

Neural Coding Computation and Dynamics

September 15–18, 2007

**Sporting Casino Hossegor
119, avenue Maurice Martin
Hossegor, France**

Programme Summary

Saturday, September 15

18:30 – 20:00 **Reception**
20:00 – 22:00 **Dinner**

Sunday, September 16

10:00 – 11:15 **Talks:** Ganguli, Ratcliff 7, 11
11:15 – 11:45 **Break**
11:45 – 13:00 **Talks:** Beck, Averbek 6, 5
13:15 – 14:30 **Lunch**
15:30 – 17:30 **Posters**
17:30 – 18:15 **Talks:** Kass 8
18:15 – 18:45 **Break**
18:45 – 19:45 **Talks:** Staude, Roxin 15, 11
20:00 – 22:00 **Dinner**

Monday, September 17

10:00 – 11:15 **Talks:** Smith, Häusler 13, 8
11:15 – 11:45 **Break**
11:45 – 13:00 **Talks:** Thiele, Tolias 16, 16
13:15 – 14:30 **Lunch**
15:30 – 17:30 **Posters**
17:30 – 18:15 **Talks:** Nirenberg 10
18:15 – 18:45 **Break**
18:45 – 19:45 **Talks:** Shlens, Kretzberg 13, 9
20:00 – 22:00 **Dinner**

Tuesday, September 18

10:00 – 11:15 **Talks:** Shenoy, Solla 12, 14
11:15 – 11:45 **Break**
11:45 – 13:00 **Talks:** Harris, Machens 7, 9
13:15 – 14:30 **Lunch**
14:30 – 15:15 **Talks:** O’Keefe 10
15:15 – 15:45 **Break**
15:45 – 16:45 **Talks:** Siapas, Frank 13, 6

Programme

Saturday, September 15

Evening

- 18:30 – 20:00 **Reception**
 20:00 – 22:00 **Dinner**

Sunday, September 16

Morning

- 10:00 – 10:45 Spikes, synapses, learning and memory in simple networks
 Surya Ganguli, University of California, San Francisco 7
- 10:45 – 11:15 Modeling behavioural reaction time (RT) data and neural firing rates
 Roger Ratcliff, Ohio State University 11
- 11:15 – 11:45 **Break**
- 11:45 – 12:30 Bayes' optimal inference, decision making, and learning with Probabilistic Population Codes
 Jeff Beck, University of Rochester 6
- 12:30 – 13:00 Data limitations on information estimates in large populations
 Bruno Averbeck, University College London 5

Afternoon

- 13:15 – 14:30 **Lunch**
 15:30 – 17:30 **Posters**

Evening

- 17:30 – 18:15 Statistical models for spike trains: some thoughts on the current situation
 Rob Kass, Carnegie Mellon University 8
- 18:15 – 18:45 **Break**
- 18:45 – 19:15 Detecting the existence of higher-order correlations in massively parallel spike trains
 Benjamin Staude, RIKEN Brain Science Institute 15
- 19:15 – 19:45 Extracting functional connectivity from simultaneous Ca²⁺ imaging of large numbers of neurons
 Alex Roxin, Pompeu Fabra University 11
- 20:00 – 22:00 **Dinner**

Monday, September 17

Morning

10:00 – 10:45	Stimulus and distance dependence of neuronal correlation in macaque primary visual cortex Matt Smith, Carnegie Mellon University	13
10:45 – 11:15	Superposition of information in large ensembles of neurons in primary visual cortex Stefan Häusler, Technische Universität Graz	8
11:15 – 11:45	Break	
11:45 – 12:30	Mechanisms and neuropharmacology of visual attention in primate V1 Alex Thiele, University of Newcastle upon Tyne	16
12:30 – 13:00	Population codes, correlations and coding uncertainty Andreas Tolias, Baylor College of Medicine	16

Afternoon

13:15 – 14:30	Lunch	
15:30 – 17:30	Posters	

Evening

17:30 – 18:15	Ruling out and ruling in neural codes Sheila Nirenberg, Weill Medical College of Cornell University	10
18:15 – 18:45	Break	
18:45 – 19:15	Spatial organization of large-scale concerted activity in primate retina Jon Shlens, Salk Institute	13
19:15 – 19:45	Multiplexing of visual information by spike rate and latency encoding in retinal ganglion cells Jutta Kretzberg, University of Oldenburg	9
20:00 – 22:00	Dinner	

Tuesday, September 18

Morning

10:00 – 10:45	Extracting dynamical structure embedded in premotor cortical activity Krishna Shenoy, Stanford University	12
10:45 – 11:15	Decoding neural signals for the control of movement Sara Solla, Northwestern University	14
11:15 – 11:45	Break	
11:45 – 12:30	How do neurons work together? Insights from auditory cortex. Ken Harris, Rutgers University	7
12:30 – 13:00	Low-dimensional network models for recordings from the prefrontal cortex Christian Machens, University of Munich	9

