, and (ii) a network with peaked tuning curves that stores the spatial location of a cue. In contrast to previous studies, we find that a number of network architectures provide good fits to the experimental data, and their connectivity structure is usually highly sparse. These networks range in structure from highly recurrent to strongly feedforward, and function at least as accurately as the biological system being modeled. Sparseness of connections is manifested by a power-law structure in the synaptic weight distribution. We emphasize that sparseness was not imposed upon the network, but rather emerged from the biological requirement that neurons maintain well-tuned persistent activity. Sparseness in this setting reflects that each neuron's firing rate, as a function of the stored variable, is well-approximated by a relatively small set of recurrent inputs. Although the reasons for this sparseness are not yet entirely clear, we speculate that this result is similar to that obtained in other smooth-function-approximation problems, in which the distributions of decomposition coefficients are often sparse. We suggest that similar sparse structures may emerge in a wide range of both memory and non-memory circuits.

[doi:](#)

II-58. Modeling firing-rate dynamics: From spiking to firing-rate networks

Evan S. Schaffer

L. F. Abbott

Columbia University

ESS2129@COLUMBIA.EDU

LFABBOTT@COLUMBIA.EDU

Firing-rate models provide an attractive approach for studying large neural networks because they can be simulated rapidly and are amenable to mathematical analysis. Traditional firing-rate models have the obvious shortcoming of using a single time constant (typically a membrane or synaptic time constant) to describe all changes in rate, and they require neurons in a network to fire asynchronously. This is likely to be violated in many cases; in fact, transient synchronization of subgroups of neurons may be an important mechanism for generating rapid behavioral responses. To address this issue without losing the advantages associated with a simple firing-rate description, we have developed a form of firing-rate model based on an approximate Fokker-Planck analysis. A Fokker-Planck equation can be used to describe the distribution of membrane potentials for a population of neurons receiving noisy input. However, most methods for approximating solutions to this type of equation lead to considerably more complex equations than are practical for large networks. For example, there is no closed-form solution describing a population of Integrate-and-Fire (IAF) neurons receiving arbitrary time-varying input. A linear approximation to the response can be calculated, yielding impressively high accuracy, but it involves cumbersome equations (e.g. Brunel & Hakim, 1999; Mattia & Del Giudice, 2002; Ostojic et al., 2009). We show here that in

a variant of this model, the Quadratic IAF, the fully nonlinear rate response can be approximated in a surprisingly simple form. Importantly, this approximate solution makes no assumptions about the shape, amplitude, or continuity of the input current. With an understanding of how dynamic external inputs drive firing rates for both asynchronous and synchronous populations of neurons, we study how units described in this way can be linked to describe the firing-rate dynamics of spiking networks with various patterns of synaptic connectivity. We find that the novel firing-rate model captures the time-varying firing-rates of the spiking network across a wide range of parameters. This holds equally well in parameter ranges where the asynchronous state is stable, and where highly synchronized firing occurs. Furthermore, the model also reproduces the dynamics of transient synchronization, which can be quite complicated. Finally, we show that the rich firing dynamics of a network of both excitatory and inhibitory neurons can be well approximated by a coupled E-I rate network. The simplicity of the model we have derived makes it highly amenable to use as the basis for network models. This will hopefully make tractable the study of the dynamics of large networks.

[doi:](#)

II-59. Optimal network architectures for short-term memory under different biological settings

Sukbin Lim
Mark Goldman

SBLIM@UCDAVIS.EDU
MSGOLDMAN@UCDAVIS.EDU

Center for Neuroscience, UC Davis

Short-term memory is thought to be maintained by patterns of neural activity that are initiated by a memorized stimulus and persist long after its offset. Because memory periods are relatively long compared to biophysical time constants of individual neurons, it has been suggested that network interactions can extend the time over which neural activities are sustained. However, the form of such interactions is currently unknown, and experimental and theoretical work has suggested a range of different network architectures that could subserve short-term memory such as feedforward networks or recurrently connected networks implementing attractor dynamics. Here, we explore the conditions under which each network may be optimal in order to gain insight into why different mechanisms might be used by different systems. For each network architecture, we characterize how the fidelity of memory is maintained in the presence of noise by calculating the Fisher information conveyed by the network activity about the stimulus strength at a previous time. Calculations are performed under several biologically relevant conditions, such as common and independent noise, variable memory durations, and additional constraints such as whether the start time of the stimulus to be memorized is known. We first consider low-dimensional ("line attractor") networks that have been suggested to occur in oculomotor and neocortical working memory systems. In the presence of uncorrelated noise, we find a paradoxical result: network performance is benefited by having an imperfect memory-holding mechanism, independent of the level of noise. We show that there is an "optimal forgetting" time constant of decay of network activity that reflects a tradeoff between having a long time constant, so that the signal does not decay, and having a short time constant so that noise does not accumulate too much. This result assumes that noise is presented continually and can build up before the signal arrives. However, if noise enters the system with the signal, or if the animal can anticipate the start of memory performance and reset its neuronal activities, then the perfect integrating mode performs better than any decaying mode. The feedforward network exhibits qualitatively different behavior: the duration of input accumulation is bounded by the number of feedforward stages, and noise flows out the system after some time. This makes the feedforward network better than the line attractor without a reset. However, the reset mechanism does not improve much the performance of the feedforward network, so it performs worse than the line attractor with reset. Together, these results suggest that there may not be a single network architecture that is optimal in all situations. Already our work has suggested how the optimal time constant of decay of activity and network architecture may be different depending upon the experimental setting, the time over which the memory must be stored, and the form in which noise arrives at the network. Currently, we are testing how correlated noise and constraints on synaptic strengths influence the information conveyed by a memory network, and are developing computational methods to find the optimal architecture in different experimental settings.

doi:

II-60. Neuroptikon: a customizable tool for dynamic, multi-scale visualization of complex neural circuits

Frank Midgley
Donald J. Olbris
Dmitri Chklovskii
Vivek Jayaraman

MIDGLEYF@JANELIA.HHMI.ORG
OLBRISD@JANELIA.HHMI.ORG
CHKLOVSKIID@JANELIA.HHMI.ORG
VIVEK@JANELIA.HHMI.ORG

Janelia Farm Research Campus, HHMI

Groups around the world are using electron and light microscopy, optogenetics and physiology to map neural connectivity in a variety of brain regions and species. Having the circuit diagram for a brain region is potentially very useful, but how should such information be represented? The visualization of neural circuits becomes difficult as the density of both a circuit's connections and the data associated with the components of the circuit increases. Visualizing the entire circuit causes fine detail to be lost, while focusing on specific components hides their context within the larger circuit. We have developed an open-source, Python-based, abstract circuit visualization package, Neuroptikon, which overcomes these problems with an expandable suite of interactive tools. Circuits are represented in Neuroptikon at both a biological and a visual level. A simple biological model allows the construction of neural circuits from regions, neurons and/or neurites and can be extended by user-defined attributes. A NetworkX (Hagberg et al. 2008) version of the circuit is maintained which allows graph-theoretic analysis. A flexible visualization layer sits above the model and allows any biological attributes (built-in or user-defined parameter-value pairs) to control the display of some or all of the circuit's components. Circuit components can be positioned in two or three dimensions either manually or automatically by one of the included layout algorithms. Visualizations can be managed independently of circuits allowing multiple styles of visualization for the same circuit (say, an anatomically accurate layout of neurons versus a layout based on wiring optimization or one chosen for maximum clarity) and re-use of the same visualization for different circuits. Neuroptikon allows interaction with circuits through its graphical and scripting interfaces. A basic set of interactive tools including local connectivity highlighting and shortest path finding are provided by the graphical interface. Both the biological model and visualization layer can be queried and modified via the scripting interface. Script commands can be executed interactively or via saved scripts. The scripting interface also allows expansion of the user interface via new layout algorithms, custom inspectors, etc. We have developed a software package that enables an abstracted representation of neural circuits suitable for conceptual analysis and experimental design. This tool can also serve as a useful front-end for repositories of neuroanatomical and neurophysiological data, and citation databases indexed by neuroanatomical features. Neuroptikon has been used to model and visualize synapse level connectivity in *C. elegans*, compartmental arborizations in *Drosophila* and region level connectivity in the primate visual cortex, and we will present these test cases as part of a demonstration of the tool. References 1. Aric A. Hagberg, Daniel A. Schult and Pieter J. Swart, "Exploring network structure, dynamics, and function using NetworkX", in Proceedings of the 7th Python in Science Conference (SciPy2008), G  l Varoquaux, Travis Vaught, and Jarrod Millman (Eds), (Pasadena, CA USA), pp. 11-15, Aug 2008

doi:

II-61. Near exact correction of path integration errors by the grid cell-place cell system

Sameet Sreenivasan¹
Ila Fiete²

SAMEET.SREENIVASAN@GMAIL.COM
FIETE@MAIL.CLM.UTEXAS.EDU

¹Center for Learning and Memory, University of Texas at Austin

²University of Texas at Austin

Path integration is the mechanism by which animals use self-motion cues to integrate their velocity and thus keep a continuously updated estimate of their position relative to a point of departure. Grid cells in the medial entorhinal cortex are widely thought to be the neural integrators involved in this computation, for many reasons including: 1) Grid cells display a spatially precise activity pattern, firing whenever the animal visits a vertex of a virtual triangular lattice tiling the floor. 2) The grid cell response is insensitive to the shape and size of the enclosure but only depends on how far the animal has moved within the enclosure. 3) The grid cell position code is generated even during navigation in the dark. 4) Grid cells project to many hippocampal subfields including CA1 and CA3, where place cells are found. It remains unclear why the brain would represent a non-periodic, local variable (location) by a highly periodic, non-local neural code. Previously, it was shown that grid cells with multiple different periods can represent position with a combinatorially large dynamic range: the capacity of the grid code for position scales exponentially with the number of different periods, and the maximum length scale greatly exceeds any single grid period. But an apparent pathology of the code is its profound noise sensitivity: small perturbations of the grid cell code translate to massive errors in represented position, typically almost as big as the exponentially large representable range. Since neural integrators are inherently noisy, each network of grid cells sharing a period accumulates errors in its integrated estimate of position. These small integration errors in grid cells would lead to large and rapidly escalating errors in estimated position. However, the observed ability of animals to follow a straight path home after foraging excursions without external sensory cues suggests they can maintain a sound estimate of position and that the brain may be employing a technique to sharply reduce these integration errors. We show how the grid cell code's exponentially large capacity and noise sensitivity endow it with properties common to a family of sophisticated error-correcting codes known in coding theory (Reed-Solomon codes, which are widely used for information storage and transmission in audio media and communications), which allow for nearly exact correction of errors, far beyond simple $1/\sqrt{N}$ improvements from averaging multiple estimates. Next, we construct a neural network model in which hippocampal place cells read out the grid cell position code to perform the error-correcting inference, then project back to the grid cells to correct the integration errors. Iteratively, this closed loop of interactions between hippocampus and entorhinal cortex enables precise internal estimation of the animal's trajectory, limited primarily by sensor errors, rather than by noise in the integrator as in conventional neural integration models. Our results demonstrate a functional role for the known connectivity between entorhinal cortex and the hippocampus. Furthermore, our model illustrates how the brain may be exploiting more sophisticated strategies than simple population averaging, through the explicit use of exact error correcting codes.

[doi:](#)

II-62. Deciding with single spikes: MT discharge and rapid motion detection

Bryan Krause¹
Geoff Ghose²

BRYANMKRAUSE@GMAIL.COM
GEOFF@CMRR.UMN.EDU

¹University of Wisconsin

²Department of Neuroscience, University of Minnesota

Sequential sampling models have been used to explain the behavioral performance and timing of reaction time tasks, but most electrophysiological studies employing these models have relied on visual tasks in which the nature of the stimulus precludes rapid performance by requiring extensive temporal integration. Behavioral data from our lab suggests that one particular assumption of these models, the even weighting of information over time, is unable to explain very fast reaction times (< 250 ms) in a well practiced task. In this task, monkeys detected a brief pulse of high contrast coherent motion embedded in noise. Electrophysiological recording of MT responses during this task demonstrates that individual neurons reliably signal the occurrence of these pulses on a time scale of tens of milliseconds. Moreover, there was a reliable relationship between activity on these time scales and the behavioral choices made by the animal. By employing a novel application of information theory, we demonstrated that the combined encoding and decoding reliabilities on particular neurons were largely sufficient to explain both behavioral reliability and timing in the task. Moreover, we found a strong correlation between encoding and decoding reliability, suggesting that the animals were basing their decisions solely on these

reliable neurons. In this study, we test whether common employed assumptions regarding neuronal coding and integration are sufficient to explain these observations. Specifically, we model MT neurons with Poisson discharge whose rate is rapidly modulated by the onset of coherent motion. Using a decay constant of 100 ms, we linearly summed the discharge from multiple neurons to produce a decision variable in an standard accumulator model. We then adjusted the decision threshold to match observed performance and studied the nature of neuronal pooling necessary to explain performance by varying both the number, selectivity, and interneuronal correlations of sampled neurons. We find that, if overall performance and reaction time are the sole constraints, a wide range of pooling models can explain our data, including ones in which hundreds of MT neurons are sampled by the animal. However, the strong correlation between encoding and decoding reliability over time scale of milliseconds observed in our recordings, places far stronger constraints on pooling models. Specifically the poor decoding performance of neurons with moderate sensory reliability can only be explained by these neurons not contributing to the decision variable. Thus, only models in which small number of reliable neurons are sampled over brief periods of time are sufficient to explain our observations. Because the model employs standard Poisson discharge and accumulation, it does not rely on complex temporal encoding or decoding schemes. Our results demonstrate the potential for very optimized neuronal pooling in the case of well-practiced tasks, in which decisions are based on small numbers of action potentials from neurons with reliable rate modulation.

[doi:](#)

II-63. Testing efficient coding: projective (not receptive) fields are the key theoretical prediction

Eizaburo Doi¹

Greg Field²

Jeffrey Gauthier²

Alexander Sher³

Martin Greschner²

Jonathon Shlens²

Timothy Machado²

Liam Paninski⁴

Debrah Gunning³

Keith Mathieson³

Alan Litke⁵

E. J. Chichilnisky²

Eero P. Simoncelli¹

EDO@CNS.NYU.EDU

GFIELD@SALK.EDU

GAUTHIER@SALK.EDU

SASHA@SCIPP.UCSC.EDU

GRESCHNER@SALK.EDU

SHLENS@SALK.EDU

TAM2138@COLUMBIA.EDU

LIAM@STAT.COLUMBIA.EDU

DGUNNING@PHYSICS.GLA.AC.UK

K.MATHIESON@PHYSICS.GLA.AC.UK

ALAN.LITKE@CERN.CH

EJ@SALK.EDU

EERO.SIMONCELLI@NYU.EDU

¹New York University

²The Salk Institute

³University of California, Santa Cruz

⁴Columbia University

⁵CERN

A fundamental principle for understanding the structure of early sensory information processing is that of efficient coding: information about the external world transmitted to the brain should be maximized, subject to limitations on resources such as firing rate and neural population size. Previous work has suggested that the receptive field structure of retinal ganglion cells (RGCs), the output neurons of the retina, is efficient in this sense [Atick & Redlich, 1990; van Hateren, 1992]. Assuming a population of identical, equally-spaced, circularly symmetric receptive fields, efficient coding principles were shown to account qualitatively for the spatial frequency responses of retinal ganglion cells at different mean illumination levels. However, previous work did not include known aspects of neural architecture, such as inhomogeneities in the photoreceptor lattice and in the spacing and structure of ganglion cell receptive fields, that can have a major impact on the theoretical predictions. Here we examine efficient coding in a framework with fewer assumptions: inhomogeneous populations of linear input and output neurons with different density and independent Gaussian noise [Campa et al., 1994; Doi et al., Cosyne08]. We

find that the efficient coding principle does not uniquely specify the receptive fields of output neurons, but instead places a strong constraint on the "projective fields" of the input neurons, that is, the strength of the connection between a given input and all the output cells. Specifically, efficient coding uniquely predicts the inner products of the projective fields of all pairs of input neurons. An experimental test of efficient coding therefore requires measurement of the pattern of connectivity between the full set of input neurons and a complete collection of output neurons. Recent advances in large-scale, high-resolution recording techniques have made such measurements possible, and consequently, allow the first quantitative test of the theory. We examine a data set in which receptive fields of complete collections of RGCs covering a region of retina were mapped at high resolution, revealing the strength of the inputs of each cone photoreceptor to each RGC, and report three principal findings. (1) The spatial pattern of the predicted inner products is similar to the measured pattern. (2) The predicted inner products explain only 16% of the variance of the measured values, in particular the norms are highly variable in the data. Adjusting these norms improves the prediction, allowing the theory to account for 51% of the variance. (3) Information transmission of the measured projective fields is approximately 90% of the theoretical limit. We conclude that the linear-Gaussian form of the efficient coding theory fails to match the data in detail, but predicts the efficiency of the retinal network for encoding natural scenes fairly accurately. A key feature suggested by this study is that the cone photoreceptors are not utilized as uniformly as the theory predicts. Nonlinear transformations and/or non-Gaussian forms of noise or image priors may explain the discrepancies.

[doi:](#)

II-64. The role the retina plays in shaping predictive information in ganglion cell populations

Stephanie E. Palmer
Michael J. Berry
William Bialek
Princeton University

SEPALMER@PRINCETON.EDU
BERRY@PRINCETON.EDU
WBIALEK@PRINCETON.EDU

We have examined how groups of retinal ganglion cells (RGCs) encode predictive information in their collective firing patterns. Predictive information is defined here as the mutual information between firing patterns across several cells in the retina at a particular time, and the firing patterns of the same neurons at a time dt in the future. Put simply, we are asking how well the firing of the retina 'now' specifies the firing of the retina in the future. We find substantial predictive information in groups of retinal ganglion cells that grows with the number of neurons pooled. This predictive information is due, in part, to the intrinsic firing properties of the ganglion cells, as well as to correlations in the stimulus. We attempt to disentangle these effects by examining responses to temporally uncorrelated white noise stimuli. We find that roughly half of the predictive information we observe can be accounted for by intrinsic properties of RGCs, while the remaining half is induced by stimulus correlations. To assess what collective properties of ganglion cell firing account for the observed predictive information, we break correlations between cells and within cells in time. We find that the predictive information in groups of ganglion cells outstretches the summed contribution from individual cells' predictive capacities, leading to substantial synergy in larger groups of RGCs. We also assess whether the way in which the retina encodes stimulus information is optimized for prediction. Preliminary evidence suggests that the retina does indeed compress information about past stimuli such that information about the future is maximally preserved.

[doi:](#)

II-65. Odour identity is represented by the pattern of activated neurons in the drosophila mushroom body

Robert Campbell¹

Glenn C. Turner¹

Kyle Honegger²

CAMPBELL@CSHL.EDU

TURNER@CSHL.EDU

HONEGGER@CSHL.EDU

¹Cold Spring Harbor Laboratory

²Watson School of Biological Sciences, Cold Spring Harbor Laboratory

The insect mushroom body (MB) is a prominent brain structure involved in olfactory learning and memory. Olfactory information is transmitted to the MB via the antennal lobe projection neurons (PNs), which receive direct input from the Olfactory Receptor Neurons (ORNs) at the sensory periphery. PN tuning curves are broad and a single mono-molecular odourant activates many of the 50 different PN classes. Electrophysiological studies have shown that the intrinsic neurons of the MB, the Kenyon cells (KCs), are highly odor-selective and that odor representations are relatively sparse in the MB. For sparse representations to be effective, a balance must be struck between stimulus-selectivity, useful for learning, and the information loss and susceptibility to noise that accompany rare responses. We are investigating these issues using two-photon Ca⁺⁺ imaging to monitor activity of up to 150-200 KCs simultaneously. Using the new genetically encoded Ca⁺⁺ indicator GCaMP 3.0, the sensitivity of which approaches single spikes, we can record detailed activity in about 8% of the total population of Kenyon cells in a single animal. Imaging from the somatic layer of the MB, we tracked the change in fluorescence of individual KC somata in response to a variety of mono-molecular odourants presented in randomised order. Response sparseness is similar to that seen using electrophysiology, with roughly 10% of neurons responding significantly to 1:100 odour dilutions. Repeated presentations of the same odour evoked similar responses, while other odours evoked different but overlapping patterns of activity. How well does the recorded KC population discriminate between different odourants and upon what basis is the discrimination made? We used multiple discriminant analysis (MDA) to construct a supervised classifier which predicts odourant identity based on the pattern of activated KCs in a single fly. Classification accuracy is about 70-95% across flies. Classifiers based upon binarised responses or mean evoked responses perform with similar accuracy. By mining the results of the classifier we were able to explore how many cells contributed to odour discrimination. Typically only 15-30 neurons are informative. Most neurons are silent. One way to derive sparse representations while minimizing information loss is if KCs integrate information from different input neurons; this would presumably produce a greater diversity of tuning curves in KCs than at earlier layers. We tested this hypothesis by imaging KCs in flies with a reduced complement of 6 odourant receptor classes with known tuning curves. We found that most KCs had tuning curves which resembled one of the 6 ORN classes but a small number of cells showed response properties that appeared to be composites of more than one ORN class. These results suggest that KCs act as feature detectors, where each cell responds selectively to information from multiple different inputs.

doi:

II-66. Efficient theta-locked population codes in olfactory cortex

Keiji Miura^{1,2}

Zachary Mainen³

Naoshige Uchida²

KMIURA@MCB.HARVARD.EDU

ZMAINEN@NEUROPHD.ORG

UCHIDA@MCB.HARVARD.EDU

¹JST PRESTO

²Harvard University

³Inst. Gulbenkian de Ciência

The olfactory system has some unique properties amongst other sensory systems. First, stimulus encoding is massively distributed and the olfactory cortices appear to lack organized spatial structures seen in the most neocortices. Second, stimulus sampling is under almost complete control by sniffing. Here, we studied how these properties shape neural representations in the olfactory cortex. Rats were trained to perform an binary odor-

mixture categorization task (using 4 mixtures of 3 odor pairs). Although rats took a variable number of sniffs during odor sampling in different trials, they achieved maximal performance with a single sniff. To investigate the neural basis of this rapid coding of odors, simultaneous recordings using tetrodes were made from multiple single units (3-21) in anterior piriform cortex (aPC). Odor-evoked responses in aPC were tightly locked to sniffing at 7-9 Hz, indicating that olfactory information is chunked in packets locked to the theta cycle. Decoding analysis shows that information contained in the neural activity during a single cycle is enough to account for accuracy of odor discriminations observed during behavior performance. Interestingly, information was nearly independent between nearby neurons both in signals and noise and, therefore, efficient. Whereas previously reported noise correlations in neocortex are substantial (mean: 0.1-0.2), in aPC they were near zero (mean: 0.0046) and independent of the similarity of odor tuning (signal correlations) and the spatial distances between simultaneously recorded neurons. These features suggest a different computational strategy than other regions of neocortex. If low noise correlations in aPC are solely due to anatomical features of the olfactory system then one might expect them to be constant. However, to the contrary, noise correlations were dynamic, being higher during non-task periods and lower during active sampling of odors, indicating that active sampling at theta frequency also play a role in the establishment of population codes in the aPC. Together, these observations show that neural coding in the olfactory cortex features extremely low noise correlations that depend on dynamical processes in widely distributed neuronal ensembles. The resulting rapid and efficient population code offers insight into the neural mechanisms of rapid odor-guided decision-making.

[doi:](#)

II-67. Temporally distributed information gets optimally combined by change-based information processing

Reza Moazzezi¹

Peter Dayan²

REZAMOAZZEZI@BERKELEY.EDU

DAYAN@GATSBY.UCL.AC.UK

¹Redwood Center for Theoretical Neuroscience, UC Berkeley

²Gatsby Computational Neuroscience Unit, UCL

Neural circuits are responsible for carrying out cortical computations. These computations consist of separating relevant from irrelevant information and noise in the incoming input (or stimuli). However, circuits do not receive all the information they need to accomplish their computations at once; rather, in many situations, the incoming input gets presented to the network over the course of a few hundreds of milliseconds. This raises the question of how the information that is provided at different times is integrated. One challenge in doing this is that the input is in general also affected by other nuisance parameters (irrelevant information) that might continually change during the time of integration and processing. Here we address this problem in the context of a hyperacuity task called the bisection task. In this task, an array of three parallel bars is presented to the subjects, who have to decide if the middle bar is closer to the right or the left bar. The signal (i.e. the relevant information) here is the sign of the deviation of the middle bar from the middle point of the array, and is fixed during a trial. Fixational eye movements, like micro-tremors, continually change the location of the array during each trial, and are therefore the source of the irrelevant information mentioned in the above paragraph. The duration of a trial is of the order of a few hundreds of milliseconds. This is therefore an excellent model task to study how neural circuits process the information that is provided gradually over time. We modelled this using a recurrent network inspired by primary visual cortical circuits. We have previously shown that in the presence of trial by trial variability of the overall location of the array of the three bars (but no variability within a trial), coding information by the early change of the network's state (which we call Change-based Processing) could lead to near optimal performance (1). This method, which is superior to conventional methods of coding by attractor states, makes its decision based on the sign of the difference between two measurements of a scalar statistic of the neural activity (in this case, its centre of mass). The timings of these two measurements could be learned, but once learned are fixed and are independent of strength of the signal. Here we show that the same method successfully combines the information that the network receives over time, still performing near optimally. Theoretical analysis of this performance indicates that an eigenmode of a sub-matrix of the recurrent weight matrix is responsible for both extracting

the relevant information, and combining it near optimally over time; this eigenmode also plays a key role in the subsequent evolution of the statistic whose change is the basis for the ultimate decision. We also demonstrate that the magnitude of the change of the statistic reflects the amount of information available in support of the decision. References: 1- Moazzezi, R. & Dayan, P. (2008). Change-based inference for invariant discrimination. *Network: Computation in Neural Systems*, 19, 236-252

[doi:](#)

II-68. Fisher information in correlated networks

David G. T. Barrett

Peter E. Latham

BARRETT@GATSBY.UCL.AC.UK

PEL@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, UCL

The information in a network of neurons depends on its correlational structure, but not in any systematic way: correlations can increase information, decrease it, or have no effect at all [1]. At least those are the theoretical possibilities. But what happens in realistic networks, where the correlational structure is not arbitrary, but is determined by recurrent connectivity and external input? We find that the parameters that maximize information also maximize correlations. Thus, for the model we consider, high correlations are synonymous with high information. The above results are based on a simple model network consisting of recurrently connected McCulloch-Pitts neurons. The network receives input from a population of neurons that code for an angular variable, denoted θ . The input consists of a noisy hill of activity centered around the true value of θ . The network connectivity has two components: a strong random one, and a weak structured one. We make the random component strong to be consistent with the observation that both excitatory and inhibitory input to a neuron is large; we make the structured component weak to prevent runaway excitation [2]. For the structured component we use Mexican hat connectivity, which matches, at least approximately, the form of the input, and thus has the potential to enhance information transmission. We have shown previously that we can compute the correlational structure analytically [3], and therefore, we can compute what is called the linear Fisher information – a bound on the inverse of the variance of θ for a linear estimator [3]. Both the correlations and linear Fisher information depend on three parameters: W and J , which determine the overall strength of the background and structured connectivity, respectively, and I_{in} , the information in the input population, which scales with the height of the noisy hill of activity. In this model, the linear Fisher information and the mean covariance are functions of J , W and I_{in} . For all realistic values of the input information, I_{in} , we find the following: 1. The linear Fisher information, denoted I_{out} , increases monotonically with J until the network becomes unstable. 2. When J is small, I_{out} is a single-peaked function of W . The position of the peak decreases as J increases, and for J large enough, it disappears altogether; in this regime (large J) I_{out} is a decreasing function of W . These results imply that maximum information transmission occurs when the structured connectivity is strong and the random connectivity weak. How do the correlations behave? They increase with J and decrease with W . Thus, because optimum information transmission occurs at high J and low W , at least for this model maximum information transmission occurs when the correlations are largest. References 1. Averbeck, Latham and Pouget. *Nature Reviews Neuroscience*, 7:358-366, 2006. 2. Roudi and Latham. *PLoS Computational Biology*, 3:1679-1700, 2007. 3. Barrett and Latham. *Frontiers in Systems Neuroscience*, doi:10.3389/conf.neuro.06.2009.03.123, 2009.

[doi:](#)

II-69. Positive reinforcement increases pooled population-coding efficacy in the auditory forebrain

James Jeanne^{1,2}
Tatyana Sharpee²
Timothy Gentner¹

JJEANNE@GMAIL.COM
SHARPEE@SALK.EDU
TGENTNER@UCSD.EDU

¹UC San Diego

²Salk Institute

To control adaptive behaviors, sensory neural circuitry must extract information from large numbers of diverse, physically complex signals in the natural world. Sensory experience helps subserve this function by evoking long-lasting changes that bias cortical circuits to represent behaviorally relevant signals. Distributing these representations across large populations of neurons further expands the coding capacity of this circuitry. However, the ability of sensory experience to modulate such population codes has not been explored. Here we test the hypothesis that experience can alter the coding properties of neural populations. Patterns of spikes across neurons in a population (a combinatorial code) can encode more information than total spike counts across that same population (a pooled code). Although dendritic apparatus may exist to accommodate a combinatorial code, the precise decoding strategies necessary to do so remain unclear. Pooled codes, in contrast, are parsimoniously accommodated by the integration properties of neural membranes. In this study, we investigate the effects of learning complex, natural birdsong stimuli on the coding efficacy of both pooled and combinatorial neural populations in an auditory cortical region of a songbird, the European Starling (*Sturnus vulgaris*). We trained starlings on a go/no-go operant song recognition task, anesthetized them with urethane, and recorded the activity of neurons in the caudolateral mesopallium (CLM) in response to training and novel songs. We partitioned the neural responses to songs into the responses to their constituent motifs (short stereotyped units of song) and computed the mutual information encoded by combinatorial and pooled codes about motif identity by pairs of neurons combined using the conditional independence approximation, as well as by single neurons. We find that, on average, learning in the context of positive reinforcement increases the percentage of information captured by the pooled code relative to that of the combinatorial code. Unlike with the combinatorial code, the increased efficacy of the pooled code comes in addition to significant learning-dependent increases in the coding capacities of single neurons. Thus, learning can modify population-coding properties in ways not attributable to coding modifications in single constituent neurons. The population-level plasticity observed here supports the biological plausibility of pooled population codes, and may help organisms meet the behavioral demands imposed by a dynamic sensory environment.

doi:

II-70. Online readout of frequency information in areas SI and SII

Adrien Wohrer¹
Ranulfo Romo²
Christian K. Machens¹

ADRIEN.WOHRER@ENS.FR
RROMO@IFC.UNAM.MX
CHRISTIAN.MACHENS@ENS.FR

¹Ecole Normale Supérieure

²Universidad Nacional Autónoma de México

How many neurons in a given area contribute to an animal's percept and behavior? Traditional experiments have suggested that a single sensory neuron can convey more information about a stimulus than an animal will use in a given task [1], raising questions about the usefulness of population-based codes. However, more recent experiments suggest that the predictive power of single neurons had been overestimated because their firing rates were computed over periods of time (~1-2 sec) much longer than what a monkey actually uses to make a decision (~200-300 msec) [2,3]. In fact, the number of neurons contributing to a code and the time scale of integration used by that code are in a natural trade-off relation. For example, the spike counts of N identical homogeneous Poisson neurons during a period of time T carry the same information as a single neuron's spikes

counted during a period NT. Here, we study quantitatively this trade-off, in somatosensory areas SI(3b), SI(1) and SII, during a two-frequency discrimination task. Instead of the traditional spike count code, we propose and test an "online readout" code, in which a sliding-window count of the population's spikes linearly provides an online estimate of stimulus value, which must be as temporally stable as possible. For each area and sliding window size w , we compute the neurometric index, i.e., the smallest variation of stimulus value which can be reliably detected from the resulting changes in the "online activity". We derive an analytical formula which allows us to predict this neurometric index as a function of the neurons' trial-averaged PSTH and cross-correlograms. Using this formula, we investigate the impact of spike-time correlations on the coding capacity of the population. Finally, we compare the neurometric indices to the monkey's psychometric index of behavioral performance. We find that: (1) Frequency discrimination from neural population activity requires readout windows of only a few tens of milliseconds to match the monkey's level of performance. (2) Admissible readout windows are markedly longer in area SI(3b) than in the higher-level areas SI(1) and SII. (3) These results still hold in the presence of various noise correlation structures that are consistent with our data. Our findings suggest the existence of a non-trivial integration of information from area SI(3b) to areas SI(1) and SII, sufficiently efficient for these areas to convey a reliable prediction of input frequency in their instantaneous population activity, computed over a few tens of milliseconds. The monkey's percept of frequency could therefore be formed in the first hundred milliseconds of neural activity (after the stimulus onset transient), with the rest of the stimulation period providing only minor changes and control. References: [1] Parker-Newsome (98), *Ann Rev Neurosc*: 2, 227-277 [2] Luna et al. (05), *Nature Neurosc*: 8, 1210-1219 [3] Cohen-Newsome (09), *J Neurosc*: 29, 6635-6648

doi:

II-71. One-dimensional dynamics of associative representations in lateral intraparietal (LIP) area

Jamie K. Fitzgerald¹

David Freedman²

Alessandra Fanini²

John Assad²

JKFITZG@FAS.HARVARD.EDU

DFREEDMAN@UCHICAGO.EDU

ALESSANDRA.FANINI@GMAIL.COM

JASSAD@HMS.HARVARD.EDU

¹Department of Neurobiology, Harvard Medical School

²Harvard Medical School

We previously showed that the lateral intraparietal area (LIP) flexibly encodes learned associations between multiple types of visual stimuli (Fitzgerald et al, Cosyne 2009). A recent recurrent network model predicts that across a population of LIP neurons, slowly varying patterns of neuronal activity-including spontaneous and delay activity during cognitively demanding tasks-are scaled versions of one another (Ganguli et al., 2008). We tested whether the LIP population activity during memory periods in shape-association and motion-categorization paradigms may be explained by a scaling of the spontaneous activity. Two monkeys performed a delayed paired association task in which they grouped six shapes into three pairs. On each trial a sample shape was presented (650 ms) followed by a delay (1500 ms) and a test shape; the monkeys released a lever if the test shape was that associated with the sample. For many neurons, the spike rate evoked by a particular shape was most similar to the activity elicited by that shape's learned associate (cells with an influence of pair during sample: 114/161, delay: 117/161, test: 77/161 neurons, ANOVA, $p < 0.05$). We previously showed that LIP neurons flexibly encode motion direction depending on how those directions are grouped or categorized (Freedman & Assad, 2006). We thus asked whether single neurons reflect associations for both motion and shape stimuli, and we found many neurons modulated by both types of associations (sample: 42/78, delay: 32/78, test 21/78, ANOVA, $p < 0.05$). For all six shape and six motion stimuli, average delay activity was correlated with spontaneous firing rates across the population of LIP neurons (correlation coefficient: shape: 0.65-0.73, motion: 0.69-0.86 monkey 1; shape: 0.59-0.73, motion: 0.60-0.68 monkey 2, $p < 0.01$), consistent with the predictions of Ganguli et al. In contrast, during the visual periods, the population activity showed markedly lower correlations (CC, shape: 0.34-0.45, motion: 0.28-0.34 monkey 1; shape: 0.36-0.44, motion: 0.28-0.33 monkey 2). The lower correlation coefficients during visual stimulation argue that the relationship between delay and spontaneous activity was not due to simple differences in

excitability among neurons. However, lower correlations during visual stimulation might be a consequence of the higher firing rates during visual stimulation. We thus examined responses from a previous experiment in which monkeys passively viewed moving stimuli that elicited a broad range of sustained responses (Fanini & Assad, 2009). We observed low correlation between these visual responses and spontaneous activity across the range of firing rates ($CC = 0.3-0.4$), arguing against the alternative excitability explanation. The one-dimensionality of delay activity also predicts that the order of selectivity for pairs of shapes or motion categories should be biased among neurons; this was indeed true for both experiments in at least one animal (chi-square, $p < 0.01$). The biases are not explained by differences in behavior among shape-pairs/motion-categories, such as performance, reaction times, and microsaccades. In conclusion, encoding of associations for multiple stimulus types during memory periods in LIP may arise from a scaling of the population spontaneous activity that is similar for associated stimuli and dissimilar for non-associated stimuli.

[doi:](#)

II-72. Modelling visual crowding of complex stimuli

Steven C. Dakin¹
Peter Bex²
John Greenwood³

S.DAKIN@UCL.AC.UK
PETER.BEX@SCHEPENS.HARVARD.EDU
JOHN.GREENWOOD@UCL.AC.UK

¹Institute of Ophthalmology, University College London

²Harvard Medical School

³University College London

We investigated how crowding - a breakdown in object recognition that occurs in the presence of nearby distracting clutter - works for complex letter-like stimuli. We first conducted a psychophysical experiment where observers reported the orientation (up/down/left/right) of a T target, abutted by a flanker composed of a randomly-positioned horizontal and vertical bar. In addition to retinotopic anisotropies (e.g. more crowding from more eccentric flankers) we report three object-centred anisotropies. First, errors included twice as many 90 degree as 180 degree target rotations than would be expected by chance. Second, flankers were twice as intrusive when they lay above or below (end-flanking) compared to left or right (side-flanking) of an upright T-target (an effect that holds under global rotation of the target-flanker pair). Third, errors induced by end-flankers resemble the flanker, but errors induced by side-flankers do not. We compared the predictions of several models of crowding to these results. In particular we describe a population coding model of spatial position - incorporating probabilistic averaging of position within contours - that can straightforwardly account for the range of psychophysical effects described. We conclude that crowding represents limits imposed by interference zones defined in retinotopic and object-centred space, and that population coding models, operating on stimulus attributes such as position and orientation, can explain the particular interactions that arise between features under crowding.

[doi:](#)

II-73. Decoding stimulus velocity from population responses in area MT of the macaque

Alan A. Stocker¹
Najib Majaj²
Chris Tailby³
J. Anthony Movshon⁴
Eero P. Simoncelli⁴

ASTOCKER@SAS.UPENN.EDU
MAJAJ@MIT.EDU
CTAILBY@GMAIL.COM
MOVSHON@NYU.EDU
EERO.SIMONCELLI@NYU.EDU

¹Department of Psychology, University of Pennsylvania

²MIT

³University of Melbourne

⁴New York University

The responses of neurons in area MT are thought to underlie the perception of visual motion in primates. However, recent studies indicate that the speed tuning of these neurons changes substantially as contrast is reduced (Pack et al., 2003; Krekelberg et al., 2006), in a way that seems inconsistent with the reduction in perceived speed seen psychophysically. To understand this apparent discrepancy, we recorded 59 MT neurons in anaesthetized macaques, and measured their responses to a broad-band compound grating stimulus presented at a broad range of velocities and contrasts. We presented the same stimuli to all neurons, adjusted only for receptive field location and preferred direction. As in previous studies with awake macaques, reducing contrast shifted the preferred velocity of most neurons toward slower speeds, as well as reducing response amplitude and tuning bandwidth. We constructed a population vector velocity decoder that operates on a neural population that includes the measured set of neurons, along with a "mirror" set tuned for the opposite direction. Using the synthetic population that represents both positive and negative velocities allows the decoder to capture the key characteristics of human velocity estimation and discrimination, including speed biases at low stimulus contrast. Specifically, we show that maintained discharge in such an MT population has an effect on the percept that is analogous to that of the slow speed prior characterized in Bayesian models of velocity perception (Stocker and Simoncelli, 2006). We also examined optimal linear decoders, and found that they produce nearly veridical percepts when operating on the full neural population, assuming that variability in the individual neuronal responses is statistically independent. Restricting these decoders to operate on a small set of model neurons whose response properties are obtained by averaging the tuning curves of similarly tuned neurons leads to qualitatively good matches to the perceptual data, but only when the decoder is optimized for stimuli drawn from naturalistic prior distributions over speed and contrast. This suggests that the response characteristics of the MT population are matched to the statistics of the natural world, in that linear decoding can approximate optimal Bayesian inference.

doi:

II-74. Neural correlates of dynamic sensory cue re-weighting in macaque area MSTd

Christopher R. Fetsch¹

Gregory C. DeAngelis²

Dora E. Angelaki¹

CRFETSCH@GMAIL.COM

GDEANGELIS@CVS.ROCHESTER.EDU

ANGELAKI@PCG.WUSTL.EDU

¹Washington University School of Medicine

²University of Rochester

Psychophysical studies have demonstrated that human observers integrate sensory information across modalities to improve perceptual sensitivity. These studies often frame multisensory cue integration as a problem in probabilistic (i.e., Bayesian) inference. In this framework, a working hypothesis is that the brain represents and combines probability distributions over stimuli, and thereby takes into account the inherent uncertainty in sensory information when making perceptual decisions. One key prediction from such optimal cue integration models is that subjects will re-weight cues according to their relative reliability (uncertainty) on a trial-by-trial basis, and indeed this has been shown in several paradigms. However, direct neurophysiological evidence for probabilistic computations during multisensory integration is scarce, and it remains unclear exactly how neuronal populations accomplish rapid re-weighting of cues based on their reliability. To address this question, we trained rhesus monkeys to perform a 2AFC fine discrimination of self-motion (heading) direction. Monkeys were seated on a motion platform facing a rear-projection screen, and on each trial were presented with a heading trajectory defined either by physical motion ('vestibular' condition), optic flow simulating observer motion ('visual' condition), or a combination of both cues ('combined' condition). Cue reliability was varied randomly across trials by changing the motion coherence of the optic flow pattern. As in previous studies, we generated optimal predictions for the cue weights by measuring performance in the single-cue conditions, then tested those predictions in the combined condition by placing the cues in conflict on a subset of trials. Our behavioral results suggest that monkeys, like humans, can dynamically re-weight cues according to their reliability. During the task, we recorded the activity of single

neurons in area MSTd, a region thought to contribute to multisensory integration for heading perception. Using ROC analysis, we quantified the behavior of an ideal observer performing the same task as the animal but using only the firing rate of an individual neuron. The majority of MSTd neurons in our sample showed near-optimal cue re-weighting with changes in reliability, similar to the monkey's behavior. We also constructed a decoding model in which a simulated observer performed the discrimination task based on MSTd population activity. On a given simulated trial, the population response (R) was generated by drawing from the individual neuron responses to a particular stimulus (s). From this response and the known tuning curves of the neurons in the sample, the model computed the likelihood $P(R|s)$ for each possible value of the stimulus, then took the maximum likelihood estimate as its choice on each trial. Even on the basis of relatively few neurons ($N = 28$), the simulated observer showed cue re-weighting that was remarkably similar to the behavior of the animal. Together with the single-neuron results, this suggests that MSTd activity implicitly encodes cue reliability on a trial-by-trial basis, and thus could contribute to the re-weighting observed behaviorally. More broadly, our results support the hypothesis that sensory populations encode the distributions used to mediate probabilistic inference, and that an explicit reliability signal is not required for optimal cue integration.

doi:

II-75. Optimal neuronal tuning curves - an exact Bayesian study of dynamic adaptivity

Steve Yaeli¹
Ron Meir²

STEVEY@TECHUNIX.TECHNION.AC.IL
RMEIR@EE.TECHNION.AC.IL

¹Department of Electrical Engineering, Technion

²Technion

Neural decoding can be fruitfully studied within an ecological framework whereby optimal performance is expected to adapt to environmental statistics. A powerful and mathematically rigorous framework can be set up within the theory of optimal point process filtering, whereby the exact posterior probability distributions of the hidden environmental states can be computed. Optimal estimators for different loss functions can be computed from these posterior distributions. For the widely used quadratic cost function, the issue of optimal neural encoding has been previously treated using information theoretic criteria and lower bounds on the minimum mean squared error (MMSE), most often based on Fisher information. Furthermore, in most cases only degenerate scenarios (e.g. fixed tuning curves, uniform population etc.) were considered. Here we employ the optimal decoding methods in a Bayesian setting in order to compute the exact MMSE, and investigate its dependence on the neuronal tuning curves (receptive fields), which constitute the encoding ensemble. Combining analytic results, obtained by using well-justified approximations to the full Bayesian solution, with numerical simulations of the exact solution, we characterize the optimal tuning curves in both space and time, and explore their relationship to the external environment. We provide precise predictions about the optimal adaptation of tuning curves to the statistical nature of the environment. Furthermore, we prove in several settings that drawing quantitative and qualitative conclusions about tuning curves based on bounds can be very misleading. In fact, the bound-based predictions may sometimes be diametrically opposed to those resulting from direct analysis of the true MMSE. This outcome is particularly significant due to the very prevalent use of approximations and bounds on the MMSE in studies of optimality. Specifically, we make the following predictions: (1) Tuning curves possess an optimal width that follows a simple relationship with decoding time and environmental statistics. (2) The optimal positions of tuning curves are related in a precisely quantifiable, yet non-trivial fashion, to their width and to the environmental statistics. (3) Tuning curves should dynamically reduce their width during the course of encoding a stimulus. (4) Narrower tuning curves are preferable in regions where stimuli are more likely. The final two predictions are consistent with empirical results, obtained in two independent physiological experiments in vivo. Our analysis explains the observed physiological phenomena in terms of performance optimization, and can thus provide insight into the possible driving mechanisms of biological sensory systems.

doi:

II-76. One rule to grow them all: A general theory of neuronal branching and its practical application

Hermann Cuntz^{1,2}
 Friedrich Forstner²
 Alexander Borst²
 Michael Häusser¹

H.CUNTZ@UCL.AC.UK
 FORSTNER@NEURO.MPG.DE
 BORST@NEURO.MPG.DE
 M.HAUSSER@UCL.AC.UK

¹University College London

²Max Planck Institute of Neurobiology

One rule to grow them all: A general theory of neuronal branching and its practical application Hermann Cuntz^{1,2}, Friedrich Forstner², Alexander Borst² and Michael Häusser¹ 1Wolfson Institute for Biomedical Research, University College London 2Max-Planck Institute of Neurobiology, Martinsried. Understanding the principles governing axonal and dendritic branching is essential for unravelling the functionality of single neurons and the way in which they connect. Nevertheless, no formalism has yet been described which can capture the general features of neuronal branching. Here we propose such a formalism, which is derived from the expression of dendritic arborizations as locally optimized graphs [1][2]. Inspired by Ramón y Cajal's laws of conservation of cytoplasm and conduction time in neural circuitry [3], we show that this graphical representation can be used to optimize these variables. This approach allows us to generate synthetic branching geometries which replicate morphological features of any tested neuron. The structure of a neuronal tree is captured by its spatial extent and by a single parameter, a balancing factor weighing the costs of conservation of cytoplasm and conduction time. This balancing factor allows a neuron to adjust its preferred electrotonic compartmentalization. In the context of the network, competitive growth and spatial tiling are governed by this same rule. Realistic large-scale synthetic networks of neurons can therefore be generated by simply defining their input and output organization. We have developed an open-source software package to implement these simulations, the "TREES toolbox", which provides a general set of tools for analyzing, manipulating and generating dendritic structure. The package includes a tool to generate artificial members of any particular cell group and a method for model-based supervised automatic morphological reconstruction of multiple cells from fluorescent image stacks. These approaches provide new insights into the constraints governing dendritic architectures. They also provide a novel framework for modelling and analyzing neuronal branching structures and for constructing realistic artificial neural networks. Acknowledgements This work was supported by the Gatsby Charitable Foundation, the Wellcome Trust, the Alexander von Humboldt Foundation, and the Max Planck Society. References [1] Cuntz H, Borst A, Segev I (2007) Optimization principles of dendritic structure. *Theor Biol Med Model* 4: 21. [2] Cuntz H, Forstner F, Haag J, Borst A (2008) The morphological identity of insect dendrites. *PLoS Comput Biol* 4: e1000251. [3] Ramón y Cajal S (1911) *Histologie du système nerveux de l'homme et des vertébrés*. Paris: Maloine.