Afternoon

13:15 – 14:30	Lunch	
---------------	--------------	--

14:30 – 15:15	The role of oscillations in the Hippocampal formation John O’Keefe, University College London	10
15:15 – 15:45	Break	
15:45 – 16:15	Decoupling through synchrony in hippocampal networks Thanos Siapas, California Institute of Technology	13
16:15 – 16:45	Coordinated reactivation of hippocampal cell assemblies during learning Loren Frank, University of California, San Francisco	6

Posters

A new class of neural encoding models for high dimensional stimuli and nonlinear contextual effects Misha Ahrens, Gatsby Unit, UCL	17
Performance of different local measures of spike time variability and estimating internal state dynamics of cortical neurons. Adrian Ponce Alvarez, INCM-UniversiteAix-Marseille II	17
Efficient supervised learning in networks with binary synapses Carlo Baldassi, I.S.I. Foundation	18
Irregular persistent activity induced by synaptic excitatory feedback Francesca Barbieri, I.S.I. Foundation	19
Near-maximum entropy models for binary neural representations of natural images Philipp Berens, Max-Planck-Institute for Biological Cybernetics	19
Population coding of categories Laurent Bonnasse-Gahot, Ecole des Hautes Etudes en Sciences Sociales	20
Interplay between a phase response curve and an activity-dependent rewiring rule of neurons leads to wireless clustering Hideyuki Cateau, RIKEN Brain Science Institute	20
Inferring Neural Firing Rates from Spike Trains Using Gaussian Processes John P. Cunningham, Stanford University	21
Studying the effects of noise correlations on population coding using a sampling method Alexander S. Ecker, Max-Planck-Institute for Biological Cybernetics	21
Attention improves object encoding in monkey area V4 Udo Ernst, Ecole Normale Superieure	22
A method for recording neural activity simultaneously from independently moveable, chronically implanted electrodes in cortical and sub-cortical areas of monkeys. Joseph Feingold, Massachusetts Institute of Technology	23
Repeating spatiotemporal spike patterns reflect functional network states in the visual cortex Kai Gansel, Max-Planck-Institute for Brain Research	24
Evaluating information transfer between neurons: a different view on causality and integration within neuron assemblies Boris Gourevitch, University of Calgary	24
Decoding responses from populations of neurons in macaque primary visual cortex Arnulf Graf, New York University	25
Extracting synfire correlation structure from massively parallel recordings Sonja Gruen, RIKEN Brain Science Institute	26
Stable, highly irregular persistent activity induced by colored noise Kosuke Hamaguchi, RIKEN Brain Science Institute	27

Laplace's method in neural decoding Shinsuke Koyama, Carnegie Mellon University	27
High frequency organization and synchrony of activity in the Purkinje cell layer of the cerebellum Clement Lena, Ecole Normale Superieure	28
Back to the future: episodic memories for control Máté Lengyel, Collegium Budapest Institute for Advanced Study	28
Decomposing neural circuits using model-based analysis: Applying the sSoTS model to fMRI data. Eirini Mavritsaki, University of Birmingham	29
Dynamics of spontaneous activity in neuronal networks Alberto Mazzoni, I.S.I. Foundation	29
STDP results from neural dynamics, and not from a specific optimality criterion Lucas Parra, City University of New York	30
The effect of single neuron properties and input diversity on the contrast response function in V1 E. Persi, CNRS	30
Time course of reactivation of memory-related cell ensembles in the rat medial prefrontal cortex during sleep. A. Peyrache, CNRS College de France	31
Sequential non-parametric change-point detection for time series models: assessing the functional dynamics of neuronal networks Fabio Rigat, University of Warwick	31
Spike train decoding without spike sorting Valerie Ventura, Carnegie Mellon University	32
Neuronal synchronization underlies effective inter-areal interactions in visual cortex Thilo Womelsdorf, Radboud University Nijmegen	32

Abstracts

Abstracts for talks appear first; those for posters next. Within each section, the abstracts are ordered alphabetically by the presenting author's last name. An index of all authors appears at the back.

Data limitations on information estimates in large neuronal populations

Bruno B. Averbeck

Sobell Institute of Motor Neuroscience and Movement Disorders, Institute of Neurology, University College London

Understanding the role of correlations among neurons in information coding is a complex problem composed of several more specific questions. We are, however, beginning to obtain answers to certain of the questions. For example, we now understand how several different measures that have been developed to assess the effects of correlations are related. Further, analysis of experimental data is beginning to demonstrate rather clearly that correlations play a limited role in information coding at the level of pairs of neurons. In an effort to move beyond pairs, theoretical studies have attempted to address the question of whether or not correlations have an effect at the population level. These studies have generally proceeded by assuming an empirically informed structure for correlations, and then using this structure to estimate information as a function of population size. These theoretical studies have shown two related facts about the effects of correlations on populations of neurons. First, effects of correlations tend to get larger as the size of the population grows. Second, small effects of correlations in pairs of neurons do not imply small effects at the population level. This implies that the empirical results that have been obtained in pairs cannot be directly extrapolated to populations. Thus, the effects of correlations at the population level will have to be assessed directly at the level of the population.

Assessing information in large populations of neurons, however, raises an additional problem, which is less important in pairs of neurons. Namely, the problem of obtaining accurate information estimates in large populations. The question I will address is, how much data is necessary to show that there is additional information in signal dependent correlations, in a large population?

We have derived analytical results which answer this question for several estimators, including a naive estimator, a regularized Bayesian estimator, and an estimator based on boosting. If we consider the specific problem of assessing the effect of signal-dependent correlations, the problem depends on several factors including the number of terms in a model that can extract information from signal dependent correlations, the number of terms in the corresponding model that can only extract linear information, the variance of the quadratic model and the variance of the linear model, which is always greater than or equal to the variance of the quadratic model. We have found that the number of trials, necessary to show that there is information in signal dependent correlations using cross validation is quadratic in the number of predictor variables. This result can be combined with a model of the effect of signal dependent correlations developed by Shamir and Sompolinsky [1] which provides us with estimates of how the variance of the quadratic model is related to the variance of the linear model as a function of the number of neurons in the population. Using this we find that we would need approximately 2000 trials to show an effect with 30 neurons recorded simultaneously. Thus, the data requirements of answering this question are large.

1. M. Shamir and H. Sompolinsky *Neural Comput.* **16**: 1105-1136 (2004).

Bayes' optimal inference, decision making, and learning with Probabilistic Population Codes

Jeff Beck
University of Rochester

Human behavior has been shown to take uncertainty into account when combining ambiguous cues and to do so in a Bayes' optimal way. This particular computation requires a neural code which represents entire probability distribution functions rather than simply estimates. Moreover, this code must be structured so that the operations available to neural circuits are, in fact, capable of implementing (and learn to implement) optimal cue combination and action selection. Here, we will show how the Probabilistic Population Coding (PPC) framework naturally links biological constraints on neural operations with an optimal form of variability that leads to specific predictions regarding stimulus conditioned neural statistics. As an example, we will show how the requirement that optimal cue combination be performed by linear operations implies that neural variability should exhibit tuning curve like behavior with arbitrary correlations and fixed (but not necessarily unit) Fano factors. We will then show that this particular form of neural variability makes it possible to optimally implement other useful probabilistic operations such as posterior diffusion/saliency, maximum likelihood estimation, and information maximization via the neural operations of divisive normalization, dynamic attraction, and delta rule learning respectively.

Coordinated reactivation of hippocampal cell assemblies during learning

Loren M. Frank and Sen Cheng
W.M. Keck Center for Integrative Neuroscience and Department of Physiology, University of California, San Francisco, CA 94143-0444

The hippocampus is essential for the acquisition of new memories for places and events. During learning the hippocampal network undergoes rapid plastic changes that establish new representations in hippocampal cell assemblies. For more permanent storage of memories, newly formed hippocampal representations are thought to be consolidated into neocortex through coordinated reactivation of cell assemblies during ripples. However, no direct link has been established between ripple reactivation and memory formation. Here we examine dynamic changes in spatial activity and reactivation of hippocampal neurons that represent novel or familiar locations. In our experiment four rats performed an alternation task in a maze containing one novel and one familiar arm. Animals clearly distinguished between the novel and the familiar arm throughout days 1-3 of novel exposure, running more slowly in the novel arm. To study the formation of spatial memories, we compared the neural representations of the novel and familiar arms in simultaneously recorded CA1 neurons. Hence, we selected neurons with place fields in the novel and familiar arms (novel arm cells and familiar arm cells, respectively).

We found that novel arm cells were initially much more likely to be reactivated (i.e., to fire at least one spike) during any given ripple than were familiar arm cells on days 1 and 2. These differences were no longer significant on day 3. For example, on day 1 novel arm cells were about two and a half times more likely to be reactivated than their familiar arm counterparts. Since individual novel arm cells were initially more active during ripples, pairs of novel arm cell should have been more co-active than their familiar arm counterparts. Indeed, on day 1 novel arm pairs were about six times more likely to be co-active than familiar arm cell pairs. This significant co-activity remained across all three days of exposure to the novel arm. Furthermore, reactivation of novel arm cell pairs was initially more coordinated: pairs were more co-active than would be expected if the two cells were firing independently during ripples.

We also measured the width (root-mean-square (RMS) time lag) of the CCG between neuron pairs including only spikes that occurred within ripples. Novel arm cell pairs were significantly more precisely coordinated than familiar arm cell pairs on days 1 and 2, but not on day 3. Novel arm cell pairs were therefore more frequently co-active in a window compatible with synaptic plasticity. This highly precise reactivation of novel arm cells was present throughout the environment, and could not be explained by proximity to the novel arm.

If, as has been proposed, reactivation is a simple replay of experience, then the level of coordinated reactivation during ripples should reflect the level of coordinated spatial activity during experience. We therefore examined the evolution of spatio-temporal structure in single cells and cell pairs. Surprisingly, reactivation was strongest and most coordinated when spatial activity was least coordinated, and reactivation became weaker as place cells came to express reliable spatio-temporal organization during exploration. These results demonstrate that coordinated reactivation of cell assemblies is strongest in novel environments where the hippocampus must form new representations and weaker in familiar environments where a stable representation is already present. Ripples originate in hippocampal area CA3, and we propose that one important role for reactivation is to replay recently learned correlations into CA1, allowing for off-line learning within the hippocampus following experience.