[doi:](#)

II-77. Columnar transformation of neural response to time-varying sounds in auditory cortex

Poppy Crum
 Xiaoqin Wang
 Johns Hopkins School of Medicine

POPPY@JHMI.EDU
 XIAOQIN.WANG@JHU.EDU

The neural representation of time-varying signals is of particular importance to our understanding of complex and biologically important sounds such as speech and music. How sub-cortical and cortical regions differentially encode temporal patterns of acoustic stimuli offers insight to the crucial properties of the stimulus being represented as well as the functional role of the neural region. It is well known that the neural representation of time-varying signals in most sub-cortical regions is time-locked to stimulus features. In contrast, previous studies of awake primate auditory cortex have shown both neurons with stimulus-synchronized phasic responses as well as neu-

rons with non-synchronized tonic responses (Lu et al, 2001). Here, we show that in the awake marmoset monkey a central stage in the cortical transformation to a firing rate-based response occurs across the cortical laminae. Neurons in upper cortical laminae show both weaker synchronization than neurons in middle thalamo-recipient laminae (layer IIIb/IV), as well as a significant shift in the percentage of neurons responsive to time-varying sounds but having no clear time-locked response features. Moreover, when the percentage of cells in primary-auditory cortex showing non-synchronized, responses is compared with previous studies of lower and higher stages in the auditory pathway, the role of within column, cross-laminar transformations in forming or enhancing the non-synchronized response pattern is evident. Preliminary models are proposed of within-column shifts in inhibition and excitation to account for cross-laminar transformations from synchronized to non-synchronized response patterns. In our experiments, laminar location of single units was identified through registration with the current-source-density (CSD) measured within the same penetration. All measurements of the CSD were generated from local-field-potential (LFP) recordings made across depth and spanning the cortical laminae using the same tungsten electrode used to record single-unit responses. A shift to a firing rate-based representation of time varying signals in upper cortical laminae suggests a change in the stimulus features critical to encoding at successively higher processing stages in auditory cortex.

doi:

II-78. Is multisensory integration Hebbian? Ventriloquism aftereffect w/o simultaneous audiovisual stimuli

Daniel Pages
Jennifer M. Groh
Duke University

DSPAGES@GMAIL.COM
JMGROH@DUKE.EDU

Visual stimuli affect the perceived location of sounds. It has been assumed that the neural mechanism supporting visual recalibration of perceived sound location involves a simple Hebbian mechanism, where simultaneously presented auditory and visual stimuli excite a common population of neurons and 'wire' the auditory stimulus to a new location. However, an alternative possibility is that visual error after auditory localization could be used to 'update' auditory space via a feedback mechanism. Under this view, what you see after you make an eye movement to a sound would play a critical role in whether/how you adjust your sense of sound location. Previous studies of the effects of vision on sound localization have allowed for both possibilities, because visual and auditory stimuli have generally been presented simultaneously, potentially permitting Hebbian associations to form, and have also been left on long enough for visual feedback to be provided following any orienting movements to the sounds, permitting plasticity to be guided by visual reinforcement. Prism adaptation experiments such as those conducted in barn owls could involve either or both mechanisms. In the present study we seek to distinguish between these possibilities by introducing a ventriloquism aftereffect - a persistent shift in the perceived location of sounds following exposure to spatially mismatched visual and auditory stimuli - using tasks permitting only one of these mechanisms to operate. Specifically, the Hebbian task involved simultaneous but short-duration visual and auditory stimuli. The visual and auditory stimuli were both turned off prior to the completion of a saccadic eye movement to the sound. In contrast, in the feedback task, the visual and auditory stimuli were never on simultaneously. Rather, the sound played first and a visual stimulus was turned on during the saccade to the sound. We tested the impact of the exposure to these two types of mismatched visual-auditory trials on the accuracy of sound localization on interleaved auditory-only trials in monkeys. We found a robust shift in auditory localization in the feedback paradigm and not in the Hebbian paradigm. The average shift in the feedback paradigm was approximately 1.2 degrees, or 20% of the 6 degree separation between the visual and auditory stimuli. Our results indicate that a feedback signal is used for visually-guided auditory plasticity in the rhesus macaque, and that coincident stimuli are not necessary. More broadly, our results show that important and behaviorally relevant interactions between sensory modalities do not require the presence of stimuli that are coincident in time.

doi:

II-79. The kinetics of fast short-term depression are matched to spike train statistics to reduce noise

William Nesse¹
 Reza Khanbabaie¹
 Andre Longtin^{1,2}
 Leonard Maler¹

WNESSE@UOTTAWA.CA
 REZAKHAN@UOTTAWA.CA
 ALONGTIN@UOTTAWA.CA
 LMALER@UOTTAWA.CA

¹Department of Cellular and Molecular Medicine, University of Ottawa

²Department of Physics, University of Ottawa

Short-term depression (STD) is observed at many synapses of the central nervous system and is important for diverse computations. Using experimental and theoretical analysis, we have discovered an entirely novel form of fast STD (FSTD) in the synaptic responses of pyramidal cells evoked by stimulation of their electrosensory afferent fibres (P-units). We find the dynamics of the FSTD are matched to the distribution of interspike interval (ISI) statistics of natural P-unit discharge. Unlike standard STD, where the kinetics of depression are typically slower than the ISI timescale that activates STD, FSTD exhibits kinetics that are fast relative to the ISIs. This makes the magnitude of the evoked EPSPs depend only on the duration of the previous ISI, tracking changes in one ISI to the next effectively instantaneously. Furthermore, when FSTD level is plotted as a function of ISI, it forms a curve similar to the cumulative distribution function of the ISI afferent input. Thus, we have discovered a biological realization of an oft-studied theory first put forward by Simon Laughlin (1981), that information transmission is done most efficiently by weighting inputs according to their input statistical distributions. Our theoretical analyses suggest that the FSTD weighting mechanism induces a noise reduction to enhance weak sensory signals. This is distinct from previously ascribed functional roles for slower STD as a high-pass filter, gain modulation, or synchrony detection.

[doi:](#)

II-80. Frequency-invariant representation of interaural time differences

Hannes L  ling^{1,2}
 Ida Siveke¹
 Benedikt Grothe¹
 Christian Leibold¹

LUELING@BIO.LMU.DE
 SIVEKE@BIO.LMU.DE
 GROTHE@BIO.LMU.DE
 LEIBOLD@BIO.LMU.DE

¹Ludwig-Maximilians-Universit  t, M  nchen

²BCCN Munich

The difference in traveling times of a sound from its origin to the two ears is called the interaural time difference (ITD). ITDs are the main cue for low-frequency-sound localization. The frequency of the stimulus modulates the ITD sensitivity of the response rates of neurons in the brain stem. This modulation is generally characterized by two parameters: The characteristic phase (CP) and the characteristic delay (CD). The CD corresponds to a difference in the temporal delays from the ears to a respective coincidence detector neuron. The CP is an additional phase offset the nature of which is still under debate. The two above characteristic quantities hence describe the best ITD at which a neuron responds maximally via $(\text{best ITD}) = \text{CD} + \text{CP}/f$. Here f describes the frequency of the pure tone stimulus. We recorded neuronal firing rates in the dorsal nucleus of the lateral lemniscus of the mongolian gerbil for pure tone stimuli with varying ITD and frequency. Interestingly, we found that CPs and CDs are strongly negatively correlated. To understand the observed distribution of CPs and CDs among the recorded population, we assessed the mutual information of firing rate and ITD in terms of these two parameters. We therefore computed signal and noise entropies from rate distributions fitted to the experiments. Our results show that the information-optimal distribution of CPs and CDs exhibits a similar negative correlation as the one observed experimentally. Assuming similar rate statistics, we make hypotheses about how CDs and CPs should optimally be distributed for mammals with various head diameters. As expected, the mutual information increases with head diameter. Moreover, for increasing head diameter the two distinct subclusters of high mutual information

(peakers and troughers) fuse into one. To reveal correlations in the neural responses, we trained support vector machines (SVMs) and analyzed the resulting weights. We trained one SVM for every ITD interval. As input we used randomly drawn population rate vectors. We found that the error increases with number of different ITDs. For the behavioral localization acuity of 20 microseconds, we find a generalization error of about 0.6 percent. Moreover, ITDs are encoded by neural populations with strongly varying coding characteristics.

[doi:](#)

II-81. Behavioral context in pigeons: motor output and neural substrates

Kimberly McArthur¹
J. David Dickman²

KMCARTH@GMAIL.COM
DDICKMAN@WUSTL.EDU

¹Department of Anatomy & Neurobiology, Washington University School of Medicine

²Washington University School of Medicine

The vestibular system generates coordinated motor responses that stabilize head orientation, gaze, and posture during motion. These motor responses depend on the animal's behavioral state, and may be enhanced, attenuated, or gated in order to optimize performance. Currently, little is known about state-dependent modification of vestibular responses regulating head orientation and posture. However, vestibular nuclei neurons provide a likely substrate for response modification, as these cells receive multisensory and other processed signals that could provide information about the behavioral state. Pigeons are an excellent model system in which to study state-dependent vestibular processing, as they display distinct head and body responses to motion under different behavioral conditions. To study state-dependent vestibular processing, we used a hydraulic motion platform to rotate the pigeon's body in space along cardinal head axes, and we compared behavioral and neural responses recorded during two behavioral states: at rest and in simulated gliding flight. We recorded head-on-body responses from adult pigeons using the 3-field search coil method. We recorded tail-in-space responses using an Optotrak Certus system. To record from the brainstem vestibular nuclei, we implanted pigeons with chronic microdrives, each driving a bundle of ten fine wires used to record single-unit neuronal data. The results of our behavioral experiments indicate that pigeons are better able to stabilize their orientation in space during simulated flight than at rest. At rest, the pigeons' head-on-body responses keep the head stable in space only at higher stimulus velocities, but these responses maintain near-perfect head-in-space stability during flight. Similarly, pigeons display a tail-on-body response to passive pitch rotation that is only present during flight. This tail response is appropriate in direction to counteract passive body pitch encountered during flight, to maintain body-in-space stability. Preliminary neural recordings indicate that a subset of neurons in the vestibular nuclear complex have vestibular sensitivity that depends on the behavioral context. Some neurons are sensitive to rotation both at rest and in flight, but the depth of firing rate modulation during sinusoidal rotation varies with context. Other neurons are consistently responsive to rotational motion exclusively during flight, thereby exhibiting context-specific vestibular motion sensitivity. Thus, behavioral evidence indicates that vestibulospinal responses in the pigeon are context-dependent: the head response is enhanced during flight, and the tail response is specific to flight. Further, initial neural recordings point to the vestibular nuclei as a substrate for these context-dependent behaviors, containing neurons that vary their response to vestibular stimuli based on whether or not the pigeon is in flight. Additional studies are needed to address the degree to which these context-dependent neural responses in the brainstem contribute to specific contextual changes in behavior.

[doi:](#)

II-82. Systematic analyses of receptive field of mammalian olfactory glomeruli

Limei Ma

Stephen Gradwohl

Qiang Qiu

Richard Alexander

Winfried Wiegraebe

Ron Yu

Stowers Institute for Medical Research

LIM@STOWERS.ORG

SCG@STOWERS.ORG

QIQ@STOWERS.ORG

RLA@STOWERS.ORG

WIW@STOWERS.ORG

CRY@STOWERS.ORG

The mouse olfactory bulb consists of ~2000 discrete glomeruli; each receives input from a single type of olfactory neurons. The glomeruli form the basis set that encode odor information. We have generated transgenic mice that express the calcium sensor G-CaMP2 to examine glomerular response to large dimensional odor stimuli in the dorsal bulb. By exposing the mice to >200 odor stimuli that varies in odor identity and concentration, we construct three dimensional receptive fields for each glomerulus to describe the responses in terms of odorant identity, concentration and response temporal structure. Most glomeruli are tuned to a variety of chemicals with distinct structure features. Odor tuning of glomeruli does not generalize to odor class, molecular feature or odor concentration, but is specific to the combination of odor identity and concentration. The receptive field evolves quickly over time and the dynamics of individual glomeruli are distinct from each other. Overall, the glomerular set represents a distributive code for different odor classes such that odors with similar chemical structures can be distinguished. We show that glomerular activity in the olfactory bulb, not chemical features, are predictive of odor discrimination in behavioral tests, suggesting odor quality is largely encoded by the spatial pattern of glomerular activity. Moreover, we observe a correlation between the odor tuning similarity and the physical distance among glomeruli. Detailed examination shows that the glomeruli form a map that is arranged according the receptive field similarity instead of odor class or molecular features. Thus, the olfactory glomeruli form hierarchical clusters according to similarity in tuning properties.

[doi:](#)

II-83. Decoding intensity-tuned neurons in the auditory system

Elisha N. Marongelli¹

Paul V. Watkins²

Dennis L. Barbour²

ENM1@CEC.WUSTL.EDU

PVW1@CEC.WUSTL.EDU

DBARBOR@BIOMED.WUSTL.EDU

¹Department of Biomedical Engineering, Washington University in Saint Louis

²Washington University in Saint Louis

Neurons whose input-output functions are peaked rather than linear or sigmoidal have been described throughout the auditory system. At the level of the auditory cortex, these intensity-tuned or nonmonotonic neurons represent the majority of neurons. While their role in auditory processing remains uncertain, one enduring theory is that they are most useful collectively for encoding sounds in an intensity-invariant manner. This theory would predict that an optimal decoder making use of intensity-tuned inputs would be able to represent a sound across intensity more accurately than an optimal decoder using only sigmoidal inputs. This prediction has not previously been evaluated, however. We tested this prediction directly by constructing optimal linear estimators (OLEs) operating on a variety of auditory input-output (rate-intensity) functions. Altogether, we recorded the responses of 544 marmoset monkey primary auditory cortex neurons when stimulated by tones fixed at each neuron's characteristic frequency and varied in intensity. Multiple stimulus repetitions yielded estimates of the neuronal noise. This population represented a wide variety of input-output function shapes and was used to seed smaller subpopulations containing various percentages of intensity-tuned and untuned neurons. OLEs were then constructed for each of these subpopulations to generate either a constant discriminability across intensity ("intensity fidelity," a linear input-output relation) or a constant representation across intensity ("intensity invariance," a constant output for any input). Because intensity-tuned neurons exhibit unique adaptation characteristics, the neuronal subpopula-

tions were modeled either in silence-adapted conditions or under dynamic conditions where fairly loud sounds were common. Estimator performance was evaluated with the root mean square error (RMSE) between the target function and the decoded function. Generally, OLEs using silence-adapted subpopulations exhibited lower RMSEs than did OLEs using dynamically conditioned subpopulations. Furthermore, the intensity fidelity condition elicited lower RMSEs than did the intensity invariance condition. Additionally, the presence of some intensity-tuned neurons in a subpopulation improved its performance in intensity invariance tasks, regardless of whether the neurons were silence adapted or dynamically adapted; the optimal proportion in these conditions was approximately 50-70% intensity-tuned. For intensity fidelity tasks, a proportion of about 30% nonmonotonic neurons improved the encoding performance of dynamically adapted populations; however, improvements were limited to the lower intensity range, -5 to 35 dB SPL. Conversely, the addition of intensity-tuned neurons provided no advantage to silence-adapted populations encoding level fidelity and in fact became detrimental to performance when they comprised over 50% of the population. These trends were consistent across population sizes ranging from 20 to 1000 neurons. Under the conditions studied, neuronal subpopulations consisting of a mixture of intensity-tuned and untuned input-output functions proved to be optimal at stimulus representation and at relative proportions that are comparable to those observed physiologically. These results indicate that intensity-tuned neurons may not be critical to intensity-invariant sound representation by themselves, but when combined with intensity-untuned neurons, the result is more accurate coding under a variety of stimulus conditions.

doi:

II-84. The structure of human olfactory space

Alexei Koulakov¹
Armen Enikolopov²
Dmitry Rinberg³

AKULA@CSHL.EDU
AE2296@COLUMBIA.EDU
RINBERGD@JANELIA.HHMI.ORG

¹Cold Spring Harbor Laboratory

²Dept. of Biological Sci., Columbia University

³Janelia Farm, HHMI

Our understanding of the sense of smell is hindered by the lack of a well-defined perceptual space and knowledge of how this space is related to the properties of odorant molecules. Here we analyze the psychophysical responses of human observers to an ensemble of monomolecular odorants. Each odorant is characterized by a set of 146 perceptual descriptors obtained from a database of odor character profiles. Each odorant is therefore represented by a point in highly multidimensional sensory space. In this work we study the arrangement of odorants in this perceptual space. We argue that odorants densely sample a two-dimensional curved surface embedded in the multidimensional sensory space. This surface can account for more than a half of the variance of the psychophysical data. We also show that only 12% of experimental variance cannot be explained by curved surfaces of substantially small dimensionality (~10). We suggest that these curved manifolds represent the relevant spaces sampled by the human olfactory system, thereby providing surrogates for olfactory sensory space. For the case of 2D approximation, we relate the two parameters on the curved surface to the physico-chemical parameters of odorant molecules. We show that one of the dimensions is related to eigenvalues of molecules' connectivity matrix, while the other is correlated with measures of molecules' polarity. We discuss the behavioral significance of these findings.

doi:

II-85. A mechanism that governs a transition from coincidence detection to integration in olfactory network

Collins Assisi
Maksim Bazhenov

University of California, Riverside

COLLINS.A@UCR.EDU
MAKSIM.BAZHENOV@UCR.EDU

Neurons can act as integrators or coincidence detectors depending on the duration over which they sum incoming action potentials. The same neuron may act as an integrator under some circumstances and a coincidence detector under others. Using a realistic computational model of the insect olfactory system we demonstrate that a switch in operational mode may be achieved by a combination of environmental contingencies and network interactions. The first olfactory relay in the locust, the antennal lobe consists of a network of excitatory projection neurons(PN) and local inhibitory interneurons(LN). Input to the antennal lobe was simulated as firing from olfactory receptor neurons to a set of PNs. Different inputs represented different odors. Individual PNs were assumed to be maximally sensitive to one odor and not others. An increase in concentration was simulated by recruiting PNs with input less than that of the preferred PNs. Odor stimulation caused the network to oscillate, a consequence of the coordinated firing of PNs and LNs; the resulting distributed, coherent population response could be measured as an oscillating local field potential. Increasing odor concentration increased the coherence of spiking across PNs and proportionally larger amplitude field potential oscillations. The Kenyon cells (KC) of the mushroom body (MB) receive direct excitatory input from PNs and delayed feedforward inhibition from lateral horn interneurons (LHIs). Coactive PNs excite KCs and compete with inhibition from LHIs creating cyclic windows of time over which KCs are sensitive to direct excitatory PN input. The durations of these windows, and therefore the operational mode of KCs, could be adaptively regulated by the concentration of the odor. We could also artificially alter the size of the integration window by truncating PN spikes that occurred after a pre-specified phase of the local field potential oscillations. To determine the effects of the KC's operational mode on the systems ability to discriminate between odors we adopted the perspective of neurons one synapse downstream from the MB. These neurons were modeled to sum inputs from randomly selected sub-populations of KCs. The resulting responses were then evaluated as a function of the size of the KC integration window. We found that for high odor concentrations the ability of KCs to discriminate between odors improved over a very short window. Larger windows proved detrimental to system function by compromising the sparseness of odor representations in the MB. Hence, for high concentrations KCs acted as coincidence detectors, a mode of operation that maximized the distance between odor representations. Low concentrations blur the distinction between similar odors. The ability of the system to discriminate odors progressively improved as the window of integration increased. The lack of coherence in PN spikes eliminated the possibility of overwhelming MB circuits thus preserving sparseness even for large integration windows. KCs therefore acted as integrators of incoming PN input. Our results suggest that the same set of neurons may operate either as integrators or coincidence detectors depending on the stimulus conditions in a manner that allows optimal separation of odor representations.

[doi:](#)

II-86. Visual features evoke reliable bursts in the perigeniculate sector of the thalamic reticular nucleus

Vishal Vaingankar¹
Cristina Soto Sanchez¹
Xin Wang¹
Amarpreet Bains¹
Friedrich T. Sommer²
Judith Hirsch¹

VAINGANK@USC.EDU
SOTOSANC@USC.EDU
XINW@USC.EDU
ABAINS05@GMAIL.COM
FSOMMER@BERKELEY.EDU
JHIRSCH@USC.EDU

¹University of Southern California

²University of California, Berkeley

Relay cells in the cat's lateral geniculate nucleus of the thalamus receive feedback inhibition from the perigeniculate sector of the reticular formation. Both relay cells and reticular neurons fire tonic spikes in addition to stereotyped bursts that must be primed by prolonged hyperpolarization. Bursts fired by relay cells have been studied intensively. Typically, they are primed by stimuli of the non-preferred polarity that occupy the receptive field for long times and are triggered just as luminance contrast returns to the preferred sign [1-4]. For example, in the case of On center relay cells, bursts signal the recent removal of a persistent dark stimulus, and vice versa for OFF center cells. We wondered what, if any, visual features might drive reticular cells to burst. Since most cells in the perigeniculate are excited by stimuli of both contrasts, it seemed unlikely that the simple, alternating bright/dark patterns that evoke bursts from relay cells would be effective. To explore this issue we made extracellular recordings from reticular cells while displaying natural scene movies. The movies elicited patterns of tonic spikes interspersed with periods of inactivity followed by bursts. To estimate the features, or spatiotemporal receptive fields (STRFs), that preceded tonic or burst spikes, we used publicly available software (STRFPak <http://strfpak.berkeley.edu>). The STRFs constructed from tonic spikes peaked shortly before the neural response, whereas the STRFs constructed from burst spikes peaked farther back in time. This is consistent with work in the lateral geniculate that suggests that stimulus sequences associated with bursts prime rather than directly elicit firing [1-4]. It was surprising to find, however, that the spatial peaks of the tonic and burst STRFs were sometimes displaced rather than overlapping. Hence, the features that evoke tonic spikes can be different from those that prime bursts. These results led us to ask if certain types of stimuli evoke bursts more reliably than others and if visually driven bursts in the perigeniculate are produced as reliably as those in the lateral geniculate [1, 2]. To address these questions, we used the Fano factor to compare trial by trial variability in bursts evoked by natural movies versus Gaussian white noise. For both reticular and relay cells, movies evoked bursts with high, subpoisson, fidelity (Fano factor < 1) whereas the Fano factor for bursts evoked by noise was typically supra-poisson (Fano factor > 1). Collectively, our results suggest that reticular bursts reliably encode specific sensory features. Thus, intrinsic thalamic circuits are capable of sophisticated visual processing. References: 1. Allitto, H.J., T.G. Weyand, and W.M. Usrey. *J. Neurosci.*, 2005. 25: p. 514. 2. Denning, K.S. and P. Reinagel, *J. Neurosci.*, 2005. 25: p. 3531. 3. Lesica, N.A. and G.B. Stanley, *J. Neurosci.* 2004. 24: p. 10731. 4. Wang, X., et al., *Neuron*, 2007. 55: p. 465.

doi:

II-87. Decoding multiple objects from populations of macaque IT neurons with and without spatial attention

Ethan Meyers¹

Ying Zhang²

Sharat Chikkerur³

Narcisse Bichot²

Thomas Serre²

Tomaso Poggio²

Robert Desimone²

EMEYERS@MIT.EDU

YZHANG00@MIT.EDU

SHARAT@MIT.EDU

BICHOT@MIT.EDU

SERRE@MIT.EDU

TP@AI.MIT.EDU

DESIMONE@MIT.EDU

¹MIT

²McGovern Inst/Dept of Brain & Cog Sci, MIT

³McGovern Institute for Brain Research, MIT

Spatial attention improves visual processing of specific objects when they are embedded in complex cluttered scenes. Computational simulations examining the limits of biologically plausible feedforward hierarchical architectures have shown that these models can successfully discriminate between objects in low levels of clutter (2-3 objects). However, under more substantial clutter (> 3-5 objects), recognition performance decreases towards chance due to interference between the different object representations at higher levels of the visual hierarchy where neurons have larger receptive fields (Serre et al., 2005). These interference effects, which have also been seen in electrophysiology (e.g., Moran and Desimone, 1985; Missal et al., 1997; Zoccolan et al., 2007) and psychophysics (Serre et al. 2007), provide a computational motivation for why attention is needed to achieve

above-chance recognition in cluttered display conditions (Serre et al. 2007; Chikkerur et al. 2009). In this work we examine this conjecture by assessing how attention affects the amount of information carried by populations of neurons in anterior inferior temporal cortex (AIT) in multiple object displays. We recorded from neurons in AIT as monkeys engaged in a spatial attention task in which three objects were displayed simultaneously at different spatial locations. A cue was shown near fixation, which pointed to the behaviorally relevant object. The monkey was rewarded for making a saccade to the cued object when it changed its color, and ignoring color changes of the distracters. Additionally, there were trials that had the same temporal sequence but only a single isolated object was displayed. We used population decoding (Meyers et al., 2008) to analyze the neural data. Using pseudo-populations of neurons (72 neuron from one monkey, 95 from another), we trained a classifier to discriminate between the objects using data from a subset of the single object trials. We then tested the classifier's predictions of what object was shown in either a different subset of single objects trials, or on the cluttered display trials, before and after the attentional cue. This allowed us to compare the amount of information present in the population when an object is shown in isolation to the amount of information about the object in a cluttered display, with and without spatial attention. When multiple objects were displayed, the classifier could predict which objects were present at a level above chance, but the accuracy was greatly reduced compared to predictions made from data from single object trials. However, after spatial attention was deployed, the decoding accuracy for the attended stimulus greatly improved, while the decoding accuracy for nonattended stimuli decreased to near chance. These findings are consistent with computational models that show AIT is capable of supporting limited representations of multiple objects (as predicted by Serre, Kreiman et al., 2007), and that attention related changes play a significant role in increasing the amount of information available about behaviorally important objects.

doi:

II-88. Contrast dependent changes in Monkey V1 gamma frequency undermine its reliability in binding/control

Supratim Ray¹

John H. R. Maunsell²

SUPRATIM_RAY@HMS.HARVARD.EDU

JOHN_MAUNSELL@HMS.HARVARD.EDU

¹HHMI & Harvard medical school

²Harvard Medical School

Electrical signals recorded from the brain reveal oscillatory behavior of the neural population. Oscillations in a frequency band between 30 and 80 Hz, called the gamma band, have been suggested to play a functional role in cortical processing such as feature binding, forming dynamic communication channels across cortical areas, or providing a temporal framework for the firing of neurons so that information could be coded in the timing of spikes relative to the ongoing gamma cycle. These hypothesized functional roles require that the neuronal assemblies processing the features of the same stimulus oscillate at the same frequency. However, the frequency of the gamma rhythm depends on simple stimulus manipulations such as size, velocity, spatial frequency and cross-orientation suppression. For stimuli with features that vary in space and time, it remains unclear whether the induced gamma rhythms in different neural assemblies that process that stimulus are stable and reliable enough to support binding, communication or coding. We tested whether increasing the stimulus contrast, which increases the level of cortical excitation, affects the frequency of the gamma rhythm in the primary visual cortex (V1) of two awake behaving rhesus monkeys. Recordings were made from a chronic array of 96 electrodes (Blackrock Systems) implanted in V1 (right hemisphere). While the monkeys fixated within a 1° window and attended to a stimulus in the opposite hemifield, we presented a static Gabor stimulus at different contrasts on the receptive fields of the neurons recorded from the microelectrodes. We found that the peak frequency of the gamma rhythm increased monotonically with stimulus contrast, from ~38 Hz at 25%contrast to ~53 Hz at 100%contrast. Changes in stimulus contrast over time caused fast and reliable gamma frequency modulation. Further, a large Gabor stimulus, whose contrast varied across space, generated gamma rhythms at significantly different frequencies in the neuronal assemblies separated by as little as 0.2° (~400 μm in cortex). Gamma oscillation frequency decreased with increasing distance between the recording microelectrode and the stimulus center, which was well accounted for by the reduction in stimulus contrast with distance. These results suggest that gamma rhythms

are generated by highly localized networks that can rapidly track the incoming excitation. In addition to having a variable peak frequency, gamma rhythms were invariably weak, on average less than a few percent of the total signal power; far weaker than the stimulus-evoked transient during the first 100 ms. The weakness and varying center frequency of gamma oscillations suggest that they are poor candidates for either a control signal or information channel. Instead, our findings are consistent with the idea that gamma rhythms are a resonant phenomenon arising from the interaction between local excitation and inhibition. Several fundamental cortical mechanisms such as divisive normalization, adaptation and gain control rely on excitatory-inhibitory interactions; thus it is not surprising that gamma rhythm is also present and is modulated during a variety of cognitive tasks such as attention, working memory or cortico-spinal interactions. Thus, it could be an important neural signature of specific cortical processes.

[doi:](#)

II-89. Contrast suppression in human visual cortex

Brouwer Gijis Joost¹
David Heeger²

GBROUWER@CNS.NYU.EDU
DAVID.HEEGER@NYU.EDU

¹Center for Neural Science, New York University

²New York University

Background The "normalization model" encompasses a linear receptive field, soft-thresholding, and divisive suppression. It has been proposed to explain stimulus-evoked responses of neurons in various visual cortical areas including V1 and MT, multi-sensory integration in MST, representation of value in LIP, olfactory processing in *Drosophila* antennal lobe, and modulatory effects of attention on visual cortical neurons. In the present study, we used fMRI to measure cross-orientation suppression and tested the normalization model in human visual cortex. Specifically, we measured the activity in each of several channels (corresponding to subpopulations of neurons) with different orientation tunings, and fit these orientation-selective responses with the normalization model. **Methods** In a training experiment, subjects ($n=4$) viewed a series of full-contrast sinusoidal gratings of 6 possible orientations. Responses of each fMRI voxel ($2 \times 2 \times 2$ mm, 24 slices sampled every 1.5 s) were modeled as a weighted sum of the responses of 6 orientation-selective channels, each with an idealized orientation-tuning curve (or basis function) with 30° orientation tuning width (half-width at half-height). Linear regression was used to estimate a matrix of weights (6 per voxel) characterizing the transformation from channel responses to voxel responses. In the main experiment, subjects viewed a series of stimuli consisting of a vertical target grating at five different contrasts (3.125%, 6.25%, 12.5% and 50%), either in isolation or superimposed with a horizontal mask grating (50% contrast). The inverse of the weight matrix (estimated in the training experiment) was used to compute channel responses from the voxel responses, separately for each visual cortical area, to each stimulus condition (different target contrasts, mask present / absent). Retinotopic visual areas, including V1, were defined by measuring the radial- and polar-angle components of the cortical retinotopic map. **Results** For the V1 channel maximally tuned to the target orientation, responses increased with target contrast, but were suppressed when the horizontal mask was added (Fig. 1A). Cross-orientation suppression was evident as a shift in the contrast gain of the channel responses. For the channel maximally tuned to the mask orientation, a constant baseline response was evoked for all target contrasts when the mask was absent (Fig. 1D); Responses decreased with increasing target contrast when the mask was present. Channels tuned for intermediate orientations exhibited intermediate effects (Fig. 1B,C). The normalization model (with 4 free parameters) provided a good fit to contrast-response functions with and without the mask, simultaneously for all 6 orientation-selective channels (48 measurements). **Conclusions** The normalization model can explain cross-orientation suppression in human visual cortex, similar to reports with single-unit electrophysiology and local field potentials in macaque and cat. The approach adopted here can be applied broadly, by assuming a basis set for neural tuning curves, to measure simultaneously the responses of each of several subpopulations of neurons (channels) in the human brain that span a particular stimulus or feature space (e.g., orientation, direction, color), and to characterize interactions (including divisive suppression) between those neural subpopulations.

[doi:](#)

II-90. Quantifying the difficulty of object recognition tasks via scaling of accuracy vs. training set size

Steven Brumby¹
 Luis M Bettencourt^{1,2}
 Craig Rasmussen^{1,3}
 Ryan Bennett⁴
 Michael Ham¹
 Garrett Kenyon¹

BRUMBY@LANL.GOV
 LMBETTENCOURT@GMAIL.COM
 CRASMUSSEN@NEWMEXICOCONSORTIUM.ORG
 RAB0146@GMAIL.COM
 MIKEH@LANL.GOV
 GARKENYON@GMAIL.COM

¹Los Alamos National Laboratory

²Santa Fe Institute

³New Mexico Consortium

⁴University North Texas

Hierarchical models of primate visual cortex (e.g. neocognitron/HMAX) have been shown to perform as well or better than other computer vision approaches in object identification tasks. However, to date only small system sizes (in numbers of neurons and synapses) have been used, commensurate with the scale of visual training sets, containing typically hundreds of images or a few minutes of video (<1 gigapixel). A rough estimate translates the size of the human visual cortex, in terms of number of neurons and synapses, to ~1 petaflop of computation, while the scale of human visual experience greatly exceeds standard computer vision datasets: the retina delivers ~1 petapixel/year to the brain (few terapixel/day), driving learning at many levels of the cortical system. This disparity of scales raises the question of how system performance may increase with larger amounts of unsupervised and supervised learning and whether it can approach human accuracy given sufficient training. Here, we describe work at Los Alamos National Laboratory (LANL) to develop large-scale functional models of visual cortex on LANL's supercomputers. We present quantitative criteria for assessing when a set of learned local representations is complete, based on its statistical evolution with the size of unsupervised learning sets. We compare representations resulting from several learning rules to well-known prototypes in primary visual cortex and known constraints in higher layers. We then quantify the difficulty of different object recognition tasks via the improvement in classification performance with the size of the supervised training set. We show that classification performance rises in a general quantitative way with the number of examples in the training set, and that these scaling coefficients quantify the difficulty of the task. Specifically we find a universal form where $\text{accuracy} = a + b \log(N)$, where a , and b are constants that depend on the details of the system architecture and layer representations and N is the number of images in the training set. Testing is performed on a fixed set of images (not seen during training), and a fit to accuracy versus N gives a , and b , and the critical set size $N^* = \exp[(100-a)/b]$. Thus, more difficult tasks correspond to larger necessary training sets N^* . We compute N^* for different standard datasets and in a variety of circumstances by varying learning rules and number of layers in the model. For example, comparing the behavior of a classifier (SVM) trained using the standard fixed Gabor V1 and imprinted V2, versus Hebbian learned V1 and V2, we see that the fully learned model starts at a lower accuracy, but eventually matches performance of the standard model as the training set grows, predicting $N^*_{\text{imprinted}} = 3996 > N^*_{\text{learned}} = 3130$ in animal/no animal object identification tasks. Finally, we discuss how scaling of performance accuracy with N provides a path to systematic improvements in system performance, regardless of initial benchmarks for small datasets, and supplies a quantitative criterion for the inclusion of new ingredients in the model, for example dealing with feed-back and lateral connectivity.

doi:

II-91. Does the visual system use natural experience to construct size invariant object representations?

Nuo Li

James DiCarlo

McGovern Inst/Dept of Brain & Cog Sci, MIT

LINUO@MIT.EDU

DICARLO@MIT.EDU

Object recognition is challenging because each object produces myriad retinal images. Responses of neurons at the top of the ventral visual stream (inferior temporal cortex, IT) exhibit object selectivity that is unaffected by the image changes. How do IT neurons attain this tolerance ("invariance")? One powerful idea is that temporal contiguity of natural visual experience can instruct tolerance (e.g. Foldiak, Neural Computation, 1991): because objects remain present for many seconds, whereas object or viewer motion cause changes in each object's retinal image over shorter time intervals, the ventral stream could construct tolerance by learning to associate neuronal representations that occur closely in time. We recently found a neuronal signature of such learning in IT: temporally contiguous experience with different object images at different retinal positions can robustly reshape ("break") IT position tolerance, producing a tendency for IT neurons to confuse the identities of those temporally coupled objects across their manipulated positions (Li & DiCarlo, Science, 2008). A similar manipulation can induce the same pattern of confusion in the position tolerance of human object perception (Cox, Meier, Oertelt, DiCarlo. Nat Neurosci, 2005). Does this IT neuronal learning reflect a canonical unsupervised learning algorithm the ventral stream relies on to achieve tolerance to all types of image variation (e.g. object size and pose changes)? To begin to answer this question, we here extend our position tolerance paradigm to object size changes. Unsupervised non-human primates were exposed to an altered visual world in which we temporally coupled the experience of two object images of different sizes at each animal's center of gaze: (e.g.) a small image of one object (P, neuronally preferred object) was consistently followed by a large image of a second object (N), rendering the small image of P temporally contiguous with the large image of N. We made IT neuronal selectivity measurements before and after the animals received ~2 hours of experience in the unsupervised, altered visual world. Consistent with our results on position tolerance, we found that this size experience manipulation robustly reshapes IT size tolerance over a period of hours. Specifically, unlike experienced controls, we found a change in neuronal selectivity (P-N) across the manipulated objects and their manipulated sizes, producing a tendency to confuse those object identities across those sizes. This change in size tolerance is specific to the manipulated objects, grew gradually stronger with increasing experience, and the rate of learning was similar to position tolerance learning (~5 spikes/s per hour of exposure). Finally, in a separate experiment, we examine how temporal direction of the experience affects the learning: do temporally-early images teach temporally-later ones, or vice-versa? We found greater learning for the temporally-later images, suggesting a Hebbian-like learning mechanism (e.g. Sprekeler & Gerstner, COSYNE, 2009; Wallis & Rolls, Prog Neurobiol, 1997). We speculate that these converging results on IT position and size tolerance plasticity reflect an underlying unsupervised cortical learning mechanism by which the ventral visual stream acquires and maintains its tolerant object representations.

[doi:](#)

II-92. Stimulus timing dependent plasticity in high- and low-level vision

David B. T. McMahon

David A. Leopold

National Institute of Mental Health

MCMAHOND@MAIL.NIH.GOV

LEOPOLDD@MAIL.NIH.GOV

Object recognition in humans is extremely plastic, as evidenced by our ability to recognize and remember new faces after a few seconds of exposure. This form of visual learning is presumably mediated by experience dependent changes in synaptic efficacy among neurons in ventral visual cortical areas that are involved in face recognition. In contrast, perceptual acuity for low-level visual features tends to be more stable, and the small changes that can occur do so only after extensive training over long periods of time. At the cellular level, synaptic strength is sensitive to the precise relative timing of pre- and post-synaptic events. This phenomenon, known as

spike timing dependent plasticity, has been linked to changes in the stimulus selectivity of V1 neurons measured physiologically, and to shifts in perceived line orientation in human subjects measured psychophysically (Yao and Dan 2001). In the current study, we tested the idea that visual responses in high-level visual cortex, and thus object recognition, can be modified by stimulus timing dependent plasticity. We hypothesized that the plasticity induced by stimulus timing manipulations would be greater in a perceptual task that depends on high-level visual cortex than for a task that can be solved based on early visual processing. To test this idea, we measured psychometric functions of subjects performing two perceptual tasks that required either a high-level (face identity) or a low-level (line orientation) judgement. We first determined the stimulus value that corresponded to the point of perceptual equivalence for both stimulus types. We then induced a shift in this perceptual midpoint using a conditioning protocol designed to induce spike timing dependent plasticity in the connections between neurons in visual cortex. In experiments measuring face identity perception, the subjects were exposed to a block of conditioning stimulation that consisted of 100 pairings of two brief (10 ms) presentations of faces A and B. In experiments measuring perceived line orientation, an identical pairing protocol was employed using gratings with slightly different orientations. In both cases, pairings presented with a brief (20 ms) stimulus onset asynchrony shifted the perceptual midpoint to the left or right depending on the order of the stimulus pairing (AB or BA). The magnitude of this perceptual shift decreased substantially or was abolished altogether when the stimulus onset asynchrony was increased beyond 60 ms. When expressed as a proportion of the slope of the psychometric function for each stimulus type, a significantly greater shift was induced for perception of face identity than for line orientation. This result supports the idea that high-level visual areas are particularly susceptible to stimulus timing dependent plasticity.

[doi:](#)

II-93. In vivo Ca2+imaging of neural activity throughout the zebrafish brain during visual discrimination

Eva A. Naumann^{1,2}

Adam Kampff¹

Florian Engert¹

¹Harvard University

²MPI

EVA@MCB.HARVARD.EDU

KAMPFF@FAS.HARVARD.EDU

FLORIAN@MCB.HARVARD.EDU

How does the brain process and combine sensory information to produce an appropriate behavior? What are the neural circuits required to discriminate different sensory stimuli and enact a simple behavioral decision? In order to understand the flow of information from sensory inputs to motor outputs, the basic process underlying a decision, it is necessary to monitor the activity of neurons throughout the brain. The larval zebrafish, a fast developing, translucent, genetically accessible vertebrate, provides a system for optically investigating the role of neural populations involved in sensory discrimination and simple decision making. We developed a visual behavioral assay to test how unambiguous and ambiguous stimuli differentially affect directed turns. Unambiguous whole field stimuli evoke consistent behavioral responses and drive neural activity in a subset of spinal cord projection neurons of the hindbrain as well as midbrain neural populations. However, when presented with ambiguous (conflicting) motion stimuli, both the behavior and neural activity provide insight into how the visual input is translated into behavior. Monocular inward motion to either eye is sufficient to evoke directed turning behavior while monocular outward motion is not, however, a conflicting converging stimulus, inward motion to both eyes, produced no turning behavior at all. We could identify neuronal populations that show activity related to each stimulus condition, with single cell resolution, throughout the brain of transgenic zebrafish expressing the Ca2+-reporter GCaMP2, using in vivo two photon microscopy. From these results, we present a "working model" for the complete visual whole field motion discrimination circuit of the zebrafish brain.

[doi:](#)

II-94. Frequency dependence of the spatial spread of the local field potential

Dajun Xing
Chun-I Yeh
Samuel Burns
Robert Shapley

XDJ@CNS.NYU.EDU
CIY@CNS.NYU.EDU
SBURNS@CIMS.NYU.EDU
SHAPLEY@CNS.NYU.EDU

New York University, Center for Neural Sci

Very slow fluctuations and oscillations in brain signals ($<10\text{Hz}$) are indices of the cortical state and possibly of mental health. However the localization of low frequency signals has been very controversial because the local field potential (LFP) at low frequency is highly coherent between distant sites. The wide spread of low frequency components in the LFP could be due to either the physical volume conductance of brain tissue or to the connectivity of neural circuitry. To decide which plays a bigger role, we estimated the visual spread of LFP in Macaque primary visual cortex in different temporal frequency bands by sparse-noise mapping experiments. A small black or white square (~ 0.2 deg) was randomly flashed for 40 ms on a grey background at different visual locations on the screen. The stimulus evoked local field potential (seLFP) to each visual location was calculated by cross-correlation between stimulus location and recorded local field potential. Then the visual spreads of the LFP at different temporal frequencies were estimated by Fourier-transforming the seLFP into different frequency bands. We predicted the visual spread of seLFP in different frequency bands for a model of the cortical tissue's impedance as a capacitive low pass filter. The model's parameters were based on the space constant of the decay of LFP coherence at different frequencies. Opposite to what is predicted by the low-pass-filter model, we found that visual spreads at different frequencies in the LFP increased with frequency from 1Hz to 30Hz, by 10%. Our result suggests that mechanisms for signal propagation in the cortex at low frequency behave as a very shallow high pass filter rather than as a low pass filter. We conclude that neural circuitry determines the propagation of signals in the cerebral cortex and that the electrical impedance of cortical tissue can be neglected for signal propagation. **ACKNOWLEDGEMENTS** This work was supported by grants from the US National Science Foundation (grant 0745253) and the US National Institutes of Health (T32 EY-07158 and R01 EY-01472) and by fellowships from the Swartz Foundation and the Robert Leet and Clara Guthrie Patterson Trust.

[doi:](#)

II-95. Network dynamic regime controls the structure of the V1 extra-classical receptive field

Daniel B. Rubin¹
Kenneth D. Miller²

DBR2107@COLUMBIA.EDU
KEN@NEUROTHEORY.COLUMBIA.EDU

¹Department of Neuroscience, Columbia University

²Columbia University

Recent work has shown that the V1 extra-classical receptive field (eCRF or "surround"), the region that surrounds the classical receptive field (CRF or "center") and exerts a subthreshold influence over the response to stimuli, has a well-defined spatial structure (Tanaka and Ohzawa 2009). To map out the spatial structure of the CRF and eCRF, these authors presented a large grating of the CRF's preferred spatial frequency (SF) and orientation (ORI), and superimposed a periodic contrast modulation. They studied the dependence of the cell's response to the modulation SF and ORI, and found that the combined CRF and eCRF typically resembled a large simple cell, with a preferred modulation ORI unrelated to the preferred luminance ORI of the CRF and a preferred modulation SF about 2-6 times lower than the preferred luminance SF of the CRF. We have previously shown (Ozeki et al. 2009) that surround suppression in V1 is a "de-amplification" rather than an inhibition - a lowering of the gain in a network with strong, destabilizing recurrent excitation stabilized by strong feedback inhibition (we term this an "inhibition-stabilized network" or ISN). Here we show that such a network, along with spatial connectivity typical of cortex, creates inhibitory activity that is periodic over retinotopic space. This network "resonates" over a narrow range of spatial frequencies, resulting in the observed tuning for modulation SF. We study a linear firing-rate-model

of V1 comprised of recurrently connected excitatory and inhibitory neuronal populations positioned on an array of retinotopic space. As in cortical networks, inhibitory connections are localized while excitatory connections are longer range. Using analytic and numeric modeling, we find that in order for inhibitory firing rates in a generic network with non-periodic connectivity to oscillate as a function of space, several requirements must be met. In particular, the network must be an ISN. Only in the ISN regime are both excitatory and inhibitory neurons tuned for a nonzero modulation SF of a contrast-modulated grating. The preferred modulation frequency is controlled by the spatial extents of the different synaptic connections ($E \rightarrow E$, $E \rightarrow I$, $I \rightarrow E$, $I \rightarrow I$). Using estimates of the extent of the lateral connections in layer II/III of cat V1, we observe preferred SF's matching closely those recorded experimentally; we also show more generally that the preferred modulation SF will be several times lower than the CRF preferred luminance SF, as observed. Lastly, once the nonlinearity of cortical networks is considered, dynamic regime (ISN or non-ISN) becomes an activity-dependent feature that changes with input strength; different contrasts drive the network into different dynamic regimes. This input-driven shift in network regime may underlie a number of contrast-dependent properties of the eCRF. In particular, this mechanism can explain the reduction in summation field size and the emergence of spatially-periodic length-tuning curves of firing rate and inhibitory conductance (Anderson et al. 2001) with increasing stimulus contrast.

doi:

II-96. Modulation of speed-sensitivity of a motion-sensitive neuron during walking

M. Eugenia Chiappe
Johannes D. Seelig
Michael B. Reiser
Vivek Jayaraman

Janelia Farm Research Campus, HHMI

CHIAPPEE@JANELIA.HHMI.ORG
SEELIGJ@JANELIA.HHMI.ORG
REISERM@JANELIA.HHMI.ORG
VIVEK@JANELIA.HHMI.ORG

Motion-sensitive neurons of the fly visual system respond selectively to specific features of the optic flow fields elicited when the animal moves through its natural environment. Flies use optic flow to guide their locomotion, and the activity of neurons in the motion pathway ultimately instructs motor commands used for course control [1]. However, motion in the extrinsic world can also cause global retinal image flow (consider, for example, a fly standing on a leaf on a windy day). This creates a potential confound between ego-motion and motion originating externally that must be resolved by incorporating feedback from other sensory and motor pathways. To examine this idea, we explored whether neurons involved in a visual-motor circuit modulate their dynamic range depending on the animal's own movement. We expressed the genetically encoded calcium indicator GCaMP3 in a subset of neurons of the lobula plate of the fruit fly, and used two-photon imaging to record from the so-called Horizontal System Equatorial (HSE) neuron. This neuron, which receives directionally selective retinotopic input from as yet unidentified motion detecting neurons, responds preferentially to motion along the horizontal axis. During the recording, tethered flies were allowed to walk on an air-suspended ball and stimulated with global optic flow fields delivered using an LED arena. We monitored the ball's rotation at high resolution and used this readout as a proxy for the tethered fly's own walking activity. We compared calcium responses of the HSE neuron to global optic flow when the fly was standing to those recorded when the fly was walking. We found that HSE neurons show significantly greater responses during walking. Moreover, peak calcium responses increased in direct correlation with the tethered fly's walking speed. Interestingly, this correlation was more evident when neurons were stimulated at higher temporal frequencies (stimuli moving at higher speeds). When the fly was standing, the temporal frequency tuning curve of the HSE neuron peaked at 1 Hz; during walking periods, the neuron responses increased in amplitude and peaked at higher temporal frequencies. Finally, the modulation of the neuron's response due to walking activity showed temporal frequency sensitivity: response gains were higher for higher temporal frequencies. By recording the activity of neurons in behaving animals, we can evaluate the performance of a circuit under more naturalistic circumstances, i.e. including sensory and mechanical feedback. Here, we show that motion-sensitive neurons quickly adjust their bandwidth depending on the behavioral state of the animal. Increasing a system's sensitivity during behavior is beneficial from an energy minimization standpoint.

The temporal-frequency- and walking-speed-dependent gain modulation we observe may additionally be used to accommodate the increased optic flow caused by a freely moving animal's own movements. We are now investigating possible mechanisms for this adaptive behavior of the circuitry, including the possibility that it is mediated by octopamine [2]. [1] Eggelhaaf et al (2002) TINS, 25: 96-102. [2] Longden and Krapp (2009). J.Neurophysiol (doi: 10.1152/jn.00395.2009).

doi:

II-97. Cross-correlation analysis reveals circuits and mechanisms underlying direction selectivity.