Spikes, synapses, learning and memory in simple networks

Surya Ganguli
UCSF

Multiple time scales underly neural computation, coding and dynamics. The interactions between biophysical dynamics occurring on these multiple time scales presents a significant challenge in understanding their collective contribution to neural computation and coding. For example how might fast spiking dynamics, intermediate range oscillations, and slower changes in synaptic efficacies collectively give rise to even longer lasting effects, such as long term memory traces? We will explore this type of question in simple model neural networks. We will first show how to explicitly compute the statistical properties of multineuronal spike trains directly in terms of such networks' synaptic efficacies and latencies. We then apply this result to generate and analyze an effective dynamics on the space of synaptic parameters that captures the slow evolution of synaptic patterns in response to spiking inputs. Such an effective dynamics can aid in bridging the time scales between fast spiking dynamics and slower changes in long term memory traces due to changes in synaptic patterns. Finally we comment on how simple models of this form might also aid in solving an inverse problem: given multineuronal spiking statistics, how can one infer underlying synaptic efficacies and latencies?

How do neurons work together? Insights from auditory cortex.

Kenneth D. Harris, Artur Luczak, Peter Bartho, Stephan Marguet, Carina Curto, Shuzo Sakata
Center for Molecular and Behavioral Neuroscience, Rutgers University

The number of spike patterns expressible by even a modest number of neurons is astronomical, particularly if precise spike times, as well as firing rates, are taken into account. In auditory cortex, neural activity can be triggered by sound stimuli, but may also occur spontaneously during silence. We found that both sensory-evoked and spontaneous activity patterns are intricately structured but highly and similarly constrained. Individual neurons responded to stimuli with consistent temporal dynamics, revealing a stereotyped spread of activity across the recorded population. The temporal sequence of neurons firing was quantitatively similar for both spontaneous and evoked events, as well as across different acoustic stimuli, suggesting that the primary difference between responses consisted of the combination of neurons participating in an event, rather than their sequential order. The set of possible neural combinations activated by presentations of any single stimulus was confined to a subset of the set of all possible combinations; the subspaces corresponding to individual stimuli lay within a larger but still constrained realm outlined by the set of spontaneous events. To investigate how such constraints could arise in neural circuits, we constructed a recurrently connected network simulation of diverse excitatory and inhibitory units, which produced similar behavior. Although constraints on spike patterns are inefficient from the point of view of energy efficient coding, we suggest they are an inevitable consequence of information processing in recurrent neuronal networks.

Superposition of information in large ensembles of neurons in primary visual cortex

Stefan Häusler¹, Wolf Singer^{2,3}, Wolfgang Maass¹ and Danko Nikolic^{2,3}

¹Inst. for Theoretical Computer Science, Graz Univ. of Technology, Austria ²Department of Neurophysiology, Max-Planck-Institute for Brain Research, Frankfurt (Main), Germany ³Frankfurt Inst. for Advanced Studies, Wolfgang Goethe Univ., Frankfurt (Main), Germany

We applied methods from machine learning in order to analyze the temporal evolution of stimulus-related information in the spiking activity of large ensembles of around 100 neurons in primary visual cortex of anesthetized cats. We presented sequences of up to 3 different visual stimuli (letters) that lasted 100 ms and followed at intervals of 100 ms. We found that most of the information about visual stimuli extractable by advanced methods from machine learning (e.g., Support Vector Machines) could also be extracted by simple linear classifiers (perceptrons). Hence, in principle this information can be extracted by a biological neuron. A surprising result was that new stimuli did not erase information about previous stimuli. In fact, information about the nature of the preceding stimulus remained as high as the information about the current stimulus. Separately trained linear readouts could retrieve information about both the current and the preceding stimulus from responses to the current stimulus. This information was encoded both in the discharge rates (response amplitudes) of the ensemble of neurons and in the precise timing of individual spikes, and persisted for several 100 ms beyond the offset of stimuli.

This superposition of information about sequentially presented stimuli constrains computational models for visual processing. It poses a conundrum for models that assume separate classification processes for each frame of visual input and supports models for cortical computation ([1], [2]) which argue that a frame-by-frame processing is neither feasible within highly recurrent networks nor useful for classifying and predicting rapidly changing stimulus sequences. Specific predictions of these alternative computational models are that i) information from different frames of visual input is superimposed in recurrent circuits and ii) nonlinear combinations of different information components are immediately provided in the spike output. Our results indicate that the network from which we recorded provided nonlinear combinations of information from sequential frames. Such nonlinear preprocessing increases the discrimination capability of any linear readout neurons receiving distributed input from the kind of cells we recorded from. These readout neurons could be implemented within V1 and/or at subsequent processing levels.

[1] D.V. Buonomano and M.M. Merzenich. *Science*. **267**:1028-1030 (1995).

[2] W. Maass, T. Natschlag and H. Markram. *Neural Computation*. **14**(11):2531-2560 (2002).

Statistical models for spike trains: some thoughts on the current situation

Rob Kass

Department of Statistics and Center for the Neural Basis of Cognition, Carnegie Mellon University

I will discuss some of the problems involved in analyzing spike train data, with an eye toward dealing with substantial numbers of simultaneously-recorded neurons. The perspective will be personal but, of course, I hope it will be useful. After giving a few simple examples to illustrate the value of the statistical modeling approach to analysis of spike train data, I will talk about non-Poisson point processes; trial-to-trial variation, adaptation, and learning; and several approaches to short and long time-scale correlation among pairs of neurons and arrays of neurons. This will provide an evaluation of where we are, and where we should be a few years from now.

Multiplexing of visual information by spike rate and latency encoding in retinal ganglion cells

Jutta Kretzberg¹, Insa Winzenborg¹, Imke C.G. Reimer¹, Leon M. Juarez Paz¹, Martin Greschner² and Andreas Thiel¹

¹University of Oldenburg, Oldenburg, Germany ²The Salk Institute for Biological Studies, La Jolla, CA, USA

The activity of retinal ganglion cells (RGCs) is the only source of information about the visual environment for the brain. Hence, the output signals of these cells must simultaneously transmit information about many dimensions of visual stimuli like light intensity, velocity etc. Moreover, several stimulus properties are changing on fast time-scales in natural environments, demanding for fast mechanisms of encoding.

Since the timing of spikes in response to stimulus changes depends strongly on the stimulus, response latency could provide a fast coding mechanism. However, unlike the experimenter, the nervous system cannot determine latency by using the information at which time the light stimulus changed. In this study, we show that transients in RGC population activity could be used as reference points to determine response latency, allowing the decoding of certain stimulus features.

We performed multi-electrode recordings from RGCs in the turtle retina. The data was analyzed with two complementary approaches, Bayesian stimulus reconstruction [1] and metric-based clustering [2].

We used two types of stimuli, i.e. slow full-field intensity modulation flicker and a moving spatial pattern, which pseudo-randomly changed speed and direction. In both situations, instantaneous stimulus changes induced transient peaks in the RGC population activity. When applying Bayesian reconstruction, several stimulus properties could be reconstructed better and faster based on response latencies of individual RGCs relative to transients in population activity than they could when spike rates were used. While relative latencies allowed better reconstruction of light intensities, temporal contrasts and changes in motion velocity, spike rates were found to be superior to estimate constant velocities.

Metric-based clustering supported these findings. This method also suggested the encoding of constant motion velocity by spike rates, because all time windows longer than 100 ms were equally well suited to estimate this property. In contrast, the estimation of velocity changes showed a clear optimum for shorter time windows, hinting towards temporal encoding.

Moreover, these findings agree with our previous result of discriminant analysis, that the combination of spike rates and relative latency improves stimulus estimation significantly [3].

In summary, we conclude that RGCs could use multiplexing of spike rates and response latencies to transmit information about several different properties of visual stimuli.

[1] Zhang et al. *J Neurophysiol* **79**:1017-1044 (1998).

[2] Victor & Purpura, *Network: Computat Neural Syst* **8**:127-164 (1997).

[3] Greschner et al., *J Neurophysiol* **96**:2845-2856 (2006).

Low-dimensional network models for recordings from the prefrontal cortex

Christian Machens¹, Ranulfo Romo², and Carlos Brody³

¹University of Munich, ²UNAM Mexico, ³Princeton

During short-term memory maintenance, different neurons in prefrontal cortex (PFC), recorded under identical conditions, show a wide variety of temporal dynamics and response properties [1]. These data are a specific example of the more general finding that neural recordings from frontal cortices often reveal that different neurons have very different response characteristics. Modeling this complexity of responses has been difficult. Most commonly, some features of the responses are focused on, and models that fit those reduced features are built. But can the full complexity of responses be easily captured? Here we attack the problem by fitting simple recurrent neural network models to the data.

Following the traditional approach, we first group neurons into different classes. When selecting neu-

rons from a single class, the estimation procedure yields a connectivity matrix with two populations of neurons coupled by mutual inhibition and self-excitation. The connectivity matrix has rank one and approximately agrees with a model we proposed earlier [2]. When selecting neurons from two classes, a connectivity matrix similar to that of the ring attractor network emerges, with a rank of two. The full complexity and richness of the observed neural dynamics, however, can only be captured when estimating a network architecture from the full set of neurons. In this case, the resulting connectivity matrix has rank four to five and its structure is dominated by randomness. Simulations of the resulting network reproduce most of the diversity and richness of neural responses found in the data. We show that several of the eigenvalues of the connectivity matrix are close zero, so that the network dynamics has either a constant or integrating flow along the respective dimensions. Finally, we discuss the consistency of the estimated connectivity matrices with the measured noise correlations.

[1] Timing and Neural Encoding of Somatosensory Parametric Working Memory in Macaque Prefrontal Cortex. C.D. Brody, A. Hernandez, A. Zainos, and R. Romo, *Cereb. Cortex* **13**:1196-1207, 2003.

[2] Flexible control of mutual inhibition: a neural model of two- interval discrimination. C.K. Machens, R. Romo, and C.D. Brody, *Science*, **307**:1121-1124, 2005.