Pamela M. Baker
Wyeth Bair

PB2001@GMAIL.COM
WYETH.BAIR@DPAG.OX.AC.UK

Department of Physiology Anatomy and Genetics, Univ. of Oxford

Cross-correlation of spike trains from pairs of neurons has proved valuable for investigating functional connectivity in the visual system. It has revealed the specific functional connections between the lateral geniculate nucleus (LGN) and simple cells in primary visual cortex (V1) [1], and exposed properties of simple cells that drive complex cells [2]. Direction selectivity is a fundamental physiological property that arises from V1 circuitry yet basic questions of how direction selective (DS) receptive fields are constructed remain unanswered. To guide experiments for uncovering cortical mechanisms of direction selectivity, we implemented a set of plausible network models of DS neurons. We demonstrate how cross-correlation analysis can reveal the circuitry and mechanisms generating DS cells, including the nature of the DS nonlinearity and the stage in the visual pathway where the time delay, inherent to motion detection, originates. Our models consisted of LGN cells, V1 simple and complex DS cells that were spiking, conductance-based integrate and fire units with physiologically-realistic intrinsic currents connected by AMPA and GABA-A synapses. We tested model architectures where DS neurons are built directly from LGN inputs or indirectly via orientation-tuned simple cells. In some models, the time delay was implemented via temporal diversity in LGN inputs, in others it was ascribed to dendritic mechanisms in the DS neuron. Both have received recent attention experimentally [3,4]. We also varied whether the DS interaction was facilitatory or suppressive. We computed shuffle-corrected cross-correlograms (CCGs) of spike trains from pairs of units that are accessible to extracellular recording in vivo. We tested the models with visual stimuli employed experimentally, including sinusoidal gratings and flashed dots. We found that cross-correlation is well-suited to determine where the DS time delay arises. In models where temporally-offset inputs synapse onto the DS unit, the peak in the CCG between any non-DS input and the DS cell occurred at the time lag typical of monosynaptically connected cells. In models where the delay was generated in dendritic subunits within the DS cell, the CCG peak shifted systematically to reflect the postsynaptic delays. Cross-correlation revealed the nature of the DS nonlinearity, with facilitatory and suppressive DS mechanisms generating distinct shapes in CCGs. We also identified key factors that determine whether connections between non-DS and DS cells are revealed in CCGs. For example, suppressive mechanisms may only become apparent in CCGs if stimuli that drive the cell in the anti-preferred direction are used. This stimulus dependence reflects the selectivity of the DS mechanism, not the input cells themselves, and is thus independent of network architecture. We explored how the magnitude of peaks and dips in the CCG change as strength and number of inputs to the DS cell vary. We address how these factors create challenges for employing this analysis in electrophysiological recording. Overall, our results suggest that CCGs from simultaneously recorded pre- and post-synaptic non-DS and DS cells, respectively, are necessary and likely sufficient for solving long-standing problems of circuitry and mechanisms underlying cortical direction selectivity.

doi:

II-98. The influence of pulvinar activity on corticocortical communication

Corey M. Ziemba¹
George R. Mangun²
W. Martin Usrey¹

CMZIEMBA@UCDAVIS.EDU
MANGUN@UCDAVIS.EDU
WMUSREY@UCDAVIS.EDU

¹Center for Neuroscience, UC Davis

²Center for Mind and Brain, UC Davis

While first-order thalamic nuclei such as the lateral geniculate nucleus (LGN) serve to transmit sensory information from the periphery to the cerebral cortex, the function of higher-order nuclei such as the pulvinar remains unclear. The pulvinar nucleus is a large and complex structure that receives input from and provides output to multiple cortical areas. Thus, the pulvinar is strategically positioned to influence neuronal communication between cortical areas, and accordingly, it has been suggested that the pulvinar is crucial for corticocortical information transmission. Here, we test this idea through simultaneous recording of visually-evoked activity in areas 17 and 18 of the anesthetized cat before and during inactivation of the pulvinar. Our results show that pulvinar inactivation has little influence on visually-evoked potentials in area 17. In contrast, the same measures in area 18 show substantial effects on late components in the evoked potential. Pulvinar inactivation also results in frequency-dependent changes in the coherence between multi-unit activity and local-field potentials in areas 17 and 18. These results demonstrate that neurons in the pulvinar nucleus make an important contribution to visual processing in higher cortical areas and support a role for the pulvinar in transmitting information between cortical areas, consistent with the notion that higher-order thalamic nuclei are important for cortico-cortical communication.

[doi:](#)

II-99. Image classification with complex cell neural networks

James Bergstra
Yoshua Bengio
Pascal Lamblin
Guillaume Desjardins
Jérôme Louradour
University of Montreal

JAMES.BERGSTRA@GMAIL.COM
BENGIOY@IRO.UMONTREAL.CA
LAMBLINP@IRO.UMONTREAL.CA
DESJAGUI@IRO.UMONTREAL.CA
JEROMELOURADOUR@GMAIL.COM

Simulations of cortical computation have often focused on networks built from simplified neuron models similar to rate models hypothesized for V1 simple cells. However, physiological research has revealed that V1 cells span a spectrum of behaviour encompassing the roles predicted by traditional simple and complex cell models. Our work looks at some of the repercussions of these modeling choices for learning. What functions can these richer V1 models learn easily that simple V1 models cannot? Are richer V1 models better at learning to categorize objects? In a first step, we implemented a few rate models of V1 neurons: affine-sigmoidal simple cell response, an energy-model, the model of [1], and the canonical circuit of [2]. We randomly initialized the filters in these models and trained a classifier on the V1 representation to categorize images from MNIST, and NORB as well as a synthetic dataset of shapes. Although we used cross-validation to permit each model family as much capacity as it could use, we found that in all cases the complex-cell models generalized significantly better from the limited number of examples used for training. In a second step, we experimented with a slow-features approach to initializing the filters instead of random initialization. This initialization yielded model neurons with additional robustness to edge position in images, and lead to still better generalization. In pursuit of this result, we introduced an approximation of the learning criterion in [3] with a computational cost that is linear (rather than quadratic) in the size of our model V1, and constant (rather than linear) in the number of training examples. In a third step, we experimented with multi-layer convolutional networks built from V1 cell model building blocks. Again, the networks with complex-like response functions outperform simple-cell-like ones. This sort of network achieved the best published test-set performance on MNIST among purely supervised algorithms. Recent complex cell rate models support an improved ability to learn to categorize objects relative to traditional simple-cell models,

without incurring extra computational expense. Additional experiments with hybrid models obtained by mixing and matching elements from the V1 models revealed two important mathematical ingredients for significantly better performance over a basic affine-sigmoid response function. Firstly, a sum of at least two squared linear filters brings a significant gain in statistical efficiency of learning. Secondly, a pointwise non-linearity that saturates polynomially rather than exponentially towards its asymptotic limit is significantly better in the majority of cases. These results underscore the importance of accurate modeling in efforts to understand learning in the visual system. References 1. Rust NC, Schwartz O, Movshon JA, Simoncelli EP (2005) Spatiotemporal elements of macaque V1 receptive fields. *Neuron* 46: 945-956. 2. Kouh, M. and T. Poggio (2008) A Canonical Neural Circuit for Cortical Nonlinear Operations. *Neural Computation* 20(6): 1427-1451. 3. Körding, KP, Kayser, C., Einhäuser, W. and König, P. (2004) How are complex cell properties adapted to the statistics of natural scenes? *Journal of Neurophysiology* 91(1):206-212.

doi:

II-100. Visual cortex unplugged: Neural recordings from rats in the wild

Agrochao Agrochao¹

Tobi Szuts¹

Vitaliy Fadeyev²

Wladyslaw Dabrowski³

Alan Litke²

Markus Meister¹

AGROCHAO@FAS.HARVARD.EDU

SZUTS@FAS.HARVARD.EDU

VF@SCIIPP.UCSC.EDU

W.DABROWSKI@FTJ.AGH.EDU.PL

ALAN.LITKE@CERN.CH

MEISTER@FAS.HARVARD.EDU

¹Harvard University

²University of California Santa Cruz

³AGH University of Science and Technology

Virtually everything we know about the function of neural systems comes from laboratory experiments in which the animal subject is either constrained in its behavior, anesthetized, or dissected into slices. The many insights gained through this process of abstraction must at some point be tested against biological reality. Do these modes of neural signaling encountered in the laboratory extend to the natural behaviors that originally motivated the research? For this purpose it will be useful to experiment on animals that are fully immersed in their natural environment, moving freely, engaging actively in multiple tasks, and generally interacting with sun, wind, and dirt. Monitoring neural signals in the wild has now become practical. We developed a wireless multi-channel recording system that acquires extracellular spikes from 64 electrodes. The device can be carried by a rat and transmits over a range of 100 m. Following our interests in visual processing, we implanted rats with multiple movable tetrodes in visual cortex. Typically, 6-8 single neurons could be recorded simultaneously. Animals participated in laboratory-bound experiments, both awake and anesthetized, to measure receptive fields. In addition, we present here initial observations during free-ranging behavior. The rat was placed in a fenced field adjacent to woods, about 10 m on a side, that included open areas, ground cover, leaf litter, and solid walls. The animal's behavior was monitored continuously by a video camera. Rats explored the enclosure, with a preference for the rim, but including the open areas. They frequently engaged in burrowing or nest building. The animals remained alert throughout the outdoor sessions, quite unlike laboratory behavior, which was commonly interrupted by bouts of sleep. In one outdoor session, neural activity was recorded from 4 visual cortex cells simultaneously on the same tetrode. Their spike trains had the following characteristics: Sparseness: The average firing rates ranged from about 1 to 10 Hz. Neurons had a tendency towards bursting, and the activity was relatively sparse. Using time bins of 100 ms, we compared the mean and variance of the spike count, n , to compute a measure of sparseness, $S = 1 - \langle r \rangle^2 / \langle r^2 \rangle$, which ranged from 0.61 to 0.95. Modulation: The activity of visual cortex neurons was strongly modulated during different behavior episodes. For example, during burrowing or nest building, when the animal's eyes were below the ground surface, these neurons ceased firing entirely for several seconds. Concerted firing: We observed substantial correlations in firing even among neurons found on different tetrodes. Pairwise correlation functions showed peaks or dips, typically within about 10 ms of zero delay, and with magnitudes around $\pm 20\%$. In one case, the central peak showed a 300% excess of synchronous spikes. These early results point to a rich set of phenomena to be explored further using insights from earlier work. For example the framework of efficient

coding theory makes predictions for the correlations among visual neurons based on their receptive fields and the structure of natural scenes. In ongoing work we are testing these predictions by combining the free-range recordings with controlled measurements of visual receptive fields.

doi:

II-101. Developing a rodent model of selective auditory attention

Chris Rodgers
Sarah Kochik
Vuong D. Vu
Michael R. DeWeese
UC Berkeley

QQRODGERS@GMAIL.COM
SKOCHIK@BERKELEY.EDU
VUONGVU@BERKELEY.EDU
DEWEESE@BERKELEY.EDU

How does the mammalian brain solve the cocktail party problem? The exigencies of life require an animal to extract the most behaviorally relevant auditory stream from a suite of simultaneous distracters. For example, a rat may have to choose whether to pay attention to the call of its mate or to the cries of its offspring. We are using a two-alternative choice paradigm to develop a rodent model of "selective auditory attention": the ability to volitionally select and process target sounds rather than distracter sounds. Crucially, the same sound can serve as target in one trial and as distracter in another trial, allowing us to control for its acoustic properties and to isolate the effect of attention. We model the selection process by training rats to respond to one sound in a mixture during one block of trials, and then, in the next block, to respond to the other sound in the same mixture. In another task, we find that rats respond more accurately to low-probability stimuli, a signature of bottom-up attention. We then increased task difficulty in a bid to titrate top-down attention against bottom-up. Preliminary results indicate that the bottom-up effect was abolished with this paradigm. In conclusion, these behavioral tasks allow us to model several components of attention: bottom-up saliency filtering, top-down control, and target selection. We next plan to record from primary auditory cortex in behaving rodents to find neural correlates of these behavioral results.

doi:

III-1. Synaptic filtering of natural spike trains in central synapses: A computational study

Umasankar Kandaswamy¹
Charles Stevens²
Vitaly Klyachko³

USK@WUSTL.EDU
STEVENS@SALK.EDU
KLYACHKO@WUSTL.EDU

¹Washington University in St Louis

²The Salk Institute

³Dept. of Cell Biology and Physiology, Washington University

Short-term synaptic plasticity (STP) is believed to function as a key neuronal mechanism of information analysis in neural circuits. Using natural patterns of neural activity recorded in behaving rodents we have recently shown that STP in excitatory hippocampal synapses serves as a high-pass filter optimized for transmission of information-carrying place-field discharges (Klyachko and Stevens, PLoS, 2006). This STP filter enables synapses to perform a highly non-linear, switch-like operation to permit the passage and amplification of signals with the place-field-like characteristics. The non-linearity of the synaptic filter is characterized by the rapid transition from the basal to a high-gain state at about 7 Hz, and by its near frequency-independence above the transition frequency in the range of 7-90 Hz. The existing models of synaptic plasticity are not able to reproduce such non-linear synaptic behavior. As a result, the synaptic mechanisms underlying this filtering paradigm remain poorly understood thus preventing incorporation of these mechanisms into larger-scale models of neural networks. Here we describe an extensive

model of STP that explains this highly non-linear synaptic behavior. Unlike most of existing models of synaptic plasticity, our current model is derived in a large part from basic principles of synaptic function. It is formulated in terms of release probability of a synapse and takes into account single-vesicle release hypothesis, calcium cooperativity of neurotransmitter release, and readily-releasable vesicle pool dynamics. In this framework, the two major forms of STP that increase synaptic strength, namely facilitation and augmentation, are formulated as transient decrease in the energy barrier for vesicle fusion leading to increased probability of release. Short-term synaptic depression is formulated as a transient reduction in the number of readily-releasable vesicles. Our computational analysis shows that this model accurately reproduces all major characteristics of the observed synaptic filtering paradigm. Our analysis suggests that facilitation is primary responsible for the rapid switching between the basal and high-gain states, while a dynamic balance between augmentation and depletion enables frequency-independence of the filter in a wide frequency range. Furthermore, the model predicts strong calcium-dependence of the filtering properties, including the transition efficacy and the amplitude of synaptic gain changes. We verified these predictions experimentally in hippocampal slices and demonstrate that indeed, the transition between the synaptic gain levels slows down an order of magnitude, while the amplitude of synaptic gain increases several times as extracellular calcium concentration decreases from 2 mM to 0.6 mM. In summary, our model of STP adequately reproduces the non-linear synaptic filtering paradigm observed during the processing of the natural spike trains. This model further proposes verifiable predictions of synaptic behavior and thus offers a useful framework to further investigate the role of STP in the information processing at the synaptic and circuit levels.

doi:

III-2. When somatic firing undermines dendritic compartmentalization

Bardia F. Behabadi

Bartlett W. Mel

BEHABADI@USC.EDU

MEL@USC.EDU

Biomedical Engineering Department, University of Southern California

Neocortical pyramidal neurons (PNs) exchange a dense plexus of interconnections, and most of these synapses are formed directly onto basal dendrites (BDs). Our working model describing the "arithmetic" of synaptic integration in these dendrites has been a 2-layer network (Archie & Mel, 2000; Poirazi et al. 2003; Polsky et al. 2004): the first layer consists of a set of separately thresholded dendritic subunits with sigmoidal input-output curves; the second layer calculates a sum of all the dendritic outflows reaching the soma, and then applies an F-I curve representing the cell's axo-somatic spike-generating mechanism. A hallmark of 2-layer models is functional compartmentalization: within-branch summation is subject to an additional layer of (dendritic) nonlinearity in comparison to between-branch summation. In an augmentation to the basic 2-layer model, we recently found evidence in both models and experiments that flexible modulation of dendritic integration, including both threshold-lowering and gain-boosting effects, and combinations thereof, can be achieved by targeting modulatory synapses non-uniformly along the proximal-distal axis of a basal dendrite (Behabadi, Polsky, Schiller & Mel, SFN 2007). This novel mechanism for flexible modulation could allow interactions between distinct pathways (vertical, horizontal, callosal, etc.) impinging on PN basal dendrites to be tailored to the modulatory needs of the specific cortical circuit. Experiments in brain slices support the augmented 2-layer model, but so far only in the subthreshold range where responses are reported as somatic depolarization (Behabadi, Polsky, Sciller & Mel SFN 2007). It remains unknown what effect the axo-somatic spike-generating mechanism may exert on dendritic integration in PN basal dendrites, though modeling studies suggest the effects can be powerful, potentially overwhelming compartment-specific dendritic nonlinearities. In this work we separate the effects of location-dependent within-branch and between-branch interactions using further compartmental modeling studies. We find that the somatic spike-generating mechanism, when allowed to have a high threshold (as in typical slice conditions), exerts a global nonlinearity that does indeed occlude within-branch nonlinearities - thus undermining the compartmentalized functioning of the dendritic tree. However, when the axo-somatic nonlinearity is neutralized by biasing the soma to near its spike threshold using background synaptic inputs, a nearly pure location-dependent within-branch nonlinearity is revealed, while between-branch nonlinearities are virtually eliminated. Our findings lend further "moral" support to the augmented 2-layer model, and point to the unexpected role of low background firing

rates in promoting compartmentalized dendritic processing.

[doi:](#)

III-3. The effect of inhibition on pyramidal cells: A conceptual model

Monika Jadi¹
Alon Polsky²
Jackie Schiller²
Bartlett W. Mel¹

JADI@USC.EDU
ALONPOL@TECHUNIX.TECHNION.AC.IL
JACKIE@TX.TECHNION.AC.IL
MEL@USC.EDU

¹Biomedical Engineering Department, University of Southern California

²Technion Medical School

Cortical computations are critically dependent on the interactions between excitatory and inhibitory influences at the single neuron level. The problem is complex given the sheer variety of inhibitory neurons (Burkhalter, 2008 and Markram 2004), but key features of excitation and inhibition can be identified that are mostly likely to influence the "arithmetic" governing E-I interactions. In this work, we have developed a conceptual framework to help unify and explain a variety of E-I effects. The framework consists of a composition of three simple biophysical functions relating to: (1) how the excitatory stimulus maps into actual delivered conductance (e.g. depressing vs. facilitating synapses) (2) the nature and location of the inhibitory inputs (e.g. somatic vs. dendritic); and (3) whether the firing threshold is affected by variability in excitation and or inhibition. Based in part on new data from slice experiments exploring the effects of inhibitory location, we show how the framework unifies and helps to explain a wide variety of seemingly diverse E-I interactions. Our framework can be extended to include the effects of inhibitory reversal potential and kinetics, and axonal inhibition.

[doi:](#)

III-4. Map dynamics of rhythmically perturbed neurons

Engelbrecht Jan¹
Kristen Loncich²
Renato Mirollo²
Michael Hasselmo³
Motoharu Yoshida³

JAN@BC.EDU
LONCICH@BC.EDU
MIROLLO@BC.EDU
HASSELMO@GMAIL.COM
MOTOHARU.YOSHIDA@GMAIL.COM

¹Physics Department

²Boston College

³Boston University

We explore patterns in the spiking of neurons receiving small perturbing rhythms with an emphasis on universal characteristics realized in both models and in vitro whole-cell recordings. Our starting point is a single-compartment model that fires periodically when stimulated with a constant current. Upon introducing a small perturbing current the dynamics of the voltage and gating variables rapidly collapse onto a 2-dimensional subspace in which the neuron's state subsequently evolves much more slowly and where successive spike times are related by a 1-dimensional map $t_{n+1} = F(t_n)$. Iterating the map yields the dynamical evolution of the spike phase (with respect to the rhythm). We then add noise to the model, and show that the spike phase evolution is characterized by a probability distribution relating consecutive phases $P(\phi_{n+1}, \phi_n)$ whose features follow from the underlying return map. We contrast the ISI and phase distributions with $P(\phi_{n+1}, \phi_n)$. Noise in the stimulus current can knock the model neuron out of entrainment, leading to a repeating cascade of spike phases as the neuron relaxes back to a periodic state. We then report on whole-cell recordings of pyramidal cells in region CA1 of rat hippocampal slices where we first inject a constant current to induce a steady firing rate and then add an additional sinewave oscillation to perturb the spike times. We show that successive phases again

evolve according to a map - strikingly similar to the model data. For appropriate currents the cell gets repeatedly knocked out of entrainment and subsequently relaxes back to entrainment through a reliable pattern of phase advances. This work demonstrates that the notion of a perturbing rhythm introducing slow dynamics characterized by a 1-dimensional return map generalizes from a single-compartment model to CA1 pyramidal cell slice recordings. Moreover, when noise frustrates entrainment, the phase dynamics can be characterized as a reproducible cascade of phase advances.

[doi:](#)

III-5. How do temporal stimulus correlations influence the performance of population codes?

Kamiar Rahnama Rad
Liam Paninski
Columbia University

KR2248@COLUMBIA.EDU
LIAM@STAT.COLUMBIA.EDU

It has long been argued that many key questions in neuroscience can best be posed in information-theoretic terms; the efficient coding hypothesis discussed by Attnave, Barlow, Atick, et al represents perhaps the best-known example. Answering these questions quantitatively requires us to compute the Shannon information rate of neural channels, whether numerically using experimental data or analytically in mathematical models. The non-linearity and non-Gaussianity of neural responses has complicated these calculations, particularly in the case of stimulus distributions with temporal dynamics and nontrivial correlation structure. In this work we extend a method proposed in [1] that allows us to compute the information rate analytically in some cases. In our approach the stimulus is modeled as a temporally correlated stationary process. Analytical results are available in both the high and low signal-to-noise (SNR) regimes: the former corresponds to the case in which a large population of neurons responds strongly to the stimulus, while the latter implies that the available neurons are only weakly tuned to the stimulus properties, or equivalently that the stimulus magnitude is relatively small. In intermediate SNR regimes, good numerical approximations to the information rate are available using efficient forward-backward decoding methods and Laplace approximations [2,3]. We find that when the number of neurons increases unboundedly, the mutual information for temporarily correlated stimuli has a simple form depending only on the Fisher information, paralleling the classical connection between Fisher and Shannon information for static stimuli [4]. For finite-size populations we are able to calculate the performance difference between a decoder based solely on temporally instantaneous observations and one which integrates temporal observations; interestingly, in both the high- and low-SNR regimes, a simple decoder which only includes temporally local neural responses is able to optimally extract information about the stimuli. References: [1] Barbieri, R. et al. (2004). Dynamic analyses of information encoding in neural ensembles. *Neural Computation*. [2] Paninski et al. (2009). A new look at state-space models for neural data. *J. Comput. Neurosci*. [3] Pillow et al. (2009). Model-based decoding, information estimation, and change-point detection in multi-neuron spike trains. Under review. [4] Brunel, N. and Nadal, J.-P. (1998). Mutual information, Fisher information, and population coding. *Neural Computation*.

[doi:](#)

III-6. The dynamic routing model of visuospatial attention

Bruce Bobier
Terry Stewart
Chris Eliasmith
University of Waterloo

BBOBIER@ENGMAIL.UWATERLOO.CA
TERRY.STEWART@GMAIL.COM
CELIASMITH@UWATERLOO.CA

The dynamic control of information routing in the brain has been most extensively studied in the context of vi-

sual attention, although the underlying mechanisms employed by the involved neural populations remain unclear. Several models have suggested that selective routing of attended information may be performed by modulating the gain of stimulus representations or by modifying synaptic connection weights. A problem with many such models is that they do not address how these processes are performed at the cellular level, or do so in a biologically implausible manner. Here we present the Dynamic Routing Model (DRM), which frames the problem of visuospatial attention as the dynamic routing information through a hierarchy of topographically organized cortical regions to an object-centered reference frame. In the DRM, the inferior and lateral pulvinar are hypothesized to encode the retinotopic location at which attention should be engaged, and project this signal to cortical control columns in each visual area. Each control column serves to guide information routing for a small group of adjacent visual columns. At each layer, the pulvinar signal is transformed by the cortical control columns to encode the location from which information should be represented. The dendrites of a neuron in a given column encode a vector containing the control signal and its inputs, and the synaptic weights of the dendrites collectively serve to approximate a function that modulates the gain of each input depending on its distance from the location specified by the cortical control signal. We begin by considering the attentional routing performed in a five layer hierarchy, where each layer is reciprocally connected with pulvinar. For each visual column, we can determine the interactions between the control and visual inputs that are required to optimally route information from the focus of attention to an object-centered reference frame at the top layer. After defining this interaction, we next investigate whether this function can be represented within the dendrites. The routing mechanism is the same for all nodes in the hierarchy, and by using the Neural Engineering Framework to model a subset of this large-scale network, we can use this high level mathematical description to analytically derive the optimal connection weights to perform the routing function. Simulation results show that in a model composed of a single cortical control population and nine visual columns, each containing 200 LIF neurons that project to a single output population, the neurons are able to accurately route visual information. We then show that the model accounts for numerous findings from neurophysiological and psychophysical studies, including attentional effects occurring earlier and more strongly in higher cortical areas, dynamic receptive field shifts, center-surround suppression, as well as impairments observed in pulvinar lesion patients. The DRM also generates several predictions concerning the character of the suppressive annulus surrounding the focus of attention, the role of pulvinar in engaging attention, and the computations performed by cortical neurons under attention.

doi:

III-7. Mechanisms recruited by attending to the time or frequency of sounds

Santiago Jaramillo

Anthony M. Zador

Cold Spring Harbor Laboratory

JARA@CSHL.EDU

ZADOR@CSHL.EDU

Valid expectations provide an advantage on behavioral performance. On tasks involving auditory stimuli, expectations can be set for different characteristics of the sounds such as spatial source, frequency, time, or other more complex features. These expectations are the basis of the behavioral and physiological effects observed in selective attention experiments. Here we have designed an experiment to test in rodents whether the mechanisms underlying auditory attention to time and frequency are distinct, by assessing whether the same neuronal population undergo attentional modulation when these stimulus characteristics are manipulated within a single animal. The stimulus in our task consists of a frequency modulated target immersed in trains of pure tones. The spatial source of the target indicates the location of reward. Temporal expectation is manipulated by presenting the target early or late within the trial with different probabilities, while expectation in frequency is controlled by the probability of the target's carrier frequency. We have first characterized the behavioral effects and changes in single cell responses when manipulating temporal expectation. Our results indicate that subjects gained an advantage in both speed and accuracy at detecting auditory targets immersed in distractors when valid expectations were available. Furthermore, responses from auditory cortex neurons suggest that the modulation of activity due to temporal expectation has similar characteristics to the changes observed during selective attention tasks in other sensory modalities. We are currently evaluating the effects of attending to frequency on behavioral sessions where the time and frequency of the target sounds are manipulated simultaneously. A comparison of the

modulation of responses for each manipulation will indicate if the advantage conveyed by valid expectations about different features of the acoustic stimulus is generated by the same neuronal mechanisms in the auditory cortex.

doi:

III-8. 252-site subdural LFP recordings in monkey reveal large-scale effects of selection attention.

Conrado A. Bosman¹

Thilo Womelsdorf²

Robert Oostenveld³

Birthe Rubehn⁴

Peter de Weerd⁵

Thomas Stieglitz⁴

Pascal Fries³

C.BOSMAN@DONDERS.RU.NL

THIWOM@IMAGING.ROBARTS.CA

ROBERT.OOSTENVELD@FCDONDERS.RU.NL

BIRTHE.RUBEHN@IMTEK.UNI-FREIBURG.DE

P.DEWEERD@PSYCHOLOGY.UNIMAAS.NL

THOMAS.STIEGLITZ@IMTEK.UNI-FREIBURG.DE

P.FRIES@DONDERS.RU.NL

¹Donders Centre

²University of Western Ontario, Canada

³Radboud University Nijmegen

⁴University of Freiburg, Germany

⁵University of Maastricht, The Netherlands

An essential mechanism during visual processing is selective attention. Selective attention lends a competitive bias to behaviorally relevant stimuli at the expense of irrelevant distracters. While information about distracters is filtered out, information about targets is routed through to be processed in depth [1]. This flexible routing of information might be subserved by a flexible pattern of synchronization among the involved brain areas [2]. In the awake monkey, several studies have demonstrated that attention enhances gamma-band synchronization inside V4 and few studies showed attentional effects on synchronization between pairs of areas. However, the pattern of synchronization among multiple areas simultaneously and their attentional modulation has not been revealed. To this aim, we trained a macaque monkey (*Macaca Mulatta*) to perform a change detection paradigm. While the monkey kept fixation, two stimuli (4 deg diameter sinusoidal gratings, drifting unidirectionally, 4 deg; of eccentricity) appeared. One of the gratings was cued to be the target stimulus. The monkey was rewarded if it reported unpredictable transient changes in the target's orientation. When reliable performance was achieved, the monkey was subdurally implanted with a flexible micromachined 252-channel ECoG (electrocorticogram) array, over large parts of the left hemisphere [3]. We obtained chronic recordings spanning from area V1, through V4, parietal and central regions up to FEF, during task performance. We quantified local and interareal synchronization during attention by obtaining frequency-resolved LFP power and coherence. We found attentional modulation of rhythmic synchronization across multiple areas. Locally, attention caused an increase in gamma power in areas V1 and V4, and a decrease in beta power over dorsal parietal and FEF electrodes. FEF was involved in two synchronization networks: 1.) a low-gamma (35 Hz) synchronization with dorsal premotor cortex that was reduced by attention, and 2.) a beta1 (18 Hz) synchronization with dorsal parietal cortex that was strongly enhanced by attention. Similarly, area V4 was involved in two synchronization networks: 1.) a beta1 (18 Hz) synchronization with dorsal parietal cortex that was not affected by attention, and 2.) a gamma (6090 Hz) synchronization with V1 and V2 that was strongly enhanced by attention. Furthermore, we found a strong attentional enhancement of beta-1 synchronization between FEF and area V1. These results further advance the notion that selective attention modulates interareal interactions by modulating interareal synchronization. Acknowledgments This research was supported by the European Community's Seventh Framework Programme (FP7/2007-2013), Grant Agreement "BrainSynch" HEALTH-F2-2008-200728 (P.F.), The Volkswagen Foundation Grant I/79876 (P.F.), and the European Science Foundation European Young Investigator Award Program (P.F.). Bibliography 1. Reynolds JH, Chelazzi L (2004) Attentional modulation of visual processing. *Annu Rev Neurosci* 27: 611-647. 2. Womelsdorf T, Fries P (2007) The role of neuronal synchronization in selective attention. *Curr Opin Neurol* 17: 154-160. 3. Rubehn B, Bosman CA, Oostenveld R, Fries P, Stieglitz T (2009) A MEMS-based flexible multichannel ECoG-electrode array. *J Neural Eng* 6: 036003.

doi:

III-9. Dynamical control of eye movements in an active visual search task: theory and experiments

Huang He
Joseph Schilz
Angela J Yu

CRANE081@GMAIL.COM
JOSEPH.SCHILZ@GMAIL.COM
AJYU@UCSD.EDU

Department of Cognitive Science, University of California, San Diego

The dynamics of cognitive control of eye movements are poorly understood. While there have been some recent analyses of "where" people fixate in a visual search task, there is a comparative paucity of understanding of "when" people saccade from one location to another. Here, we propose an optimality framework, using stochastic control theory, to examine both questions. We cast the visual search task as a sequential, dynamic decision-making problem, in which sensory processing and motor control, as well as their dynamic interplay, are optimized with respect to behavioral objectives such as speed and accuracy. We also present human behavioral data from a novel visual search task, in which we carefully control the level of stimulus noise and the relative costs of inaccuracy and search delay. We manipulate key computational parameters in the experiment, such as the spatial distribution of reward and information, which are the key determinants of the trade-off between exploitation and exploration, respectively. We measure subjects' sequential choices of fixation locations as well as fixation duration, and demonstrate that, relative to the optimality model, subjects efficiently learn the statistics of the environment and utilize them to optimally control their eye movements to achieve quantitatively specified task objectives. We conclude from this study that the active component of vision is highly optimized and adaptive with respect to environmental statistics and task demands, and that stochastic control theory provides valuable tools for the formulation of an optimality framework for sequential decision-making problems such as active sensory processing.

doi:

III-10. Decision-making dynamics and behavior of a parietal-prefrontal loop model

David Andrieux
Xiao-Jing Wang

DAVID.ANDRIEUX@YALE.EDU
XJWANG@YALE.EDU

Yale University School of Medicine

Single-units studies with behaving monkeys have revealed that the posterior parietal cortex area LIP and the prefrontal cortex (PFC) are critically involved in cognitive functions such as working memory and decision making, and similar neural signals are often found in these areas. Computational models propose that the observed activity patterns, such as persistent or ramping activity, are generated by a strongly recurrent (multistable attractor) network, but it remains unclear whether such neural signals originate within each area or arise from the reciprocal loop interactions. It is also poorly understood what the differential circuit properties and functions of these regions are. To explore these questions, we considered a model of mutually interacting parietal and prefrontal modules. Each module contains excitatory populations selective for different choice options that compete with each other through a local inhibitory pool. By varying the strength of local recurrent excitation, each module can operate in different regimes, i.e. exhibiting attractor dynamics or not. The modules are connected via long-range excitatory projections to both excitatory (E) and inhibitory (I) populations, determining the strength and balance of E and I connections. This framework allows us to explore the different patterns of local and long-range connections in a systematic way, and to characterize the resulting networks behaviors. We find new emergent behaviors when two

modules are reciprocally connected. First, when none of the modules individually behaves as an attractor network, the long-range interactions (providing recurrent excitation between the two modules) can give rise to slow ramping activity underlying time integration and self-sustained activity for working memory storage. Second, we find that coupling between the two modules can enlarge their domain of bistability. Thus, a local area can be shifted in and out of the attractor regime by a change in the effective inter-areal connectivity, providing a mechanism to control and gate the dynamical states of the different areas. Third, the two modules can reach different choices (one module selects option A, whereas the other module selects B). Such conflict states are observed when both modules are strongly recurrent and the long-range projections are biased toward targeting inhibitory cells, in hard decisions with noisy sensory information—a situation possibly relevant in complex decision tasks. These results reveal that novel computational behaviors can result from synaptic interactions between different areas in the brain. Importantly, the global behavior is not predetermined by the individual states (attractors or not) of the two local circuits in isolation. Furthermore, we envisage that different global functional states can be achieved depending on task-specific control inputs to the different modules. The coupling between LIP and PFC thus provide them with rich gating mechanisms. In this way the system can show flexible and adaptive behavior in accordance with the task at hand.

doi:

III-11. Time-varying gain modulation on neural circuit dynamics and performance in perceptual decisions

Ritwik K. Niyogi^{1,2}
KongFatt Wong-Lin^{3,4}

RITWIK.NIYOGI@GATSBY.UCL.AC.UK
KFWONG@MATH.PRINCETON.EDU

¹Gatsby Computational Neuroscience Unit, UCL

²CSBMB, Princeton University

³Program in Applied & Computational Mathematics

⁴Princeton Neuroscience Institute, Princeton U

Recent studies have shown that urgency signals may play an important role during the temporal integration of sensory evidence in perceptual decision-making. Furthermore, it has been shown that neurons in the lateral intraparietal area (LIP) can exhibit temporal integration properties in perceptual decision-making tasks, as well as multiplicative gain modulation properties. Our theoretical work connects this growing body of research areas. We investigate how time-varying single-cell gain modulation affects the network dynamics of a biological neural circuit model (Wong, Huk, Shadlen and Wang, 2007) for motion-discrimination tasks. Since we are interested in two-alternative forced-choice tasks, we implement a three-population network model. Two populations of excitatory pyramidal cells are selective to the coherent motion directions in a random dot motion-discrimination task. A third population of interneurons mediates pooled inhibition. We use dynamical systems analyses to study how the network dynamics evolves over different epochs of a trial. In order to enable categorical choice, the two selective populations need to remain at a common stable high firing-rate during the target-period prior to the motion stimulus onset. In reaction time (RT) tasks, the firing-rates typically diverge at a higher firing-rate shortly after motion stimulus onset (e.g. Roitman and Shadlen, 2002; Huk and Shadlen, 2005). Our analyses reveal that increasing the gains of both excitatory and inhibitory neurons is necessary to permit such neuronal dynamics to occur; an unstable saddle steady-state (that causes the divergence) has to be higher than the stable firing-rate during the target-period. In addition, the network is required to operate near a critical bifurcation point. Simulating our model, we are able to strikingly capture the full characteristics of the LIP neuronal dynamics in motion-discrimination tasks more accurately than past models, while accurately reproducing the behavioral data. Our model is also able to reproduce the timecourse of neuronal firing-rates in the fixed-viewing duration (FD) version of this task (e.g. Shadlen and Newsome 2001; Roitman and Shadlen 2002). It has been observed that in the FD task, the firing-rates of the selective populations are lower, and diverge from a level lower than during the target-period, contrary to that in the RT version. Firing-rates are also typically lower than the decision threshold during the delay-period in FD task. These phenomena can be accommodated in the same model by having lower gains than in the RT task. This approach is consistent with the intuition that decision-making that aims to

maximize reward-rate (total number of correct responses over total time taken), by optimizing a speed-accuracy tradeoff in a RT task, would be more demanding and therefore involve more attentional resources than decision-making in a FD task, where the reward is dependent only on accuracy. In fact, our simulations predict that a short time-constant of gain modulation, e.g. indicating fast attentional modulation, enables the maximum reward rate to be attained. Our model thus provides an integrative and coherent understanding of the interplay among separate neuronal processes to enable flexible and optimal decision performance.

doi:

III-12. An optimality framework for understanding the psychology and neurobiology of inhibitory control

Pradeep Shenoy¹

Rajesh P. N. P. N. Rao¹

Angela J. Yu²

PSHENYOY@CS.WASHINGTON.EDU

RAO@CS.WASHINGTON.EDU

AJYU@UCSD.EDU

¹University of Washington

²Department of Cognitive Science, UCSD

Inhibitory control, defined as the ability to withhold or modify actions that may no longer be appropriate, is a critical aspect of human behavior. Deficits in behavioral inhibition have been implicated in a variety of psychiatric conditions such as ADHD and addiction. The classic stop signal or countermanding task, where an initial movement command is occasionally countered by a "stop" signal instructing response inhibition, has been used to study many aspects of behavioral and neural processing that may relate to a more general notion of inhibitory control. We propose an optimality based framework for explaining behavior and neural mechanism in the stop signal task. Two characteristic behavioral outcomes are the lower reaction times on trials where response inhibition fails, and higher rates of inhibition failure with later stop signals. The prevalent theoretical model (the race model; Logan & Cowan, 1984) provides a mechanistic description of this behavior as being the outcome of a race between two independent processes, but does not directly explain behavior as a consequence of task demands and experimental parameters. The race model, therefore, is silent on how behaviour should change as a consequence of experimental manipulations such as long-term and short-term statistics of the relative frequency of stop signals (Emeric et al., 2007). We overcome these limitations with the application of two principled computational tools. First, we use Bayesian probability theory and a model of subjective sensory processing to explicitly compute and track the state and history of the environment. Second, we define an objective function that directly encodes experimentally enforced constraints such as reward, punishment and deadlines. Together, this formal framework enables us to concretely test the hypothesis that subjects' behavior in this task is optimal in terms of maximizing the objective function given the estimated state of the environment. We show that our model captures classic behavioral results in the countermanding paradigm, by comparing it to monkey behavioral data (Hanes & Schall, 1995). We also successfully model aspects of behavior unexplained by the race model, such as the effect of the fraction of stop signal trials and immediate trial history on reaction times and success at inhibition (Emeric et al., 2007). Thus, on a behavioral level, our framework motivates and generalizes the assumed features of the race model, while providing insight into the computational and behavioral import of empirical measures such as the stop signal reaction time (SSRT). Various neural populations, including the frontal eye field and superior colliculus, have been implicated in countermanding (Hanes et al., 1998, Pare & Hanes, 2003). Neurons in these regions show differential activity after a delay when comparing movement and successfully withheld movements. The race model offers a mechanistic connection between behavior and such activity, since the SSRT approximates the delay before the neural activities diverge. We show that computational variables in our model show a very similar temporal profile to neural activity in the FEF. Our optimality model therefore provides a unified framework for understanding behavior, optimal computation and neural processing underlying inhibitory control.

doi:

III-13. Discounting as task termination, and its implications

Schrater Paul¹
Constantin A. Rothkopf²

SCHRATER@CS.UMN.EDU
ROTHKOPF@FIAS.UNI-FRANKFURT.DE

¹University of Minnesota

²FIAS

Discounting is a natural part of the formulation of most sequential decisions problems and has been reported to underlie human and animal behavior in numerous empirical studies. Yet, typically discounting is treated as an arbitrary parameter needed to bound the expected future reward, and is treated as a fixed constant for all states. However, rather than an arbitrary parameter, discounting can also be interpreted as the probability of task termination. By assigning individual discount rates to separate states, we show it provides a framework for exploration bonuses, a rational basis for bounded computation and provides a basis for the automatic construction of stochastic options, or macro actions. Our results capitalize on work by Sonin (2008), who shows a generalization of the Gittins index to a Markov chain that allows for state dependent discount rates. This generalized index shows that the exploratory bonus for an option is the reciprocal of the probability of task termination. Moreover, there is a recursive algorithm to compute the generalized index, which produces an ordering of states in terms of the obtainable reward. Furthermore states that are traversed on the way to high-reward states can be eliminated, producing abstract states with new corresponding transition dynamics, and the probability of task-termination on these abstract states. We show how to use this elimination method to construct stochastic options (sub-policies) that find these abstract states, and that the choice between options can be computed via an index function. Using the termination probability formulation, we can derive exploratory incentives for each option and show that the effect of transition uncertainty is to reduce exploration incentive - in particular, incentive is reduced by the uncertainty over the set of next states - effectively a branching factor on the look ahead. This result provides a rational basis for bounding computation in model-based reinforcement learning. We apply the framework to a well-known reinforcement learning problem that is challenging for exploration - the "chain game." Our analysis decomposes the problem into a simple binary choice between two options, given enough experience with the transition probabilities, and we can quantify the difficulty in learning the better option. Humans placed in this environment fell into one of two distinct groups; one group performed enough unrewarded exploratory actions to find the better option while the second group under-explored and found the worst option. In debriefing, subjects in the former group reported finding worse option quickly, but believed that higher rewards were possible and thus continued exploration. Conversely, subjects in the latter group typically reported an initial exploratory phase, but upon finding the worse option believed the rest of the chain wasn't worth exploring. We believe these results provide new insight into the relationship between abstraction, exploration and the overall chance of task termination, and give a computational basis for understanding why exploration should be tied to competence or effectiveness, i.e. the ability to complete the task.

[doi:](#)

III-14. Optimal decision-making in multisensory integration

Jan Drugowitsch¹
Alexandre Pouget¹
Gregory C. DeAngelis¹
Dora E. Angelaki²

JDRUGOWITSCH@BCS.ROCHESTER.EDU
ALEX@CVS.ROCHESTER.EDU
GREGD@CABERNET.WUSTL.EDU
ANGELAKI@PCG.WUSTL.EDU

¹Department of Brain & Cognitive Sciences, University of Rochester

²Washington University

Combining multiple sources of information is essential for robust decision making under the influence of noise. For fixed duration tasks and constant flow of information over time it has been previously shown that both humans and animals can perform this combination optimally or near-optimally in the context of multisensory integration. However, it is more natural in everyday situations for the flow of sensory information to vary over time, and for

subjects to make decisions whenever they feel confident (reaction time tasks). In such a setting, it is unclear what constitutes behavior that is Bayes-optimal, and if subjects exhibit such behavior. The aim of this study is to address both questions by ideal observer modeling and by conducting experiments with human subjects performing a reaction-time version of a heading discrimination task. Optimality in multisensory integration is commonly assessed by testing if cues are weighted optimally according to their reliabilities. Such a procedure ignores the process of accumulating evidence over time and is therefore unsuitable to investigate multisensory integration in reaction time tasks. In contrast, bounded drift-diffusion models (DDMs) are able to explain both speed and accuracy of a decision, but only for single modalities. Here we show how to extend the DDM approach to combine multiple sensory modalities optimally by weighting them according to their respective reliability even when these reliabilities change over time. This combination rule leads to what we call local optimality, i.e., the resulting model combines the momentary evidence optimally. There is a second type of optimality, which we call global optimality, which concerns whether accuracy (percentage of correct response) in the multisensory conditions is the Bayes-optimal combination of the accuracy in the unisensory conditions. This latter form of optimality depends on how the decision bound is set. We show how to adapt this bound to achieve global optimality. We tested human multisensory integration in a reaction-time version of a fine heading discrimination task to test for both local and global optimality. In this task, subjects were seated on a motion platform, experienced motion with a Gaussian velocity profile, and reported whether their heading direction was left or right of straight ahead. In each trial, heading was indicated by vestibular cues related to platform movement, visual cues (optic flow) presented in a 3D random-dot stimulus, or both cues in combination. The reliability of the visual modality was controlled by presenting three different coherence levels of visual motion, and the subjects were free to report their perceived heading at any time after stimulus onset. Considering only the accuracy of their decisions (global optimality), subjects' cue combination weights are found to be sub-optimal, contrary to what is widely reported in fixed-duration experiments (as opposed to a reaction time task). However, if we consider both the speed and accuracy of their decisions, subjects' behavior is largely consistent with optimal combination and accumulation of evidence (local optimality). Specifically, subjects trade off some accuracy in the multimodal condition to achieve lower reaction times, as compared to the slower (visual) unimodal condition.

[doi:](#)

III-15. Role of anterior cingulate cortex in patch-leaving foraging decisions

Benjamin Hayden
Michael Platt
Duke University

HAYDEN@NEURO.DUKE.EDU
PLATT@NEURO.DUKE.EDU

One major problem faced by foraging animals is when to leave the patch in which they are currently feeding and move to a richer, more distant, patch. Optimal decisions in such situations are provided by an elegant and powerful mathematical framework, Charnov's marginal value theorem. The existence of these formal guidelines makes patch-leaving decisions an attractive model for studying the neural basis of richer, more complex and naturalistic decisions than are typically studied. We designed a novel behavioral task that recapitulates the essential aspects of patch-leaving decisions. On each trial, monkey subjects chose between two targets that corresponded to decisions to remain at and leave a patch respectively. The 'remain in patch' target provided a juice reward after a short delay, but its value diminished each time it was chosen, while the 'leave the patch' target led to a longer delay, provided no reward, and reset the value of the forage target at a high starting value. We systematically varied the duration of the delay associated with leaving - a factor that is analogous to travel time in natural foraging decisions. We found that subjects' choices closely matched the guidelines set by Charnov's theorem, endorsing the notion that this task captures the critical elements of foraging decisions. We measured responses of neurons in dorsal anterior cingulate cortex (dACC) while two subjects performed this task. We found that firing rates rose to a constant threshold as a patch was depleted, that upon reaching this threshold, subjects switched to a new patch. The rate at which this firing rate rose, and the corresponding value of this threshold was affected by travel time between patches, suggesting that ACC responses incorporate environmental parameters that affect patch-leaving decisions. Collectively, these data provide a possible neuronal basis for complex and natural economic decisions, and provide insight into the role of ACC in economic decision-making.

[doi:](#)

III-16. A neuronal model for context-dependent change in preference

Alireza Soltani¹
Benedetto De Martino²
Antonio Rangel²
Colin Camerer²

SOLTANI@BCM.EDU
BMARTINO@CALTECH.EDU
RANGEL@HSS.CALTECH.EDU
CAMERER@HSS.CALTECH.EDU

¹Baylor College of Medicine

²Caltech

One of the basic assumptions underlying classical rational choice theories is that the preference between given options does not depend on the presence or absence of irrelevant or inferior options. Nevertheless this assumption, known as the independence of irrelevant alternatives, does not hold true in many cases and its violation is instantiated in behavioral effects such as background context and tradeoff contrast. Moreover, a behavioral phenomenon known as asymmetric dominance (i.e. the effect of an asymmetrically dominated new option on the preference of two equally preferable options) has been shown to violate regularity and similarity hypotheses in economic theories (Huber et al 1982). Recent neurophysiological evidence has indicated that neural representation of an option value in the orbitofrontal cortex does not depend on the other available option (Padoa-Schioppa and Assad 2008). However, the same value representation is modulated by the range of values in a given session (Padoa-Schioppa 2009). While these neural findings seem to provide evidence against the behavioral data mentioned above, there are well-established results in sensory areas which support contextual modulations. For example, the neural representation of a visual stimulus is strongly modulated by the attributes of other stimuli in the scene (e.g. background effect) and their relationship to that stimulus (e.g. orientation contrast effect). These contrasting observations raise questions regarding the function of context-dependent neural representation and its relationship to the behavioral effects mentioned above. In order to investigate these questions and reach to a neurally-plausible model, we designed an experiment to study change in the preference between two target options (monetary gambles with a given magnitude and probability) due to the introduction of a new option (decoy gamble). To detect even a small change in the preference, we constructed the two target gambles such that they were equally preferable for each subject. By presenting the decoy at different regions of the gamble space, we found systematic changes in the preference between the target gambles depending on the location of the decoy at the single subject level (in contrast to previous studies which showed similar modulations between groups of subjects). We showed a model, which assumes dependence of neural representation on the range of option values on each trial, can capture the change in preference observed in our experiment. In this model, the value representation for a given attribute (i.e. magnitude or probability) is modulated by the inverse of the range of values on that attribute, independently of other attribute values. We compared the results of our model with competing models, discussed its similarity and differences to those models, and described its implications for value computations. Finally, we argue that the proposed neural mechanism is crucial for our evaluation system in order to deal with different ranges of values which occur in decision making in everyday life.

[doi:](#)

III-17. What, when and how of target detection in visual search.