Ruling out and ruling in neural codes

Sheila Nirenberg

Weill Medical College of Cornell University

The subject of neural coding has generated much debate. A key issue is whether the nervous system uses coarse or fine coding. Each has different strengths and weaknesses and, therefore, different implications for how the brain computes. For example, the strength of coarse coding is that it’s robust to fluctuations in spike arrival times. Downstream neurons don’t have to keep track of spike train structure. The weakness, though, is that individual cells can’t carry much information, so downstream neurons have to be able to pool signals across cells and/or time to obtain sufficient information. With fine coding, individual cells can carry a great deal of information, but downstream neurons have to be able to resolve spike train details to obtain it. Here we set up a strategy to determine what the neural code can and can’t be and used it at the level of the retina. We recorded from essentially all the retinal output cells an animal uses to solve a task, evaluated the cells’ spike trains for as long as the animal evaluates them, and used optimal, i.e., Bayesian, decoding. This makes it possible to obtain an upper bound on the performance of codes and thus eliminate those that aren’t viable. Our results show that standard coarse coding is insufficient; finer, more information-rich codes are necessary.

The role of oscillations in the Hippocampal formation

John O’Keefe

University College London, London, UK

The electrical activity of the hippocampal formation is characterized by several oscillatory potentials of different frequencies: theta (6 - 11 Hz), gamma (30 - 60 Hz), and high-frequency ripples (100- 200 Hz). In this talk I will discuss the functions of theta activity, one of the best-known of these oscillations. I will argue that it has at least three functions: 1) to bind together the activity of cells in different parts of the hippocampal formation and more widely distributed brain regions; 2) to set up the conditions for LTP or LTP depending on the relationship of synaptic inputs to a given cell and the phase of theta; and finally, 3) to provide one of the two oscillators of different frequencies which interact to produce interference patterns within the principal cells of the hippocampal formation and which determine the temporal firing patterns in these cells and their relationship to that of other cells. I will concentrate primarily on the latter function. The first evidence for interference patterns came from the phase precession effect originally observed in CA1 pyramidal cells by O’Keefe and Recce in 1993, and now known to exist in CA3 pyramids, layer 2 stellate cells of the entorhinal cortex and, to a lesser extent, dentate granule cells. We found that hippocampal pyramidal cells fired in bursts at approximately the theta frequency as the rat ran through the place field but that the preferred phase of firing did not remain constant but instead moved to earlier phases with each successive burst. Furthermore the phase of firing correlated

had a higher correlation with the animal's location within the field than with any other variable, for example, time since entry into the field. We speculated that this phase precession effect might be due to the interaction of two oscillators of slightly different frequencies summing within the pyramidal cell. This model would explain many of the observed phenomena and in particular the fact that the phase precession often continued in the second half of the field despite the fact that the firing rate decreased at same time. An alternative was that the precession effect was generated in the input structures to the hippocampus and propagated into the hippocampus. Evidence in favor of this latter view came from an experiment by Zugaro and Buzsaki who showed that temporary silencing of pyramidal cells as the animal ran through the field did not alter the relationship between phase and location. When the cell resumed firing, it did so at the appropriate phase for the animal's location in the field. The recent discovery of grid cells in the medial entorhinal cortex by Hafting and Fyhn in the Moser's lab opens up the possibility that they may be the source of the phase precession. Grid cells fire with multiple place fields located at the vertices of a triangular grid. Cells located more ventrally in the entorhinal cortex have larger spacing between firing fields than those in the dorsal entorhinal cortex. We now know that the grid cells located in layer 2 also show the phase precession effect. In a recent computational model, Neil Burgess, Caswell Barry and I have proposed a modification of the dual oscillator interference model which accounts for the firing patterns observed in grid cells. One feature of the model has already been tested and supported. This is that the resonant frequency of the intrinsic oscillator within the EC layer 2 stellate cells should vary systematically from dorsal to ventral entorhinal cortex in order to account for the change in scale of the grid size in that dimension. This has now been demonstrated by Hasselmo's group who have also shown that the time constants of the H current vary systematically from dorsal to ventral in a way which would explain the changes in natural frequency.

Modeling behavioural reaction time (RT) data and neural firing rates

Roger Ratcliff
Ohio State University

There has been progress in both psychological modeling and neurophysiology in attempting to understand simple decision processes. The kinds of models used are sequential sampling models of the class of diffusion processes. We present a leaky accumulator model with negatively correlated starting points and discuss how it accounts for standard RT phenomena including accuracy, correct and error RTs and their distributions. This model assumes evidence is accumulated in two separate accumulators to separate decision criteria. We then show how simulated paths in the two accumulators mimic neural firing rate data recorded from buildup cells in the superior colliculus of rhesus monkeys in a brightness discrimination task. I will focus on the criteria for behavioral modeling, including fitting RT distributions and error RTs. Much work in the neurophysiology domain currently ignores these dependent variables, but in the psychological domain, they have proved decisive in deciding among models.

Extracting functional connectivity from simultaneous Ca^{2+} imaging of large numbers of neurons

Alex Roxin,¹ Nicolas Brunel² and Vincent Hakim³

¹Universitat Pompeu Fabra, Barcelona, Spain. ²CNRS, Paris, France. ³ENS, Paris, France.