Vidhya Navalpakkam
Pietro Perona
Caltech

VIDHYA@CALTECH.EDU
PERONA@CALTECH.EDU

We often search for targets in cluttered scenes, such as a friend in the crowd. When and how do we decide whether the target is present in the scene or not? This question has been surprisingly ignored, although reac-

tion time (RT) has been used since 1980s to infer search difficulty. While several quantitative models exist for predicting search accuracy, fewer models predict RT, and these are either qualitative or descriptive. We propose a generative model of what, when and how decisions about target presence in a scene are made. The model consists of four steps. First, the observer obtains information (e.g. via cortical V1 mechanisms) about the stimulus at each location and time instant. We assume that these stimulus observations are corrupted by white noise of variance σ^2 . Second, information across scene locations is combined using optimal Bayesian inference to compute at each instant t evidence for target presence in the scene, (i.e., the log likelihood ratio of target presence vs. absence at time t). Third, the observer accumulates evidence over time. Fourth, the observer decides 'yes' if the accumulated evidence exceeds criterion c , 'no' if the evidence falls below criterion $-c$, and otherwise continues to observe the scene. c is a shift in criterion, which depends on the prior probability of target presence (frequency), and the reward associated with the different responses. This model has only 2 free parameters (noise σ^2 , criterion c), compared to many parameters in drift diffusion models. The model can explain several phenomena in search including how RT and accuracy are affected by set-size (number of items in the scene), target-distractor discriminability and distractor heterogeneity. For example, the model shows that the slope of RT vs. set-size in target present trials is zero in pop-out tasks (i.e., time to find the target is independent of set-size), but the slope increases as target-distractor discriminability decreases. It explains the roughly 2:1 ratio in slopes of RT-set-size in target absent and present trials, and shows why the time to find the target is typically faster than the time to quit search when the target is absent. A recent study (Wolfe et. al, 2005) reported that rare targets, such as bombs in simulated airline passenger bags, are often missed. Our model replicates this behavior qualitatively, by showing how varying target frequency (e.g., from 50% to 10% to 2%) shifts the criterion such that the starting point of the decision process is closer to 'no' than 'yes' criterion, leading to high miss error rates, and fast abandoning of search. In addition, our model predicts that adjusting rewards by increasing penalties on miss errors will decrease these errors and increase RTs. We validated these predictions through psychophysics experiments with 4 human subjects. To summarize, we have proposed a generative model of visual search that with only 2 free parameters, can explain several search phenomena including how accuracy and RT are affected by set-size, target and distractor features, frequency and reward.

doi:

III-18. Beyond the edge: Amplification and temporal integration by recurrent networks in the chaotic regime.

Taro Toyozumi

L. F. Abbott

TARO@NEUROTHEORY.COLUMBIA.EDU

LFA2103@COLUMBIA.EDU

Department of Neuroscience, Columbia Univ.

Randomly connected networks exhibit a transition from fixed-point to chaotic activity as the variance of their synaptic connection strengths is increased. At the transition point, known as the edge of chaos, networks display a number of desirable features (Bertchinger and Natschläger, 2004), including large gains and integration times for responses to external inputs. Away from this edge, in the fixed-point regime that has been the focus of most models and studies, gains and integration times fall off dramatically, which implies that parameters must be fine tuned with considerable precision if high performance is required. Here we show that the fall-off in gains and integration times is slower on the chaotic side of the transition, meaning that good performance can, under appropriate conditions, be achieved with less fine tuning. Our study is based on a dynamical mean-field calculation of the Fisher information provided by a large recurrently connected network about its external input. The Fisher information bounds decoding accuracy for both static and dynamical stimuli, and can also be used to quantify the integration time or memory lifetime of a network (Ganguli et al., 2008). To quantify the behavior of the Fisher information near the transition point, we evaluate the critical exponents of the Fisher information under a simple case. For any antisymmetric smoothly saturating response nonlinearity, the Fisher information diverges more rapidly on the chaotic side than the non-chaotic side of the edge of chaos. We show that, with observation noise, the Fisher information is maximized at the edge of chaos, where the network time-constant shows a critical

slowdown and small inputs are highly amplified. Furthermore, at a given distance away from the transition point, the chaotic state is often more informative than the non-chaotic state, and it provides a longer lasting memory. The analytical expression for the Fisher information provides an intuitive picture of the trade-off between increasing the signal due to a larger gain and increasing chaotic "noise". The presence of observation noise emphasizes the importance of increasing the signal over decreasing the internally generated noise, providing an advantage to the chaotic state. N. Bertchinger, and T. Natschläger (2004) Real-time computation at the edge of chaos in recurrent neural networks. *Neural Comput.* 16, 1413-1436. S. Ganguli, D. Huh and H. Sompolinsky (2008) Memory traces in dynamical systems. *Proc. Natl. Acad. Sci (USA)* 105, 18970-18975.

[doi:](#)

III-19. Control of persistent spiking activity by background correlations.

Mario Dipoppa
Boris Gutkin

MARIO.DIPOPPO@ENS.FR
BORIS.GUTKIN@ENS.FR

Group for Neural Theory, LNC, DEC, ENS

A telltale feature of working memory (WM) is the sustained neural activity associated with holding necessary information on-line (e.g., Romo et al. 1999 *Nature*: 399, 470-473). A key unresolved question is how the cortical machinery manipulates this sustained activity in order to perform WM tasks. Specifically, a unified spike-based process that can "gate-in" the information at the task initiation, control the activity dynamics during the memory period and "gate-out" the memory trace at task completion remains to be specified. Here we show that it is possible to control multiple aspects of the persistent activity at once via variable correlations in the background noise and suggest that this control depends on transient synchronization of the spiking. As a paradigmatic model of sustained activity we consider a network of QIF neurons coupled with all-to-all recurrent synaptic excitation and background noise in which we can modify the global amount of correlations among neurons. In accordance with previous studies, above a critical synaptic strength and time-scale the network exhibits multi-stable behavior with sustained activity states and quiescence. We find that in the absence of noise, the sustained activity state with the lowest rate shows a specific spike-time structure referred to as "splay-state". This spike-time structure is characterized by maximum anti-synchrony among the active neurons (the spikes of $N-1$ neurons constituent in a network sized N are equally splayed out within an interspike interval of the N th neuron). We show this splay-state spike arrangement to be robust to variations in PSP shape. We develop analytic conditions for the stability of the splay-state using event-driven spike-time maps. Furthermore the network responses and the robustness of the sustained splay-state are examined under noisy conditions. For a given noise strength, we find that noise correlations robustly promote transitions from the splay-state to the quiescent state while suppressing spurious reactivation of the sustained activity. We find that these transitions are due to noise induced increases in spike-time coherence among the active neurons. This effect of noise-induced synchronization becomes dominant as the network size grows. We thus find a stochastic analogue of synchrony-induced turn-off of sustained activity as proposed by Gutkin et al., (2001 *J Comp Neurosci*: 11, 121-134). We further find that for strongly correlated noise no splay-state activation is possible. Hence by tuning correlations in the background memory-unrelated activity we can control 1. the ability of an incoming stimulus to initiate sustained activity, 2. spurious activations of sustained activity, 3. the mean life-time of the persistent state, 4. the ability of excitatory transients to read-out and turn-off the activity rapidly. This mechanism could be used to conceive a new model of working memory where the basin of attraction of the persistent state is controlled by a correlation parameter. In conclusion, correlated stochastic input can gate persistent activity during a working memory task and slow variations of the background correlations may embed task-timing information directly into the working-memory trace.

[doi:](#)

III-20. Multiple routes to functionally feedforward dynamics in cortical network models.

Edward Wallace¹
 Marc Benayoun²
 Wim van Drongelen³
 Jack Cowan¹

EWJWALLACE@GMAIL.COM
 MARCB@UCHICAGO.EDU
 WVANDRON@PEDS.BSD.UCHICAGO.EDU
 COWAN@MATH.UCHICAGO.EDU

¹Dept. of Mathematics, University of Chicago

²CNS Graduate Program, University of Chicago

³Pediatric Neurology, University of Chicago

Networks with closely balanced inhibition and excitation exhibit dynamics of particular interest to computational neuroscientists. They may respond quickly to sudden changes in afferent input, or presented with slow changes in input they may move between synchronous and asynchronous dynamics. This has been attributed to a hidden functionally feedforward structure in the network, meaning that a co-ordinate change to the dynamical variables reveals an underlying unidirectional flow of activity. It is therefore of interest to determine which combinations of feedforward structure and input obtainable under physiological constraints produce which network dynamics, and what this means for information processing. We generate random sparse matrices with strong hidden feedforward structure, by adjusting the spectral properties of sub-blocks of the connectivity matrix. This may be achieved by fixing the mean values of the sub-blocks of the matrix. Less intuitively, it may also be achieved by adjusting their variances. Then, the interplay between the connectivity matrix and the input, depending on their mutual correlations, causes the network to move between dynamical regimes. A special case of this in a stochastic spiking network is a transition from asynchronous firing to aperiodic synchronous firing, with spikes grouped into so-called neuronal avalanches. The spectral properties necessary for functionally feedforward dynamics may hold in networks with a variety of connection topologies. In particular, we show how to generate networks with small-world connectivity and underlying feedforward dynamics. This is relevant to the problem of inferring a network's connectivity from its dynamics: experimental measurements of the dynamical repertoire obtainable from functionally feedforward networks are not enough to show that the network under observation does or does not have small-world topology.

[doi:](#)

III-21. Visualizing classification decisions of hierarchical models of cortex

Will A. Landecker¹
 Steven P. Brumby²
 Mick Thomure¹
 Garrett T. Kenyon²
 Luis M. A. Bettencourt²
 Melanie Mitchell¹

LANDECKW@CS.PDX.EDU
 BRUMBY@LANL.GOV
 THOMURE@CS.PDX.EDU
 GARKENYON@GMAIL.COM
 LMBETT@LANL.GOV
 MM@PDX.EDU

¹Portland State University

²Los Alamos National Laboratory

Recent work on computer vision systems using hierarchical feed-forward models of visual cortex has reported accuracy exceeding 80% on certain binary object category discrimination tasks using natural image sets. The relatively good performance of these models compared to other computer vision systems creates a need for visualizing how the hierarchically constructed representations are being utilized, and whether correct classification is due to relevant learned features or to artifacts arising from the limited number of images used in these experiments. In particular, it is important to ask if the model is indeed classifying features from the object (e.g., animal), rather than the background (e.g., forest). Previous work has focused on visualizing the hierarchically constructed features learned by the model in early stages of visual cortex (V1, V2) in order to show that salient patterns have been captured from the training images, but has not demonstrated that these patterns are used to classify relevant

portions of the image, i.e., foreground instead of background. We now propose a method of extracting information from a supervised classifier and tracing back through the hierarchical model in order to determine the degree to which each region of the image contributed to the image's overall classification. Using a naïve Bayes classifier and calculating the posteriors associated with each high-level image feature, we trace down through the hierarchy to create an informative visualization of which low-level image features were associated with the positive (foreground) and negative (background) classes. This method of associating low-level image features (in our case, Gabor functions at specific locations, orientations, phases, and scales) with image classes (positive and negative) also suggests a new way to quantitatively evaluate a model's recognition of relevant image regions. To this end, we evaluate the statistical overlap between the image regions associated with the positive class and the pixels belonging to the object being classified. We propose this "relevance recognition" as a useful measure of the extent to which classifications were caused by the foreground or background of images. We compare the performance of different learning rules used to train PANN (Petascale Artificial Neural Network), an implementation of a hierarchical feed-forward model developed by the Synthetic Visual Cognition team at Los Alamos National Laboratory. Results show that the use of different learning rules to construct prototypes in V2 lead to higher classification accuracy. We determine how this happens in terms of relevant regions of the image for object classification. The visualizations and the measure of relevance recognition are applicable to a large class of hierarchical feed-forward models of visual cortex, and provide valuable information about the aggregation of low level features that enable high level classification. In this way, they go beyond the standard measures of classification accuracy to provide analytical tools for cognitive processes and are a step towards the implementation of feed-back connectivity aimed at disambiguating high level object classification.

[doi:](#)

III-22. Hidden structures detection in nonstationary spike trains

Ken Takiyama¹
Masato Okada²

TAKIYAMA@MNS.K.U-TOKYO.AC.JP
OKADA@K.U-TOKYO.AC.JP

¹The University of Tokyo.

²The University of Tokyo / RIKEN

Elucidating neural encoding from irregular neural activities is one of the most important issues in neuroscience. Firing rates abundantly include information encoded by neurons and are often estimated averaging neural activities across trials. We however need to estimate firing rates using only one or the least spike trains especially in the case that neural activities represent internal processes such as decision making and motor planning because the process differs from trial to trial. Firing rate estimation using only one spike train hence plays a crucial role to elucidate neural encoding. Many studies have been proving probabilistic models can estimate firing rates using one spike train in recent years. These studies assume stationarity of mean and temporal correlation within a trial. Stationarity means time series has temporally uniform statistical property. We call what does not have stationarity as nonstationary. Many studies, using hidden markov model (HMM), have been showing neural network activity transit among neural states within a trial. Since statistical properties of firing rates differ drastically in all neural states, firing rates include nonstationarity with in a trial in many cases. Recent studies have been indicating that neural states reflect the properties of inputs encoded by neurons. Because the studies also suggest the transitions show trial-by-trial variation, neural state estimation using only one spike train plays a crucial role for elucidating neural encoding. We construct the algorithm that can simultaneously estimate nonstationary firing rates, neural state transition timings and neural state numbers using only one spike train. Our algorithm consists of Switching State Space Model(SSSM). SSSM defines more than one prior distributions one of which generates observation data. SSSM can also estimate neural state transition timings where a prior distribution that generates observation switches. The HMMs assume firing rates are constant in each neural state. Firing rates however naturally showing temporal variation, Excluding the time variation probably obscures the neural state information. SSSM enables to estimate neural state transitions with considering the temporal variation. Learning and estimation algorithm are constituted of Variational Bayes method. Automatic relevance determination induced by Variational Bayes enables to estimate neural state number at a time. Synthetic data analysis shows our algorithm can simultaneously estimate nonstationary firing rates, neural state transition timings and neural state numbers using one spike train

with high accuracy. Applying to real data, which is available on Neural Signal Archive, reveals our algorithm can estimate neural state information from area MT data. These neural states probably correspond to transient and sustained states which have been detected heuristically in area MT data. We confirm the HMM can not detect these neural states. These results suggest our algorithm has the versatility in neural state estimation and the effectiveness on neural encoding elucidation.

[doi:](#)

III-23. Microcircuits of stochastic neurons

Stefano Cardanobile

Stefan Rotter

BCCN Freiburg

CARDANOBI@BCCN.UNI-FREIBURG.DE

STEFAN.ROTTER@BIOLOGIE.UNI-FREIBURG.DE

Multiplicatively interacting point processes and applications to neural modelling Stefano Cardanobile & Stefan Rotter BCCN Freiburg Mathematical analysis of complex neural network dynamics is both challenging and important for research in neuroscience. Most current approaches, though, rely on mean-field approximations, which have difficulties to evaluate the influence of network structure on its spiking dynamics. We exploit the stochastic nature of neuronal firing and set up a point process framework, based on the observation that the escape noise of real neurons is exponential with respect to their membrane voltage [1]. Assuming linear integration of inputs, this translates into a multiplicative interaction rule on the level of instantaneous firing rates: each incoming spike effectively multiplies the instantaneous firing rate by a fixed "synaptic weight". This approach is in contrast to Hawkes' linear model [2], where the instantaneous firing rate is given by a convolution of the input spike rate with a linear temporal filter. This effectively prevents the implementation of inhibition in this model. We proved that the equations governing the dynamics of the expected firing rates in our multiplicative system are of Lotka-Volterra type, if one ignores covariances [3]. Based on numerical simulations, we show that this approximation works quite well under very general conditions. Asymptotically, the observed firing rates coincide with the solutions of the associated rate equations even in cases where the rates do not converge to a fixed point, but exhibit transient dynamics. Multiplicatively interacting point processes offer an interesting novel framework for the study of complex neural network dynamics. To illustrate this claim, we finally describe some structured networks that are able to process information, and discuss specifically competing neural populations to describe experiments where rivaling features are perceived. Our model qualitatively replicates the unimodal distribution of dwell times as observed in experiments, and it leads to an intuitive explanation of the switching dynamics. The project has been supported by BMBF grant 01GQ0420 to the BCCN Freiburg. References: [1] Jolivet et al. (2006) Predicting spike timing of neocortical pyramidal neurons by simple threshold models. J. Comp. Neurosci. [2] Hawkes (1971) Spectra of some self-exciting and mutually exciting point processes. Biometrika [3] Cardanobile & Rotter (2009) Multiplicatively interacting point processes and applications to neural modeling. J. Comp. Neurosci., conditionally accepted. Available on arXiv.org 0904.1505

[doi:](#)

III-24. Bayesian methods for intracellular recordings: electrode artifact compensation and noise removal

Yongseok Yoo

Jonathan W. Pillow

University of Texas at Austin

YYOO@MAIL.UTEXAS.EDU

PILLOW@MAIL.UTEXAS.EDU

Intracellular recording techniques provide a fundamental tool for understanding the functional properties of single neurons. Typical experiments involve the injection of a time-varying current while recording a neuron's membrane

potential. Unfortunately, injecting and recording through a single electrode gives rise to substantial measurement artifacts. The resistance and capacitance of the recording electrode make it difficult to know either the exact time-varying current entering the cell or the voltage across the membrane. Classically, this problem has been addressed using "bridge" compensation and capacitance neutralization, which model the electrode as a simple RC circuit, and subtract the effects of this circuit from the recording. However, the electrodes used in most intracellular experiments do not closely obey this description. Recent work from Brette et al (2008) suggested an elegant alternative, known as Active Electrode Compensation (AEC), which can be described in three steps: first, in a hyperpolarized regime, inject a small-amplitude white noise current and perform reverse correlation to obtain a "composite" linear kernel, reflecting both membrane and electrode filtering properties; second, separate the electrode kernel from the membrane kernel (which should have the shape of a decaying exponential) by fitting an exponential to the tail of the composite kernel and fitting the electrode kernel "Ke" to the remainder; third, actively compensate by subtracting the injected current convolved with Ke from the recorded voltage. This gives an accurate estimate of membrane potential; the time-varying current across the membrane can be obtained in a similar manner. Here we develop a recursive Bayesian filtering approach to electrode artifact compensation, which extends the AEC framework by incorporating explicit noise in both the measurements and the voltage dynamics. As in AEC, we assume that sub-threshold dynamics are linear and that the electrode behaves as an arbitrary linear filter on the injected current. If both noise sources are Gaussian, this setup is a linear dynamical system with Gaussian noise, which has a well-known, optimal solution known as the Kalman filter (or the Kalman smoother, for off-line applications). This problem allows for efficient maximum-likelihood estimation of the model parameters via the expectation-maximization (EM) algorithm. Using simulated data, we show that EM gives estimates of the electrode kernel that are more accurate than those obtained by AEC, particularly in the case of large measurement noise. This model can also be used to remove noise, but we cannot straightforwardly apply the model—which assumes linear, sub-threshold dynamics—to nonlinear, spiking voltage data. Instead, we fit a separate model of the membrane dynamics to a short segment of spiking data. Then, using (1) measurement noise and an electrode kernel estimated from sub-threshold recordings and (2) a dynamical model estimated from spiking data, we can accurately compensate for electrode artifacts and measurement noise. The resulting model allows for online filtering, which is highly relevant to dynamic clamp recordings, and offline "smoothing" of voltage traces. We apply these methods to simulated data from a Hodgkin-Huxley neuron, and show that fitting is robust to noise level and accurate with small amounts (~1s) of data.

doi:

III-25. A normative theory of short-term synaptic plasticity

Jean-Pascal Pfister¹

Peter Dayan²

Máté Lengyel¹

JEAN-PASCAL.PFISTER@ENG.CAM.AC.UK

DAYAN@GATSBY.UCL.AC.UK

M.LENGYEL@ENG.CAM.AC.UK

¹University of Cambridge, Dpt. of Engineering

²Gatsby Computational Neuroscience Unit, UCL

Synapses are highly complex dynamical elements whose efficacies undergo perpetual change over several time scales. In particular, due to short-term plasticity (STP), a sequence of presynaptic spikes can elicit a sequence of increasing (facilitation) or decreasing (depression) postsynaptic potentials. Facilitation and depression can even co-occur in the same synapse. The regularity and magnitude of these changes argue against their being treated only as wanton variability. However, existing computational suggestions are restricted to select neural subsystems or types of change, despite the ubiquity of STP. Here, we suggest that the two forms of STP collectively solve a computational problem that is faced by any neural population in which neurons represent and compute with analogue quantities, but communicate with discrete spikes. Almost all neurons transmit only a partial, digital account (spike or no spike) of the full analogue signal (membrane potential) associated with computation. Synapses thus face an inverse problem, i.e. estimating the presynaptic membrane potential based only on the observed spikes. Given a statistical (generative) model of presynaptic membrane potential fluctuations and spiking, the solution to this inverse problem is an optimal Bayesian estimation filter. We show that under a simple model of membrane potential dynamics involving a mixture of different time scales, and a non-linear Poisson model of spike genera-

tion, the optimal estimator features non-stationarities like those exhibited by postsynaptic potentials. Specifically, analysis of the steady-state properties of the filter shows that its dynamics generally match those of STP, a relationship that can be shown analytically in a simple limiting case via a reduction to a canonical dynamical model of STP. This close analogy yields a functional account of STP, under which the local postsynaptic potential and the level of recruited synaptic resources track the (scaled) mean and variance of the estimated presynaptic membrane potential. We also show that a dynamical synapse with STP performs significant better at this estimation task than a static synapse whose efficacy is fixed. Our theory is readily testable, since it suggests a precise relationship between quantities that have been subject to extensive, separate, empirical study - namely the 'natural statistics' of a neuron's membrane potential dynamics, the form of its spiking non-linearity, and the form of STP it expresses in its efferent synapses. This work was supported by the Wellcome Trust and the Gatsby Charitable Foundation.

doi:

III-26. Dynamic Bayesian network model on two opposite types of sensory adaptation

Yoshiyuki Sato¹
Kazuyuki Aihara²

YOSHI@SAT.T.U-TOKYO.AC.JP
AIHARA@SAT.T.U-TOKYO.AC.JP

¹Institute of Industrial Science, The University of Tokyo

²The University of Tokyo

We can adjust our sensory and motor systems to fit the changes in our bodies and the surrounding environment. Such adaptational phenomena have been found to play an important role in facilitating an understanding of the nervous system. In recent times, a wide variety of studies have shown that human perception and action can be regarded as optimal computations to compensate for the stochastic nature of our sensory and motor systems and the environment. Specifically, some studies have shown that adaptation can also be regarded as Bayesian inference of relevant parameters. Two opposite types of adaptational effect have been observed in psychophysical experiments. For example, after a subject is exposed to audio-visual stimuli with biased temporal difference repeatedly, the subject tends to perceive the repeated stimuli as simultaneous, and this effect is called "lag adaptation". This type of adaptational effect has been observed in a very wide range of domains. Recently, an opposite type of effect was found in tactile perception, which was called "Bayesian calibration" [1]. In the experiment, repeated stimuli were more unlikely to be judged as simultaneous. However, the mechanism and function of this new type of adaptational effect is not clear. The two types of adaptational effects, the lag adaptation type and the Bayesian calibration type, can be explained as the adaptive learning of the likelihood functions [2] and the prior distributions [1], respectively, in the Bayesian inference of the stimulus properties. Therefore, these two effects are complementary to each other from, both, the viewpoint of their phenomenal properties and that of their Bayesian models. Although we constructed a model which included changes in, both, the likelihood functions and the prior distributions [3], the model is unable to provide a concrete interpretation of the determinants of the adaptational effect. Here, we construct a dynamic Bayesian network model of perceptual adaptation to provide a unified explanation of the adaptational effects. The model includes hidden parameters that determine the mean values of the likelihood functions and the prior distribution functions, which are assumed to be Gaussian functions. We assume that the adaptational effect is the result of the estimation of these parameters. We show that this model can reproduce both types of adaptational effects depending on the model parameters. With some modification to the optimal inference in this model structure, we analytically derive the parameter condition to determine the type of adaptation. By analyzing the result, we derive an interesting relationship between the adaptational type and the stimulus presentation method. Our model predicts that if adapting stimuli are random around a constant mean value, the effect would be more like the lag adaptation type, and if adapting stimuli have a random walk nature, the effect would be more like the Bayesian calibration type. Reference [1] M. Miyazaki et al. *Nature Neuroscience*, 9(7):875-7, 2006 [2] Y. Sato et al. *Neural Computation*, 19(12):3335-3355, 2007 [3] Y. Sato and K. Aihara. *Artificial Life and Robotics*, 14, 2009 (in press)

doi:

III-27. Methods for neural circuit inference from population calcium imaging data

Joshua T. Vogelstein¹

Timothy A. Machado²

Yuriy Mishchenko²

Adam M. Packer²

Rafael Yuste²

Liam Paninski²

¹Johns Hopkins University

²Columbia University

JV.WORK@JHU.EDU

TIM.MACHADO@GMAIL.COM

YURIY.MISHCHENKO@GMAIL.COM

AP2416@COLUMBIA.EDU

RMV5@COLUMBIA.EDU

LIAM@STAT.COLUMBIA.EDU

Calcium imaging techniques have become an important and ubiquitous tool for studying neural circuits. Advances in fluorescence microscopy and calcium sensors allow researchers to obtain data with increasingly better spatial and temporal resolution—and provide them with the ability to observe ever finer details of population dynamics. However, many fundamental questions about neural coding and circuit connectivity are not directly approachable with raw fluorescence data. Here we present results from the application of novel inference techniques to calcium imaging data: a fast spike inference algorithm [1], and a Monte Carlo expectation maximization algorithm for inferring neural connectivity [2]. We have developed a faster than real time algorithm that infers the approximately most likely spike train for each neuron in an imaged population. A simple generative model of calcium dynamics was used to formulate a concave objective function. In order to ensure that negative spikes are not inferred, while preserving our ability to use standard gradient ascent methods, a barrier term was imposed. Since the Hessian term in our objective function is a tridiagonal matrix, we can implement the Newton-Raphson method in linear time by using standard banded Gaussian elimination methods. By generalizing our model of calcium dynamics to an entire population, we can infer the probability of a single neuron spiking given the fluorescence activity of the entire network. A maximum a posteriori estimate of the model parameters is fit to the observed data through the use of a Monte Carlo expectation maximization algorithm. The sufficient statistics are computed using a spike inference algorithm [1,3] and a hybrid blockwise Gibbs sampler. On simulated noisy calcium data, connectivity matrices (even for more than 100 neurons) can be accurately reconstructed using this approach. In order to refine these methods for use on real data, we used calcium indicators in vitro to image spontaneous neural activity in mouse cerebral cortex. To verify the accuracy of our spike inference methods, we recorded from individual neurons intracellularly during imaging. Across a variety of preparations, our fast spike inference algorithm outperformed the optimal linear deconvolution method (a Wiener filter) and also accurately inferred the timing of most action potentials detected intracellularly. These data are now being used to test our connectivity inference algorithm. We can validate our model by comparing the neurons that were inferred to create negative connections with genetically labeled inhibitory interneurons present in the tissue. Inferred synaptic connections can also be directly verified using paired whole cell recordings. [1] Vogelstein, JT, et al. Online nonnegative deconvolution for spike train inference from population calcium imaging. In preparation. [2] Mishchenko, Y, et al. A Bayesian approach for inferring neuronal connectivity from calcium fluorescent imaging data. *Annals of Applied Statistics*. In press. [3] Vogelstein, JT, et al. Spike inference from calcium imaging using sequential Monte Carlo methods. *Biophysical Journal*, 97(2), 636-655 (2009).

[doi:](#)

III-28. Coincidence detection in active neurons

Cyrille Rossant

Romain Brette

Ecole Normale Supérieure

CYRILLE.ROSSANT@ENS.FR

ROMAIN.BRETTE@ENS.FR

Neuronal synchronization is ubiquitous in the nervous system, yet its functional role for information processing is still unclear. Because of the leak current, two input spikes are more likely to make a postsynaptic neuron fire

when they are synchronous. This coincidence detection property has been demonstrated in vivo for thalamo-cortical processing (Usrey, Alonso and Reid (2000) J Neurosci 0(14)), but the theory of coincidence detection in active neurons is still sparse. We tried to answer the following question: in an active neuron (with background synaptic activity), what is the extra probability of firing a spike in response to two input spikes, as a function of the delay between them? Our approach is based on several properties of cortical neurons in vivo: the membrane time constant is short compared to typical inter-spike intervals; the membrane potential is stochastic, with a mean well below threshold (balanced regime); the autocorrelation time constant is dominated by slow inhibitory fluctuations. Incorporating these properties in a stochastic spiking model allowed us to obtain quantitative estimates of coincidence detection properties. We obtained an approximation of the extra firing probability as a function of the delay between two pre-synaptic spikes, from which we derived two quantities: the strength and timescale of coincidence detection. The strength of coincidence detection was defined as the relative increase in firing in response to coincident spikes compared to distant ones. We found that coincident spikes can be more than twice as efficient as distant spikes in realistic situations. We extended our results to calculate the extra firing rate induced by a pool of correlated input spike trains and found that the postsynaptic firing rate was very sensitive to the strength of correlations. The timescale of coincidence detection was defined from the decay of the extra firing function and corresponds to the temporal window of interaction between two input spikes. We found that the timescale of coincidence detection can be expressed as the product of the time constant of post-synaptic potentials (PSPs) and of a quantity defined by the membrane potential distribution and the synaptic weight (maximum of the PSP). This quantity is always smaller than 1, it approaches 1 at high noise level and 0 at low noise level. Our estimates were consistent with numerical simulations. Our modeling results show that, in realistic situations, 1) fine correlations can have a strong impact on the response of a post-synaptic neuron and 2) spiking models are sensitive to finer correlations than expected from the value of the membrane time constant.

doi:

III-29. A delta-rule approximation to Bayesian inference in change-point problems

Robert C. Wilson¹

Matthew Nassar²

Joshua Gold²

RCW2@PRINCETON.EDU

NASSARR@MAIL.MED.UPENN.EDU

JIGOLD@MAIL.MED.UPENN.EDU

¹Princeton University

²University of Pennsylvania

Whether it is a stock market crash, falling in love, or task switching in controlled behavioral experiment, life is full of "change-points" that cause the rules of the game to change abruptly, often without warning. To thrive in these circumstances, individuals must recognize these events quickly and adapt their behavior accordingly - for example by pulling out their money, putting on cologne, or adopting the appropriate task rules. As with other inference problems, ideal-observer models can help shed light on the computational demands of, and solutions to, change-point problems. Online Bayesian solutions to this problem have been reported [Adams & MacKay, Technical report Cambridge University, 2007; Fearnhead & Liu J. R. Statist. Soc. B 69(4):589, 2007], which can capture many aspects of human behavior on a simple change-point task [Nassar et al., in preparation]. However, despite this success, the computations implied by the Bayesian models seem biologically implausible. These computations require the brain to maintain and update a constantly growing probability distribution (the "run-length distribution") over all possible locations of the last change-point. In this work we address this shortcoming by systematically reducing the full Bayesian solution to show that it can be well approximated by a form of delta rule. This reduced Bayesian algorithm updates its current beliefs about the world based on the difference between the model's current predictions about the world and the observed reality. Given the substantial evidence for this kind of prediction-error signal in the brain [e.g. Schultz et al., Science 275:1593, 1997], the model is more plausible than the full Bayesian solution. This approach also yields a massive reduction in computational cost, going from $O(t)$ computations per time step at time t , to $O(1)$. Our approach rests on drastically simplifying the representation of the run-length distribution. In particular, we introduce a reduced distribution that is not maintained over all

possible run-lengths but instead on just two run-lengths: the first a run-length of zero, which assumes that a change-point just occurred; and the other a non-zero run-length, which represents the expected time since the last change-point. Using this reduced distribution, we derived update equations for the parameters of the model that, for the mean of the predictive distribution, take the form a delta rule. We find that the learning rate depends on both average run-length and the probability of a change-point at the current time, given the latest data. We also show that an extension of this work to deal with hierarchical change-point models [Wilson et al., submitted] leads to a hierarchy of delta-rules. This hierarchical model can adjust to increasingly complex change-point dynamics, including cases not considered by the reported online Bayesian solutions, in which the frequency of change-points is unknown. This work establishes a formal relationship between a biologically plausible algorithm based on the delta rule and ideal-observer Bayesian models. The results allow us to characterize in detail the strengths and limitations of both kinds of model in performing on-line inference for a range of complex change-point problems.

doi:

III-30. Cellular mechanisms that may contribute to prefrontal dysfunction in psychosis

Vikaas S. sohal¹

Karl Deisseroth²

VIKAAS@STANFORD.EDU

DEISSERO@STANFORD.EDU

¹Dept. of Psychiatry and Behavioral Sciences, Stanford University

²Stanford University

Prefrontal cortical dysfunction is a hallmark of psychotic illnesses such as schizophrenia. Hyperdopaminergic states are hypothesized to contribute to psychosis, and psychotomimetics including phencyclidine also reproduce prefrontal dysfunction. However, unifying neural mechanisms that may be shared by diverse causes of prefrontal dysfunction remain unknown. Here, we studied how phencyclidine (PCP) and the D2 agonist quinpirole affected the function of pyramidal neurons in layer 5 of the prefrontal cortex. We made whole-cell patch-clamp recordings from layer 5 pyramidal neurons in prefrontal cortical slices from 6-10 week old mice. We stimulated neurons with depolarizing current pulses, trains of simulated excitatory post-synaptic currents (sEPSCs), and light flashes which activated channelrhodopsin-2 (ChR2) expressed transgenically in layer 5 pyramidal neurons under the Thy1 promoter. We characterized the input-output properties of individual pyramidal neurons by computing the mutual information between the rates of input sEPSCs and output spikes. We found that bath application of (+/-) Quinpirole at doses of 10 microM (n = 5 cells), or (-) Quinpirole at doses of 5 microM (n = 9 cells), produced periods of prolonged depolarization outlasting injection of depolarizing current to in 9/14 cells. The prolonged depolarization elicited spiking in 5/9 cells, and could be reversed by application of the D2 antagonist and antipsychotic haloperidol (1-2 microM; n = 3 cells). Quinpirole also occasionally produced firing following trains of light flashes that could be reversed by haloperidol (1-2 microM). 5 microM (-) Quinpirole also decreased the amount of information that pyramidal cells transmitted about the rate of input sEPSCs, from 1.10 +/- 0.05 bits to 0.96 +/- 0.08 bits (information per 100 msec window; p < 0.05; n = 5 cells). At a concentration of 5 microM, PCP frequently produced similar periods of prolonged depolarization outlasting depolarizing current injection (5/8 cells), and occasionally produced depolarization that was concurrent with the injection of depolarizing current resulting in wide spikes and doublets (additional 1/8 cells). These depolarizations produced plateau potentials during responses to EPSC trains, resulting in periods of time during which neurons were non-responsive to input. Layer 5 prefrontal neurons are known to maintain persistent activity, and send output signals which may contribute to the role of the prefrontal cortex in working memory and other cognitive functions. Here we show that both quinpirole and phencyclidine evoke activity-dependent depolarizations in layer 5 prefrontal neurons, in some cases interfering with the transmission of inputs. These effects may support persistent activity while reducing the sensitivity of the prefrontal cortex to incoming inputs. The pathological persistence of abnormal cognitions, and perseverative behavior in tasks that depend on the prefrontal cortex are major features of psychosis, and may relate to these neural changes. Thus our results suggest a common cellular mechanism that may contribute diverse forms of psychosis.

doi:

III-31. In vivo multi-single-unit extracellular recordings from identified neural populations in fruit flies

Misha B. Ahrens¹
 Mladen Barbic²
 Brian Barbarits²
 Brian G. Jamieson³
 Vivek Jayaraman²

MBA22@CAM.AC.UK
 BARBICM@JANELIA.HHMI.ORG
 BARBARITS@SBMICROSYSTEMS.US
 BRIANJ@SBMICROSYSTEMS.US
 VIVEK@JANELIA.HHMI.ORG

¹Cambridge University

²Janelia Farm Research Campus, HHMI

³SB Microsystems

The fruit fly's genetic toolbox makes it a promising model system for systems neuroscience. Neural recordings from its brain have so far been made using single-channel glass electrodes or fluorescence imaging with genetically encoded calcium indicators. The former has high temporal resolution but is limited to single neurons; the latter can capture multi-neuron activity with high spatial resolution, but has poor temporal resolution. Multi-unit recordings with high temporal specificity can be achieved with multi-electrode probes, but existing probes are an order of magnitude too large for the fruit fly brain. Here, we describe the development of a new set of miniature multi-electrode probes with pad sizes that are significantly smaller than currently available probes, and yet of low enough impedance to capture spiking activity in the fruit fly brain with single neuron resolution. Using a focused ion beam system to both cut through and deposit platinum on commercially available multi-electrode probes from NeuroNexus, we designed probes which are ~20 microns wide, and contain up to five shanks. Each shank terminates in a recording pad with a surface area of roughly 25 square microns. With four- or five-channel probes we could spike-sort up to five independent units from central brain recordings. Simultaneous loose patch recordings from identified neurons confirmed that the waveforms detected belonged to single units. The power of *Drosophila* is in its well-developed genetics, with the free availability of thousands of fly lines that allow the targeted expression of exogenous proteins to selected neural sub-populations. Genetic targeting in this manner allows responses in an entire brain region, such as the central complex (CC), to potentially be mapped cell-type by cell-type in small groups of a few neurons each. We used the Gal4-UAS system to express channelrhodopsin-2 (ChR2) in small populations of identified neurons in sub-regions of the CC. Loose-patch recordings confirmed that a brief pulse of blue light elicited one or more spikes in ChR2-expressing neurons with a latency low enough to discriminate directly activated neurons from those activated through monosynaptic connections with ChR2-positive neurons. In a manner similar to that employed in a recent paper where the authors used viruses to target ChR2 to specific neurons in the rodent auditory cortex [2], we could then identify a subset of the spike-sorted units as belonging to the genetically specified neural population based on their responses to blue light. We are now using this approach to target probes to areas that maximize our yield of identified neurons. Our miniaturized multi-site probes allow us to record from small and densely packed neurons. We have used these probes to make the first multi-single-unit extracellular recordings in the fruit fly, but probes with similar site sizes and spacing may also be useful for recordings in packed areas of the vertebrate brain. Finally, this technique can be combined with genetically targeted expression of ChR2 to perform in vivo recordings of multiple identified neurons in the fruit fly central brain. 1. Lima, Znamenskiy, Hromadka, and Zador (PLoS ONE, 2009)

doi:

III-32. Layered sparse associative network for soft pattern classification and contextual pattern completion

Evan Ehrenberg
 Pentti Kanerva
 Friedrich Sommer

EVANEHRENBURG@BERKELEY.EDU
 KANERVA@SICS.SE
 FSOMMER@BERKELEY.EDU

Redwood Center for Theoretical Neuroscience, University of California, Berkeley

Traditional models of associative networks [1, 5] have impressive pattern completion capabilities. They feature high storage capacity and can reconstruct incomplete sensory input from stored memory in a context-sensitive fashion. However, associative networks are not useful as pattern classifiers beyond simple toy examples where the classes form well-separated clusters in pattern space. Conversely, multi-layer feed-forward neural networks [4] can be trained to solve challenging classification tasks, but they cannot perform pattern completion. Here we tested the ability of sparse two-layer associative networks to learn to 1) recognize patterns in real world data and 2) use memory content to reconstruct noisy input in a context-sensitive fashion. Specifically, we used the memory models with supervised learning on handwritten characters (NIST database [3]) and assessed how well the memory could predict the class of unknown input and reconstruct the input pattern. We started with the Kanerva memory model [2], which has a two-layer structure. The first neural layer has fixed random synapses and contains many more neurons than input fibers, which are sparsely activated. This stage maps the input space into a high-dimensional space of sparse patterns. The synapses of second layer neurons store associations between the high-dimensional sparse patterns and desired output patterns via Hebbian plasticity. We trained the model with the NIST data by storing for each input pattern the given digit interpretation as well as the input pattern in autoassociative fashion. When cross-validated with (noisy) inputs, the model frequently performed pattern recognition and reconstruction successfully. However, ambiguous inputs could not be classified, and the reconstruction was a mixture pattern not corresponding to a valid handwritten digit. To overcome these limitations, we constructed a new model with a two-layer structure similar to the described Kanerva model but with two important new features, enabling soft classification and pattern completion in two subsequent phases. One new feature is labeling first layer neurons during supervised training according to how their activity is correlated with the occurrence of certain classes. Thus, in the classification phase when a new input drives first layer neurons, estimates of class membership are encoded in active populations of first layer cells. The second feature is a phase of selective pattern completion performed after classification. In this phase first layer neurons are also influenced by lateral interactions and feedback from the second layer. This competitive dynamical process accomplishes pattern completion that is contingent on a specific interpretation rather than creating a mixture of the entire memory content, even if the input is ambiguous. We compare the model to current methods of classifying the NIST data. We found that the model exhibits competitive classification performance and adds a valuable explanatory component not provided by ordinary classification algorithms. Finally, we discuss how the proposed memory architecture would map onto dentate gyrus and CA3 of the hippocampus. 1. Hopfield, J., PNAS, 1982. 9:p.2554. 2. Kanerva, P., Sparse Distributed Memory. MIT Press. 1988. 3. LeCun, Y., C. Cortes, <http://yann.lecun.com/exdb/mnist>. 4. Rumelhart, D.E., et al., Nature, 1986. 323:p.533. 5. Willshaw, D.J., et al., Nature, 1969. 222:p.960.

doi:

III-33. Spike-timing theory of working memory

Botond Szatmáry^{1,2}
Eugene M. Izhikevich^{1,2}

BOTOND.SZATMARY@BRAIN CORPORATION.COM
EUGENE.IZHKEVICH@BRAIN CORPORATION.COM

¹The Neurosciences Institute, San Diego, CA

²Brain Corporation, San Diego, CA

Working memory (WM) is part of the brain's memory system that provides temporary storage and manipulation of information necessary for cognition. Although WM has limited capacity at any given time, it has vast memory content in the sense that it acts on the brain's nearly infinite repertoire of lifetime memories. Existing models, however, fail to explain how WM functionality emerges in the brain's vast memory content. We show that large memory content and WM functionality emerge spontaneously if we take the spike-timing nature of neuronal processing into account. This is in contrast with previously suggested mechanisms of WM, where spike-timing is ignored and the models' explanatory power is limited to systems having small repertoires of long-term memories represented by, e.g., carefully selected non-overlapping populations of neurons. In our model, memories are represented by extensively overlapping groups of neurons that exhibit stereotypical time-locked spatio-temporal spike-timing patterns, called polychronous patterns. This mechanism, for example, allows for a set of neurons with two distinct patterns of synaptic connections with appropriate axonal conduction delays to form two distinct

polychronous neuronal groups (PNGs). In other words, PNGs are defined by distinct sets of synapses, and not by the neurons per se, which allows neurons to take part in multiple PNGs and enables the same set of neurons to generate several different distinct stereotypical precise spatio-temporal spike-timing patterns. PNGs arise spontaneously in simulated realistic cortical spiking networks shaped by spike-timing dependent plasticity (STDP). In our model, these polychronous patterns are the basis for the large memory content, and the activation of a PNG is the underlying currency of information. Activation of PNGs in spiking networks happens spontaneously due to stochastic synaptic noise. These reactivations, however, can be biased by short-term changes in synaptic efficacies, which, in our model, are implemented in the form of short-term STDP, where short-term synaptic changes depend on the conjunction of pre- and post-synaptic activity. Using simulations, we show how such associative synaptic plasticity can select externally cued PNGs into WM by temporarily strengthening the synapses of the selected PNGs: This strengthening increases the spontaneous reactivation frequency of the selected PNGs, resulting in irregular, yet systematically changing elevated firing activity patterns of intra-PNG neurons, consistent with those recorded in vivo during WM tasks. Note that despite the fact that PNGs share neurons amongst each other, activity of one PNG does not spread to the others, therefore, frequent reactivation of a selected PNG does not initiate uncontrollable activity in the network. Hence, our WM mechanism can work in a network with large memory content. Our theory explains the relationship between precise spikes and slowly changing firing rates of neurons engaged in active maintenance of WM, and it points to the connection between WM and perception of elapsed time on the order of seconds. It also predicts that polychronous structures are essential for cognitive functions like WM, and such structures may be the basis for memory replays involving, for example, prefrontal cortex, visual cortex, and hippocampus.

doi:

III-34. Cue-based feedback enables remapping in a multiple oscillator model of place cell activity

Joseph D. Monaco¹

Kechen Zhang²

Hugh T. Blair^{3,4}

James J. Knierim¹

JMONACO@JHU.EDU

KZHANG4@JHMI.EDU

BLAIR@PSYCH.UCLA.EDU

JKNIERIM@JHU.EDU

¹Krieger Mind/Brain Institute, Johns Hopkins

²Biomedical Engineering, Johns Hopkins

³Psychology, UCLA

⁴Brain Research Institute, UCLA

Place fields in rat hippocampus consist of both a firing-rate component [6] and a temporal component defined by spike-phase precession relative to local theta [7]. Previous models based on oscillatory phase interference [e.g., 5] can account for phase precession, but not for the remapping that can occur when an animal is exposed to novel spatial information. A novel room may elicit complete remapping in which the population spatial code becomes statistically independent. However, subtler cue manipulations can induce partial remapping and other more graded spatial recoding effects in which some degree of coherence with previous representations is retained. Double-rotation experiments, in which sets of local and distal cues are rotated relative to each other around a circular track, have shown that activity in the CA3 subregion is significantly more coherent than in CA1 [4]. Thus, it is critical to our understanding of hippocampal function to have models of spatial coding that can explain graded remapping as well as all-or-none complete remapping. While somato-dendritic dual-oscillator models have been examined closely [5], it is not clear how to couple them with environmental cues to explore these sorts of effects. We demonstrate a more recent generalization of oscillatory interference models featuring multiple oscillator inputs [1]. Each oscillator's phase is modulated by the velocity vector of the trajectory such that the population phase code provides stable path integration. First, we show that arbitrarily connected output units can produce spatially-modulated activity. Further, due to the combinatoric rarity of synchronizing N oscillators, there is a lower bound on the number of theta inputs to achieve sparse responses. Second, we demonstrate a cue-based phase-code feedback that represents learned fixed-points of the trajectory. This makes spatial representations robust to noise,

but also allows cue manipulations similar to double-rotation experiments. Simulating double-rotation using actual trajectories, we found that the diversity of remapping behavior among the output population depended on the number of cues, the feedback gain and the relative contributions of path integration and phase-code feedback. We found a diversity of both cue-following and ambiguous outputs qualitatively similar to the experimental data using moderate overall feedback gain and a small number of low-spatial-extent cues. Recent intracellular recordings of place cells demonstrated increased theta power within-field and intracellular phase precession relative to extracellular theta [3], both of which result from this model. Notably, this model enables complete remapping with a phase reset of the sort that may occur upon introduction to a novel environment and provides a possible common-input explanation for the concurrency of hippocampal remapping and entorhinal grid realignment [2]. Thus, the multiple oscillator model provides insight into phase code mechanisms that may underlie a wide array of rate and temporal coding effects and remapping phenomena in hippocampus. [1] Blair, Zhang. (2009). SfN Abstract, 192.27; [2] Fyhn, Hafting, Treves, Moser, Moser. (2007). *Nature*, 446(7132):190-194; [3] Harvey, Collman, Dombeck, Tank. (2009). *Nature*, 461(7266):941-946; [4] Lee, Yoganarasimha, Rao, Knierim. (2004). *Nature*, 430(6998):456-459; [5] O'Keefe, Burgess. (2005). *Hippocampus*, 15:853-866; [6] O'Keefe, Dostrovsky. (1971). *Brain Research*, 34(1):171-175; [7] O'Keefe, Recce. (1993). *Hippocampus*, 3(3):317-330.

doi:

III-35. Independent snapshot memories in hippocampus: Representation of touch- and sound-guided behavior.

Pavel M. Pavel¹
Ekaterina Vinnik¹
Christian Honey²
Jan Schnupp²
Mathew E Diamond¹

ITSKOVPA@GMAIL.COM
LULSWINNICK@GMAIL.COM
CHRISTIAN.HONEY@GMAIL.COM
JAN@PHYSIOL.OX.AC.UK
DIAMOND@SISSA.IT

¹SISSA

²Oxford University

Understanding the mechanisms by which sensory experiences are stored is a longstanding challenge for neuroscience. Up to date there is no evidence about how neurons represent the behavioral significance of tactile stimuli, or how tactile events are encoded in memory. We have previously shown that different aspects of a tactile categorization task are readily encoded by hippocampal neurons. The focus of the current work was to characterize interactions between object-related coding location-related coding. To investigate these issues, we recorded single-unit firing and local field potentials from the CA1 region of hippocampus while rats performed a task in which tactile stimuli informed the rat of reward location. On each trial, the rat touched a textured plate with its whiskers and then chose to turn either left or right to obtain a reward. Two textures were associated with each reward location. Over one-third of the 750 sampled neurons encoded the identity of the texture: their firing differed for two stimuli associated with the same reward location. Around 80 percent of the neurons encoded the behavioral significance of the contacted textures: their firing differed according to the reward location with which the stimulus was associated. The responses of neurons to a given stimulus in different locations were independent. This was not the case for reward location signals: neurons that carried a signal in one location were more likely to carry a signal in the other location. Of the neurons encoding reward location in both locations, half represented reward location in allocentric coordinates and half in the egocentric coordinates. Although some neurons encoded both reward location and texture, the number of these neurons did not exceed that expected from 2 partially overlapping but independent populations of neurons. Additional experiments were carried out, on another set of rats, to generalize some of the above findings from the tactile to the auditory modality. On each trial, the rat leaned into the gap and heard one of four sounds which were distributed along a vowel continuum from "A" to "I". Two sounds were associated with each reward location, and the experiment was repeated in 2 locations. As in the tactile task, more than 80 percent of neurons represented reward location and more than 20 percent of neurons represented the identity of the sound (the vowel). The role of context on the stimulus and reward location signals was the same as in the tactile experiments. As in the tactile task representations of sounds were

independent across 2 platforms and there was no systematic relationship between the representation of sound and reward location. Responses to the same sound were tested during passive listening. About 10 percent of hippocampal neurons distinguished between the sounds although their coding and sound preference could not be predicted from their response during the discrimination task. The results from the tactile and auditory experiments converge on the idea that each feature (stimulus, or place) is represented as a snapshot that engages a neuronal population independently of other snapshots.

[doi:](#)

III-36. Contextual information for maintaining coherent egocentric-allocentric maps

colin Jimenez Rezende*
 Danilo Molter*
 Wulfram Gerstner
 EPFL - LCN

REZENDE.DANILO@EPFL.CH
 COLIN.MOLTER@EPFL.CH
 WULFRAM.GERSTNER@EPFL.CH

To navigate in an environment without relying on landmarks based response strategies, the central nervous system must be able to integrate the sensory information perceived in an egocentric coordinate system and to transform it into an allocentric coordinate system. Neurophysiological evidences demonstrated that the allocentric coordinate system is maintained by hippocampal place cells, firing at specific orientation-independent locations in an environment, and by head direction cells, firing according to specific directional headings. The computation underlying this egocentric-allocentric transformation has been recently enlightened by the observation that entorhinal cortical cells (EC cells), located one layer upstream hippocampal place cells, fire at multiple locations in the environment, these locations forming a regular hexagonal grid. Different computational models have been proposed to explain how hippocampal place fields can emerge from the conjunctive activity of entorhinal grid fields. These models successfully explain how an environment can be transformed into an allocentric coordinate system by relying only on idiothetic cues. However, when moved from one environment to another, grid fields are realigned and reoriented leading to incoherent place cells activity. Facing that situation, the entorhino-hippocampal circuitry should first decide the novelty of the environment. Then, if the environment is recognized, the grid fields should be reset to the configuration leading to the previously learned allocentric system. To address this problem, idiothetic cues alone are not sufficient. Here, based on neuroanatomical and neurophysiological data, we propose a minimal computational model of the egocentric-allocentric transformation which relies both on idiothetic cues and on preprocessed visual information. During movement, preprocessed visual information is gathered by layer III EC cells while grid cells firing activity in ECII is mainly mediated by idiothetic cues and head direction cells (see Supplemental for a detailed Figure). During exploration the hebbian learning of connections between grid cells and CA1 cells lead to the emergence of CA1 place fields. Simultaneously, backward connections (simulating the CA1-parasubiculum-ECII connections) are learned using the same hebbian rule and ECIII-CA1 connections are learned providing associations between views and place fields. Similarly, head direction cells activity is updated by idiothetic cues while associations between these cells and preprocessed views are learned. We hypothesized the existence of 'context cells' located in the ventral hippocampus that encode environment specific information and that can be rapidly triggered based on limited visual information. When grids are randomly realigned and reoriented, based on visual information, context cells can either recognize the environment or can launch a novelty signal. In the former case, context cells can reorient the grid fields. Based on same visual information, ECIII cells can reactivate the place cells coding for the rat's location which in turn successfully reset the position of the grid cells activity. As a consequence, when tested on pre-learned environment, our model is able to successfully maintain a correct reference frame for the egocentric-allocentric transformation.

[doi:](#)

III-37. Attention selects informative populations

Preeti Verghese
Alexander Wade

Smith Kettlewell Eye Research Institute

PREETI@SKI.ORG
WADE@WADELAB.NET

Motivation: Ideal observer theory predicts that observers use the most informative neural population for any particular task and that the identity of this neural population differs for different tasks. For orientation discrimination, Fischer information is highest in populations that have a preferred orientation tilted away from the target orientation as these neurons have orientation tuning functions that are most sensitive (steepest) in the region of the target orientation. For contrast discrimination, information is highest at the target orientation although the tuning is expected to be much weaker. Here we used source-imaged EEG to measure neural responses in humans while they performed two different tasks on the same stimuli. For each task, we asked whether our subjects attended to the most informative neural population as theory predicts. **Methods:** Observers performed either a contrast or an orientation discrimination task on one of two static, vertical targets, located 5° to the left and right of fixation. A cue indicated both the task type (contrast or orientation) and the location of the change (left or right). In order to probe the responses of neurons that were modulated attentionally but not driven directly by the target, we used frequency-tagged gratings within annuli that surrounded the target and flickered at 15 and 20 Hz on the left and right sides. These annuli were present throughout each 2-second trial. We used the frequency-tagged response amplitudes generated by these gratings to estimate the amplitude of the attentional modulation in the surrounded target region. Theory predicts that attention should preferentially modulate neurons tuned 20° away from the target in the orientation discrimination task, but not the contrast discrimination task. To test this, we used two different annulus grating orientations: one with the same orientation as the central, vertical target, and the other tilted 20° away. For each observer we collected high-density EEG data while they performed these tasks. We then used a minimum norm inverse procedure combined with realistic MR-derived head models and retinotopically-mapped visual areas to estimate cortical activity due to the grating annuli. To quantify the attentional modulation in each subject's primary visual cortex (V1) we computed an attention modulation index (AMI): the difference between the attended response and unattended response normalized by the unattended response. **Results:** Attention to a spatial location clearly increased the amplitude of the response to the annulus surrounding that location. We observed these spatially-dependent amplitude modulations in both the contrast and orientation discrimination tasks. More importantly, the pattern of modulation depended on the task. For contrast discrimination both the 0° and 20° neural populations showed similar enhancements in area V1 consistent with the prediction that orientation tuning is weak in this task. For orientation discrimination, only the 20° population exhibited attentional modulation. These findings indicate that humans are able to attend selectively to the most informative neural population even at the level of primary visual cortex and that these populations change depending on the nature of the psychophysical task.

[doi:](#)

III-38. How well does local neuronal activity predict cortical hemodynamics?

Yevgeniy B. Sirotnin
Aniruddha Das

Columbia University, Dept of Neuroscience

YSIROTNIN@GMAIL.COM
AD2069@COLUMBIA.EDU

We had earlier reported a hitherto-unknown hemodynamic signal in primary visual cortex (V1) of alert monkeys performing a periodic visual task. This novel signal entrains to anticipated task onsets but is independent of visual input or any measurable local neuronal activity. The overall hemodynamic response to periodic visual tasks combines this trial-related signal with the more familiar visually-evoked response (Sirotnin & Das, Nature 2009). Here we follow up with three questions further analyzing the physiological underpinnings of trial-related and visually-evoked hemodynamics in alert behaving animals. They are: 1: Do the visually-evoked and trial-related hemodynamic signals add linearly? 2: How linear is the relationship between visually-evoked hemodynamics and neural

activity? 3: Is it possible to identify ongoing hemodynamic fluctuations outside the context of a periodic task - in particular, in a 'resting state'? How do such ongoing fluctuations relate to neural activity? To answer these questions we used our published technique of intrinsic-signal optical imaging with simultaneous electrode recordings in V1 of alert behaving macaques. The animals either performed a periodic fixation task with visual stimulation; or alternately, they remained in a state of restful alertness in a dimly lit or dark room with no visual stimulus. 1: In tasks with visual stimulation, we found that the neuronally-predicted and trial-related signals added linearly with each other. The trial-related signals accompanying trials with visual stimulation closely matched the trial-related signal that we had found earlier for periodic fixation tasks undertaken in total darkness: using a hemodynamic kernel to regress away 'neurally-predictable' hemodynamic responses in visual tasks revealed a residual hemodynamic signal essentially identical to the trial-related signals found earlier in the dark room. 2: We found that the stimulus-evoked components of neural and hemodynamic signals were linearly related to each other over the entire dynamic range of stimulus contrast used (5 log units of contrast of a full-field drifting grating). The correlation with hemodynamics was strongest for spiking (multi-unit activity: MUA) followed by high-gamma (66-130 Hz) local field potentials (LFP). 3: With the animals in a state of alert restfulness we found robust hemodynamic signals in the 'resting state' V1. These signals were very poorly related, however, to neuronal activity. By using multivariate linear regression where we combined, as independent regressors, the simultaneously recorded MUA and multiple LFP bands (ranging from 2 Hz to 130 Hz) we could explain only a small fraction of the hemodynamic signal. The MUA accounted for the largest fraction of the hemodynamic signal while low-frequency LFP (delta: 2-4 Hz) accounted for very little. We found that fluctuations in the heart rate independently accounted for an additional fraction of the hemodynamic signal that was comparable with the fraction related to the MUA. But even on including heart rate fluctuations in our multivariate analysis we could account for a maximum of ~50% of the resting state hemodynamic fluctuations. This suggests that a significant portion of the ongoing hemodynamic signals may be driven by neuromodulatory or other distal inputs that are not reflected in local neuronal activity.

doi:

III-39. Decrease in synaptic variance improves perceptual ability

Robert C. Froemke
Michael M. Merzenich
Christoph E. Schreiner

RFROEMKE@PHY.UCSF.EDU
MERZ@UCSF.PHY.EDU
CHRIS@PHY.UCSF.EDU

University of California, San Francisco

Receptive fields in the adult nervous system are highly structured. In the primary auditory cortex (AI), neurons are tuned to both sound frequency (pitch) and intensity (loudness), and these tuning properties are largely determined by the organization of synaptic inputs onto AI neurons. Cortical synapses are highly plastic, even in the adult brain, and thus receptive fields and the functions that they support have the potential to be rapidly modified to improve perception and behavioral performance. Here we studied the functional consequences of AI synaptic receptive field plasticity, combining in vivo whole-cell recording with behavioral training in adult rats. We focused on intensity tuning of excitatory inputs, as previously we have examined frequency tuning of inhibitory inputs (Froemke et al., *Nature* 2007), and asked how changes in AI synaptic intensity tuning would affect auditory discrimination behavior. As a naturalistic method for modifying synaptic receptive fields, we paired specific auditory stimuli with activation of the cholinergic basal forebrain (Goard and Dan, *Nat. Neurosci.* 2009), simulating the effects of heightened arousal or directed attention to behaviorally-relevant percepts. Whole-cell voltage-clamp recordings were made in vivo from AI neurons in anesthetized animals, while basal forebrain activation in behaving animals was performed with a custom-built neuroprosthetic device (Froemke et al., *Keck Futures Initiatives* 2007). We first asked if AI neurons and networks could be retuned to prefer low-intensity tones. We paired basal forebrain stimulation with presentation of relatively quiet pure tones, usually 30-50 dB from best level and 1-3 octaves from best frequency. We measured frequency-intensity synaptic receptive fields before and after pairing, and found that responses at paired stimuli were greatly enhanced, while responses at the original best stimuli were reduced. This specific reduction in responses to the original best stimuli were activity-dependent, such that for a period of approximately 10-20 minutes following pairing, whichever stimuli evoked the largest excitatory responses were selectively depressed, via a mechanism associated with Ca²⁺ release from intracellular stores.

These synaptic modifications preserved the net excitatory drive onto AI neurons. The increase in response at the paired stimuli was precisely matched by the decrease in response at the original best stimuli. As a consequence, average synaptic strength was unaltered by basal forebrain pairing. However, as large responses became smaller, and small responses became larger, the variance of synaptic receptive field responses decreased. Information theoretic analysis suggested that this decrease in variability might improve the detection of weak signals such as lower-intensity sounds, at the expense of discrimination of the pitch identity for higher-intensity sounds. Finally, we tested and confirmed this hypothesis by training a separate set of adult rats to detect and discriminate between several pure tones of different frequencies and intensities. In summary, basal forebrain pairing sets in motion a dynamic set of changes to AI synaptic receptive fields, acting to enhance responses to paired stimuli while decreasing originally larger responses. These changes preserve excitatory drive and reduce response variance, improving behavioral performance for some perceptual abilities, but sometimes at the expense of performance for other tasks.

doi:

III-40. Spike-based Expectation Maximization

Bernhard Nessler
Michael Pfeiffer
Wolfgang Maass

NESSLER@IGI.TUGRAZ.AT
PFEIFFER@IGI.TUGRAZ.AT
MAASS@IGI.TUGRAZ.AT

Graz University for Technology

There exists solid experimental evidence that synaptic weights are subject to spike-timing-dependent plasticity (STDP). However it is not clear how STDP contributes to the emergence and maintenance of powerful computations in network of neurons. We report a surprising theoretically founded connection between STDP in the context of an ubiquitous motif of cortical networks of neurons, winner-take-all (WTA) circuits [Douglas et al., *Annu. Rev. Neurosci.*, 27:419-451, 2004], and theoretically optimal methods for unsupervised learning. More specifically, we prove that STDP approximates in such context stochastic online expectation maximization (EM) for the discovery of hidden causes for complex input patterns. More precisely, each application of STDP to a neuron on the competitive layer that fires can be understood as an approximation of the M-step in stochastic online EM. The E-step consists simply of the application of the WTA-circuit with the resulting slightly changed competitive balance to the next spike inputs. At the heart of this theoretical approach lies the observation that for certain forms of STDP the weights converge to values that can be interpreted as log of the conditional probability that the presynaptic neuron has fired just before, given that the postsynaptic neuron has fired. This principle provides a direct link between STDP and EM. On the basis of this principle one can achieve in computer simulations of STDP in networks of spiking neurons surprising network learning results. For example, we demonstrate that they can learn without supervision to discriminate handwritten digits after having seen a few thousand examples from the MNIST database (transformed into high-dimensional spike patterns through standard population coding), and to detect and discriminate repeating spatio-temporal patterns of spikes within continuous high-dimensional spike input streams. Furthermore STDP induces in the weight vectors internal models for characteristic prototypes of input patterns, such as prototypes of handwritten digits, as predicted by our theoretical analysis. Our results also show that STDP is able to learn an optimal bayesian inference if both inputs and outputs are realized as probabilistic population codes [Ma et al., *Nature* 6:11:1432-1438, 2006]. Our theoretical framework predicts that unsupervised learning with STDP works best if weight increases are additive, but depend negative exponentially on the current value of the weight. Furthermore it predicts that the size of weight decreases is independent of the current weight size. Both predictions have been confirmed by experimental data (see Fig. 1 in [Montgomery et al., *Neuron* 29: 691-701, 2001] and Fig. 5c as well as the text on p. 1153 of [Sjöström et al., *Neuron* 32:1149-1164 2001]).

doi:

III-41. Evaluation of memories through synaptic tagging

Marc Papper¹
Richard Kempter²
Christian Leibold³

MARCPAEPPER@GOOGLEMAIL.COM
R.KEMPTER@BIOLOGIE.HU-BERLIN.DE
LEIBOLD@BIO.LMU.DE

¹University of Munich

²Humboldt-Universitat zu Berlin

³Ludwig-Maximilians-Universitat, Munchen

Long-term stability of synaptic changes requires the synthesis of proteins in the cell soma. Synaptic tagging has been suggested as a cellular mechanism to achieve presynaptic specificity of somatic protein synthesis, as required by learning theories (Frey and Morris, 1997). The presynaptic activation is thereby assumed to tag the synapse such that it is able to capture plasticity-related proteins whereas untagged synapses cannot capture them. Here we provide a computational model to test the tagging hypothesis on its potential to provide reasonable memory lifetimes in an online learning paradigm. The model is based on a discrete linear dynamics of synaptic state distributions as proposed by Amit and Fusi 1994 (see also Leibold and Kempter 2008). We find that tagging is only beneficial if it is used to evaluate memories, in which important ones evoke protein synthesis and unimportant ones do not. We find that then our tagging model exhibits a volatile equilibrium configuration where most synapses are in a neutral state that can easily be modified, and only the comparatively few synapses that are essential for the retrieval of the important memories are more stable against plasticity stimuli. The model also provides a parameter regime in which it solves the distal reward problem where the initial exposure of a memory item and its evaluation are temporally separated. We furthermore derive optimal transition probabilities from the stable synaptic states and show that those have to be at least as stable in time as we require the memories to be. References: Frey U, Morris RGM (1997) *Nature* 385:533. Amit D, Fusi S (1994) *Neural Comput* 6:957. Leibold C, Kempter R (2008) *Cereb. Cortex* 18:67.

[doi:](#)

III-42. Optimal architectures for fast-learning, flexible networks

Vladimir Itskov¹
Anda Degeratu²
Carina Curto¹

VITSKOV2@MATH.UNL.EDU
DEGERATU@AEI.MPG.DE
CCURTO2@MATH.UNL.EDU

¹University of Nebraska-Lincoln

²Max Planck Institute, Potsdam

New memories (patterns) in some brain areas, such as hippocampus, can be encoded quickly. Irrespective of the plasticity mechanism (or learning rule) used to encode patterns via changes in synaptic weights, rapid learning is perhaps most easily accomplished if new patterns can be learned via only small modifications of the initial synaptic weights. It may thus be desirable for fast-learning and flexible neural networks to have architectures which enable large numbers of patterns to be encoded by only small perturbations of the synaptic efficacies. What kinds of network architectures have this property? We define the perturbative capacity of a network to be the number of memory patterns that can be learned under small perturbations of the (effective) synaptic weights. We propose that candidate architectures for fast-learning, flexible networks should be networks with high perturbative capacity. What are optimal architectures that maximize a network's perturbative capacity? We investigate this question for threshold-linear networks. Here the memory patterns encoded by the recurrent network are groups of neurons that co-fire at a stable fixed point for some external input. We prove that for an arbitrary matrix of effective synaptic weights, the network's memory patterns correspond to stable submatrices. This enables us to study the perturbative capacity of a network by analyzing the eigenvalues of submatrices under small perturbations of the synaptic weights. In the case of symmetric threshold-linear networks, we find (analytically) the set of optimal network architectures that have maximal perturbative capacity. For these networks, any of the possible memory patterns can be selectively encoded via small perturbations of the synaptic weights. We show that these architec-

tures correspond to a highly restricted set of possible sign patterns governing the effective interactions between principal neurons, and we completely describe these patterns. In particular, we find that at least one half of the effective weights must be inhibitory, and thus the optimal architectures reflect inhibition-stabilized networks with a significant level of inhibition. Finally, we study a larger class of threshold-linear networks (where the matrices are no longer assumed symmetric), and find that our qualitative results continue to hold. The optimal architectures we discover provide a benchmark for comparison to experimentally obtainable estimates of the effective interactions between pyramidal neurons in fast-learning networks. They also give clues as to the differences we might expect between recurrent networks in areas such as the hippocampus, and more rigid networks in areas such as primary sensory cortices, where it may be undesirable for the allowed response patterns to be sensitive to small perturbations in synaptic weights.

[doi:](#)

III-43. Disconnection of monkey orbitofrontal and rhinal cortex impairs assessment of motivational value

Andrew M. Clark
Sebastien Bouret
Elisabeth A. Murray
Barry J. Richmond

CLARKAND@MAIL.NIH.GOV
BOURETS@MAIL.NIH.GOV
MURRAYE@MAIL.NIH.GOV
BARRYRICHMOND@MAIL.NIH.GOV

Laboratory of Neuropsychology, NIMH, NIH

Associating stimuli with the rewards that they signal and appropriately adjusting motivation in proportion to the desirability of these outcomes is an important component of goal directed behavior. In Rhesus monkeys, it has previously been shown that bilateral lesions of two structures commonly thought to play a role in either reward valuation - orbitofrontal cortex (OFC) - or multi-modal associations - rhinal cortex (Rh) - disrupt monkeys' ability to adjust their motivation in response to cue-reward associations. Given the strong reciprocal connections between these areas, we hypothesized that network level interaction of OFC and Rh is necessary for assessing the motivational value of visual stimuli. Accordingly, we examined the impact of crossed lesions of OFC and Rh on monkeys' behavior in a task in which they learned to associate a visual cue with the amount of reward earned for performing an instrumental action. On each trial, animals earned a variable amount of liquid reward for releasing a lever after a change in the color of a visual target; trials began with a visual stimulus that signaled the magnitude of the forthcoming reward. Consistent with previous results, we found that monkeys' adjusted their performance as a function of reward size and accumulated reward; error rates decreased hyperbolically with increasing reward size and increased logarithmically with an increase in accumulated reward. This suggests animals were less motivated to work both for smaller rewards and after having consumed more rewards. Monkeys were trained on 4 sets of 4 cue-reward pairings before we performed a unilateral aspiration lesion of either OFC or Rh. After receiving this first stage of the crossed lesion monkeys were tested on the last pre-operative cue set (old) and a new cue set (new); we then performed the remaining lesion in the contralateral hemisphere prior to a further round of old/new testing. We used ANOVA to test for effects of reward size, accumulated reward, treatment and testing order, estimating significance at the $p < 0.05$ level. There were inconsistent effects of the unilateral first stage lesions (2 out of 4 animals, OFC first stage). After receiving the first stage OFC lesion both animals showed a further impairment after the second stage Rh ablation, with the effect of the crossed lesion being greater for the old (tested first) versus the new cue set. To determine whether this was due to an effect of the crossed lesion on retention as opposed to new learning, rather than an effect of testing order, we conducted a further round of old/new testing. We observed a significant interaction between testing order and old/new status, suggesting a gradual post-lesion recovery rather than a specific retention deficit. We were able to reinstate the deficit by simply doubling the number of cue-reward pairings, even several months post-lesion, suggesting that the initial impairment was a specific effect of the treatment. In sum, our results thus far indicate that the interaction between OFC and Rh is necessary to form the associations required for normal assessment of the motivational value of visual stimuli.

[doi:](#)

III-44. Reward-modulated spike timing-dependent plasticity requires a reward-prediction system

Nicolas Frémaux
Henning Sprekeler
Wulfram Gerstner
LCN, EPFL

NICOLAS.FREMAUX@EPFL.CH
HENNING.SPREKELER@EPFL.CH
WULFRAM.GERSTNER@EPFL.CH

Spike-timing dependent plasticity (STDP) has been shown to perform unsupervised learning tasks such as receptive fields development. However, STDP fails to take behavioral relevance into account, and as such cannot learn behavioral tasks. Recent publications have suggested to extend STDP by conditioning the induction of plasticity on the contingency of the pre-post pairing of classical STDP with a neuromodulatory "reward signal". We call this reward-modulated STDP model R-STDP (Izhikevich, 2007; Legenstein et al., 2008). In this study, we show that R-STDP (and, in general, any learning rule based on such a combination of an unsupervised learning rule with reward-modulation) suffers from a bias problem, which will in most cases impede learning. This problem can be solved only if the average of the reward signal is zero for each task. To learn a response to a single stimulus, it is enough to subtract a baseline reward value (the mean). However, in the case where the post-synaptic neurons have to learn different tasks for different input stimuli, this requires an external reward-prediction system (or "critic"), to subtract the stimulus-dependent or task-dependent baseline from the reward. Reward-modulated learning rules derived analytically from the policy-gradient framework of reinforcement learning (Pfister et al., 2006; Florian, 2007; Baras and Meir, 2007), called R-max in the following, don't suffer from this bias problem. We illustrate our findings with two learning paradigms. First, we teach a simulated 1-layer feed-forward network to respond to a 1 second input pattern of spike trains with a precise pattern of output spike trains. At the end of each trial, a scalar reward signal is broadcast to the network, representing how well the output spike train matched the target, with a running average subtracted so that the average reward is zero. Both R-STDP and R-max can learn the task. However, as soon as the average reward is not zero or that two or more spike-train response tasks have to be learned, R-STDP fails spectacularly, whereas R-max suffers only a modest decrease in performance. Second, we learn a hand movement trajectory controlled by spiking neurons using a population vector coding similar to that found in motor areas (Schwartz et al., 1988). We use the same network structure as before, albeit with more units. Each output neuron codes for a particular direction of motion in 3 dimensions. The output spike trains produced in a single trial are transformed into a motion sequence through population vector coding. The reward signal is calculated by comparing the motion produced by the network to a target motion, and subtracting a baseline reward. Again, both learning rules can learn the task, but R-STDP fails to learn more than one different motion without subtracting a task-dependent baseline from the reward signal. In summary, we show that to learn behavioral tasks, reward-modulated STDP needs a reward-prediction system that can infer future expected rewards from current stimuli. This is a strong requirement, but dopaminergic neurons in the primate VTA could be plausible candidates (Schultz, 2007).

[doi:](#)

III-45. Tagging and capture: a bridge from molecular to behavior

Lorric Ziegler
Wulfram Gerstner
EPFL - LCN

LORRIC.ZIEGLER@EPFL.CH
WULFRAM.GERSTNER@EPFL.CH

Synaptic plasticity is widely accepted as a substrate for learning and memory. The maintenance of memory traces over extended periods depends upon a late phase of plasticity which exhibits complex dynamics. Tagging and capture has been thoroughly analyzed in rat hippocampal slices (Frey & Morris, 1997). It is part of the late phase of plasticity and determines what synapses are going to be maintained among those that have undergone potentiation or depression during the induction phase. This mechanism depends on protein synthesis, triggered via neuromodulators (here dopamine), and is mainly thought to be non local. As a consequence associative phenom-

ena can be observed where tagged synapses take advantage of other sources of protein production occurring in a certain time window. A behavioral analog has recently been studied (Moncada & Viola, 2007). An inhibitory avoidance training paradigm was used to investigate long term memory in rats and its consolidation either by strong stimulus or by exploration of a novel environment. This was shown to depend on protein production in dorsal hippocampus. We exploit this striking similarity by building a simple neural network implementing fear memory formation and its effect on action triggering process in rats. We use TagTriC (Ziegler et al., 2008), a model of induction as well as maintenance of synaptic plasticity, as learning rule in parts of the connectivity scheme and observe its action at the system level. This allows to refine the link between long term potentiation and long term memory, hence nicely connecting behavioral observables and their molecular substrates.

[doi:](#)

III-46. Predicting the task specificity of learning

Jacqueline M. Fulvio

Paul Schrater

University of Minnesota

FULVJ002@UMN.EDU

SCHRATER@UMN.EDU

A recurrent finding in the field of perceptual learning is that the learning is often highly-specific to the trained task and stimuli. Using the language of reinforcement learning, such results would suggest that subjects have learned stimulus-specific policies rather than more general world models. Recently, two bases for action selection have emerged from reinforcement learning: (i) policies and value functions computed from experience; (ii) planning via predictive models of future outcomes. The former strategy is model-free, but suffers from inflexibility as it is only useful for determining actions from states that have been experienced. The latter strategy uses a model to compute a policy on the fly, which is computationally-intensive and subject to the propagation of errors due to model inaccuracy, but allows for simulation of outcomes over a range of conditions, and thus transfers more broadly. We hypothesize that the brain will adopt the learning strategy that is expected to maximize performance in a particular domain. The current study uses a motion extrapolation task. An object moves along a path of variable curvature before disappearing and a prediction of where the object will reemerge at a specified distance beyond the point of occlusion is made. In the absence of feedback, we find that subjects naturally adopt one of two predictive process models in carrying out the task: a constant acceleration model or a constant velocity model depending on the subject's acceleration sensitivity. Kalman filter simulations of subject performance revealed that the choice of model was determined by a transition in relative model performance as a function of the likelihood of the model fit to the curve path information and the model's predicted extrapolation. We also provided feedback to test whether a similar transition occurs between policy and predictive model learning as a function of the relative performance between the strategies. The number of individual state action pairs that need to be learned in order to perform the task was hypothesized to play a role in this transition. Few state-action pairs means the learner will experience many exposures to the states, which enhances the reliability of the policies. As the number of state-action pairs increases, the number of exposures decreases, reducing the reliability of the policies. Furthermore, the computational load in policy retention increases as the number of pairs increases. Thus, the transition was expected to occur at the point at which policies become less reliable than generative model predictions. Simulations of extrapolation performance under Q-learning and the Kalman filter on action selection in the motion extrapolation task quantified these predictions. As the number of policies needed for successful extrapolation increased, extrapolation performance by the Kalman filter model is more reliable than the policies learned via Q-learning. Combined with human data, these results contribute to evidence that performance criteria determine learning strategies, specifically those that optimize the reliability of performance.

[doi:](#)

III-47. Eye position modulation of visual responses in the lateral intraparietal area lags the eye movement

Yixing Xu¹
Carine Karachi²
Michael Goldberg³

YIXING.XU@GMAIL.COM
CARINE.KARACHI@GMAIL.COM
MEG2008@COLUMBIA.EDU

¹Columbia University

²CR-ICM U679 Pitie Salpetriere Hospital

³Dept. of Neuroscience, Columbia University

Gain fields, eye position modulated neuronal responses to behaviorally relevant visual stimuli, are commonly found in cells of the posterior parietal cortex (Andersen and Mountcastle, 1983; Andersen et al., 1990). Neurons in areas 7a and LIP are thought to facilitate the transformation of retinal and eye position information into head-centered coordinates (Zipser and Andersen, 1988). Head-centered neurons are noticeably absent in these brain areas, so the gain fields fulfill an important role in the generation of a stable, body-centric coordinate frame for performing visually guided movements. The gain fields are thought to derive their eye position information from corollary discharge (Andersen and Mountcastle, 1983), and coordinate transformation models implicitly assume that gain fields are updated when the saccade is executed, so that the gain signal is always reliable. The timing of gain field modulation of neuronal responses, however, has never been studied. We have designed a set of experiments to test the timing of visual gain fields in LIP. Two monkeys have been trained in a task in which they perform two sequential saccades. The first saccade is made along the gain field, either from an excitatory to inhibitory eye position, or vice versa. A 100 ms visual probe is flashed in the cell's receptive field at a variable latency, ranging from 0 to 600 ms, after the conclusion of the first saccade. The monkeys are rewarded when they make a memory-guided saccade to the location of the flashed stimulus. The flashed stimulus probes the cell's visual gain field response during a period, after the first saccade, in which the gain fields should be stable, according to the current literature. Our results in two monkeys show that visual responses to briefly flashed stimuli are unstable immediately after this saccade, and do not reflect the execution of the saccade. In more than 90% of LIP neurons recorded with visual gain fields, the response to the visual probe does not update until 100 to 200 ms after the first saccade. Thus, visual gain field timing is slow, similar to the timing of the proprioceptive eye position signal in area 3a of primary somatosensory cortex (Wang et al., 2007). Proprioception is an alternative source of cortical eye position information that complements corollary discharge and provides feedback after an eye movement has been executed. Neck proprioception drives LIP head-on-body position gain fields (Snyder et al., 1998), so we hypothesize that a slow proprioceptive signal may drive LIP eye position gain fields. Proprioception is unnecessary for online control of double step saccades (Guthrie et al., 1983) and open-loop pointing (Lewis et al., 2000). If gain fields utilize proprioception as their source of eye position information, then they would also be unnecessary and too slow for online coordinate transformations for action. Our conclusions are that the gain fields are updated after an eye movement, that they may be too slow to be used for online motor planning, and that they are slow enough that they could derive their input from oculomotor proprioception.

[doi:](#)

III-48. Oscillatory spiking activity in primate superior colliculus is related to spatial working memory

Lovejoy Lee
Richard J. Krauzlis
Salk Institute for Biological Studies

LOVEJOY@SALK.EDU
RICH@SALK.EDU

We report the presence of oscillatory spiking activity associated with memory-guided saccades in the primate superior colliculus (SC). Behavior contingent on spatial memory is associated with persistent activity in regions of the brain related to the planning and execution of movement. How this activity persists in the absence of a sensory

stimulus is currently a matter of debate, and its resolution would be of great significance to understanding both working memory and memory-contingent control of behavior. One potential explanation is that spatial memory is stored as reverberant oscillatory activity (Goldman-Rakic, 1995). The oscillatory activity is triggered upon the disappearance of spatially specific sensory input; in turn, it drives persistent activity elsewhere in the circuit in the absence of visual input and thus serves as local working memory. Indeed, oscillatory activity associated with working memory has been observed in frontal and parietal cortical areas in both human and non-human primates performing working memory tasks. For example, persistent spiking activity can be observed prior to memory-guided movements in both brainstem and cortical areas, including the superior colliculus (SC), frontal eye fields (FEF), and lateral intraparietal area (LIP); corresponding oscillatory spiking activity in the gamma-band (30-80 Hz) has been observed in both FEF (Buneo et al, 2003) and LIP (Pesaran et al, 2002) during the same period. We examined persistent activity recorded in the intermediate and deep layers of the SC while monkeys performed memory-guided saccades. Using multi-taper spectral analysis, we detected the presence of a narrow peak in the power spectrum of spiking activity in a subset of SC neurons. For most such cells, the peak appeared in the gamma-band (30-80 Hz), while the remainder appeared in the beta band (15-30 Hz) with no clear distinction between them. The peak only appeared when the memory-guided saccades were directed into the response field of the neuron and only after the visual cue was extinguished. Therefore the oscillatory activity was temporally and spatially specific to the remembered saccade goal, and could thus serve as local storage of working memory. These results show that some superior colliculus neurons could have dynamic memory fields like those seen in cortex. These dynamic memory fields appear to be in register with the well-known retinotopic motor map for saccades found in the SC and could form a "memory map" which is a component of a distributed circuit for spatial working memory. Buneo, C., Jarvis, M., Batista, A., and Andersen, R. (2003). Properties of spike train spectra in two parietal reach areas. *Exp. Brain Res.* 153, 134-139. Goldman-Rakic, P. (1995). *Cellular Basis of Working-Memory.* *Neuron* 14, 477-485. Pesaran, B., Pezaris, J., Sahani, M., Mitra, P., and Andersen, R. (2002). Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat. Neurosci.* 5, 805-811.

doi:

III-49. Ensemble activity underlying movement preparation in prearcuate cortex

Rachel Kalmar¹
John Reppas²
Stephen Ryu³
Krishna Shenoy⁴
William Newsome²

KALMAR@STANFORD.EDU
JBR1@STANFORD.EDU
SEOULMAN@STANFORD.EDU
SHENOY@STANFORD.EDU
BNEWSOME@STANFORD.EDU

¹Stanford University

²HHMI, Dept. of Neurobiology

³Dept. of Neurosurgery

⁴Dept. of Electrical Engineering

Movement preparation allows the rapid and accurate execution of voluntary movements, and can be influenced by factors that may change from moment to moment, such as attention and differences in stimulus properties. Consequently, movement preparation unfolds differently across many repetitions of the same movement. Averaging neural responses across many repetitions is necessary to interpret single-cell recordings, but diminishes our ability to characterize the dynamics of the underlying process. However, simultaneous recording from populations of neurons allows dynamics of movement preparation to be estimated on single trials. Our goal is to characterize these dynamics, to gain insight into the process underlying movement preparation. In the oculomotor system, the dynamics of movement preparation in individual cells are typically described using a rise-to-threshold model. As activity at the location on the cortical map corresponding to the target location increases, the saccade to that target is considered more "prepared", and the saccade can be initiated once activity crosses a threshold (Hanes & Schall, 1996). Reaction times are shorter when cells increase their firing rates to this threshold more rapidly. However, this relationship has only been studied in individual cells with saccade targets in their response field. During

behavior, the brain must combine input from populations of cells with wide ranges of response field positions relative to the upcoming saccade. To study the dynamics of movement preparation, we took an alternative approach that permits us to examine movement preparation by integrating information from a larger population of neurons with heterogeneous response fields. Here, we recorded peri-saccadic activity from ensembles of neurons in the prearcuate cortex in two monkeys. While monkeys performed visually-guided eye movements, we measured firing rates of a population of oculomotor neurons using a 96-electrode "Utah" array (Blackrock Microsystems, Salt Lake City, UT). Using the simultaneity of the neural responses recorded, we correlated trial-by-trial variations in population activity with the monkey's latency to saccade. In our population analysis, each neuron's firing rate describes a dimension in a high-dimensional firing rate space. The activity of the population evolves over the course of a trial, tracing out a trajectory within this high-dimensional space. These trajectories vary from trial-to-trial, but follow a stereotyped path. On individual trials, we found that the further along this path the population activity is when the monkey receives the go cue, the shorter the saccade latency. In both monkeys, this population measure accounted for more variance in reaction time than a measure of firing rate increase in individual neurons. Further, this relationship between prearcuate responses and saccadic reaction times was similar to that observed for reach reaction times and population dynamics in PMd (Afshar et al., SfN 2008). This suggests that some aspects of the neural strategy underlying movement preparation may be common to both systems. Our framework for analyzing neural population activity and dynamics should permit new extensions of single-neuron-level models, and may offer further insight into general mechanisms of movement preparation across motor systems.

doi:

III-50. Network mechanisms for the modulation of gamma spike phase by stimulus strength and attention

Paul Tiesinga^{1,2}
Terrence Sejnowski³

P.TIESINGA@SCIENCE.RU.NL
TERRY@SALK.EDU

¹Radboud University Nijmegen

²University of North Carolina at Chapel Hill

³Salk Institute; HHMI; U. of Cal. San Diego

Recent experiments demonstrate that in prefrontal cortex spikes at a specific gamma (30-80 Hz) phase relative to the local field potential are most informative about stimulus identity, where the specific value of the phase depends on the order of the stimulus within the stimulus sequence (Siegel, M., Warden, M.R., and Miller, E.K. PNAS, in press). Within the context of a model cortical circuit this corresponds to changes in the relative phase between excitatory (E) and inhibitory (I) cells. Gamma oscillations can emerge via the interneuron-gamma (ING) or pyramidal-interneuron gamma (PING) mechanism. We use theoretical arguments and simulation to determine how the relative phase between E and I cells can be modulated in cortical circuits. In the ING mechanism, I cells synchronize in the gamma frequency range when they are sufficiently depolarized. The resulting synchronous inhibition in turn synchronizes the E cells. Our simulations indicate that there are two scenarios for phase modulation. First, when the E cells have a firing rate much less than the oscillation frequency and when in addition they receive an incoherent excitatory drive, they will fire near the troughs of the inhibitory conductance. Increasing the depolarization of the E cells will increase their firing rate, but will not significantly change their mean phase relative to I cells. Second, when the E cells are sufficiently depolarized they will fire at a rate similar to the oscillation frequency, at a high precision (about 1 ms) and at a specific phase. Increasing E cell depolarization changes the E phase because it decreases the delay between E and I cells. The first scenario is more realistic in the absence of a strong stimulus, whereas the second might hold for circuits directly driven by a stimulus. In the PING mechanism, gamma oscillations emerge when the E cells are sufficiently depolarized and when a synchronized E cell volley recruits a synchronous I cell volley, which inhibits the E cells for approximately a gamma cycle. When the E cells have recovered, the oscillation cycle starts over. We find that hyperpolarizing the I cells will delay the I cell volley, thus increasing the phase of the I cells relative to the E cells. By contrast, when the depolarization of the E cells is increased, they recover faster from inhibition, which increases the oscillation frequency, but does not strongly alter the delay between I and E cells. E or I cells are depolarized by neuromodulators, such as, for

instance, acetylcholine released by projections from the basal forebrain or dopamine. These projections affect E and I cells differently, which allows for a division of labor because the preceding results show that targeting E or I cells has different effects on spike rate, oscillation frequency and phase. Under the hypothesis that in visual cortex feedforward inputs drive the E cells and top-down inputs target I cells, the PING mechanism predicts that contrast increases gamma power and oscillation frequency and attention alters the relative phase between I and E cells.

doi:

III-51. A model of vPFC neurons performing a same-different task: an alternative model of working memory

Jung Lee¹
Yale Cohen²

JLEE@MAIL.MED.UPENN.EDU
YCOHEN@MAIL.MED.UPENN.EDU

¹Department of Otorhinolaryngology, University of Pennsylvania School of Medicine

²University of Pennsylvania School of Medicine

Working memory is a fundamental component of cognition. Classically, the neural correlates of working memory is thought to involve persistent or sustained activity that is mediated by strongly recurrently connected groups of neurons. Recently, we recorded from neurons in the ventrolateral prefrontal cortex (vPFC), while monkeys reported whether sequentially presented auditory stimuli were the same or different and found that vPFC neural activity is predictive of the monkeys' decisions. Since this task required the monkeys to compare sequential stimuli, vPFC neurons should have access to a memory trace of the previous stimulus. However, we could not identify such a trace (i.e., persistent activity) in the firing rates of vPFC neurons between stimulus presentation nor could we identify a memory trace in neurons of the superior temporal gyrus (STG), a cortical region which projects to the vPFC. Instead, we found that memory traces appeared to be stored in the local-field potentials (LFPs) of the vPFC. Here, we present a model of the vPFC that incorporates both LFPs (membrane potentials) and spiking activity. This circuit has three layers: an input, middle, and output layer; the output layer is the read-out layer. The middle layer contains two groups of neurons each of which is tuned to a different stimulus and a third group which is "untuned". The two tuned neural groups are also bilaterally connected with a group of inhibitory inter-neurons. All of the synaptic weights are static except for the synapses between the tuned neurons and the inter-neurons and the untuned neurons and the output (read-out neurons). These dynamic synapses are updated by a simple rule that captures only qualitative properties of realistic dynamic synapses as proposed by Tsodyks et al.: action potentials potentiate the synapse; otherwise, it is depressed. As a consequence of this architecture, a negative feedback loop is established between the tuned neurons and the inter-neurons. When the same stimulus is presented sequentially, this feedback loop suppresses activity in those neurons that are tuned to this stimulus, but it does not affect the activity levels in the other tuned group. Consequently, when a different stimulus is presented, the responses of these neurons are robust. As a result, this circuit can discriminate between sequential stimuli. Furthermore, due to the connectivity of the untuned neural group, the spiking activity of the output neurons (1) does not habituate when the same stimulus is presented sequentially and (2) is enhanced when a different stimulus is presented. Whereas the spiking activity does not habituate, an analysis of the network's membrane potential indicates that the average membrane potential does habituate. This overall pattern of spiking activity and membrane-potential changes mimics that seen in our vPFC data. It is possible to build a circuit-level model of working memory based on dynamic synapses that does not produce persistent activity. This model raises the possibility of alternative mechanisms underlying working memory.

doi:

III-52. Designing optimal stimuli to control neuronal spike timing

Yashar Ahmadian

Adam M. Packer

Rafael Yuste

Liam Paninski

Columbia University

YA2005@COLUMBIA.EDU

ADAMPACKER@COLUMBIA.EDU

RMV5@COLUMBIA.EDU

LIAM@STAT.COLUMBIA.EDU

We develop fast computational methods for optimally designing a natural or artificial stimulus to make a neuron emit a desired spike train. We consider three specific examples of artificial stimulation methods: extracellular electrical stimulation (Salzman et al. 1990), two-photon uncaging of caged neurotransmitters (Nikolenko et al., 2008), and optical activation of genetically implanted light-sensitive ion channels (Han et al. 2007). We also consider the case of optimizing a sensory stimulus (e.g., the spatiotemporal modulation of visual contrast) for this purpose. We adopt a model based approach, using relatively simple biophysical models which describe how, in the case of each stimulation method, the input affects the spiking activity of the neuron. For example, in the case of photo-stimulation of light-sensitive ion channels, we model how laser light interacts with the ion channels in a neuron and how the latter affect its membrane potential and hence its spiking activity. Depending on the type of the neuron in question, we have used both a conductance based leaky integrator model, and a resonator model inspired by a certain linearization of the Hodgkin-Huxley equations, to describe the membrane potential dynamics. Finally, in the case of sensory stimuli, we use a generalized linear model to effectively capture how the whole upstream sensory network encodes the stimulus in the spike train of the neuron in question. Based on these models, we solve the reverse problem of finding the best time-dependent modulation of the input that makes the neuron emit a spike train which with highest probability will be close to a target spike train. However, this problem as stated is ill-posed: for example, if we can inject any arbitrary current into a cell, we can simply make the cell fire any desired pattern. Instead, we need to impose constraints on the set of allowed stimuli, as there are limitations on the stimuli we can safely apply in any physiological preparation without damaging the cells or causing other unwanted results. Thus, the task becomes a constrained convex optimization problem. We have developed fast methods for solving such optimization problems (Paninski et al. 2009). These methods can be implemented in real time and are also potentially generalizable to the case of many cells without losing tractability. This makes them suitable for neural prosthesis applications. Our simulations show that our methods provide an automatic, fast, and stable way of constructing the best possible input. These simulations can also be used to gauge how precisely spike trains can be induced in practice, given realistic values for the constraints such as the maximum allowed magnitude of current, light intensity, etc. We are in the process of experimentally testing these methods on neurons in cortical slices. Our work is motivated by several possible applications in neuroscience. As an example, our method can be used to produce desired spike trains in a number of selected neurons in some network. Observing the effect of the produced spikes on the subsequent activity in the whole network can then help study its connectivity patterns.