It seems likely that many aspects of higher brain function involve the concerted activity of large numbers of interacting neurons. Inspired by this insight, much theoretical effort has been put into understanding the collective behavior of coupled neurons in computational models. However, such models have proven difficult to constrain given the paucity of large-scale, simultaneous recordings of neuronal activity, largely due to technical impediments. With the advent of voltage-sensitive dye and Ca^{2+} imaging techniques, it is now possible to record the activity of large numbers of neurons simultaneously, opening the doors to quantitative studies of network-wide neuronal interactions in brain tissue.

Here we discuss the analysis of calcium-imaging data from slices of primary visual cortex and medial prefrontal cortex in mice from the laboratory of Rafael Yuste. Intracellular recordings confirmed that

calcium-signal transients in individual cells in these slices corresponded to transitions between cortical up- and down-states. Earlier work on these data revealed the presence of reliably repeating spatio-temporal sequences of calcium-transients [1,2], which appeared more often than chance compared to selected surrogate time series. We have revisited these data and compared the frequency of occurrence of such patterns to several possible models. We analyzed 15 data sets totaling over an hour of recordings in over 2000 cells.

Of the 15 data sets we find that the occurrence of repeating spatio-temporal patterns in 10 of them is consistent with each neuron behaving as an independent Poisson process with a refractory period. In four of the data sets, the patterns are best described by a probabilistic network model for which we extract a transition probability matrix describing interactions between neurons. We study the structure of the resulting connectivity matrices and find, consistent with data from paired intracellular recordings [3], highly nonrandom features including an over-representation of doublet and triplet interactions. Interestingly, however, we find that the statistical property of the connectivity which most directly influences the firing patterns of the cells is the degree distribution, i.e. the number of links into and out of each cell. We find a large number of ‘hubs’, cells with an especially large number of links.

[1] R. Cossart, D. Aronov and R. Yuste *Nature* **423**:283-288 (2003).

[2] Y. Ikegaya et al., *Science* **304**:559-564 (2004).

[3] S. Song et al., *PloS Biology* **3**:507-515 (2005).

Extracting dynamical structure embedded in premotor cortical activity

Krishna V. Shenoy¹, Byron M. Yu¹, Afsheen Afshar¹, Mark M. Churchland¹, John P. Cunningham¹, and Maneesh Sahani²

¹Stanford University, Stanford, CA ²Gatsby Computational Neuroscience Unit, University College London

Our seemingly effortless ability to reach out and swat a fly or grab a cup belies the sophisticated neural computations at work in our nervous system. It has long been recognized that, before moving, we somehow prepare neural activity such that, when called upon, the desired movement unfolds. But the goals of movement preparation and the underlying neural mechanisms remain poorly understood. I will describe some of our recent electrophysiological investigations of how premotor cortex prepares and helps execute movements. Our results suggest that the brain is attempting to optimize preparatory neural activity [1] and can delay movement until this activity is sufficiently accurate [2]. We note that the spiking activity during motor preparation exhibits dynamics beyond that driven by external stimulation, presumably reflecting the extensive recurrence of neural circuitry. We have been investigating methods for capturing the dynamics from (96 channel) simultaneous neural recordings using low-dimensional non-linear dynamical systems models, with underlying recurrent structure and stochastic point-process output. I will present recent work using latent variable methods that simultaneously estimate the system parameters and the trial-by-trial dynamical trajectories [3], and recent evidence that low-dimensional trajectories may be able to better predict the reaction time of the upcoming movement. Characterizing these dynamics may reveal important features of neural computation, and may be useful for further increasing the performance of neural prostheses [4,5].

[1] Churchland MM, Yu BM, Ryu SI, Santhanam G and Shenoy KV *J. Neurosci.* **26**:3697-3712 (2006).

[2] Churchland MM and Shenoy KV *J. Neurophysiol.* **97**:348-359 (2007).

[3] Yu BM, Afshar A, Santhanam G, Ryu SI, Shenoy KV and Sahani M *NIPS* **18**:1545-1552 (2006).

[4] Santhanam G, Ryu SI, Yu BM, Afshar A and Shenoy KV *Nature* **442**:195-198 (2006).

[5] Yu BM, Kemere C, Santhanam G, Afshar A, Ryu SI, Meng TH, Sahani M and Shenoy KV *J. Neurophysiol.* **97**:3763-3780 (2007).

Spatial organization of large-scale concerted activity in the primate retina

Jonathon Shlens¹, Greg D. Field¹, Jeffrey L. Gauthier¹, Martin Greschner,¹ Matthew I. Grivich¹, Dumitru Petrusca,² Alexander Sher², Alan M. Litke² and E.J. Chichilnisky¹

¹Salk Institute for Biological Studies, ²UC Santa Cruz

All visual signals in the brain originate in the electrical activity of retinal ganglion cells (RGCs). Standard models implicitly assume that RGCs signal information independently of one another. However, several studies have demonstrated significant concerted activity in nearby pairs of RGCs which may fundamentally alter visual signals. Very little is known about the spatial structure of this concerted activity in large populations of RGCs, particularly in the primate retina.