[doi:](#)

III-53. Hidden Markov models for the stimulus-response relationships of multi-state neural systems

Sean Escola¹

Liam Paninski²

GSE3@COLUMBIA.EDU

LIAM@STAT.COLUMBIA.EDU

¹Center for Theoretical Neuroscience, Columbia University

²Columbia University

Recent experimental results suggest that neural networks are associated with multiple firing regimes, or states, such as tonic and burst modes in thalamus (for review, see Sherman, Trends in Neuroscience, 2001) and UP and DOWN states in cortex (e.g. Anderson et al., Nature Neuroscience, 2000). It is reasonable to speculate that neurons in multi-state networks that are involved in sensory processing might display differential firing be-

haviors in response to the same stimulus in each of the states of the system, and, indeed, Bezdudnaya et al. (Neuron, 2006) showed that temporal receptive field properties change between tonic and burst states for relay cells in rabbit thalamus. Motivated by these results, we previously presented a general framework for estimating state-dependent neural response properties from paired spike-train and stimulus data assuming that neuronal assemblies transition between several discrete hidden states (Escola and Paninski, Cosyne, 2007). We modified the traditional hidden Markov model (HMM) framework to permit point-process observables such as spike-trains, and, for maximal flexibility in our model, we allowed an external, time-varying stimulus, if present, and the neurons' own spike histories to drive both the spiking behavior in each state and the transitioning behavior between states. We showed that an appropriately modified expectation-maximization algorithm could be constructed to learn the model parameters and gave preliminary results with simulated data. Although HMMs have been used previously to analyze neuronal data (e.g. Abeles et al., Proceedings of the National Academy of Sciences, 1995; Chen et al., Neural Computation, 2009), our model is an extension to the stimulus and history-dependent regime. In this poster, we review this previous work and then apply our model to a recently published data set of known multi-state neuronal ensembles (Jones et al., Proceedings of the National Academy of Sciences, 2007). We show that inclusion of spike-history information significantly improves the fit of the model compared to the analysis given in Jones et al. We then show that a simple reformulation of the state-space of the HMM's underlying Markov chain allows us to implement a hybrid half-multi-state/half-histogram model which captures more of the neuronal variability than either a simple HMM or a simple peri-stimulus time histogram (PSTH) model alone. This hybrid model learns firing-rate histograms that are triggered by the state-transition times rather than the trial start-times (i.e. these are state-dependent peri-transition time histograms or PTTHs as opposed to traditional PSTHs), and uncovers interesting and unexpected transition-locked dynamics in the data such as oscillations that are phase-locked to the transition times. We believe that techniques such as these may allow for the identification of data as multi-state that could not have been so identified by earlier methods, particularly data derived from neural systems where it is the network dynamics that are state-dependent rather than simple features such as firing rate, inter-spike interval distribution, or resting membrane potential.

doi:

III-54. Dopamine-modulated dynamic cell assemblies generated by the GABAergic striatal microcircuit

Mark Humphries^{1,2}

Ric Wood²

Kevin Gurney²

M.D.HUMPHRIES@SHEFFIELD.AC.UK

RIC.WOOD@SHEF.AC.UK

K.GURNEY@SHEF.AC.UK

¹Ecole Normale Supérieure

²University of Sheffield

The striatum, the principal input structure of the basal ganglia, is crucial to both motor control and learning. It receives convergent input from all over neocortex, hippocampal formation, amygdala and thalamus, and is the primary recipient of dopamine in the brain. Within the striatum is a GABAergic microcircuit that acts upon these inputs, formed by the dominant medium-spiny projection neurons (MSNs) and fast-spiking interneurons (FSIs). There has been little progress in understanding the computations it performs, hampered by the non-laminar structure that prevents identification of a repeating canonical microcircuit. We sought to solve this problem by searching for dynamically-defined computational elements within a full-scale model of the striatum. In the process, we have made significant progress in large-scale modelling of this structure. We constructed a new three-dimensional model of the striatal microcircuit's connectivity, implemented at 1:1 scale, neuron-for-neuron. The anatomical model was instantiated with our new dopamine-modulated neuron models of the MSNs and FSIs. A new model of gap junctions between the FSIs was introduced and tuned to experimental data. Finally, we developed a novel spike-train clustering method, suitable for large-scale models; applying this to the outputs of the model allowed us to find groups of synchronised neurons at multiple time-scales. We found that, with realistic in vivo background input, small assemblies of synchronised MSNs spontaneously appeared, consistent with experimental observations. The number of assemblies and the time-scale of synchronisation were strongly

dependent on the simulated concentration of dopamine. Such small cell assemblies, forming spontaneously only in the absence of dopamine, may contribute to motor control problems seen in humans and animals following loss of dopamine cells. We dissected the contributions of the circuit elements to the formation of the cell assemblies, and found that the FSI input was crucial in desynchronising the MSN activity. We also showed that feed-forward GABAergic input from the FSIs counter-intuitively increases the firing rate of the MSNs. Our interpretation of these results is that, in healthy striatum, localised, phasic changes of FSI activity switch the type of computations performed by MSNs. A phasic increase in local FSI activity would, in turn, increase the local MSN responses to ongoing cortical input, performing a striatum-wide "selection" computation on cortical inputs without using winner-takes-all. By contrast, a phasic decrease in local FSI activity would, for the same MSNs, promote competition between them through their network of inhibitory local collaterals. Thus, striatal FSIs seem able to set the scale and type of MSN computation.

doi:

III-55. Stationary envelope synthesis (SES): A universal method for phase coding by neural oscillators

Hugh T. Blair^{1,2}

Adam Welday¹

I Gary Shlifer³

Matthew Bloom³

Kechen Zhang⁴

BLAIR@PSYCH.UCLA.EDU

AWELDAY@UCLA.EDU

GARYSHLIFER@GMAIL.COM

MATTHEW.L.BLOOM@GMAIL.COM

KECHEN.ZHANG@GMAIL.COM

¹Psychology, UCLA

²Brain Research Institute, UCLA

³Dept of Psychology, UCLA

⁴Biomedical Engineering, Johns Hopkins Univ

The rat hippocampus and limbic cortex contain spatially tuned neurons-including place cells, grid cells, and boundary cells-which exhibit firing rate maps with different geometries, but their spike trains are all similarly modulated by 4-12 Hz theta rhythm. These neurons are widely believed to be components of a neural system for path integration, which tracks the rat's position by integrating the velocity of its movements over time. Burgess et al. (2005) introduced an oscillatory interference model of path integration based on the principle that if an oscillator's frequency is linearly modulated by movement velocity, then its phase will encode a position signal. This principle has been exploited to show how phase interference between velocity-controlled oscillators (VCOs) can synthesize envelope waveforms that mimic the periodic spatial firing rate maps of grid cells (Burgess et al., 2007; Giocomo et al., 2007; Hasselmo et al., 2007; Burgess 2008). Here we show that these oscillatory interference models of grid cells represent a special case of a more general principle-referred to as stationary envelope synthesis (SES)-whereby phase interference among VCOs can synthesize stationary envelope waveforms that approximate any desired function in a vector space of arbitrary dimensions. A neural architecture for implementing the SES principle is proposed, consisting of a ring oscillator matrix (ROM) which contains a bank of VCOs from which stationary envelope functions can be synthesized (Blair et al., 2008). Simulations are presented to show how a target neuron that sums transiently synchronous inputs from the ROM can synthesize spatial envelopes that mimic the firing rate maps not only of grid cells (as in prior models), but also of place cells and boundary cells. An essential prediction this model architecture is that burst frequencies of some theta cells in the rat brain must be modulated by the cosine of the rat's movement direction, and we present experimental data showing that this indeed appears to be true for theta cells recorded from the anterior thalamus of freely behaving rats. Based on this theory and evidence, it is conjectured that spatially tuned neurons in hippocampus and cortex sum feedforward inputs from ROM circuits in the thalamus, and then return feedback connections to the thalamus to complete a stable attractor loop which facilitates accurate path integration that is robust to noise. Spatial path integration can be generalized from 2D environments into higher dimensions for tracking the momentary value of any sensory, motor, or memory state that traces a continuous trajectory through an M-dimensional Euclidian vector space. According to the SES principle, a bank of oscillators with frequencies that are modulated by the time derivative of such a trajectory

(regardless of its path) should be able generate any desired envelope function in a vector space with any number of dimensions. It is speculated that this capability might be broadly useful for solving problems of invariant pattern recognition, by making it possible to synthesize envelope functions that pick out bounded subsets of a state space that contain all possible transformations of a given recognition target.

[doi:](#)

III-56. Representation of environmental statistics by neural populations

Deep Ganguli¹
Eero P. Simoncelli²

DGANGULI@CNS.NYU.EDU
EERO.SIMONCELLI@NYU.EDU

¹Center for Neural Science, NYU

²New York University

The efficient coding hypothesis proposes that sensory systems have evolved so as to efficiently represent signals experienced in natural environments (Barlow, 1961). Quantitatively, this hypothesis can be expressed as the desire to maximize mutual information between environmental signals (as characterized by a prior probability density), and the response of a neural population, subject to resource constraints. Here we consider both the perceptual and physiological implications of efficient coding constrained by both the size of the population, and the average firing rate. Information is difficult to optimize for a multi-dimensional neural system, so we rely on two approximations. First, we use the Fisher bound on mutual information (Brunel et al., 1998), which can be rewritten in terms of the KL divergence between the square root of Fisher information and the prior (McDonnell, 2008). The bound becomes tighter as the signal to noise ratio of the population code increases (e.g. with a larger number of neurons and/or higher firing rates). Second, we solve for the optimum of this objective function, ignoring the normalization that would be necessary to make the root Fisher information a proper density. Under these approximations, an efficient population code should achieve a Fisher information proportional to the square of the prior, while satisfying resource constraints. Since the inverse of root Fisher information also provides a lower bound on the discriminability of stimuli (Series et al., 2009), this solution provides a direct perceptual prediction. Specifically, efficient coding implies that the inverse of perceptual discriminability should look like the prior. We tested this prediction for two different stimulus attributes: orientation, and speed. In the first case, we show that the inverse of orientation discriminability (data from Girshick et al., VSS2009) closely resembles the prior distribution of orientations we measured from natural images. In the speed case, we examined both discrimination data (Stocker et al., 2006) and the responses of a population of Macaque MT cells to stimuli of different speeds (data from Majaj et al., Cosyne2007). The latter data set was collected explicitly for the purpose of assessing the embedding of prior information in MT populations (Stocker, VSS2009). Thus, we can use the efficient coding result to make two independent predictions for the prior on speed. The first prediction is obtained by inverting the discrimination thresholds, as in the orientation case. A second prediction is obtained by computing the normalized root Fisher information of the MT population, which we accomplish by assuming independent Poisson noise and parameterized tuning curves fit to the data. Remarkably, the two predicted priors are consistent with each other and with the finding that human visual speed perception is consistent with a Bayesian observer using a prior that favors slower speeds (eg., Stocker et al. 2006). The efficient coding hypothesis allows us to establish a quantitative link between environmental statistics, neural processing, and perception. Although we have resorted to several approximations to achieve these results, we are currently exploring reformulations that can eliminate the need for such approximations while reaching similar conclusions.

[doi:](#)

III-57. Origins of contrast gain control in isolated cortical neurons: deriving the code from the dynamics

Michael Famulare
Rebecca Mease
Adrienne Fairhall
University of Washington

FAMULARE@U.WASHINGTON.EDU
BECKIN@U.WASHINGTON.EDU
FAIRHALL@U.WASHINGTON.EDU

A neuron's coding strategy is defined by the relationship between its inputs and its spiking output. Much of the computation performed by single neurons can be captured by a linear-nonlinear cascade (LN) model: a linear filter extracts the relevant feature in the stimulus that drives spiking, and a nonlinear decision function determines the probability of a spike for a particular value of the filtered stimulus. For many systems, the LN model varies in response to changes in stimulus statistics. This form of adaptive coding often serves to maximize information transmission. In a recent set of experiments, we found that single neurons in the developing mouse sensorimotor cortex exhibit nearly perfect contrast gain control in parallel with the expression of basic spike-generating ion channels. Via whole-cell recording, we stimulated synaptically-isolated neurons with white Gaussian noise current with fixed mean but different variance and calculated LN models using reverse correlation techniques. By nearly perfect contrast gain control, we mean that the nonlinear decision function is essentially identical for all stimulus ensembles with the same mean, after the standard deviation of the stimulus ensemble is scaled out. In other words, the neurons normalize out the overall context in which stimuli appear before making decisions about spiking. Over the same developmental timecourse, the ratio of the maximal sodium and potassium conductances for each cell converges to a common population value, suggesting that this quantity is controlled. These experiments demonstrate that single cortical neurons can perform contrast gain control using a minimal set of voltage-gated channels. Given the apparent simplicity of these results, we were motivated to address two questions. The first is to develop a mathematical understanding of how the dynamics of a single cell can give rise to contrast gain control. The second is to determine to what extent the parameters of a neuronal system must be finely tuned in order to express this property. Conductance-based modeling shows that the expressed ratio of Na/K is favorable for the appearance of gain control. To further analyze this system, we chose to investigate the posed issues within the framework of a reduced, analytically tractable yet experimentally plausible spiking neuron model, the exponential integrate-and-fire model (EIF). We show how to tune this model to show contrast gain control in a manner consistent with the behavior seen in the experimental data. To do so, we make use of Fokker-Planck techniques coupled with our previously-developed technique of subthreshold stochastic linearization to derive from first principles the reduction of a simple nonlinear dynamical system to a static LN model. This provides insight into the more general issue of the limitations of that reduction. The EIF model is an effective reduction for neurons like the ones studied experimentally that are leak-dominated below threshold with rapid-onset sodium kinetics. While we focus here on this simple model, the conceptual framework extends to more biophysically-plausible models. This work constitutes a step toward deriving analytical connections between specific ion channel dynamics and properties of the resulting functional computation performed by a neuron.

[doi:](#)

III-58. Single neuron dynamics determine the strength of chaos in the balanced state

Michael Monteforte¹
Siegrid Löwel²
Fred Wolf¹

MONTE@NLD.DS.MPG.DE
SIEGRID.LOEWEL@UNI-JENA.DE
FRED@NLD.DS.MPG.DE

¹MPI for Dynamics and Self-Organization, BCCN

²Friedrich-Schiller-University Jena

The asynchronous irregular firing activity of neurons in the cortex results from strongly fluctuating synaptic in-

puts, originating from a dynamical balance of excitation and inhibition. This network effect has been termed the balanced state. Its statistical properties are relatively well understood. Its dynamical nature, however, is not. Originally Vreeswijk and Sompolinsky reported extremely strong chaos, whereas, more recently, Zillmer et al. and Jahnke et al. found stable dynamics. To clarify these contradictory results we comprehensively investigated the dynamics of balanced neural networks for different neuron models. In all networks neurons were sparsely coupled by relatively strong synapses without synaptic delays. We chose analytically solvable single neuron models to allow fast event-based simulations and the analytical calculation of the single spike Jacobian. Based on the single spike Jacobian we numerically calculated all Lyapunov exponents and derived the Kaplan-Yorke attractor dimension and Kolmogorov-Sinai entropy production rate of the entire network. First, we examined networks of two canonical neuron models, leaky integrate and fire (LIF) and theta-neurons. Intriguingly, the dynamics of the balanced state was extremely sensitive to the choice of single neuron model. LIF neurons led to stable dynamics, whereas theta-neurons led to extensive deterministic chaos. The latter was characterized by largest Lyapunov exponents of 20 to 40Hz, attractor dimensionality of 20 to 60% of the phase space and entropy production rates of 0.5 to 0.9bit per spike. In the LIF networks, weak chaos occurred when changing from delta-pulse coupling to correlated synaptic currents. It was characterized by largest Lyapunov exponents of 0 to 10Hz, attractor dimensionality up to 0.4% and entropy production of 0 to 0.001bit per spike. Thus, the chaotic dynamics of the LIF networks was up to 3 orders of magnitude weaker than that observed in theta-neuron networks. A qualitative difference between the LIF and theta-neuron model is an active action potential (AP) initiation in the latter. Our results thus suggest that the AP initiation dynamics can strongly influence chaos in the balanced state. To directly address this we then investigated networks of a new neuron model with variable AP onset rapidness r , called rapid-theta-neuron. This model is able to interpolate between theta-neurons ($r=1$) and LIF neurons ($r=\infty$). Gradually increasing the AP onset rapidness monotonously decreased the attractor dimensionality and entropy production, which vanished at a critical value of $r>100$. Interestingly, the largest Lyapunov exponent showed a maximum at AP onset rapidness of order 10. This fits standard conductance-based neuron models like the Wang-Buzsaki model. The actual AP onset rapidness of cortical neurons may, however, be much bigger [Naundorf et al. 2006, see also McCormick et al. & Naundorf et al. 2007]. The largest Lyapunov exponent and entropy production might then be much lower in realistic cortical networks. These results demonstrate that the dynamics of single neuron spike initiation critically determine the occurrence and strength of chaos in balanced neural networks. They suggest that cortical neurons with their fast AP onsets might be tuned to reduce information loss by the chaotic network dynamics.

doi:

III-59. A non-stationary copula-based spike count model

Arno Onken^{1,2}

Steffen Grünewälder³

Matthias H. J. Munk⁴

Klaus Obermayer¹

AONKEN@CS.TU-BERLIN.DE

STEFFEN@CS.UCL.AC.UK

MATTHIAS.MUNK@TUEBINGEN.MPG.DE

OBY@CS.TU-BERLIN.DE

¹Technische Universität Berlin

²BCCN Berlin

³University College London

⁴MPI for Biological Cybernetics

Recently, detailed dependencies of spike counts were successfully modeled with the help of copulas [1, 2, 3]. Copulas can be used to couple arbitrary single neuron distributions to form joint distributions with various dependencies. This approach has so far been restricted to stationary spike rates and dependencies. It is known, however, that spike counts of recorded neurons can exhibit non-stationary behavior within trials. In this work, we extend the copula approach to capture non-stationary rates and dependence strengths which are on the order of several hundred milliseconds. We use Poisson marginals for the single neuron statistics and several copula families with and without tail dependencies to couple these marginals. The rates of the Poisson marginals and the dependence strengths of the copula families are time-dependent and fitted to overlapping 100 ms time bins using the inference for margins procedure. To reduce the model complexity we then use regularized least-squares fits

of polynomial basis functions for the time-dependent rates and dependence strengths. The approach is applied to data that were recorded from macaque prefrontal cortex during a visual delayed match-to-sample task. Spike trains were recorded using a micro-tetrode array and post-processed using a PCA-based spike sorting method. We compare the cross-validated log likelihoods of the non-stationary models to the corresponding stationary models that have the same marginals and copula families. We find that taking non-stationarities into account increases the likelihood of the test set trials. The approach thereby widens the applicability of detailed dependence models of spike counts. This work was supported by BMBF grant 01GQ0410. [1] A. Onken, S. Grünwälder, M. H. J. Munk, and K. Obermayer. Analyzing short-term noise dependencies of spike-counts in macaque prefrontal cortex using copulas and the flashlight transformation. *PLoS Computational Biology*, in press. [2] A. Onken, S. Grünwälder, M. H. J. Munk, and K. Obermayer. Modeling short-term noise dependence of spike counts in macaque prefrontal cortex. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 1233-1240, 2009. [3] P. Berkes, F. Wood, and J. Pillow. Characterizing neural dependencies with copula models. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 129-136, 2009.

[doi:](#)

III-60. Multiple spike time patterns occur at bifurcation points of membrane potential dynamics

J. Vincent Toups¹
 Jean-Marc Fellous²
 Peter J. Thomas³
 Terrence J. Sejnowski⁴
 Paul H. Tiesinga⁵

TOUPS@PHYSICS.UNC.EDU
 FELLOUS@EMAIL.ARIZONA.EDU
 PJTHOMAS@CWRU.EDU
 TERRY@SALK.EDU
 TIESINGA@PHYSICS.UNC.EDU

¹Univ. North Carolina Chapel Hill

²University of Arizona

³Department of Mathematics, Case Western Reserve University

⁴Salk Institute; HHMI; U. of Cal. San Diego

⁵Radboud University Nijmegen; UNC Chapel Hill

The response of a neuron to fluctuating current injected in vitro typically elicits a reliable and precisely timed sequence of action potentials when the same waveform is repeated. The effects that different stimulus conditions occurring in vivo might have on the reproducibility of output spike times remains an open question. To address this question, we somatically injected an aperiodic current into cortical neurons in vitro and varied the amplitude of the fluctuations and the constant pedestal or offset. As the amplitude of the fluctuations was increased the spike reliabilities increased and the spike times remained stable over a range of values. However, at exceptional values called bifurcation points, large shifts in the spike times were obtained in response to small changes in the stimulus amplitude. At such bifurcation points, multiple spike patterns were revealed with an unsupervised method. Increasing the current offset, which mimicked an increase in network activity, also increased the spike time reliability, but the spike times shifted earlier with increasing offset. Although the reliability was reduced at bifurcation points, the information about the stimulus time course was increased because each of the spike time patterns contained different information about the input.

[doi:](#)

III-61. Information scaling, efficiency and anatomy in the cerebellar granule cell layer.

Guy Billings¹
 Andrea Lorincz²
 Pdraig Gleeson¹
 Zoltan Nusser²
 Angus Silver¹

G.BILLINGS@UCL.AC.UK
 LORINCZ@KOKI.HU
 P.GLEESON@UCL.AC.UK
 NUSSE@KOKI.HU
 A.SILVER@UCL.AC.UK

¹University College London

²Institute of Experimental Medicine, Hungarian Academy of Sciences

A key challenge in neuroscience is to understand the relationship between the structure of neural circuits and their function. However, relatively little is known about how the anatomy of neural networks affects the information theoretic properties of neural coding. An ideal system in which to study this problem is the cerebellar granule cell layer. This input-layer neural network can be approximated as feedforward and has a highly conserved and relatively simple bipartite structure in which excitatory mossy fibers synapse onto granule cells within glomeruli. Two anatomical features of this network are at once striking and mysterious: Firstly, on average each granule cell is synaptically innervated by 4 different mossy fibers (typical range 3-7). Secondly, the ratio of the densities of granule cells and glomeruli is 2.8:1. Here we consider the effect of these parameters on the energy efficiency, redundancy and information content of coding within this network. We calculate the mutual information between glomeruli and granule cells in a simple binary linear threshold model of the cerebellar granule cell layer. This approach was combined with a linear cost model that penalizes unnecessarily large or active networks. To quantify redundancy, we compute the log of the number of possible output codes minus the actual entropy of the outputs. We then tuned the parameters of the network, including the fraction of the inputs activated, so as to minimize the cost per bit of the resulting encoding. We find that the energetically most efficient encodings are strongly redundant, sparse and bestow a pseudo-metric structure on the input code space. But these encodings are also lossy: expressing only around 70% of the input information. For this reason we also examined lossless encodings that minimize the cost per bit. These encodings are less strongly redundant, are less sparse and bestow a metric structure on input code space. Over the range of granule cell to glomerulus ratios explored we find that between 3 and 5 connections per granule cell maximizes the energy efficiency of the most efficient encodings. For a ratio of 2.8:1 this was achieved with a threshold of 2 active inputs and a glomerulus activation probability of 0.1. Lossless encoding can also be achieved by adjustment of the probability of activation of glomeruli to around 0.7 and the threshold to equal the number of connections. We also examined the effect of altering the granule cell to glomerulus ratio. Upon reducing the ratio to 1:1 we find that the network can transmit information with an energy-efficient but lossy code, but that there is no possible lossless encoding. This suggests that around 4 connections per granule cell combined with their relative abundance might enable a dynamic tradeoff between energy efficiency and lossless encoding under the control of thresholding. This is the first information theoretic hypothesis about the of the number of granule cell connections. Funded by BBSRC, MRC and the Wellcome Trust.

[doi:](#)

III-62. Dynamic changes in single cell and population activity during the acquisition of task behavior

Joshua Swearingen¹
 Marcelo Reyes¹
 Catalin V. Buhusi²

SWEARINJ@MUSC.EDU
 REYE@MUSC.EDU
 BUHUSI@MUSC.EDU

¹Medical University of South Carolina

²Dept. Neurosciences, Medical University of South Carolina

Previous studies indicate rather abrupt changes in individual animals undergoing task learning which is detectable through analysis of behavior [1, 2] and gross population activity through electro-corticographs [3], though it is unclear how this information is acquired, represented, and processed in the brain. Advances in multi-channel electrophysiology make it possible to record simultaneously from many neurons, bringing the power to address such questions, though such an analysis requires more sophisticated methods to handle the growing complexity of the data. We seek to investigate whether we can detect complex changes during animal learning, in both single cell and population level activity, using unsupervised classification techniques. Data were collected from nine rats, implanted with two 2x8 electrode arrays directed at dorsal striatum and medial prefrontal cortex. The specific neuronal changes that may underlie task-oriented learning are unknown. While individual neurons emit spikes of action potentials in series that contain information in their firing rate and temporal pattern, they may also switch between a number of general patterns of activity, reflecting the nature of the larger neural network they exist within. In this study we use two methods of unsupervised classification to help identify such changes within the activity of single neurons, as well as population ensembles, in the units recorded in both pre-frontal cortex and striatum in single sessions. The first method, based on work by Fellous et al [4], classifies changes within neurons by creating trial by trial similarity matrices and uses this as a basis for k-means clustering. The second is suitable for the classification of population activity, creating dimension reduced spaces that illustrate gross network changes over trials and time using principle components analysis. Using these analyses we show that neural activity evolves from early trials when the animals have not yet learned the task, and transitions in neural properties are typically synchronized to the observed behavioral changes. A number of distinct features develop over trials in sub-populations of single cells, and these changes lead to a significant increase in the classification of ensemble activity in post-acquisition trials. We conclude that these results can help to understand the way information is distributed and embedded in the dynamics of the frontal and striatal populations during task acquisition, and the methods provide a way to successfully classify complex functional dynamics within neurons. References: [1] Gallistel et al. "The learning curve: implications of a quantitative analysis." *Proceedings of the National Academy of Sciences of the United States of America* 101, no. 36 (September 7, 2004): 13124-13131. [2] Smith et al. "Dynamic analysis of learning in behavioral experiments." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 24, no. 2 (January 14, 2004): 447-461. [3] Ohl et al "Change in pattern of ongoing cortical activity with auditory category learning." *Nature* 412, no. 6848 (2001): 733-736. [4] Fellous et al. "Discovering Spike Patterns in Neuronal Responses." *J. Neurosci.* 24, no. 12 (March 24, 2004): 2989-3001.

doi:

III-63. A stimulus-dependent maximum entropy model of the retinal population neural code

Einat Granot-Atedgi¹

Gasper Tkacik²

Ronen Segev³

Elad Schneidman¹

EINAT.GRANOT@WEIZMANN.AC.IL

GTKACIK@PRINCETON.EDU

RONENSGV@BGU.AC.IL

ELAD.SCHNEIDMAN@WEIZMANN.AC.IL

¹Department of Neurobiology, Weizmann Institute of Science

²Department of Physics, University of Pennsylvania

³Ben-Gurion University

The nature of the code by which information is represented by the joint activity patterns of groups of neurons, depends on the stimulus-response properties of each of the single cells, and the dependencies among neurons. However, most models of neural encoding, describing the selectivity and stochastic nature of neural response to various stimuli, have focused on single neurons or small groups of cells, using simplified, artificial stimuli. Recent results, in several different neural systems, have used maximum entropy pairwise models to show that the typically weak pairwise correlations between neurons add up to dominate the population activity patterns. While reflecting the strong effect of pairwise correlations on the collective behavior of the population, these models have not addressed the stimulus dependent properties of neural population code. Here we aim at elucidating

how large neural populations encode sensory information, and uncovering the functional role of the interactions between neurons. We present a novel stimulus-dependent maximum entropy model, which captures both the pairwise correlations between neurons and the stimulus-dependent firing rates of single cells. We apply this model to the recordings of large groups of retinal ganglion cells responding to artificial and naturalistic stimuli, and show that it significantly outperforms conditionally independent models. In particular, our model captures the time-dependent activity of single cells and the population activity patterns for both classes of stimuli, where the classical receptive-field models perform poorly. Finally, we find that the pairwise interaction map that underlies the population response is similar under different stimuli, suggesting that the functional interactions between cells could encode a prior over the binary words a neural population may emit. Our results provide a framework for combining single cell receptive field models with the maximum entropy description of network states. We show that this approach allows accurate modeling of large neural population responses to non-stationary and rich stimuli, and suggests a biologically plausible way for the neural systems to implement population decoders.

doi:

III-64. Effects of spike-driven feedback on neural gain and pairwise correlation

John Bartels¹
Brent Doiron²

JBARTELS@ALUM.RPI.EDU
BDOIRON@PITT.EDU

¹University of Pittsburgh

²Dept of Mathematics, University of Pittsburgh

Effects of spike-driven feedback on neural gain and pairwise correlation John D. Bartels¹, Brent Doiron^{1,2}; ¹: Department of Mathematics, University of Pittsburgh. ²: The Center for the Neural Basis of Cognition, Pittsburgh. Both single neuron and neural population spiking statistics, such as firing rate or temporal patterning, are critical aspects of many neural codes. There has been tremendous experimental and theoretical effort devoted to understanding how nonlinear membrane dynamics and ambient synaptic activity determine the gain of firing rate responses. This is primarily motivated by in vivo recordings, from a variety of sensory systems, where response gain is shifted by, for instance, stimulus contrast, anesthetic state, and subject attention. However, there is increasing experimental evidence that the same manipulations that affect firing rate gain also modulate the pairwise correlation between neurons. In this study, we explore how spike-driven intrinsic and synaptic feedback co-modulate firing rate gain and spike train correlation. We consider a pair of excitable leaky integrate-and-fire neurons where each neuron receives feedback inputs driven from past spike events. To distinguish the effects of intrinsic versus network influences, we study two "network" architectures. The first is a "self-coupled" network comprising of a pair of neurons, which have internal, spike driven dynamics (depolarizing or hyperpolarizing after potentials) yet lack synaptic coupling [1]. The second is a "cross-coupled" network, where the intrinsic dynamics are replaced by synaptic activity between the neurons. We find that while the relative firing rate gain is sensitive to feedback polarity, it is invariant to both the network architecture and timescale over which gain is computed. We next investigate the correlation transfer properties of these networks by driving the neurons with weakly correlated, Gaussian white noise sources [2]. Our correlation results contrast to our gain results; spike train correlation shows a strong dependence on polarity, architecture, and timescale. For each of these scenarios, we apply linear response theory to derive expressions for both dynamic gain and correlation between spike trains for the paired neurons in the weakly coupled limit. We find good agreement between our theoretical results and the corresponding spike-count statistics computed from simulated spike trains. Our results suggest that commonly employed intrinsic and synaptic modulations of firing rate gain will have strong consequences for network correlation transfer. [1] de la Rocha J, Doiron B, Shea-Brown E, Josic K, Reyes A. Correlation between neural spike trains increase with firing rate. *Nature*. 448: 802-806, 2007. [2] Doiron B, Oswald A-M, and Maler L. Interval Coding II: Dendrite Dependent Mechanisms. *J. Neurophysiol*. 97: 2744-2757, 2007.

doi:

III-65. A neuronal population measure of attention predicts behavioral performance on individual trials

Marlene R. Cohen
John H. R. Maunsell
Harvard Medical School

MARLENE_COHEN@HMS.HARVARD.EDU
JOHN_MAUNSELL@HMS.HARVARD.EDU

Internal states such as attention have dramatic effects on behavior. On blocks of trials when an animal directs attention to a particular area of space, perception at that location is greatly improved compared to blocks of trials when attention is directed elsewhere. Like all neuronal and behavioral processes, attention varies from moment to moment, but the behavioral consequences of fluctuations in attention within an attentional condition have been impossible to track using either psychometric performance, which by definition is measured over large numbers of trials. Previous studies have therefore only studied the effects of attention averaged over long periods. Because attention modulates the responses of sensory neurons, it should in principle be possible to detect fluctuations in attention by measuring neuronal responses. However, obtaining single trial estimates of attention is not possible from responses of a single neuron. Many sensory neurons are likely to contribute to any perceptual process, so trial-to-trial changes in single neuron responses should not be expected to strongly correlate with behavioral fluctuations caused by attention. More importantly, it is impossible to dissociate variability in single neuron responses due to attention from variability due to other factors. These problems may be solved by basing a single trial measure of attention on the responses of many simultaneously recorded neurons. We obtained a single trial metric of attention that quantified the similarity of responses of dozens of neurons in visual area V4 to mean responses in two attention conditions. We were able to show for the first time that the response of a modestly-sized population of sensory neurons can reliably predict an animal's performance on a single trial. We found that within an attentional condition, this attentional modulation of a population of neurons varies substantially within and between trials, and that the monkey's ability to detect a given small change in a stimulus varies from near perfect performance to near zero depending on the amount of modulation of the population. We also used this single trial measure to test the validity of the metaphor of an attentional "spotlight" that has been long used to describe visual attention, which suggests that attention is a limited resource that can be directed to a specific location or small set of locations. Implicit in this idea is the assumption that increasing attention to one location reduces attention at other locations, so the amount of attention allocated to two locations should typically be anticorrelated. By simultaneously recording the responses of populations of neurons in opposite hemispheres, we were able to test this assumption directly. Unexpectedly, we found that on a trial-to-trial basis, the amount of attention allocated to two locations in opposite hemifields was uncorrelated, suggesting that at any moment attention is allocated to each spatial location independently. Consistent with recent proposals that attention acts through the same mechanisms that control response normalization, these results suggest that attention to locations in opposite hemifields is governed not by a single top-down control mechanism, but by local groups of neurons whose variability is independent.

[doi:](#)

III-66. A nearly optimal correlation-independent readout of population activity

Ching-Ling Teng¹
Peter Latham²
Jonathan W. Pillow³

TENG@VIRGINIA.EDU
PEL@GATSBY.UCL.AC.UK
PILLOW@MAIL.UTEXAS.EDU

¹University of Virginia

²Gatsby Computational Neuroscience Unit, UCL

³University of Texas at Austin

Perceptual decision-making requires the brain to read out, or decode, population spiking activity of sensory neurons. If neurons were independent, decoding would be relatively straightforward. However, it is well known that spiking activity of cortical neurons is not independent: trial-to-trial fluctuations in activity are shared among neu-

rons. Moreover, these correlations can change, when driven by other factors like attention, arousal, and task demands. In this study, we examine how such changes in inter-neuronal correlations affect performance of a fixed decoder. We consider two types of decoders: one that exploits knowledge of the correlations between neurons in the sensory representation, and one that does not. We analyze the effects of noise correlation on perceptual decision-making using a simple model with two stages of sensory processing. The first stage represents the population encoding of sensory information. The second stage consists of a linear pooling mechanism that read out, or interpret, sensory information. Specifically, we assume that a scalar-valued stimulus is encoded by a population of neurons with Gaussian-shaped tuning curves corrupted by additive noise. The covariance of the noise is taken to be either stimulus-independent or stimulus-dependent, and we consider a variety of correlational strengths and structures. For each type of noise correlation, we determine the optimal linear readout rule for different tasks: stimulus detection, identification, and fine and coarse discriminations. In addition, we analytically derive the relationship between single-neuron and behavioral variability (choice probability) for different tasks and noise correlation; this provides experimentally testable predictions. Our analysis shows that if noise correlations are present, the optimal decoder is a center-surround linear filter. If noise correlations are absent, the optimal linear decoder has no surround inhibition. Moreover, the optimal decoder for independent responses performs poorly when applied to responses with correlated noise. Thus, it would seem that correlations are important and a decoder - presumably implemented in downstream networks of neurons - needs to know about them. But does a decoder really need to know about the correlation structures? More specifically: is there a decoder that can provide accurate readout regardless of the details of the correlational structure, including no correlations at all? If so, correlations would become unimportant because a decoder would not need to know them to perform near-optimally. Indeed, that is what we found: a center-surround readout is remarkably robust to various correlational structures, including changes in correlational lengths, strengths, and stimulus dependency. In conclusion, a center-surround readout provides a flexible and effective way to decode both independent and correlated neural populations, including those whose correlation structures can change when driven by factors like attention, arousal, and task demands. Our results suggest that the problem of decoding correlated population activity may be easier than was thought.

doi:

III-67. Sampling based inference with linear probabilistic population codes

Jeff Beck¹
 Alexandre Pouget²
 Peter Latham¹

BAYESIAN.EMPIRIMANCER@GOOGLEMAIL.COM
 ALEX@CVS.ROCHESTER.EDU
 PEL@GATSBY.UCL.AC.UK

¹University College London

²University of Rochester

An ever increasing corpus of behavioural data from animals as simple as insects and those as complex as humans indicates that probabilistic reasoning is the rule when it comes to tasks that range from low level sensory-motor transformation to high level cognition and decision making. This remarkable set of behaviours require a neural code which (1) represents probability distributions over task relevant latent variables (stimuli), (2) is consistent with observed neural statistics, and (3) can be used to implement the operations of probabilistic inference using biologically plausible neural operations. Over the past few years we have worked to develop such a probabilistic population code or PPC and have utilized such a code to model tasks which involve probabilistic operations such as cue combination, evidence accumulation, prior implementation, and maximum a posteriori estimation. However, to some extent, these are the easy problems probabilistic inference, or at least, these are the problems of probabilistic inference for which the specific type of PPC which we have proposed, the linear PPC, is ideally suited. Unfortunately, as is often the case, a representation which makes a particular problem easy can make another problem difficult. In this case, the probabilistic operation made difficult by this choice of code is the marginalization operation. Here, marginalization involves taking a complex generative model for neural activity, r , which is conditioned upon many latent variables (i.e. some $p(r|s_1, s_2, \dots, s_n)$) and then inverting that model to obtain a posterior distribution over only a few task relevant latent variables. This involves a potentially high dimensional integral over say $s_2 \dots s_n$ to obtain a marginal posterior distribution $p(s_1|r)$. Previously, we had some success in

generating biologically plausible networks which perform the relatively low dimensional marginalization operations needed to implement computations such as non-linear coordinate transforms for sensory-motor information integration, Kalman filters for motor control and object tracking, explaining away in infants (a.k.a backward masking), and auditory localization. This was accomplished using networks which implement a quadratic non-linearity and divisive normalization, operations which are observed throughout cortex. This was intriguing as, previously, divisive normalization had been implicated in gain control, attention and redundancy reduction, and these results suggest a more generic and possibly unifying computational role. Unfortunately, it was not clear from this previous work that these results would generalize to higher dimensional marginalization operations which can only be made tractable by utilizing sampling methods. Indeed, it was not initially clear that sampling methods were compatible with the probabilistic population coding approach at all. In this work, we address this issue by first demonstrating that samples can be generated naturally within the PPC framework in a variety of ways, some of which require nothing more than stochastic synapses and dynamic attractor networks comparable to those used to generate maximum a posteriori estimates. Moreover, we show that the samples generated in this way can be easily adapted to the purpose of implementing the marginalization operation in a way that, once again, indicates that divisive normalization performs a critical computation role.

doi:

III-68. Decoding multiscale word and category-specific spatiotemporal representations from intracranial EEG

Alexander M. Chan^{1,2}

Eric Halgren³

Chad Carlson⁴

Orrin Devinsky⁴

Werner Doyle⁴

Ruben Kuzniecky⁴

Thomas Thesen⁴

Chunmao Wang⁴

Donald Schomer⁵

Emad Eskandar⁶

AMCHAN@MIT.EDU

EHALGREN@UCSD.EDU

CCARLSO2@GMAIL.COM

OD4@NYU.EDU

WERNER.DOYLE@NYUMC.ORG

RUBEN.KUZNIECKY@NYUMC.ORG

THOMAS.THESEN@MED.NYU.EDU

WANGCH@NIDCD.NIH.GOV

DSCHOMER@BIDMC.HARVARD.EDU

EESKANDAR@PARTNERS.ORG

¹Harvard-MIT Health Sciences & Technology

²Massachusetts General Hospital

³Department of Radiology, UC San Diego

⁴NYU Comprehensive Epilepsy Center

⁵Beth Israel Deaconess Medical Center

⁶Neurosurgery, Massachusetts General Hospital

Current communication prostheses do not directly decode language intent but instead utilize indirect information, such as imagined movement or the P300 potential, to choose letters, words or symbols. A more powerful and intuitive approach for such systems would decode language-related neural activity directly. To investigate the possibility of extracting language content from neural activity, we examined multiscale spatiotemporal representations of semantic category as well as individual words via intracranial EEG in 13 patients undergoing invasive monitoring for epilepsy. Subjects were required to recognize a set of visually or acoustically presented nouns which were either animals or non-living objects. Using support vector machines, a class of high-dimensional machine learning algorithms, we demonstrate robust decoding of both semantic category and individual word representations from these recordings despite high intersubject variability in type and placement of electrodes. Information was found at multiple spatial scales including subdural macroelectrode recordings, microelectrode multiunit activity (MUA) and current source density (CSD) measurements, and single unit firing rates. Achieved accuracies were as high as 91% using macroelectrode features and 90% using microelectrode features when decoding semantic category (chance = 50%). When decoding between individual words, macroelectrode features yielded accuracies of up to 62% and microelectrode features yielded accuracies of up to 52% (chance = 20%). Combining features from multi-

ple spatial scales resulted in increased decode performance, suggesting the information provided at the different spatial resolutions is not identical. Examination of the most informative features indicates both medial temporal structures and diverse areas of lateral neocortex, including perirhinal, temporal, and lateral occipital areas, contribute to the representation of individual words and animal/object category. Utilizing a hierarchical decision tree to discriminate specific word properties sequentially utilizing appropriate features at each stage further improves performance over a single multiclass decoder. These results suggest that both word and category-specific information is present in intracranially recorded neural activity and that the information at different spatial scales, from population recordings to single unit firing rates, is not entirely redundant. Furthermore, decoder weights suggest that both medial temporal structures and lateral cortex contribute to the representation of animal and non-living objects categories as well as individual words. These signals may eventually serve as the basis for a rapid and intuitive language prosthetic device.

doi:

III-69. Exploring the statistical structure of large-scale neural recordings using a sparse coding model

Amir Khosrowshahi¹

Jonathan Baker²

Roger Herikstad³

Shih-Cheng Yen³

Christopher J. Rozell⁴

Bruno A. Olshausen⁵

AMIRK@BERKELEY.EDU

JBAKER@CNS.MONTANA.EDU

ROGER.HERIKSTAD@GMAIL.COM

SHIHCHENG@NUS.EDU.SG

CROZELL@GATECH.EDU

BAOLSHAUSEN@BERKELEY.EDU

¹Redwood Center for Theoretical Neuroscience, University of California, Berkeley

²Weill Cornell Medical College, New York

³National University of Singapore

⁴Georgia Institute of Technology

⁵University of California, Berkeley

We present a method for exploratory data analysis that attempts to learn the underlying statistical structure of large-scale neural recordings via a sparse coding model (Olshausen and Field, 1996; Olshausen 2003). In this model, a vector time series of filtered electrode potentials is represented as the sum of a convolution of latent sparse coefficients with a set of kernels. The set of kernels are learned from the data by maximizing the log-likelihood of the model through stochastic coordinate-wise ascent separately in the sparse coefficients and kernel parameters. At each step, a random batch of data is used to infer sparse coefficients using the current kernel basis, then the kernels are updated to minimize residual error. The algorithm was implemented in parallel with the inference step scaling linearly in the number of processes. The efficiency of the learning step was improved by the use of a 2nd-order stochastic gradient method. The model was applied to data recorded from visual cortex of an anesthetized cat viewing full-field natural movies. The recording device was a single shank polytrode with 54 contacts staggered vertically in two columns inserted perpendicular to the cortical surface and spanning all layers. The model was applied separately to bandpass filtered data (500Hz-10kHz) and to LFP (0-150Hz) from a single penetration. For the bandpass data, a subset of learned kernels consisted of localized waveforms that corresponded closely with spike waveforms estimated through a separate cluster-based spike sorting algorithm. For the LFP, the learned kernels recovered the dominant oscillation frequencies in the gamma range that were a characteristic feature of this particular recording. While the biophysical underpinnings of the learned kernels in some cases are unknown, the method provides a data-driven parsing of the multi-channel recordings into different groups according to the spatiotemporal statistics of neural activity. These groups, or "virtual units," thus serve as additional candidates for regressing with features of the stimulus, along with single-unit activity, to help unravel the ensemble coding and circuit dynamics that underlie cortical computations.

doi:

III-70. Modulation of STDP by the structure of pre-post synaptic spike times

Gordon Pipa¹
 Marta Castellano¹
 Raul Vicente¹
 Bertram Scheller²

PIPA@MPIH-FRANKFURT.MPG.DE
 CASTELLANO@MPIH-FRANKFURT.MPG.DE
 VICENTE@MPIH-FRANKFURT.MPG.DE
 SCHELLER@EM.UNI-FRANKFURT.DE

¹MPI for Brain Research

²Clinic for Anaesthesiology, J. W. Goethe Univ

Spiking activity recorded in vitro and in vivo preparations often contains rich auto structure. This auto structure can range from bursty to regular renewal processes, and can be even non renewal involving very complex temporal spiking pattern. Despite the fact that such auto-structure has been discussed to be involved in encoding information it is mostly ignored in modeling studies. Instead, modeling often assumes that spiking activity can be described by Poissonian firing. Here we studied the impact of non Poissonian renewal activity on structure formation by spike timing dependent synaptic plasticity. To this end we simulated a conductance based integrate and fire neuron that received input from 200 to 2500 inhibitory and excitatory neurons. This presynaptic activity was modeled by renewal processes with gamma distributed inter-spike interval (ISI) distributions. Using such a gamma process allowed us to systemically vary the regularity, ranging from Poissonian firing for a gamma processes with a shape factor on 1 (coefficient of variation of the ISI distribution, $CV=1$) to extremely regular firing with a shape factor of 100 ($CV=0.1$). In a first step we show that the temporal structure of post synaptic firing depends critically on the auto structure of the presynaptic activity even if the presynaptic population contain on the order of a couple of thousands mutually independently firing neurons. This clearly counter argues the assumption that a large enough presynaptic population can be modeled as Poissonian population activity, and raises the question to what degree this dependence of the post synaptic firing onto the pre-synaptic auto structure also modulated synaptic plasticity. In a second step we investigate this impact. We show that the auto-structure of the pre-synaptic activity first changes the steady state distribution of synaptic weights. Second, that the learning rate, that is the speed with which the distribution of weights is changing over time, is substantially higher than in the case of regular firing neurons than in the case of Poissonian firing. Third, that the impact of non-Poissonian firing is also modulated by the rate distribution of the presynaptic population, as well as by the rate relation between the populations of excitatory and inhibitory neurons. If both the excitatory and inhibitory population fires with a regular gamma process the learning rate is highest when the rate of the excitatory and inhibitory scales as $n:1$ with n being an integer number. Our findings give rise to new modulatory effects of synaptic learning due to STDP. Such that the learning rate of STDP can be regulated by a modulation of the autostructure of the spike trains, i.e. the regularity of presynaptic firing and its rate (in case that firing is not Poissonian). Both effects are frequently observed in neuronal recording and had been associated with attention, learning, short term memory and other cognitive tasks. In this light our findings might establish a link between neuronal activity changes and task or behavioral modulation of learning and structure formation in recurrent networks. Acknowledgements: Research was funded by the European Community GABA project (FP6-NEST No. 043309)

[doi:](#)

III-71. Spatio-temporal credit assignment in population learning

Johannes Friedrich
 Robert Urbanczik
 Walter Senn

FRIEDRICH@PYL.UNIBE.CH
 URBANCZIK@PYL.UNIBE.CH
 SENN@PYL.UNIBE.CH

Department of Physiology, University of Bern

Learning by reinforcement is important in shaping animal behavior, and in particular in behavioral decision making. Such decision making is likely to involve the integration of many synaptic events in space and time. However, using a single reinforcement signal to modulate synaptic plasticity, as suggested in classical reinforcement learning algorithms, a twofold problem arises. Different synapses will have contributed differently to the behavioral deci-

sion, and even for one and the same synapse, releases at different times may have had different effects. Here we present a plasticity rule which solves this spatio-temporal credit assignment problem in a population of spiking neurons. The learning rule is spike-time dependent and maximizes the expected reward by following its stochastic gradient. Synaptic plasticity is modulated not only by the reward, but also by a population feedback signal. While this additional signal solves the spatial component of the problem, the temporal one is solved by means of synaptic eligibility traces. In contrast to temporal difference (TD) based approaches to reinforcement learning, our rule is explicit with regard to the assumed biophysical mechanisms. Neurotransmitter concentrations determine plasticity and learning occurs fully online. Further, it works even if the task to be learned is non-Markovian, i.e. when reinforcement is not determined by the current state of the system but may also depend on past events. The performance of the model is assessed by studying three non-Markovian tasks. In the first task, the reward is delayed beyond the last action with non-related stimuli and actions appearing in between. The second task involves an action sequence which is itself extended in time and reward is only delivered at the last action, as it is the case in any type of board-game. The third task is the inspection game that has been studied in neuroeconomics, where an inspector tries to prevent a worker from shirking. Applying our algorithm to this game yields a learning behavior which is consistent with behavioral data from humans and monkeys, revealing themselves properties of a mixed Nash equilibrium. The examples show that our neuronal implementation of reward based learning copes with delayed and stochastic reward delivery, and also with the learning of mixed strategies in two-opponent games.

doi:

III-72. Single-neuron spike timing depends on global brain dynamics

Canolty Ryan
Karunesh Ganguly
Steven Kennerley
Kilian Koepsell
Charles Cadieu
Jonathan Wallis
Jose Carmena

University of California, Berkeley

RCANOLTY@GMAIL.COM
KARUNESH.GANGULY@UCSF.EDU
SKENNERLEY@BERKELEY.EDU
KILIAN@BERKELEY.EDU
CADIEU@BERKELEY.EDU
WALLIS@BERKELEY.EDU
CARMENA@EECS.BERKELEY.EDU

Brain rhythms have been suggested to play a role in both local cortical computation and long-range communication between areas, with a dynamic hierarchy of neuronal oscillations modulating activity at a variety of levels within the multi-scale brain network. Evidence supporting this view includes the micro-scale dependence of neuronal spiking upon the proximal local field potential (LFP), the meso-scale dependence of spiking upon spatial voltage patterns, and the macro-scale dependence of spiking upon phase coherence between distinct cortical areas. However, few previous studies examined multiple scales simultaneously and therefore the relative importance of these distinct spatiotemporal scales upon spike timing remains unknown. To investigate this issue, multiple single-unit spike times and LFPs were recorded via electrode arrays implanted bilaterally in either motor or prefrontal regions (primary motor and dorsal premotor sites in two subjects; cingulate, orbitofrontal, dorsolateral and ventrolateral prefrontal sites in two subjects). Here we show that many neurons exhibit a preference for distinct, frequency-specific multi-scale patterns of LFP phase coupling. In particular, using a probabilistic multivariate phase model, we show that the large-scale phase coupling structure present during the spike times of a given neuron is similar over repetitions of the same experimental task but is significantly different from the baseline phase coupling structure. Interestingly, however, the preferred phase coupling structure for two different neurons located in the same cortical area can be quite dissimilar, and the pattern of phase coupling for one given neuron can shift as a function of the behavior task. We discuss the implications these findings have for different models of neuronal computation and long-range communication, and consider ways in which future work may use this multi-scale dependence to improve decoding in brain-machine interfaces.

doi:

III-73. Temporal precision of the olfactory system

Roman Shusterman¹

Matt Smear¹

Thomas Bozza²

Dmitry Rinberg¹

¹Janelia Farm, HHMI

²Northwestern University

SHUSTERMANR@JANELIA.HHMI.ORG

SMEARM@JANELIA.HHMI.ORG

BOZZA@NORTHWESTERN.EDU

RINBERGD@JANELIA.HHMI.ORG

Olfaction is traditionally considered a 'slow' sense, but recent evidence demonstrates that rodents are capable of making extremely difficult odor discriminations rapidly, in as little as a single sniff. But can olfaction operate at even faster time scales? Odors can vary on a time scale finer than the period of a sniff cycle (roughly 100-500 ms) - odors are temporally structured by air turbulence, the sniff waveform, and the complex geometry of the nasal turbinates. Does the olfactory system preserve this precise timing information? And can an animal use this faster variation to guide behavior? To understand the temporal aspects of olfactory information processing, we combine electrophysiological, optogenetic, and behavioral approaches. We have analyzed the temporal structure of mitral cell activity in awake mice. We find very precise spiking patterns relative to the sniffing cycle, with jitter as small as 10 ms during odor presentation. This precise timing may carry information about the stimulus during odor presentation: different odors evoke different patterns in the same cell and different cells respond differently to the same odors. Moreover, we observe that the spontaneous activity of mitral cells (in the absence of odorants) reaches a minimum during the inhalation phase of the sniffing cycle. Interestingly, we find that odor responses begin during this period of decreased spontaneous activity, and have precise temporal structure during this phase of the sniffing cycle. The data suggest that the suppression of mitral cell activity during the inhalation phase modulates responses to odorants. To better control the timing of our stimulus, we have generated transgenic mice that express Channelrhodopsin2 (ChR2) in all olfactory sensory neurons (OSNs) from the OMP locus, which allows us to decouple sniffing from stimulus delivery by stimulating sensory neurons with light. In our ongoing work we are measuring the difference of mitral cell responses to the same stimulus at different phases of the sniffing cycle. Our preliminary data indicate that light pulses given during inhalation and exhalation evoke different mitral cell responses. To ask whether mice can perceive the timing of OSNs activation relative to sniff, we trained the mice in a head-fixed go-no-go discrimination paradigm. After OMP-ChR2 mice learn to report light detection, we next train them to discriminate identical light pulses solely on the basis of when in the sniff cycle they occur. We find that mice can discriminate temporal differences as small as 10 ms relative to the sniff cycle. Precise temporal neuronal responses to odors and the ability of mice to discriminate spatially-identical OSN stimuli on the basis of sub-sniff-cycle timing differences provide strong evidence that timing plays as crucial a role in olfactory neuronal coding and perception as it does in other sensory modalities. Olfaction may not be such a 'slow' sense.

[doi:](#)

III-74. Complementary encoding of sound features by local field potentials and spikes in auditory cortex

Stephen V. David

Nima Mesgarani

Serin Atiani

Shihab A. Shamma

University of Maryland, College Park

SVD@UMD.EDU

NIMA@GMAIL.COM

SATIANI@UMD.EDU

SAS@UMD.EDU

The local field potential (LFP) is believed to represent synchronous inputs to neurons near the site in the brain where it is measured. Therefore the functional relationship between the LFP and spiking activity of single neurons at the same site should describe computations performed in that brain area. To study this relationship, we simultaneously recorded single unit spike and LFP activity from the primary auditory cortex (A1) of awake, passively listening ferrets during the presentation of rippled noise and speech stimuli. We used the method of linear stimu-

lus reconstruction to compare stimulus spectrograms encoded by the spikes and LFP. For both rippled noise and speech, spectrograms were reconstructed more accurately from joint spike-LFP signals than from either spikes or LFP alone. Thus the different signals encoded complementary spectro-temporal stimulus features. This difference was most striking for temporal modulations. The LFP encoded high rate (16-50 Hz) modulations more accurately while spikes encoded low rates (4-16 Hz) more accurately. In order to explore the specific relationship between spikes and LFP at individual recording sites, we compared the spectro-temporal receptive field (STRF) for each signal measured from the rippled noise stimuli. We found a basic correspondence between spike and LFP STRFs (best frequency, latency, etc.). However, tuning to temporal modulations paralleled the results of the reconstruction analysis. LFP STRFs tended to be tuned to higher temporal modulation rates than spike STRFs. These results indicate that synchronous inputs to A1 emphasize higher temporal modulations than the spiking activity of single neurons. Spike responses appear to transform the representation of faster modulations into a rate code that cannot be recovered with linear reconstruction. In addition, these results illustrate that the LFP in A1 can be modulated by stimuli at rates up to 50 Hz, and possibly higher. This stands in contrast to visual cortex, where activity in this gamma frequency range is usually attributed to intrinsic oscillations rather than stimulus-evoked responses. Support: NIH R01DC005779, K99DC010439

doi:

III-75. Modeling peripheral auditory processing in the precedence effect

Jing Xia
Barbara Shinn-Cunningham
Boston University

JXIA@BU.EDU
SHINN@BU.EDU

The "precedence effect" (PE) is a perceptual phenomenon whereby a pair of clicks presented with a brief inter-stimulus delay (ISD) are typically heard as a single "fused" sound image whose perceived direction is near the location of the leading part of the stimulus. The PE is used to explain robust perception of sound location in reverberant environments. Specifically, localization is made reliable because listeners attribute greater perceptual weight to the localization information reaching the ears first, minimizing the influence of later-arriving reflections or echoes. The vast majority of past studies of the PE explored situations in which listeners perceived two temporally close, wideband clicks of equal amplitude. For such clicks, peripheral, monaural interactions between the phases of the lead and the lag in the cochlea result in internal, effective interaural time differences (ITDs) that may differ from those imposed on the external stimuli. Moreover, the onset adaptation of the inner hair cell - auditory nerve (IHC-AN) synapses may contribute to the localization dominance of the leading click by reducing the responses to the lagging click, which arrive soon after the initial activation. In the current study, subjects were asked to match the intracranial lateral position of an acoustic pointer to that of PE stimuli. To explore the peripheral interactions in the PE, localization was measured for wideband clicks as well as narrowband clicks centered at 500 Hz, with 1- and 2-ms ISDs. Stimuli were specifically chosen to produce very different internal ITDs, based on preliminary modeling results. To explore the effects of onset adaptation in the responses, the intensity of the leading click was attenuated parametrically. As controls, in some trials a single click was presented from the leading or lagging location. A neural model was used to predict behavioral results. Basilar membrane (BM) responses in the cochlea were simulated by bandpass filtering the acoustic stimuli. The BM outputs fed a biologically plausible AN model that included a three-compartment diffusion model to simulate adaptation of the IHC-AN synapses. The left- and right-ear AN outputs were used to compute the interaural cross-correlation (IACC). The IACC functions of the AN outputs were summed across frequency for each interaural delay. Finally, the peak of the across-frequency sum gave the estimate of the perceived ITD of the stimuli. We show that dominance of the lead diminished gradually as the lead attenuation increased. For narrowband clicks, subjects perceived the PE stimuli at locations predicted by the internal ITDs resulting from peripheral interactions in the cochlea, not at the "true" ITDs in the external stimuli. The perceived location was more lateral (to either the leading or the lagging side) for paired clicks with the 1-ms ISD than the 2-ms ISD. For wideband clicks, the lead dominance was stronger for the 2-ms ISD than the 1-ms ISD. Behavioral results were well explained by the model that included peripheral interactions and firing-rate adaptation. Results demonstrate that BM filtering, combined with nonlinearities in the cochlea, help lead to the perceptual dominance of a preceding sound on localization.

doi:

III-76. Auditory cortex neuronal tuning is sensitive to statistical structure of early acoustic environment

Hania Koever
Yi-Ting L. Tseng
Kirt Gill
Shaowen Bao
UC Berkeley

HANIAKOEVER@BERKELEY.EDU
LESLIE_YTTSENG@BERKELEY.EDU
GHANI003@GMAIL.COM
SBAO@BERKELEY.EDU

Whereas sensory stimuli vary along a continuum, behavioral responses to these stimuli are discrete. The perceptual system can help translate sensory inputs into behaviors by parsing continuous stimuli into discrete perceptual categories. Some forms of categorical perception are innate and may reflect intrinsic characteristics of sensory processing; others, however, depend on learning and experience. In humans, for example, native English speakers perceive a distinct difference between the speech sounds /la/ and /ra/, whereas native Japanese speakers do not. Likewise, some species of songbirds categorically perceive note duration in a manner that depends on experience with their specific native song dialect. Despite evidence for experience-dependent categorical sound perception in a range of species from humans to rodents, it is not clear what rules neurons use to segment acoustic inputs into categories. Evidence from linguistics suggests that the statistical structure of acoustic inputs during an early stage of development affects categorical perception. For example, the distribution of speech sounds plays a role, with bimodal, but not unimodal distributions of speech sounds leading to categorical perception in infants. Additionally, infants are sensitive to the transitional probabilities between speech sounds, and perceptually group sounds that repeatedly occur together. To investigate the effect of the statistical structure of early acoustic inputs on neural tuning in auditory cortex, we exposed young rats (postnatal day (p)9 to p45) to sequences of pure tones that were grouped into two categories based on transitional probabilities alone. Pure tones were organized into one-second long sequences of six tones, in which all tones within a sequence were drawn from either a low frequency or high frequency category of tones. In this design, tones within a category had a high probability of occurring successively within a sequence (high transitional probability), whereas tones from different categories never occurred within a sequence (low transitional probability). We used extracellular recording in anaesthetized rats to map primary auditory cortex, and compared neural tuning in exposed rats to that of naïve rats, as well as rats exposed to control tone sequences with homogeneous statistics. Neurons in primary auditory cortex of the categorical sequence group differed from those in the naïve and control groups in terms of the shape of their frequency response tuning curves. In particular, neurons tuned to frequencies near category boundaries had steeper slopes at the category boundary. In addition, receptive fields exhibited a tendency to code for either the low or high category, rather than straddling two categories. In conclusion, neuronal tuning in primary auditory cortex is sensitive to the transitional probabilities of sounds encountered during an early period of development. We are currently performing behavioral experiments to determine whether the altered neuronal tuning has a perceptual correlate.

doi:

III-77. Up-states are rare in awake auditory cortex

Tomas Hromadka¹
Michael DeWeese²
Anthony M. Zador¹

HROMADKA@CSHL.EDU
DEWEESE@BERKELEY.EDU
ZADOR@CSHL.EDU

¹Cold Spring Harbor Laboratory

²University of California, Berkeley

Up and down states—depolarized and hyperpolarized plateaus of membrane potential—have attracted a considerable attention, as they have been suggested to underlie persistent activity in cortical networks. The persistent activity, in turn, has been hypothesized to play a role in such interesting processes as short-term (working) memory, or attention. Large, persistent changes in membrane potential can also occur in a response to behavioral state changes, or sensory stimulation. Up and down states have been described in frontal cortical areas, somatosensory, visual, olfactory areas, striatum, etc. However, the nature and prevalence of up states have not been well studied in the auditory cortex. We have previously found that membrane potential in the auditory cortex consists of bumps, i.e. brief, stereotyped excursions of membrane potential. By contrast, 'canonical' up-states are usually thought of as stereotyped plateaus at which the membrane potential remains for a prolonged period (hundreds of milliseconds to seconds). Answering the question of prevalence of up-states is complicated by the fact that there is no universally agreed upon definition of what constitutes up and down states. To characterize the nature and prevalence of up-states in the awake auditory cortex, we used whole-cell patch-clamp recording techniques to record intracellular activity of neurons (n=20) in auditory cortex of awake head-fixed rats (Sprague-Dawley, postnatal day 24-30). We have generalized the definition of up-states, and defined several membrane potential thresholds. We found that long up-states were rare in awake auditory cortex, with only 4 % of up-states longer than 200 ms. Most neurons displayed only brief up-states (bumps) and spent most of the time near their resting potential, typically spending less than 4 % of recording time in up-states longer than 200 ms. It is unclear why auditory cortex would differ from other sensory areas in terms of subthreshold dynamics. Whether the apparent lack of up-states in the auditory cortex reflects an inherent difference in cortical networks between auditory and other sensory areas, or is simply caused by difference in nomenclature (and bumps are simply brief up-states) remains an open question.

[doi:](#)

III-78. Hearing the song in noise

Moore Richard¹
Patrick R. Gill²
Frederic E. Theunissen³

CHANNING_MOORE@BERKELEY.EDU
PRG56@CORNELL.EDU
THEUNISSEN@BERKELEY.EDU

¹UC Berkeley

²Cornell University

³UC Berkeley Department of Psychology

One major task of the auditory system is to pick out and process relevant sounds against a background of noise. Some degree of background rejection can be achieved by having cells tuned selectively for properties found in sounds of interest but not present in background noise. This scheme fails, however, to separate a relevant sound from a background of like sounds. We therefore propose that more complicated mechanisms may be at work. We recorded from auditory fields L and NCM in urethane-anaesthetized zebra finches while playing zebra-finch songs and modulation-limited noise. We presented each stimulus alone, and we also combined each song with several different noise stimuli. This resulted in three conditions: noise alone; song alone; and song plus noise. We compared the PSTH for the combined trials with the respective PSTHs for each stimulus alone. We also computed receptive fields (STRFs) from the noise-alone and song-alone trials, and compared the predicted response to the combined stimuli with our observed result. In both cases, we found that responses to sounds masked by noise were better predicted by song-alone responses, indicating that these areas employ a more complex form of stream segregation than matched filters.

[doi:](#)

III-79. A normalization model of multi-sensory integration

Tomokazu Ohshiro¹
Dora E. Angelaki²
Gregory C. DeAngelis¹

TOSHIRO@BCS.ROCHESTER.EDU
ANGELAKI@PCG.WUSTL.EDU
GDEANGELIS@CVS.ROCHESTER.EDU

¹University of Rochester

²Washington University School of Medicine

Many physiological studies have examined how multisensory neurons respond when inputs from two different modalities (e.g. visual, auditory) are presented together and separately. From these studies, a set of empirical principles have emerged, including the principle of inverse effectiveness, the spatial principle, etc. However, there has been no general computational framework that accounts for these empirical properties. We propose that a multi-sensory version of divisive normalization, inspired by the work of Heeger and colleagues, can account for a wide variety of findings on multisensory integration. We also describe a physiological test of the normalization model. Each multisensory neuron in our model receives inputs from two different sensory modalities, and performs a linear weighted sum of those inputs followed by a static nonlinearity (e.g., rectification and squaring) that models the transformation of membrane potential to firing rate. The response of each model neuron is then divided by the summed activity of all other neurons in the pool, including neurons with a wide variety of receptive field locations and tuning properties. An optional component is synaptic depression of the unisensory inputs to the multisensory neuron. Our normalization model accounts for many of the most salient properties described in the physiology literature, including: 1) Inverse effectiveness. Model neurons exhibit super-additive interactions when both inputs have low intensities, and sub-additive interactions when one or both inputs have high intensity. The expansive nonlinearity (i.e., squaring) drives super-additivity at low intensities, whereas normalization produces sub-additivity for high intensities; 2) The spatial principle. When one sensory input is made non-optimal, multisensory neurons often show responses that are suppressed relative to the stronger input. The normalization model naturally explains this because the suboptimal stimulus contributes little to the underlying linear response of the neuron, but still contributes strongly to the normalization signal; 3) Uni-sensory vs. multi-sensory interactions. Physiology studies report that two stimuli presented to the same sensory modality (e.g., vision) interact sub-additively whereas two stimuli of different modalities can produce super-additive interactions. Incorporating a saturating nonlinearity (e.g., synaptic depression) into the unisensory inputs to the model naturally accounts for this finding. To test for the existence of a multisensory stage of normalization, we recorded from neurons in the medial superior temporal area (MSTd) of macaque monkeys, where visual (optic flow) and the vestibular cues to self-motion are integrated. By manipulating visual and vestibular stimulus directions, we tested a key prediction of the model which is analogous to the spatial principle: a non-optimal stimulus for one modality, which is excitatory on its own, should be able to suppress a near-optimal response to the other modality when the cues are combined. This prediction was confirmed for neurons in MSTd. This finding is readily explained by normalization at the network level, but cannot be explained by any simple form of nonlinearity in the unisensory processing pathways prior to multisensory integration or within the multisensory neuron itself. In closing, normalization may provide a unifying computational framework to account for many aspects of multisensory integration in single neurons.

doi:

III-80. Odor trail tracking by rats in surface and air borne conditions

Adil Khan
Urvashi Raheja
Upinder Bhalla

ADIL@NCBS.RES.IN
URVASHI@NCBS.RES.IN
BHALLA@NCBS.RES.IN

National Centre for Biological Sciences

Rats, being nocturnal, rely heavily on their sense of smell. They have long been known to be proficient at odor-based tasks in the lab. Precise measurements by ours and other groups have revealed abilities to both identify and localize odors within the timescale of a single sniff (~150ms). They have also been known to be capable of locating

food in large arenas such as households and fields. These feats presumably use air-borne and surface borne odor trails. We have attempted to study the nature of this trail tracking in both these conditions. In this process we have tried to establish behavioral methods which study odor guided navigation in more natural behavioral tasks yet with well controlled stimulus conditions. To study surface borne odor tracking, we constructed a treadmill which had a sheet of paper as the running surface. The paper had an odor trail in the form of a chocolate piece rubbed in a meandering line on the paper. Rats learned to track this trail while running on the treadmill. Video analysis which tracked their nose movement revealed their tracking strategies. Since this primary task provides visual and tactile cues also, the task was compared to a condition with purely odor cues also. To study air borne plume tracking, we built an arena where rats were trained to locate the source of an air borne odor emerging from one of multiple compartments. Rats learnt this task and showed a range of different localization strategies. We further characterized the nature and degree of laminarity of the odor plume, and correlated different plume characteristics with different strategies used by the rats.

doi:

III-81. The flow of expected and unexpected sensory information through the distributed forebrain network

Maolong Cui¹
Jozsef Fiser¹
Don Katz¹
Alfredo Fontanini²

MLCUI@BRANDEIS.EDU
FISER@BRANDEIS.EDU
DBKATZ@BRANDEIS.EDU
ALFREDO.FONTANINI@STONYBROOK.EDU

¹Department of Psychology, Brandeis University

²Dpt Neurobiology&Behavior, SUNY Stony Brook

Forebrain taste information processing is accomplished mainly by three reciprocally connected forebrain regions - primary gustatory cortex (GC), (basolateral) amygdala (AM), and orbitofrontal cortex (OFC)- loosely characterized as the neural sources of sensory, palatability-related, and cognitive information, respectively. It has been proposed that the perception of complex taste stimuli involves an intricate flow of information between these regions in real time. However, empirical confirmation of this hypothesis and a detailed analysis of the multidirectional flow of information during taste perception have not yet been presented before. We have simultaneously recorded local field potentials from GC, AM, and OFC in awake behaving rats under two conditions as controlled aliquots of either preferred or not preferred taste stimuli were placed directly on their tongues via intra-oral cannulae. Half of the deliveries were "active", as the rat pressed a bar to receive the taste upon receiving an auditory 'go' signal, the other half of deliveries were "passive" when the rat received a tastant at random times. Peri-delivery signals from the three areas were analyzed by computing transfer entropy, a method that measures directional information transfer between coupled dynamic systems by assessing the reduction of uncertainty in predicting the current state of the systems based on their previous states. The results of this analysis reveal the complexity and context specificity of perceptual neural taste processing. Passive taste deliveries caused an immediate and strong flow of information that ascended from GC to both AM and OFC ($p < 0.001$). However, within the 1.5-2.0 sec in which our rats typically identified and acted on (swallowing or expelling) the tastes, feedback from AM to GC became a prominent feature of the field potential activity ($p < 0.001$). This finding confirms and extends earlier single cell results showing that palatability-related information appears in AM single-neuron responses soon after taste delivery, and that there is a sudden shift in the content of both GC and AM single-neuron responses at ~1.0 sec following delivery, as palatability-related information appears in GC and subsides in AM. The neural response to active taste deliveries differed from that to passive deliveries in important ways. The massive immediate GC to AM/OFC flow was greatly decreased and delayed. Instead, there was an increased and lasting information flow from OFC to GC ($p < 0.01$) immediately after the tone. The likely reason for this reduction was obvious: tone onset led to an anticipation of taste delivery that activated a descending flow of information from the "cognitive centers" in OFC to the primary sensory cortex, which greatly changed the actual neural processing of the stimulus itself in GC. These results place earlier single-neuron findings into a functional dynamic framework, and offer an explanation of how the parts of the sensory system work together to give rise to complex perception. They suggest

that perception is not a simple bottom-up process in which a stimulus is coded by progressively "higher" centers of the brain, rather various bottom-up and top-down effects jointly define and greatly alter stimulus processing even in the primary sensory areas.

[doi:](#)

III-82. Sparse coding of natural stimuli in the midbrain

Maurice J. Chacron

MAURICE.CHACRON@GMAIL.COM

Department of Physiology

Sparse neural codes (i.e. codes in which neurons respond selectively to few sensory stimuli) have been observed widely across animal taxa including the human hippocampus, monkey visual cortex, locust olfactory system, and songbird HVC. Theoretical studies suggest that the generation of sparse neural codes critically depends on non-linear mechanisms. Systems with easily characterized natural stimuli and anatomy are expected to yield significant insight into the nature of these non-linear mechanisms. Weakly electric fish provide an attractive model system for studying sparse coding due to their well-characterized anatomy and physiology. We recorded from neurons located within the midbrain Torus Semicircularis (TS, equivalent to the inferior Colliculus) using the patch clamp technique in vivo. We used transient natural communication stimuli called chirps that must be distinguished from a background stimulus. We found that some TS neurons fired either a single action potential (type A) or a burst of action potentials (type B) only in response to the chirp and not to the background. In contrast, afferent neurons have been shown to respond to both the background and the chirp. The TS neurons thus provide a neural correlate of segregating a transient signal from a background. Most surprisingly, some neurons were able to distinguish between background and chirp even when the two stimuli had similar temporal frequency content. We used a combination of mathematical modeling and in vivo manipulations of neural activity via pharmacology and current injection to reveal the mechanisms that make TS neurons selective for chirp stimuli. We found that, in type A neurons, the background caused a hyperpolarization in the membrane potential via shunting inhibition below the spiking threshold. The increased membrane conductance promoted spiking in response to coincident activity elicited by the chirp stimulus and activation of high-threshold potassium channels prevents further spiking in response to the chirp. In contrast, we found that the mechanism by which type B neurons selectively responded to chirps involved integration of input from both ON-type and OFF-type afferent neurons. These two neuron types are shown to be approximately out of phase in their linear responses to the background but approximately in phase in their nonlinear responses to the chirp. Simple integration of both of these inputs by TS neurons is therefore sufficient to explain why they only respond to the chirp. These results show a mechanism by which neurons can segregate streams of information and may have applications to other systems such as detecting transients occurring over a background in both the visual and auditory systems.

[doi:](#)

III-83. 2D encoding of concentration and concentration gradient in *Drosophila* ORNs

Anmo J. Kim

Aurel A. Lazar

Yevgeniy Slutskiy

Columbia University

AK2263@COLUMBIA.EDU

AUREL@EE.COLUMBIA.EDU

YS2146@COLUMBIA.EDU

The lack of a deeper understanding of how olfactory receptor neurons (ORNs) encode odors has hindered progress in understanding olfactory signal processing in higher brain centers. We investigate the encoding of time-varying odor stimuli by *Drosophila* ORNs and their spike domain representation for further processing by the

network of glomeruli. We built a novel low-turbulence odor delivery system that enables a precise control and reproducible delivery of airborne odorants. The system provides exact control of odor concentration and concentration gradient on a millisecond time scale. Using a photo-ionization detector for monitoring odor concentration values in real-time, we found that odor waveforms reaching the antennae of fruit flies can be reproduced to within a tolerance of 1%. We augmented the odor delivery system with the capability of simultaneously recording the in-vivo extracellular activity of two ORNs. A wide range of time-varying odor waveforms were used in in-vivo recordings of ORNs expressing the same receptor. Spiking activity of single ORNs activated by essentially the same odor waveforms could be evaluated from repeated experiments for a wide range of concentration and concentration gradient value pairs. In order to evaluate the spike-timing precision, we simultaneously recorded from a single fly the activity of two neurons expressing the same receptor. Overall, we recorded the spiking activity of (i) neurons expressing different receptors in response to the same odorant, and (ii) neurons expressing the same receptor in response to different odorants. Our analysis of the recordings demonstrates that ORNs respond to a given time-varying stimulus with a high degree of spike-timing precision. This precision is conserved across multiple repetitions of the same time-varying odor waveform for ORNs expressing the same receptor. Further, we report a qualitatively stereotyped response to a given waveform across a population of excited ORNs in a single fly and across different flies. Based on an extensive analysis of the recordings, we propose a novel two-dimensional encoding paradigm for the representation of excitatory time-varying odor stimuli by fly ORNs. We identify odor concentration and its rate of change as the predominant odor characteristics determining the response of ORNs. Using concentration and concentration gradient as input variables, we construct a novel 2D encoding manifold in a three-dimensional space that characterizes the response of a neuron. We quantitatively show how to predict the response of an ORN to an excitatory time-varying stimulus by choosing an appropriate trajectory embedded in the 2D encoding manifold. Our work demonstrates an adaptive two-dimensional encoding mechanism for *Drosophila* ORNs. At very low odorant concentrations, ORNs encode positive concentration gradients. Conversely, at high concentrations ORNs encode the odorant concentration. The 2D encoding manifold clearly shows that *Drosophila* ORNs encode both odor concentration and concentration gradient and provides a quantitative description of the neural response with a predictive power not seen before. Acknowledgements. The work presented here was supported by NIH under grant number R01DC008701-01 and was conducted in the Axel laboratory at Columbia University. The authors would like to thank Dr. Richard Axel for insightful discussions and his outstanding support.

doi:

III-84. Interactions of rat whiskers with air currents: implications for flow sensing

Venkatesh Gopal¹
Minwoo Kim²
Charles Chiapetta²
Joel Russ²
Michael Meaden³
Mitra Hartmann⁴

VGOPAL@ELMHURST.EDU
MKIM212@ILLINOIS.EDU
CCHIAPP2@ILLINOIS.EDU
JOELW.RUSS@GMAIL.COM
MEADENM@ELMHURST.EDU
M-HARTMANN@NORTHWESTERN.EDU

¹Department of Physics, Elmhurst College

²University of Illinois at Urbana-Champaign

³Elmhurst College

⁴Northwestern University

Many species of mammals have a regular array of facial whiskers (mystacial vibrissae) that emerge from sensory follicles embedded in the cheek. Each whisker-follicle pair constitutes a highly sensitive mechano-transducer, and the whiskers are often used for the tactile exploration of objects. Rats, for example, actively sweep their whiskers against objects between 5 - 12 Hz in a behavior called "whisking." Using only whisking movements, a rat can tactually extract object features such as size, shape, orientation and texture. However, many animals with large and prominent vibrissae, such as dogs, do not actively whisk. Why, then, are the vibrissae so prominent, and so regularly arranged, even in species that do not actively whisk? We reasoned that animal nervous systems have

evolved around the need to find food, water, and mates. For most mammals, the search for these resources is largely olfactory, requiring the detection of airborne odorants which are usually present at very low concentrations. The dispersal of an airborne odorant is strongly dependent on local air currents, which are generally turbulent and have complex flow patterns. Olfactory localization thus requires an animal to detect the chemical odorant, to sense wind direction, and then to integrate these chemical and fluid dynamic cues to make a decision on where to move next. We therefore hypothesized that in addition to their direct tactile function, whiskers may also serve as sensitive detectors of air currents, and that the rat's vibrissal array can be used to measure information about local airflow. To test this hypothesis, we measured the interactions of vibrissae with air currents in anesthetized rats. Turbulent air streams were blown onto the vibrissae of an anesthetized rat at various angles. Whisker deflections were measured using high-speed video cameras at a frame rate of 1KHz. Whisker kinematic parameters were measured and compared across vibrissae on both sides of the face. Our preliminary results show significant differences in response frequency and amplitude across the right and left sides of the vibrissal array depending on the orientation of the air stream. We discuss these results in the context of odor localization behaviors. This work was supported by NSF awards IIS-0613568 and IOS-0818414 (MJZH) and an Elmhurst College Summer Research Collaboration Grant (VG, MM).

doi:

III-85. Re-testing the energy model: identifying features and nonlinearities of complex cells.

Timm Lochmann¹
Joseph N. Stember²
Tim Blanche³
Daniel A. Butts¹

TIMM.LOCHMANN@GMAIL.COM
JOESTEMBER@GMAIL.COM
TJB@BERKELEY.EDU
DAB@UMD.EDU

¹Department of Biology, University of Maryland

²Weill Cornell Medical College, New York

³Helen Wills Neuroscience Institute

Complex cells in the primary visual cortex are a prominent example for which the classical (linear) receptive field models fail. In particular, complex cells are defined by having orientation selectivity that is invariant with respect to spatial phase. This property of complex cells is explained by the energy model, whose output is calculated through input filters comprised of spatially orthogonal inputs that form a quadrature pair, which are then squared to result in phase invariance. While this model provides a possible mechanism to reproduce the observed phase invariance, it is not clear the degree to which either assumption is strictly true, and how they are implemented by neuronal circuitry. The assumptions of the energy model can be tested using appropriate models applied to extracellular recordings from complex cells in the primary visual cortex (V1). Previous studies have used spike-triggered covariance (STC) to identify directions in the stimulus space that V1 neurons are sensitive to. STC provides an orthogonal coordinate system for the stimulus subspace that the neuron is selective to, but has two main drawbacks: First, the features relevant to a cell do not have to be orthogonal and therefore STC does not allow for their identification directly. Second, especially in cases of correlated inputs corresponding to non-orthogonal features, STC analysis is not sufficiently sensitive and fails to identify all relevant directions for limited amounts of data. Given that many cells in the visual cortex receive inputs from multiple sources, this is of practical importance. We have developed a General Nonlinear Modeling (GNM) approach to estimate the relevant features and non-linearities underlying the spiking responses from extracellular data recorded in V1. Our approach makes only weak assumptions about the form of the nonlinearities and temporal dependencies, as it estimates their shape from data. This makes our approach well suited to test the assumptions underlying the energy model in an unbiased way, and to identify alternative solutions if they provide a better account of the data. Consistent with the predictions of the energy model (and previous findings using STC), we find two nearly orthogonal input features with symmetric, bowl shaped nonlinearities resembling the combined inputs of quadrature pairs. However, we show that the orientation of these orthogonal features in the subspace identified by STC does not affect the model predictions, suggesting that they represent the principle components of network inputs comprised of many similar

elements tuned to different phases. In addition to these filters, we also robustly detect a third input direction outside of the subspace delimited by STC. This input has a monotonic non-linearity, suggesting that it could stem from thalamic inputs or another simple cell. Including the third feature yields significantly better goodness of fit, raising the question of its role in stimulus processing in the context of phase invariance assumed in the energy model. Our results shed light on the functioning of complex cells, and more generally, they highlight how appropriate nonlinear modeling can add significant insight into what determines the response of sensory neurons.

[doi:](#)

III-86. Human versus machine: comparing visual object recognition systems on a level playing field.

Nicolas Pinto¹
 Najib J. Majaj²
 Youssef Barhomi²
 Ethan A. Solomon²
 David D. Cox³
 James J. DiCarlo²

PINTO@MIT.EDU
 MAJAJ@MIT.EDU
 BARHOMI@MIT.EDU
 ESOLOMON@MIT.EDU
 COX@ROWLAND.HARVARD.EDU
 DICARLO@MIT.EDU

¹MIT

²McGovern Inst/Dept of Brain & Cog Sci, MIT

³Rowland Institute, Harvard

It is received wisdom that biological visual systems easily outmatch current artificial systems at complex visual tasks like object recognition. But have the appropriate comparisons been made? Because artificial systems are improving every day, they may surpass human performance some day. We must understand our progress toward reaching that day, because that success is one of several necessary requirements for "understanding" visual object recognition. How large (or small) is the difference in performance between current state-of-the-art object recognition systems and the primate visual system? In practice, the performance comparison of any two object recognition systems requires a focus on the computational crux of the problem and sets of images that engage it. Although it is widely believed that tolerance ("invariance") to identity-preserving image variation (e.g. variation in object position, scale, pose, illumination) is critical, systematic comparisons of state-of-the-art artificial visual representations almost always rely on "natural" image databases that can fail to probe the ability of a recognition system to solve the invariance problem [Pinto et al PLoS08, COSYNE08, ECCV08, CVPR09]. Thus, to understand how well current state-of-the-art visual representations perform relative to each other, relative to low-level neuronal representations (e.g. retinal-like and V1-like), and relative to high-level representations (e.g. human performance), we tested all of these representations on a common set of visual object recognition tasks that directly engage the invariance problem. Specifically, we used a synthetic testing approach that allows direct engagement of the invariance problem, as well as knowledge and control of all the key parameters that make object recognition challenging. We successfully re-implemented a variety of state-of-the-art visual representations, and we confirmed the high published performance of all of these state-of-the-art representations on large, complex "natural" image benchmarks. Surprisingly, we found that most of these representations were weak on our simple synthetic tests of invariant recognition, and only high-level biologically-inspired representations showed performance gains above the neuroscience "null" representation (V1-like). While in aggregate, we found that the performance of these state-of-the-art representations pales in comparison to human performance, humans and computers seem to fail in different and potentially enlightening ways when faced with the problem of invariance. We also show how our synthetic testing approach can more deeply illuminate the strengths and weaknesses of different visual representations and thus guide progress on invariant object recognition.

[doi:](#)

III-87. Model of visual target detection applied to rat behavior

Philip Meier
 Pamela Reinagel
 University of California, San Diego

PMEIER@UCSD.EDU
 PREINAGEL@UCSD.EDU

The perception of visual features is influenced by the nearby spatial context of an image. Such contextual processing may be useful for natural vision, and can be experimentally revealed. In previous work, we have shown that rat's behavioral performance on a visual task is influenced by the relationship of features: specifically, collinear flankers make it more difficult to detect the presence of an oriented target. It is known that the contrast of the center target and surround flankers influence human and monkey behavior as well as cat and monkey neurophysiology. Here we present new rat behavioral data on the influence the contrast of collinear flankers at multiple target contrasts. Our main finding is that increasing flanker contrast impoverishes target detection and biases rats to report the presence of the target. Interestingly, surround contrast does not mask target detection which would bias the rat to report the absence of a target more often. This runs counter to predictions made by a divisive normalization model. Additionally, the presence of flankers decreases performance. This opposes the hypothesis that the primary influence of flankers is to reduce the uncertainty of the target location. Contextual effects in visual tasks could have multiple causes including normal perceptual processing, the task-specific allocation of attention, optical blurring, neural integration over visual space, and cognitive confusion. Perceptual processing is likely to be responsible for effects of the flankers that are sensitive to the precise geometric configuration, but here we address the influence of contrast within a single configuration by using an uncertainty model with a max decision rule (Pelli, 1985) to account for rats' behavior. Our model includes an attention term for the suppression of the task irrelevant distracters, and a spatial integration term that represents feature pooling. With a small number of parameters (attention spatial focus, feature pooling magnitude, channel noise, number of channels, contrast sensitivity function) the model qualitatively captures the trends of the flanker's influence on bias and performance. If a fixed decision threshold is used, the model predicts that subjects would have a variable hit rate and a constant false alarm rate for a given flanker contrast. However, rats display a false alarm rates that increase with flanker contrast. Rats probably change their decision criteria for different stimulus conditions because conditions are blocked into 150 trials with the same stimulus parameters. A model with an adaptive decision criterion is proposed to better fit the hit rates and false alarms. We note that the model only contains knowledge of contrast, and does not attempt to explain results that depend on interactions of orientations or stimulus configuration. Which stage of the rat's visual system is responsible for the influence of the flankers? To narrow down the cause of lateral interaction, we have begun to present collinear stimuli of varying contrasts while recording extra-cellular spikes from the lateral geniculate nucleus of the rat.

[doi:](#)

III-88. Do quantal dynamics of graded synaptic signal transfer adapt to maximise the rate of information?

Xiaofeng Li¹
 Shiming Tang¹
 Mikko Juusola^{2,1}

XIAOFENG.LI@SHEFFIELD.AC.UK
 TANGSHM@SUN5.IBP.AC.CN
 M.JUUSOLA@SHEFFIELD.AC.UK

¹Beijing Normal University

²University of Sheffield

In synapses, information in incoming voltage changes is converted to quantal bursts of neurotransmitter, released from vesicles to a synaptic cleft that separates neighbouring neurones. Changes in the neurotransmitter concentration are then picked up by specific receptor-complexes in the post-synaptic membrane, thereby channelling information back to voltage changes in the post-synaptic neurone. In the classic view, synaptic vesicles are uniform in size, with each carrying similar doses of neurotransmitter, enabling transmission of pulsatile messages. However, the rate of information transfer is much higher in graded potential synapses^{1, 2}, with some findings

suggesting that quantal vesicle release (output) may change with the transmitted messages (input). It remains an open question whether the dynamics of quantal vesicle release in graded potential synapses adapt to optimise the rate of information transfer. Flies have modular eye structure and characteristic layout, which enables reliable intracellular electrophysiology from individual neurones in their first synaptic processing layer, the lamina. The lamina contains a system of neurons consisting of photoreceptors (R1-R6) and interneurons: large monopolar cells (LMCs) and an amacrine cell that co-process visual information. While photoreceptors depolarize and LMCs hyperpolarize to light, owing to a web of feedback and feedforward synapses their graded voltage responses are shaped together³. Here we exploit in vivo fly preparations (*Calliphora* and *Drosophila*) to investigate quantal histaminergic transmission from photoreceptors to LMCs. We recorded voltage responses of photoreceptors and LMCs [i] to light backgrounds, [ii] to repeated pseudorandomly modulated light contrast patterns and [iii] to naturalistic light contrast series. [iv] We also recorded voltage noise in LMCs when synaptic output from photoreceptors was silenced by massive Na⁺-K⁺-exchanger driven hyperpolarisation, following intense light pulsation⁴. By analysing the signal and noise properties of synaptic throughput for the different stimulus conditions in the same neurones, we show that quantal transmitter release changes with changing light inputs. With low SNR inputs, the mean post-synaptic unitary events are large and slow, leading to low-passing responses with high gain. But with high SNR inputs, we find more, smaller and faster synaptic quanta, which sum up band-passing voltage responses with lesser gain. These results, together with our anatomical findings with electron microscopy, support the idea that the quantal vesicle release in photoreceptor-LMC synapses adapts to ongoing light inputs. By dynamically adjusting the size and numbers of the transmitted quanta, the photoreceptors-LMC synapse seems to strive to maximise the flow of visual information from the ever changing visual world. 1. Juusola, Uusitalo & Weckström. *J Gen Physiol* 105, 117-48 (1995). 2. de Ruyter van Steveninck & Laughlin. *Nature* 379, 642-645 (1996). 3. Zheng et al. *J Gen Physiol* 127, 495-510 (2006). 4. Uusitalo et al. *J Neurophysiol* 74, 470-3 (1995).

doi:

III-89. Bursts and visual encoding in LGN during natural state fluctuations in the unanesthetized rat

Erik D. Flister
 Pamela Reinagel
 UCSD

ERIK.FLISTER@GMAIL.COM
 REINAGEL@BIOMAIL.UCSD.EDU

We implanted chronic wells in adult Long-Evans rats over the stereotaxic coordinates of LGN. In multi-hour sessions over subsequent weeks, we recorded well-isolated extracellular spiking activity and LFP while the rats were headfixed but unanesthetized. Once a visually responsive unit ventral to hippocampus was isolated, a CRT or LED array presented spatially uniform flickering visual stimuli. We held cells as long as possible (average 1 hour) while displaying two 30 sec repeated flicker patterns. In the first pattern, frame luminances were drawn independently at 100Hz from a gaussian distribution with contrast chosen such that the central 99% of the values fell in the luminance range of the display; the 0.5% tails were clipped to the extrema. In the second pattern, the flicker reprised a 30 sec contiguous segment from a photodetector recording taken by a person walking through a forest in daylight [1]. The recording was resampled from 1200 to 100 Hz. In the case of the CRT, the values were scaled so that the brightest 1% were clipped, so that periods of dim flicker still resolved in the CRT's reduced dynamic range (8 bits) relative to the recording (15 bits). Units were well isolated, as verified by spike shape and refractory analysis. Single unit responses were not sparse, but showed strong spike triggered average stimuli. We identified putative low threshold calcium bursts by the criterion of 100 ms quiescence followed by interspike intervals of < 4 ms [2]. Most cells fired a mixture of both burst and tonic (non-burst) spikes. In many cases, bursts also showed a strong average triggered stimulus, corroborating earlier reports of visually driven LGN bursts in anesthetized cats [3] and demonstrating that these results were not merely an artifact of anesthesia. These data support the claim that visual stimuli sometimes hyperpolarize LGN neurons enough to deinactivate T-channels, priming a calcium burst to be triggered upon subsequent superthreshold stimuli, even in natural physiological states. Because the rats were unanesthetized, both LFPs and firing characteristics were highly nonstationary, varying along a continuum of states that likely correspond to alertness levels. Many recordings contained brief epochs (< 10 min)

of enhanced LFP power at < 20 Hz, coinciding with a high burst rate and low tonic firing rate. During other epochs, the same cell responded with a higher tonic firing rate, lower burst rate, and reduced rhythmicity. Here, we characterize these states in terms of their LFP spectral power, spike-field coherence, firing rate, burstiness, reliability, precision, sparseness, and stimulus encoding using reverse correlation and information measures. [1] van Hateren, JH. 1997. Processing of natural time series of intensities by the visual system of the blowfly. *Vision Res.* [2] Lu SM, Guido W, Sherman SM. 1992. Effects of membrane voltage on receptive field properties of lateral geniculate neurons in the cat: contributions of the low-threshold Ca_2 conductance. *J Neurophysiol.* [3] Denning KS, Reinagel P. 2005. Visual Control of Burst Priming in the Anesthetized Lateral Geniculate Nucleus. *J Neurosci.*

[doi:](#)

III-90. Velocity coding and octopamine in an identified optic flow-processing interneuron of the blowfly

Kit D. Longden
Holger G. Krapp

KIT@IMPERIAL.AC.UK
H.G.KRAPP@IMPERIAL.AC.UK

Dept of Bioengineering, Imperial College London, UK

Flying generates predictably different distributions of optic flow compared to walking, or resting. In the blowfly, for instance, flight doubles the magnitude of the optic flow experienced during mean yaw rotations compared to walking. Flight is also extremely energy-intensive, raising the blowfly's metabolic rate by an order of magnitude from the resting state. A sensorimotor system adjusted for rapid responses and the high bandwidth of optic flow experienced during flight could allow the fly to avoid wasting energy through imprecise motor action. However, neural processing that covers a higher input bandwidth itself comes at higher energetic costs which would be a bad investment when the animal was not flying. We were interested to know if the blowfly adjusts the performance of its optic flow-processing neurons to its current locomotor state. Octopamine (OA) is a biogenic amine that is central to the initiation and maintenance of the flight states in insects. It is released during flight in local tissue and into the haemolymph, affecting the metabolism, muscle properties and sensory processing. We used an OA agonist chlordimeform (CDM) to simulate the widespread OA release during flight, and recorded extracellularly the spiking activity of the identified H2 cell in the third visual neuropil, the lobula plate. The H2 cell is a wide-field optic flow-processing interneuron, mainly sensitive to horizontal motion, and is involved in yaw stabilisation reflexes. CDM doubled the spontaneous spike rate of the H2 cell, from 7.9 ± 1.7 Hz to 15.6 ± 2.7 without having an impact on the cell's preferred temporal frequency when presented with constant velocity gratings. It did, however, alter the sensitivity and dynamics of the H2 responses. In particular, the initial responses, 20-60 ms after stimulus onset, were significantly increased over a broad range of velocities. The adaptation of the response was also significantly reduced over the 0.5-4 s interval. Upon stimulation in the anti-preferred direction, the CDM-induced elevation of the spontaneous activity increased the inhibitory signalling range. Across nearly all velocities the response latencies were reduced, on average by 4%. Consistent with the greater sensitivity after the application of CDM, the mutual information of the responses to a grating moved at a uniform, zero-symmetric velocity distribution was increased by 35%. This change was not merely a consequence of the greater inhibitory signalling range because CDM also increased the mutual information of the responses to a uniform, positive velocity distribution by 40%. The mean spike rate in these experiments also increased, such that the information per spike decreased by 20%. These findings suggest that OA can modulate the sensitivity and dynamics of the responses in blowfly optic flow-processing interneurons to adjust to the higher stimulus bandwidth experienced during flight. The increased signalling range, more rapid and longer lasting responses employ more spikes and thus consume greater amount of energy. It appears that for the fly to invest more energy into sensory processing during flight is more efficient than wasting energy on under-performing motor control.

[doi:](#)

III-91. Exact statistical analysis of visual inference amid eye movements

Eran A. Mukamel¹
 Yoram Burak¹
 Markus Meister¹
 Haim Sompolinsky²

ERAN@POST.HARVARD.EDU
 YBURAK@FAS.HARVARD.EDU
 MEISTER@FAS.HARVARD.EDU
 HAIM@FIZ.HUJI.AC.IL

¹Harvard University

²Harvard University and Hebrew University

Sensory information about objects of interest in the world is generally corrupted by neural noise, but also by confounding signals. For example, research on high-acuity visual perception has focused on the limits imposed by variability of neural responses from the retina. Here, we consider a second important limit that results from random movement of the eyes during fixation. By introducing a simplified model of neural variability and stochastic eye movement, we find an analytic expression for the optimal estimator in a common psychophysical task. Our exact analysis shows that estimator performance is limited by both noise sources. The model also reveals a critical role for spike timing in the optimal decoding strategy. For concreteness, we analyze the estimation of a gap between two lines in a Vernier hyperacuity task. The absolute position of the image drifts across the retina owing to the random-walk trajectory of eye position. The brain must estimate both eye position and gap using exclusively the spike trains of retinal ganglion cells. We calculate the exact joint probability distribution for the eye position and the gap in our model. We find simple iterative update rules for the mean and covariance matrix as spikes are emitted. These rules show how the accuracy of gap estimation improves over time. We also express the optimal estimate of the gap as an explicit function of the previous history of spiking. The optimal decoding strategy depends on a single dimensionless parameter that characterizes the two sources of uncertainty about the position of the lines: the root mean squared displacement of the eyes between subsequent spikes in any two ganglion cells, divided by the width of a ganglion cell receptive field. In the limit of very slow eye movements, the optimal decoder makes equal use of all observed spikes to estimate the position of each line and find their separation. In the opposite limit of fast eye movements, the decoder uses only near-synchronous spikes arising from each of the lines. Such nearly synchronous spikes can arise due to Poisson spike variability. Our analysis provides the full phase diagram for the optimal decoding strategy as a function of parameters. We further determined the fundamental performance limits for image feature estimation by calculating the mean squared estimation error. We find that the performance of the optimal estimator after a fixed viewing period is, to a good approximation, limited by the larger of the two sources of uncertainty, eye movement amplitude between spikes and the width of a receptive field. By exactly solving a statistical model of early visual processing, our work reveals the interplay of neural variability and eye movements in high-acuity vision. Such exact analysis, though reliant on simplifying modeling assumptions, provides insight into the principles constraining neural decoding in sensory systems when confounded by unknown signals such as observer motion. A similar framework might be useful for analyzing other high-precision sensory tasks, such as auditory localization in the barn owl or tactile perception in the rodent vibrissal system.

[doi:](#)

III-92. Encoding stereomotion with neural populations using IOVD and CD mechanisms

Qiuyan Peng
 Bertram E. Shi
 Dept. of ECE, HKUST

EEPENG@UST.HK
 EEBERT@EE.UST.HK

Stereomotion refers to the motion-in-depth (MID) that is either approaching or receding from the observer. Two possible cues have been proposed for MID perception: changing disparity (CD) and interocular velocity differences (IOVD). The CD mechanism combines binocular information to estimate disparity first, and then estimates disparity changes over time. The IOVD mechanism estimates monocular motion first, and then estimates the

difference between the motion in the two eyes. Although CD and IOVD are mathematically equivalent for natural stimuli, there is both physiological and psychophysical evidence for and against both mechanisms in the visual system. Currently, it is thought that both mechanisms contribute to the perception of motion in depth, but they may be used to different extents on different tasks. There has been little past work in constructing neurally plausible models for stereomotion that can be directly applied to pairs of image sequences. What work has been done has considered only single or pairs of neurons. Sabatini et al. have described how to construct single stereomotion selective neurons using the IOVD mechanism. We have described how to construct pairs of stereo-motion selective neurons tuned to the CD cue. However, single or even pairs of neurons are not enough to estimate stereomotion quantitatively, since neural responses depend upon multiple stimulus dimensions, e.g. orientation and spatial/temporal frequency. Thus, stimulus variables are thought to be encoded by the joint responses from a population of neurons displaying a diversity of tuning parameters. We describe two populations of physiologically plausible neurons encoding stereomotion: one using the IOVD mechanism and the other using the CD mechanism. Both are two-stage models cascading populations of motion and disparity energy-like models, but differing in the order of the cascade (Fig. 1). We refer to these as the IOVD and the CD energy models. For the IOVD energy model, populations of motion energy neurons for each eye are combined by a population of disparity energy-like neurons that encodes differences in the motion between the two eyes. For the CD energy model, populations of disparity energy neurons combine information from each eye resulting in a distributed representation of disparity. This representation is fed into a population of motion energy-like neurons, which encode temporal changes in disparity. These models enable us for the first time to examine stereomotion estimates from populations of physiologically plausible neurons using the same visual stimuli used in physiological or psychophysical experiments, such as random dot stereograms. Our results are consistent with the existing psychophysical literature. The increased sensitivity to the IOVD cue for stimuli with relatively high temporal frequency stimuli is consistent with our finding that the IOVD model has higher MID resolution for large MID speed. The decrease in sensitivity to MID with increasing disparity pedestals is consistent with our finding that MID speed estimation variance increases for stimuli at non-zero disparity pedestals. Our models also generate the testable prediction that neurons selective for CD should exhibit higher selectivity than neurons selective to IOVD.

doi:

III-93. Sparseness is not actively optimized in V1

Pietro Berkes
Benjamin Lee White
Jozsef Fiser

BERKES@BRANDEIS.EDU
BLWHITE@BRANDEIS.EDU
FISER@BRANDEIS.EDU

Brandeis University

Sparse coding is a powerful idea in computational neuroscience referring to the general principle that the cortex exploits the benefits of representing every stimulus by a small subset of neurons. Advantages of sparse coding include reduced dependencies, improved detection of co-activation of neurons, and an efficient encoding of visual information. Computational models based on this principle have reproduced the main characteristics of simple cell receptive fields in the primary visual cortex (V1) when applied to natural images. However, direct tests on neural data of whether sparse coding is an optimization principle actively implemented in the brain have been inconclusive so far. Although a number of electrophysiological studies have reported high levels of sparseness in V1, these measurements were made in absolute terms and thus it is an open question whether the observed high sparseness indicates optimality or simply high stimulus selectivity. Moreover, most of the recordings have been performed in anesthetized animals, but it is not clear how these results generalize to the cell responses in the awake condition. To address this issue, we have focused on relative changes in sparseness. We analyzed neural data from ferret and rat V1 to verify two basic predictions of sparse coding: 1) Over learning, neural responses should become increasingly sparse, as the visual system adapts to the statistics of the environment. 2) An optimal sparse representation requires active competition between neurons involving recurrent connections. Thus, as animals go from awake state to deep anesthesia, which is known to eliminate recurrent and top-down inputs, neural responses should become less sparse, since the neural interactions that support active sparsification of responses are disrupted. To test the first prediction empirically, we measured the sparseness of neural

responses in awake ferret V1 to natural movies at various stages of development, from eye opening to adulthood. Contrary to the prediction of sparse coding, we found that the neural code does adapt to represent natural stimuli over development, but sparseness steadily decreases with age. In addition, we observed a general increase in dependencies among neural responses. We addressed the second prediction by analyzing neural responses to natural movies in rats that were either awake or under different levels of anesthesia ranging from light to very deep. Again, contrary to the prediction, sparseness of cortical cells increased with increasing levels of anesthesia. We controlled for reduced responsiveness of the direct feedforward connections under anesthesia, by using appropriate sparseness measures and by quantifying the signal-to-noise ratio across levels of anesthesia, which did not change significantly. These findings suggest that the representation in V1 is not actively optimized to maximize the sparseness of neural responses. A viable alternative is that the concept of efficient coding is implemented in the form of optimal statistical learning of parameters in an internal model of the environment. This work has been supported by the Swartz Foundation and the Swiss National Science Foundation.

[doi:](#)

III-94. Motion and reverse-phi stimuli that do not drive standard Fourier or non-Fourier motion mechanisms

Qin Hu

Jonathan D. Victor

Weill Cornell Medical College

QIH2002@MED.CORNELL.EDU

JDVICTO@MED.CORNELL.EDU

Detection of motion is a crucial component of visual processing, and is generally considered to consist of two stages: an early stage in which local motion is extracted and a later stage at which local motion signals are combined into object motion or flows. Early motion processing is generally considered to be carried out by first-order (Fourier) and second-order (non-Fourier) mechanisms. Fourier motion mechanisms extract motion when the pairwise spatiotemporal correlation of luminance signal is present. Non-Fourier mechanisms are thought to work via local nonlinear pre-processing, such as flicker detection or extraction of unsigned contrast, followed by a spatiotemporal correlation of the resulting signals. To probe the computations underlying motion perception, we created a new class of non-Fourier motion stimuli: binary movies characterized by their 3rd- and 4th-order spatiotemporal correlations. As with other non-Fourier stimuli, they lack second-order correlations, and therefore their motion cannot be detected by standard Fourier mechanisms. Additionally, these stimuli lack pairwise spatiotemporal correlation of edges or flicker - and thus, also cannot be detected by extraction of one of these features, followed by standard motion analysis. Nevertheless, our psychophysical results showed that many of these stimuli produced apparent motion in human observers. The pattern of responses - i.e., which specific spatiotemporal correlations led to a percept of motion - was highly consistent across subjects. Moreover, for many of these stimuli, inverting the overall contrast of the stimulus reversed the direction of apparent motion. This "reverse phi" phenomenon, as well as the high-order-only spatiotemporal correlation of the stimulus, challenge existing models, including models that correlate low-level features (e.g., the Reichardt model and spatiotemporal energy models) and gradient models. Simple augmentations of those models - for example, spatiotemporal filters followed by non-quadratic nonlinearities - can account for some aspects of the percepts, but not for others. This suggests that a full account of motion percepts driven by high-order spatiotemporal correlations will lead to a more complete understanding of the computations underlying early motion processing.

[doi:](#)

III-95. Temporal integration of motion and cortical normalization in macaque V1

Douglas McLelland
 Pamela M. Baker
 Bashir Ahmed
 Wyeth Bair

DOUGLAS.MCLELLAND@DPAG.OX.AC.UK
 PAMELA.BAKER@DPAG.OX.AC.UK
 BASHIR.AHMED@DPAG.OX.AC.UK
 WYETH.BAIR@DPAG.OX.AC.UK

Department of Physiology Anatomy and Genetics, Univ. of Oxford

There is an increasing awareness that the spatio-temporal receptive field properties of sensory neurons are not necessarily fixed but can vary dynamically, e.g., in a stimulus dependent manner. For example, in complex direction selective (DS) cells in macaque primary visual cortex (V1), the window of temporal integration for motion stimuli is short for fast motion but extends for slower motion [1], a phenomenon referred to as adaptive temporal integration (ATI). As a result, sensitivity over a range of stimulus velocities is higher than could be achieved with a fixed spatio-temporal filter. We investigated the site and possible mechanisms of ATI. One possibility is that ATI is inherited from changes in the integrative properties of cells earlier in the visual processing stream, such as non-orientation tuned cells in the retina, lateral geniculate nucleus, or geniculo-recipient layers of the cortex. An alternative hypothesis is that ATI is linked to the phenomenon of cortical normalization. With increased stimulus speed, thalamic input to V1 increases, as does the activity of many cells within V1, which should engage normalization mechanisms. This could be associated with an increase in synaptic input (excitatory and perhaps inhibitory) to DS cells, leading to a reduction in membrane time constant, which has been proposed to play a role in changes in temporal dynamics in V1 [2,3]. We devised a masking paradigm which would test for both of these mechanisms. We made single unit extracellular recordings of complex DS cells ($n = 21$) in the primary visual cortex of anesthetized, paralysed macaques. We tested ATI using the same dynamically moving stimulus as used previously [1]. From these results, we selected a fast and a slow stimulus speed which yielded clear and distinct temporal kernels in a spike-triggered average. We then re-tested integration at these speeds, but included a superimposed, dynamically moving orthogonal grating, likewise at fast and slow speeds. This mask should engage mechanisms of ATI if they are present in non-orientation tuned cells or arise from cortical normalization. This would be revealed most obviously as a characteristic change in the integration kernel (reduction in width, and rightward shift of the peak) for slow target motion when the fast mask was included. No such characteristic change in integration kernel was found. Thus we conclude that ATI arises in cortex once orientation tuning has already been established, and is not dependent on a broadly-tuned mechanism such as cortical normalization. We hypothesize that ATI may be intrinsic to the mechanism responsible for direction selectivity. Further, although the mask did not yield the change in integration kernel characteristic of ATI, it nonetheless yielded a reduction in kernel amplitude. This change is characteristic of a decrease in stimulus selectivity of the response, that is an increase in the proportion of spikes not driven by the "preferred" stimulus. We discuss this result in the context of cortical normalization. References [1] Bair and Movshon (2004) *J Neurosci* 24:7305-7323 [2] Reid, Victor and Shapley (1992) *Vis Neurosci* 9:39-45 [3] Carandini, Heeger and Movshon (1997) *J Neurosci* 17:8621-8644

[doi:](#)

III-96. Relationship of contextual modulations in V1 and V2 revealed by non-linear receptive field mapping

Anita M Schmid
 Jonathan D. Victor
 Weill Cornell Medical College

AMS2031@MED.CORNELL.EDU
 JDVICTO@MED.CORNELL.EDU

Responses of neurons in primary visual cortex (V1) are modulated by context, due to processes such as surround suppression and facilitation. These modulations possibly serve as a mechanism for detecting texture borders. Neurons in secondary visual cortex (V2), on the other hand, respond more directly to texture borders. It

is as yet unclear how these processes relate to each other: are they independent of each other and serve different purposes or are they connected and conjointly serve the purpose of detecting texture borders? To gain more insight into this relationship, we investigated the dynamics of both types of nonlinearities. We recorded from single neurons in V1 and V2 of anesthetized monkeys. The stimulus consisted of a 4 by 5 or 6 by 6 grid of adjacent rectangular regions, covering the classical and non-classical receptive field. Each region contained sinusoidal gratings with either the preferred orientation or the non-preferred orthogonal orientation controlled by an m-sequence with frame durations of either 20 or 40 milliseconds. In V1, for frame durations of 20 milliseconds, only positive interactions were observed; responses were larger when the same orientation was displayed in two neighboring patches than when two orientations were different from each other. This interaction was only seen for regions aligned along the preferred orientation axis and is consistent with iso-orientation facilitation along the axis of the receptive field. For frame durations of 40 milliseconds, however, positive as well as negative interactions were observed in V1. Both types occurred for neighboring regions aligned along the preferred orientation as well as neighboring regions aligned orthogonal to the preferred orientation. These nonlinear responses are consistent with known center-surround interactions, but interestingly, they occurred most often between regions that individually elicited a positive response and therefore can be considered within the "center" of the receptive field. In V2, negative interactions leading to larger responses to texture borders were very robust for transient V2 neurons [1] for frame durations of both 20 and 40 milliseconds. Also, these interactions in V2, which are consistent with a spatial differentiation operation, started earlier than the contextual interactions in V1. This study shows, firstly, that contextual modulations in V1 are dynamic processes and do not merely take place between a static receptive field "center" and "surround". Positive interactions, which indicate preference for continuous orientations, occur with a shorter latency than negative interactions, which yield larger responses for orientation texture borders. Interestingly, very short frame durations of 20 milliseconds do not elicit negative interactions in V1, suggesting that they need more time to build up, whereas the positive interactions act faster. Most strikingly, nonlinear interactions in V2 occur earlier than those in V1. We conclude that the fast V2 responses to texture borders arise independent of the slower contextual modulations in V1. Support: NIH R01EY09314. References: 1. Schmid AM, Purpura KP, Ohiorhenuan IE, Mechler F and Victor JD (2009) Subpopulations of neurons in visual area V2 perform differentiation and integration operations in space and time. *Front. Syst. Neurosci.* 3:15.

doi:

III-97. Black' dominance measured with different stimulus ensembles in macaque primary visual cortex V1

Chun-I Yeh
Dajun Xing
Robert Shapley

New York University, Center for Neural Sci

CIY@CNS.NYU.EDU
XDJ@CNS.NYU.EDU
SHAPLEY@CNS.NYU.EDU

Most neurons in output layer 2/3 (not input layer 4c) of the primary visual cortex (V1) have stronger responses to 'black' (negative contrast) than to 'white' (positive contrast) when measured by reverse correlation with sparse noise (randomly positioned dark and bright spots against a gray background; Jones and Palmer 1987). The black-dominant response of monkey V1 neurons may serve as the neuronal substrate for human perception: in many perceptual tasks black is much more salient than white. Furthermore, the degree of the black dominance in V1 depends on the stimulus ensemble - the black dominance in layer 2/3 of V1 is much stronger when neuronal responses are measured with sparse noise than with Hartley subspace stimuli (a series of gratings at different orientations, spatial frequencies and spatial phases; Ringach et al 1997). Sparse and Hartley stimuli differ significantly in many ways. First, the individual stimulus size of the sparse noise is much smaller and therefore activates fewer V1 neurons simultaneously than that of the Hartley dense noise. Second, the dark and bright pixels of the sparse noise are present separately in time while those of the Hartley dense noise are shown simultaneously. Third, there are spatial correlations along the long axis of the Hartley stimuli that are not present in sparse noise. Which of these differences might contribute to the disparity in black dominance? We addressed this question by introducing a third stimulus ensemble - a binary checkerboard white noise (Reid et al 1997) - to measure the black-

dominant responses in monkey V1. The binary white noise and Hartley dense noise share several properties in common: both activate a larger population of neurons than the sparse noise, and dark and bright pixels appear simultaneously under both conditions. However, unlike Hartley dense noise, neighboring pixels of binary white noise are uncorrelated. To make fair comparisons across different stimulus ensembles, only those neurons with significant responses (signal-to-noise ratio > 1.5) to all three ensembles were included ($n=32$). Consistent with previous findings (Yeh et al 2009), black-dominant neurons largely outnumbered white-dominant neurons in layer 2/3 of V1 with all three stimulus ensembles (percentage of black-dominant neurons: 76~82%), while the numbers of black and white-dominant neurons were nearly equal in input layer 4c (percentage of black-dominant neurons: 40~60%). The degree of the black-dominant response was significantly stronger for binary white noise than for Hartley dense noise ($p < 0.02$, Wilcoxon signed rank test). These results indicate that the degree of black dominance depends on the spatial structure of the stimulus ensemble. Acknowledgements: This work was supported by NIH-EY001472, NSF-0745253, the Robert Leet and Clara Guthrie Patterson Trust Postdoctoral Fellowship, and the Swartz Foundation.

[doi:](#)

III-98. The adapting receptive field surround of a large ON ganglion cell

Karl Farrow
Botond Roska

KARL.FARROW@FMI.CH
BOTOND.ROSKA@FMI.CH

Friedrich Miescher Institute

A classic example of the context dependant nature of retinal ganglion cell receptive fields, first noted by Stephen Kuffler (1953), is their the light-level dependent narrowing and widening. Here we detail the receptive field adaptation of a genetically targeted large ON ganglion cell in the mouse retina at different dark adapted states and start to unravel its input structure. Using a combination of two-photon microscopy and a mouse expressing YFP in a subset of its ganglion cells (PVCr $\times$ Thy1Stp-EYFP) we performed targeted patch clamp recordings from a large ON ganglion cell type, termed PV1. The spatial receptive field properties were explored using spots and annuli of different sizes. These stimuli were presented at different background light intensities ranging from ~10 to 107 photon/ $\mu\text{m}^2/\text{s}$. The morphology of each cell was reconstructed from the two-photon scans. Recording the spiking output of PV1 cells, we found that they have much larger receptive fields in dark adapted conditions as compared to light adapted conditions. To assess the relative contribution of excitatory and inhibitory inputs we calculated their combined reversal potential during the different stimuli. We found that the reversal potential of the cell decreased towards that of inhibition as the adapting light level was increased, suggesting that the level of surround inhibition increases with increasing background light level. Pharmacological experiments suggest this relative increase in inhibition is carried by GABAergic wide field amacrine cells as the application of picrotoxin, but not strychnine shifts the reversal potential towards excitation. These experiments demonstrate that the surround of the PV1 cell gets stronger as one moves from scotopic to mesopic or photopic light conditions and that this surround is likely mediated by wide field amacrine cells.

[doi:](#)

III-99. Functional dissection of optomotor pathways in the Drosophila optic lobe.

Saskia E. J. de Vries
Thomas R. Clandinin
Stanford University

SASKIAD@STANFORD.EDU
TRC@STANFORD.EDU

Visual motion cues are critical for flies to navigate through their environment, yet precisely how the nervous

system uses these signals to guide behavior is only incompletely understood. In freely walking flies such optomotor behaviors are shaped by differences in the speed and density of the visual motion stimulus, and reflect separable changes in the translation or rotation of the animal. Using a forward genetic screen in which random sub-populations of cells were inactivated, a specific group of tangential cells in the lobula plate, the Foma-1 neurons, were shown to play a critical role in linking motion stimuli to rotational responses. Here, we record the electrophysiological activity of these Foma-1 cells to determine their role in visual motion behavior. We target these neurons for recording by driving the expression of GFP in the cells, and record their spiking activity using a loose patch configuration. Visual stimuli are generated on a high speed CRT and focused onto the fly's eyes using two coherent fiber optic bundles. We find that at least one of the Foma-1 neurons is directionally selective, preferring motion along the longitudinal axis of the fly, in the dorsal part of its visual field. Interestingly, the preferred direction of motion changes based on specific parameters of the visual stimulus. When tested with a sparse stimulus, the cell prefers motion from back-to-front. On the other hand, when probed with a dense stimulus, the cell prefers motion from front-to-back. This reversal of the Foma-1 neuron's direction selectivity with stimulus density mirrors the changes Foma-1 inactivation has on the behavioral responses of flies viewing either sparse or dense optomotor stimuli. Based on these results, we propose that the Foma-1 neuron enhances the turning behavior of the fly when it encounters a motion stimulus.

[doi:](#)

III-100. Retention of perceptual categorization following bilateral removal of area TE in rhesus monkeys

Narihisa Matsumoto^{1,2}

Richard Saunders²

Katalin Gothard³

Barry Richmond²

XMATUMO@NI.AIST.GO.JP

RICHARDSAUNDERS@MAIL.NIH.GOV

KGOTHARD@EMAIL.ARIZONA.EDU

BJR@LN.NIMH.NIH.GOV

¹AIST

²NIMH/NIH

³University of Arizona

Bilateral ablation of inferior temporal (IT) cortex impairs pattern discrimination in monkeys [1,2]. Single neurons in area TE of IT cortex have marked tuning to visual stimuli like objects or faces [3-5]. Physiological recordings (single neuron and fMRI) show activity grouped into categorical-like functions (face and non-face patches) [6-10]. These findings lead to the hypothesis that area TE plays a critical role in categorizing objects or faces. To test this hypothesis, monkeys were tested before and after bilateral removals of area TE on a task that requires the association of 2 perceptual categories (e.g. human faces vs. monkeys faces) of visual stimuli with different incentive values. In the task the monkey grasps and holds a bar, a visual cue appears for 400 ms, and then a red dot appears in the center of the stimulus, 500-1500 ms later the red dot changes to green. If the monkey releases the bar within 3 seconds after the dot changes to green, the color of the dot changes to blue and in high incentive trials a liquid reward is also delivered. In low incentive trials the trial ends. After an intertrial interval a new cue is chosen pseudorandomly for the next trial. We used 20 dogs/20 cats, or 20 human/20 monkey faces for the visual cues. Before TE ablations, the monkeys made significantly more errors in trials with low incentive stimuli than in trials with high incentive stimuli ($p < 0.01$, chi-square test); the monkeys distinguished dogs vs. cats and human vs. monkey faces. After TE removals, the monkeys continued to make significantly more errors in low incentive trials than in high incentive trials ($p < 0.01$). To test that they could generalize on one testing day, we used 240 trial-unique exemplars of dogs and cats and 240 exemplars of human and monkey faces. The monkeys (with ablations) again made significantly more errors in the low incentive trials than in the high incentive trials ($p < 0.05$). Our results so far seem to show the possibility that ablations of area TE fail to interfere with perceptual categorization. Funding: AIST, IRP/NIMH/NIH References [1] Iwai, E., Mishkin, M. (1969). *Exp Neurol*, 25, 585-594. [2] Cowey, A., Gross, C.G. (1970). *Exp Brain Res*, 11, 128-144. [3] Desimone, R., Albright, T.D., Gross, C.G., Bruce, C. (1984). *J Neurosci*, 4, 2051-2062. [4] Richmond, B.J., Optican, L.M., Podell, M., Spitzer, H. (1987). *J Neurophysiol*, 57, 132-146. [5] Fujita, I., Tanaka, K., Ito, M., Cheng, K. (1992). *Nature*, 360, 343-346.

[6] Hadj-Bouziane, F., Bell, A.H., Knusten, T.A., Ungerleider, L.G., Tootell, R.B. (2008). *Proc Natl Acad Sci USA*, 105, 5591-5596. [7] Wang, G., Tanaka, K., Tanifuji, M. (1996). *Science*, 272, 1665-1668. [8] Sugase, Y., Yamane, S., Ueno, S., Kawano, K. (1999). *Nature*, 400, 869-873. [9] Sigala, N., Logothetis, N.K. (2002). *Nature*, 415, 318-320. [10] Kiani, R., Esteky, H., Mirpour, K., Tanaka, K. (2007). *J Neurophysiol*, 97, 4296-4309.

doi:

III-101. A model of efficient change detection through the interaction of excitation and inhibition

Nabil Bouaouli^{1,2}
Sophie Deneve³

NABIL.BOUAOULI@ENS.FR
SOPHIE.DENEVE@ENS.FR

¹Ecole Normale Supérieure, Paris, France

²Radboud University, Nijmegen, Netherlands

³Group for Neural Theory, LNC, DEC, ENS Paris

Psychophysical experiments have shown that visual perception requires constant retinal motion induced by small eye movements. On the cellular level, "ON" sensory neurons respond more strongly to the appearance of their preferred stimulus and remain mainly quiescent during its presence. Moreover, depending on the context of the stimulus, these neurons can respond remarkably precisely in time. This marked preference for transients, also called adaptation, was reported in visual, auditory, and somato-sensory pathways and was observed at various levels of sensory pathways. However, many mechanisms can lead to such adaptation. Foremost are spike based adaptation, delayed inhibition, and short term synaptic depression (STD). Thus, it seems likely that a major role of early sensory processing is detecting sudden changes in sensory scenes and that some or all of these mechanisms are crucial for the underlying computations. To clarify this issue, we used an inference-based probabilistic framework to explore the hypothesis that sensory neurons are tuned to optimally detect, as quickly and reliably as possible, a sudden transient in their preferred stimuli. We derived a Bayesian model of change detection parameterized by the stimulus temporal statistics and the strength of the input. The model suggests a neural implementation by a "minimal" circuit that computes the probability of stimulus onset. This circuit links two pyramidal cells through direct excitation and delayed feed-forward inhibition, respectively involving STD and facilitation, and is widespread throughout sensory cortices. Consequently, we found strong correlations between the stimulus temporal statistics and some biophysical synaptic parameters. For instance, the time constant of depression is strongly related to the temporal dynamics of the stimulus: fast stimuli predict fast recovery from depression, while slower stimuli induce slower recovery time constants. Similarly, the absolute synaptic efficacy is related to the stimulus strength: stronger inputs, corresponding to easier-to-detect transients, yield larger EPSPs, because each spike is more "informative". We next explored the predictions of our model for sensory coding and adaptation, particularly in early stages of visual processing. The model predicted the biphasic "ON" or "OFF" temporal receptive fields (tRFs) reported in the retina and LGN. However, as observed experimentally, the predicted responses to time varying stimuli are also sparser and far more temporally precise than would be predicted by the tRF alone. Moreover, as found in early visual stages, response gains and tRFs adapt well to stimulus variance, adjusting consequently the temporal precision of the response. We then investigated cellular explanations of such adaptation and found that excitation and inhibition interact differently at low and high stimulus variance, yielding different firing statistics of the neuron. Inhibition is always delayed relative to excitation which is, contrarily, "time-locked" to the sharp transients of the stimulus. Strikingly, at high stimulus variance, inhibition becomes more important and less delayed than for small variance, suggesting that stronger stimuli recruit more inhibition to ensure more precise firing. In conclusion, our model shows that efficient change detection is implemented by coupling feedforward inhibition and STD and relies on how the balance of excitation and inhibition is modulated by the stimulus.

doi:

Author Index

Ölveczky B., [82](#)

K., [81](#)

R., [44](#)

Abbott L., [93](#)

Abbott L. F., [149](#), [157](#), [201](#)

Acuna D., [124](#)

Adesnik H., [118](#)

Agrochao A., [188](#)

Ahmadian Y., [133](#), [140](#), [227](#)

Ahmed B., [259](#)

Ahrens M., [64](#)

Ahrens M. B., [211](#)

Aihara K., [207](#)

Aimone J. B., [144](#)

Aksay E., [156](#)

Aladini O., [60](#)

Alexander R., [175](#)

Alonso J.-M., [91](#)

Amari S.-i., [88](#)

Amodei D., [112](#)

Andken B. B., [112](#)

Andrieux D., [195](#)

Angelaki D., [96](#)

Angelaki D. E., [169](#), [198](#), [247](#)

Arleo A., [147](#)

Arthur J., [26](#)

Asaad W. F., [76](#)

Asari H., [105](#)

Assad J., [167](#)

Assisi C., [177](#)

Atiani S., [243](#)

Avermann M., [63](#)

Büsing L., [145](#)

Babadi B., [149](#)

Baccus S., [106](#)

Badel L., [135](#)

Bains A., [177](#)

Bair W., [186](#), [259](#)

Baker J., [240](#)

Baker P. M., [186](#), [259](#)

Ballard D., [121](#)

Bao S., [100](#), [245](#)

Baraniuk R., [93](#)

Barbarits B., [211](#)

Barbic M., [211](#)

Barbour D. L., [175](#)

Barhomi Y., [252](#)

Barlow H. B., [32](#)

Barrett D. G. T., [165](#)

Bartels J., [236](#)

Battaglia F. P., [70](#)

Bavelier D., [34](#)

Bazhenov M., [177](#)

Beck J., [238](#)

Behabadi B. F., [190](#)

Behrens T., [47](#)

Behrens T. E., [125](#)

Beierholm U., [149](#)

Benayoun M., [203](#)

Bengio Y., [187](#)

Bennett R., [181](#)

Bennur S., [127](#)

Bergstra J., [187](#)

Berkes P., [101](#), [257](#)

Bernacchia A., [49](#)

Berry II M. J., [105](#)

Berry M., [112](#)

Berry M. J., [162](#)

Bethge M., [30](#)

Bettencourt L. M., [181](#)

Bettencourt L. M. A., [203](#)

Bex P., [168](#)

Bhalla U., [247](#)

Bhattacharyya A., [85](#)

Bialek M., [49](#)

Bialek W., [155](#), [162](#)

Bichot N., [178](#)

Biessmann F., [85](#)

Billings G., [234](#)

Billings S., [108](#)

Billings S. A., [36](#)

Black M. J., [91](#)

Blair H. T., [213](#), [229](#)

Blanche T., [251](#)

Bloom M., [229](#)

Boahen K., [26](#)

Bobier B., [192](#)

Boerlin M., [95](#)

Bonin V., [104](#)

Borst A., [171](#)

- Bosman C. A., 194
 Bossaerts P., 150
 Bouaouli N., 263
 Bouret S., 220
 Bozza T., 243
 Brainard M., 72
 Brette R., 208
 Brincat S. L., 143
 Brody C. D., 34, 49
 Brumby S., 181
 Brumby S. P., 203
 Brunel N., 56
 Brunton B. W., 34
 Buckley C. L., 80
 Buehlmann A., 87
 Buhusi C. V., 234
 Burak Y., 82, 92, 256
 Burge J., 36
 Burns S., 184
 Butts D. A., 99, 111, 251
 Buzsáki G., 68
 Buzsaki G., 140

 Cadieu C., 242
 Calabrese A., 98
 Callaway E. M., 23
 Camerer C., 200
 Campbell R., 163
 Cardanobile S., 205
 Carlson C., 239
 Carmena J., 242
 Carr M. F., 41
 Castellano M., 241
 Casti A. R. R., 111
 Cazé R., 43, 122
 Chacron M., 134
 Chacron M. J., 94, 249
 Chan A. M., 239
 Chang S. W. C., 153
 Cheng K., 75
 Chestek C., 154
 Chiapetta C., 250
 Chiappe M. E., 67, 106, 185
 Chichilnisky E. J., 133, 161
 Chiel H. J., 83
 Chikkerur S., 178
 Chklovskii D., 65, 159
 Chou W.-C., 60
 Chow S. F., 132
 Churchland M. M., 156
 Clandinin T. R., 261
 Clark A. M., 220
 Clopath C., 145
 Coca D., 36, 108
 Coen-Cagli R., 57

 Cohen M. R., 237
 Cohen Y., 226
 Conte M., 130
 Cools R., 126
 Costa G., 54
 Costello G., 51
 Cottureau B. R., 114
 Cotton R. J., 50
 Cowan J., 203
 Cox D. D., 252
 Cronin B., 63
 Crum P., 171
 Cui M., 248
 Cullen K., 134
 Cullins M. J., 83
 Cumming B. G., 52
 Cummins G., 103
 Cunningham J., 154
 Cunningham J. P., 156
 Cuntz H., 171
 Curto C., 219

 Dabrowski W., 188
 Dakin S. C., 168
 Das A., 216
 David S. V., 99, 101, 243
 Daw N., 55
 Dayan P., 57, 124, 126, 149, 164, 206
 De Martino B., 200
 de Vries S. E. J., 261
 de Weerd P., 194
 DeAngelis G., 96
 DeAngelis G. C., 169, 198, 247
 Deco G., 87
 Deger M., 116
 Degeratu A., 219
 Deisseroth K., 141, 210
 Deneve S., 95, 263
 Desbordes G., 91
 Desimone R., 178
 Desjardins G., 187
 Devinsky O., 239
 DeVries S. H., 117
 DeWeese M., 245
 DeWeese M. R., 189
 DeWitt E. E. J., 73
 Diamond M. E., 214
 DiCarlo J., 182
 DiCarlo J. J., 112, 252
 Dickman J. D., 174
 Diesmann M., 116
 DiMattina C., 61
 Dipoppa M., 202
 Doi E., 90, 161
 Doiron B., 25, 57, 94, 236

- Dolan R., 149
Dolan R. J., 126
Dotson N. M., 47
Doyle W., 239
Drover J., 130
Drugowitsch J., 198
Dubnau J., 43
Dubreuil A., 82
Duzel E., 149
Dyer E., 93
- Ehrenberg E., 211
Eliasmith C., 192
Engert F., 183
Enikolopov A., 176
Erlich J. C., 49
Escola S., 227
Eshel N., 124
Eskandar E., 239
Eskandar E. N., 76
Eule S., 46
Everling S., 120
- Fadeyev V., 188
Fairhall A., 231
Fairhall A. L., 27
Falkner A., 53
Famulare M., 231
Fang G., 144
Fanini A., 167
Farinella M., 24
Farrow K., 261
Fee M. S., 39
Fellous J.-M., 233
Fellows M., 79
Ferster D., 115
Fetsch C. R., 169
Field D. J., 30
Field G., 161
Fiete I., 159
Finn I., 115
Fiser J., 93, 101, 248, 257
Fisher D., 156
Fitzgerald J. K., 167
Flister E. D., 254
Fok S., 96
Fonseca A., 54
Fontanini A., 248
Forstner F., 171
Frémaux N., 221
Frank L., 153
Frank L. M., 41, 141
Freedman D., 167
Freeman J., 38
Friedel E., 126
- Friederich U., 36
Friedrich J., 241
Fries P., 66, 194
Fritz J. B., 101
Froemke R. C., 217
Frolov R., 109
Fulvio J. M., 222
- Gabbito B. M., 81
Gage F. H., 144
Ganguli D., 230
Ganguli S., 68
Ganguly K., 242
Gardner J., 75
Gardner T., 101
Gauthier J., 161
Geisler W. S., 31, 36
Gentner T., 166
Gerstner W., 63, 71, 135, 145, 215, 221
Gerwinn S., 30
Ghebreab S., 115
Ghose G., 160
Ghose K., 110
Gijs Joost B., 180
Gilja V., 154
Gill K., 245
Gill P. R., 246
Girshick A. R., 139
Gleeson P., 24, 234
Glimcher P., 73, 128
Goddard C. A., 117
Goelzer M., 126
Gold J., 35, 127, 209
Goldberg J. H., 39
Goldberg M., 223
Goldberg M. E., 53
Goldman M., 156, 158
Goncalves P., 131
Goodman P. H., 69
Gopal V., 250
Gothard K., 262
Grünewälder S., 232
Gradwohl S., 175
Granot-Atedgi E., 235
Gray C. M., 47
Green C. S., 55, 124
Greenwood J., 168
Greschner M., 161
Groh J. M., 28, 172
Grothe B., 173
Gruen S., 88
Gu Y., 96
Guitart Masip M., 149
Gunning D., 161
Gupta A., 144

- Gurney K., [228](#)
 Gutkin B., [43](#), [202](#)
 Gutmann M., [58](#), [59](#)
- H.Steven H. S., [115](#)
 Häusser M., [171](#)
 Haefner R. M., [30](#)
 Hairston D., [27](#)
 Haiss F., [102](#)
 Halgren E., [239](#)
 Ham M., [181](#)
 Hamilton L. S., [100](#)
 Harasawa N., [75](#)
 Hartmann M., [250](#)
 Haruno M., [75](#)
 Hasan M. T., [102](#)
 Hasenstaub A., [23](#)
 Hasselmo M., [191](#)
 Hayden B., [199](#)
 He H., [195](#)
 He S., [55](#)
 Heasley B., [35](#)
 Heeger D., [180](#)
 Hehrmann P., [136](#)
 Heinz A., [126](#)
 Helias M., [116](#)
 Helmchen F., [102](#)
 Hennequin G., [71](#)
 Herikstad R., [240](#)
 Herz A., [119](#)
 Hillar C., [84](#)
 Hinton G. E., [32](#)
 Hirsch J., [177](#)
 Histed M. H., [104](#)
 Hofer S., [107](#)
 Holy T. E., [112](#)
 Honegger K., [163](#)
 Honey C., [214](#)
 Hooks D., [81](#)
 Horwitz G., [37](#)
 Hosoya H., [137](#)
 Hromadka T., [245](#)
 Hu Q., [258](#)
 Hu T., [65](#)
 Huber Y.-X., [81](#)
 Huguenard J., [117](#)
 Humphries M., [228](#)
 Humphries M. D., [43](#)
 Hunt L., [125](#)
 Hussar C., [109](#)
 Huys Q. J. M., [124](#), [126](#)
 Hyvärinen A., [59](#)
 Hyvärinen A., [58](#)
- Ichinohe N., [75](#)
- Ignatova I., [86](#)
 Immonen E.-V., [109](#)
 Isely G., [84](#)
 Issa E. B., [112](#)
 Itskov V., [219](#)
 Izhikevich E. M., [212](#)
- Jadi M., [191](#)
 Jamieson B. G., [211](#)
 Jan E., [191](#)
 Jaramillo S., [193](#)
 Jayaraman V., [67](#), [106](#), [159](#), [185](#), [211](#)
 Jayet L., [69](#)
 Jeanne J., [166](#)
 Jimenez Rezende* c., [215](#)
 Jin J., [91](#)
 Johnson A., [141](#)
 Johnson D., [93](#)
 Johnston K., [120](#)
 Joshi P., [142](#)
 Juusola M., [36](#), [108](#), [253](#)
- Kalmar R., [224](#)
 Kalwani R., [35](#)
 Kampff A., [183](#)
 Kandaswamy U., [189](#)
 Kanerva P., [211](#)
 Karachi C., [223](#)
 Karlsson M. P., [41](#)
 Katz D., [248](#)
 Kaufman M., [156](#)
 Kaveri S., [74](#), [75](#)
 Kayser C., [29](#)
 Kemere C., [141](#)
 Kemp A., [129](#)
 Kempter R., [219](#)
 Kennerley S., [242](#)
 Kenyon G., [181](#)
 Kenyon G. T., [203](#)
 Kerr C., [129](#)
 Kersten D., [55](#)
 Khan A., [247](#)
 Khanbabaie R., [173](#)
 Khosrowshahi A., [240](#)
 Kim A. J., [249](#)
 Kim H., [62](#)
 Kim M., [250](#)
 Kira S., [52](#)
 Klyachko V., [189](#)
 Knierim J. J., [213](#)
 Knoblauch A., [146](#)
 Knudsen E., [117](#)
 Ko H., [107](#)
 Kochik S., [189](#)
 Koepsell K., [242](#)

- Koever H., 245
 Kohn A., 89
 Koriyama T., 81
 Kording K., 63
 Koster U., 59
 Koulakov A., 176
 Krapp H. G., 255
 Krause B., 160
 Krauzlis R. J., 223
 Kretz R., 85
 Krumin M., 134
 Kulkarni J., 133
 Kumar A. L., 94
 Kuzniecky R., 239

 Löwel S., 231
 Lüling H., 173
 Lai D., 120
 Lamblin P., 187
 Lamme V., 115
 Landecker W. A., 203
 Landy M. S., 139
 Latham P., 237, 238
 Latham P. E., 131, 165
 Laudano A., 50
 Laughlin S., 119
 Laughlin S. B., 31
 Lazar A., 70
 Lazar A. A., 249
 Lee D., 49
 Lee J., 28, 226
 Lee L., 223
 Legenstein R., 146
 Leibold C., 173, 219
 Lengyel M., 101, 206
 Leopold D. A., 182
 Lerchner A., 131
 Lesica N., 107
 Li J., 153
 Li N., 182
 Li X., 253
 Lim S., 158
 Lim Y., 101
 Lisberger S., 153
 Lisman J., 41
 Liston D., 122
 Litke A., 161, 188
 Liu S., 96
 Lochmann T., 251
 Logotheitis N., 29
 Loncich K., 191
 Long L. N., 144
 Longden K. D., 255
 Longtin A., 173
 Lorincz A., 24, 234

 Losonczy A., 140
 Lott G. K., 67
 Louie K., 128
 Louradour J., 187
 Lu H., 83
 Luk C.-H., 54

 Müller K.-R., 85
 Ma L., 175
 Ma W. J., 60
 Maass W., 146, 218
 Machado T., 161
 Machado T. A., 208
 Machens C. K., 131, 166
 Macke J. H., 30
 Magee J., 140
 Magnasco M., 100
 Mainen Z., 54, 163
 Majaj N., 168
 Majaj N. J., 252
 Maler L., 173
 Mangun G. R., 187
 Mante V., 128
 Manu M., 106
 Marder E., 24
 Margoliash D., 72
 Margolis D., 102
 Marongelli E. N., 175
 Marre O., 112
 Massoglia D., 51
 Mathieson K., 161
 Matsumoto N., 262
 Matthijs v. d. M., 122
 Maunsell J. H. R., 110, 179, 237
 Mayrhofer J., 102
 Mazor O., 92
 McArthur K., 174
 McDermott J. H., 96
 McGuire L., 79
 Mckee S. P., 114
 McLelland D., 259
 McMahon D. B. T., 182
 McManus J. M., 83
 Meaden M., 250
 Mease R., 231
 Medina J., 153
 Meier P., 253
 Meinecke F., 85
 Meir R., 170
 Meir Y., 44
 Meister M., 92, 105, 188, 256
 Mel B. W., 190, 191
 Mensi S., 63
 Merzenich M. M., 217
 Mesgarani N., 243

- Meyers E., 178
 Middleton J. W., 57
 Midgley F., 159
 Milford M., 83
 Miller E. K., 143
 Miller J., 103
 Miller K., 140
 Miller K. D., 184
 Millner S., 26
 Mirollo R., 191
 Mishchenko Y., 208
 Mitchell M., 203
 Miura K., 163
 Moazzezi R., 164
 Molter* D., 215
 Monaco J. D., 213
 Monteforte M., 231
 Moore T., 39, 76
 Movshon J. A., 168
 Mrsic-Flogel T., 107
 Mukamel E. A., 256
 Mulder-Rosi J., 103
 Munk M. H. J., 232
 Munoz D., 77, 78
 Murakami M., 54
 Muralidharan K., 65
 Murray E. A., 220
 Murray I., 27

 Nagel K., 42
 Nakahara H., 74, 75
 Nassar M., 35, 209
 Natarajan R., 27
 Nategh N., 106
 Nathe A. R., 41
 Naud R., 63
 Naumann E. A., 183
 Navalpakkam V., 200
 Neimark Geffen M., 100
 Nesse W., 173
 Nessler B., 218
 Newsome W., 224
 Newsome W. T., 128
 Nienborg H., 52
 Niven J., 119
 Niyogi R. K., 196
 Nolan C. R., 83
 Nolan M. F., 45
 Norcia A. M., 114
 Nowotny T., 80
 Nusser Z., 24, 234
 Nuyujukian P., 154

 O'Connor D. H., 81
 O'Donnell C., 45

 O'Reilly J., 47
 Obermayer K., 232
 Ohshiro T., 247
 Okada M., 204
 Olbris D. J., 159
 Olshausen B. A., 240
 Omar C., 57
 Onken A., 232
 Oostenveld R., 194
 Ostojic S., 56
 Otani S., 147
 Otchy T., 82
 Otte S., 23
 Oxenham A. J., 96

 P. N. Rao R. P. N., 197
 Päpper M., 219
 Packer A. M., 208, 227
 Pages D., 172
 Paik S.-B., 138
 Palmer S. E., 162
 Paninski L., 45, 90, 98, 133, 161, 192, 208, 227
 Panzeri S., 29
 Papanastassiou A., 112
 Pastalkova E., 68
 Pasternak T., 109
 Paul S., 198
 Pavel P. . M., 214
 Peng Q., 256
 Pennartz C., 66
 Perona P., 200
 Petersen C., 63
 Peyrache A., 70
 Pfeiffer M., 218
 Pfister J.-P., 71, 206
 Pichler B., 107
 Pillow J. W., 133, 138, 205, 237
 Pinto N., 252
 Pipa G., 70, 241
 Pitkow X., 140
 Platt M., 199
 Platt M. L., 29
 Poggio T., 178
 Polsky A., 191
 Pouget A., 198, 238
 Preuschoff K., 150
 Priebe N., 115

 Qiu Q., 175
 Quoy M., 69

 Rad K. R., 192
 Raheja U., 247
 Rainer G., 85
 Rangel A., 123, 200

- Rasmussen C., 181
 Ratliff C. P., 117
 Ray D., 123
 Ray S., 179
 Read J., 88
 Rebecca V. D. H., 113
 Reid R. C., 21, 104
 Reinagel P., 253, 254
 Reiser M. B., 67, 106, 185
 Rennie C., 129
 Reppas J., 224
 Reyes A. D., 22
 Reyes M., 234
 Richard M., 246
 Richardson M. J. E., 135
 Richmond B., 262
 Richmond B. J., 62, 220
 Riecke H., 132
 Rinberg D., 176, 243
 Ringach D., 138
 Robinson P., 129
 Rodemann T., 151
 Rodgers C., 189
 Roiser J. P., 124
 Romo R., 166
 Roska B., 261
 Rossant C., 208
 Rothkopf C. A., 60, 151, 198
 Rotter S., 116, 205
 Royer S., 140
 Rozell C. J., 240
 Rubehn B., 194
 Rubin D. B., 184
 Rushworth M., 125
 Russ J., 250
 Ryan C., 242
 Ryu S., 154, 224
 Ryu W., 155

 Sabes P., 79
 Sadagopan S., 115
 Sahani M., 64, 136
 Salazar R. F., 47
 Salinas E., 51
 Sato T. R., 81
 Sato Y., 207
 Saunders R., 262
 Savin C., 148
 Scanziani M., 118
 Schaefer S. Y., 40
 Schaffer E. S., 157
 Scheller B., 241
 Schiff N., 130
 Schiff N. D., 40
 Schiller J., 22, 191

 Schilz J., 195
 Schinkel-Bielefeld N., 99
 Schmid A. M., 259
 Schneider A., 134
 Schneider D., 98
 Schneidman E., 235
 Schnupp J., 214
 Schomer D., 239
 Schrater P., 55, 124, 141, 222
 Schreiner C. E., 217
 Schumacher J., 98
 Schwartz G., 105
 Schwartz O., 57
 Seelig J. D., 67, 185
 Segev R., 235
 Sejnowski T., 225
 Sejnowski T. J., 233
 Sengupta B., 119
 Senn W., 241
 Seo H., 49
 Serre T., 178
 Shadlen M. N., 52
 Shamir M., 89
 Shamma S. A., 99, 101, 243
 Shankar S., 51
 Shapley R., 184, 260
 Sharpee T., 166
 Shaw K. M., 83
 Shelly I. L., 40
 Shenoy K., 154, 156, 224
 Shenoy P., 197
 Sher A., 161
 Sheynikhovich D., 147
 Shi B. E., 256
 Shimazaki H., 88
 Shimokawa T., 86
 Shinn-Cunningham B., 244
 Shinn-Cunningham B. G., 101
 Shinomoto S., 62, 86
 Shlens J., 133, 161
 Shlifer I. G., 229
 Shoham S., 134
 Shriki O., 89
 Shusterman R., 243
 Silver A., 24, 234
 Simoncelli E. P., 38, 90, 96, 133, 139, 161, 168, 230
 Simons D. J., 57
 Singer A. C., 41
 Sirotnin Y. B., 216
 Siveke I., 173
 Slutskiy Y., 249
 Smear M., 243
 Smeulders A., 115
 Snyder L. H., 153
 sohal V. S., 210

- Solomon E. A., 252
 Soltani A., 200
 Sommer F., 84, 211
 Sompolinsky H., 68, 256
 Song Z., 108
 Soo F., 105, 112
 Soto Sanchez C., 177
 Soudry D., 44
 Sprekeler H., 221
 Sreenivasan S., 159
 Sridharan D., 26, 117
 Stanford T. R., 51
 Stanley G. B., 91
 Steinmetz N., 76
 Stember J. N., 251
 Stemmler M., 119
 Stephens G., 155
 Stevens C., 189
 Stevenson I. H., 63
 Stewart T., 192
 Stieglitz T., 194
 Stocker A. A., 168
 Stone L., 122
 Sunkara A., 96
 Sur M., 63
 Suzuki S., 75
 Svoboda M., 81
 Swearingen J., 234
 Szatmáry B., 212
 Szuts T., 188
- T. Sommer F., 177
 Tailby C., 168
 Taillefumier T., 100
 Takalo J., 86
 Takiyama K., 204
 Tang S., 253
 Tankus A., 134
 Teng C.-L., 237
 Thesen T., 239
 Theunissen F. E., 246
 Thomas P. J., 83, 233
 Thomure M., 203
 Thoroughman K. A., 40
 Tiesinga P., 225
 Tiesinga P. H., 233
 Tkacik G., 235
 Tolias A. S., 50
 Touns J. V., 233
 Toyozumi T., 201
 Trappenberg T., 77, 78
 Triesch J., 60, 70, 148, 151
 Tseng Y.-T. L., 245
 Tsigankov D., 46
 Tumer E., 72
- Turner G. C., 163
 Turner R. E., 136
 Tuthill J. C., 106
 Tzounopoulos T., 25
- Uchida N., 163
 Ueno K., 75
 Urbanczik R., 241
 Usrey W. M., 187
- Vähäsöyrinki M., 86
 Vahasoyrinki M., 109
 Vaingankar V., 177
 van der Meer M., 152
 van Drongelen W., 203
 van Rossum M. C. W., 45
 van Wingerden M., 66
 Varberg Z., 141
 Vasconcelos N., 65
 Vasilaki E., 145
 Veit J., 85
 Verghese P., 216
 Vervaeke K., 24
 Vicente M., 54
 Vicente R., 241
 Victor J., 130
 Victor J. D., 258, 259
 Vidne M., 133
 Vincent B., 48
 Vinck M., 66, 120
 Vinnik E., 214
 Vogelstein J. T., 107, 208
 Vu V. D., 189
- Wade A., 216
 Wallace E., 203
 Wallis J., 242
 Wallis J. D., 54
 Wang C., 239
 Wang X., 171, 177
 Wang X.-J., 49, 195
 Wang Y., 91
 Warren T., 72
 Watanabe M., 77, 78
 Watkins P. V., 175
 Weber B., 102
 Weckström M., 86
 Weckstrom M., 109
 Wei W., 90
 Weisswange T. H., 151
 Welday A., 229
 Wenger M., 113
 Wessel R., 120
 White B., 93
 White B. L., 257

Wick S. D., [132](#)
Wiegraebe W., [175](#)
Wiles J., [83](#)
Williams S. T., [130](#)
Wilson R., [42](#)
Wilson R. C., [35](#), [209](#)
Wimmer G. E., [152](#)
Wohrer A., [166](#)
Wolf F., [90](#), [231](#)
Womelsdorf T., [66](#), [120](#), [194](#)
Wong-Lin K., [196](#)
Wood R., [228](#)
Woolley S., [98](#)
Wyeth G., [83](#)

Xia J., [244](#)
Xing D., [184](#), [260](#)
Xu Y., [223](#)

Yaeli S., [170](#)
Yang Y., [33](#)
Yeh C.-I., [184](#), [260](#)
Yen S.-C., [240](#)
Yoo Y., [205](#)
Yoshida M., [191](#)
Yu A. J., [195](#), [197](#)
Yu B., [154](#)
Yu R., [175](#)
Yuste R., [208](#), [227](#)

Zador A. M., [33](#), [97](#), [193](#), [245](#)
Zaraza D., [72](#)
Zemel R., [27](#)
Zemelman B., [140](#)
Zhang F., [141](#)
Zhang K., [61](#), [213](#), [229](#)
Zhang P., [55](#)
Zhang Y., [178](#)
Zhao Y., [25](#)
Ziegler L., [221](#)
Ziemba C. M., [187](#)
Znamenskiy P., [97](#)