We performed large-scale multi-electrode recordings of the electrical activity of several hundred ON and OFF parasol (magnocellular-projecting) RGCs in isolated peripheral macaque monkey retina. The regular mosaic organization of the recorded cells indicated that we recorded from nearly every cell of both types in a 4x8 degree region of the visual field. In the presence of steady, spatially uniform photopic illumination, pairs of RGCs fired synchronously (within 10 ms) several-fold more often than expected by chance, indicating significant network interactions. Synchrony declined with distance between cells, and was universal among nearby cells of the same type, indicating that it arises from local and highly stereotyped circuitry.

To probe the spatial structure and scale of network interactions, we measured the total number, contiguous spatial extent, and number of adjacencies in the collection of cells firing within each 10 ms time bin. To test whether concerted firing can be explained by known interactions between immediate neighbors in the mosaic, we used a maximum entropy approach borrowed from statistical mechanics to predict the above measurements based purely on measured pairwise correlations between neighboring cells. The predictions of this model accurately reproduced the data. This suggests that the spatial structure of spontaneous activity in parasol cells of primate retina can be understood entirely on the basis of single cell activity and nearest neighbor interactions.

[1] Shlens et al, *Journal of Neuroscience***26**(32):8254-8266 (2006).

Decoupling through synchrony in hippocampal networks

Thanos Siapas, Eugene Lubenov
California Institute of Technology

The level of synchronization in distributed systems is often controlled by the strength of the interactions between individual elements. In brain circuits the connection strengths between neurons are modified under the influence of spike timing-dependent plasticity rules (STDP). We show that when recurrent networks exhibit population bursts, STDP rules exert a strong decoupling force that desynchronizes activity. The decoupling force of STDP may be engaged by the synchronous bursts occurring in the hippocampus during slow-wave sleep. We discuss experimental evidence that these bursts may lead to the selective erasure of information from hippocampal circuits as memories are established in neocortical areas.

Stimulus and distance dependence of neuronal correlation in macaque primary visual cortex

Matthew A Smith¹ Adam Kohn² Ryan C Kelly¹ Tai Sing Lee¹

¹Center for the Neural Basis of Cognition, Carnegie Mellon University ²Department of Neuroscience, Albert Einstein College of Medicine

Common synaptic input is thought to give rise to both correlated trial-to-trial response variability and synchronous spiking in nearby cortical neurons. As a way to determine the source of this common input, we measured how correlation between neurons depends on the similarity of their tuning properties, the cortical distance between them, and the orientation of the stimulus.

We recorded from the primary visual cortex (V1) of anesthetized, paralyzed macaque monkeys using two techniques: (1) pairs of electrodes spaced 300 to 500 microns apart, and (2) A 100-electrode array. The array consisted of a 10 x 10 grid with 400 micron spacing between adjacent electrodes. Electrode length was 1.0 mm and the array was implanted 0.6 mm into cortex, resulting in superficial layer recordings. We presented sinusoidal grating stimuli drifting in different directions to elicit neuronal responses. From each electrode, we recorded all responses which exceeded a threshold and sorted these waveforms offline. Simultaneously, we recorded the local field potential (LFP) from each electrode, a signal considered to be related to the sum of all synaptic input to a region of tissue.

We measured correlation on two different time scales from the spiking data. Spike count correlation (rsc, or "noise" correlation) was calculated over the course of the entire stimulus presentation (1.28 seconds). Synchrony (precise temporal correlation) was measured by examining the spike-train cross-correlogram (CCG) peak height near zero time lag. We determined the coherence of the raw LFP in seven frequency bands, ranging from delta (1-4 Hz) to very high gamma (100-150 Hz). We found that spike count correlation did not change with stimulus orientation. Although rsc was smaller for widely separated neurons, it remained significantly above zero even at separations greater than 4 mm. The peak height of the CCG was strongest for the orientation that produced the best response for the pair of neurons. Similar to spike count correlation, CCG peak height decreased with distance, but at a faster rate. In addition we found that the coherence of high frequency LFP components falls off more precipitously with distance than that of low frequency LFP components (similar to a previous study in awake behaving macaques [1]). We also calculated the time scale of spiking correlation and compared it to the LFP. Pairs of single units which had their strongest correlation on a time scale of roughly 200 ms tended to have spiking activity that was tightly locked to the LFP.

Together, our results suggest that sharp synchrony and spike count correlation on long timescales arise from different mechanisms. The dependence of synchrony on distance and stimulus orientation suggests it is primarily (but not exclusively) driven by common input arising from short-range cortical circuitry. The spatial extent of spike count correlation and its lack of dependence on stimulus orientation suggest it is mediated by connections which cover large regions of cortex. Furthermore, different frequency components of the LFP are related to measures of spiking correlation on different time scales. Pairs of neurons which are correlated on a broad time scale appear to be the most influenced by the LFP.

[1] D.A. Leopold, Y. Murayama, N.K. Logothetis *Cereb. Cortex.* **13**:422-433 (2003).

Decoding neural signals for the control of movement

Sara A. Solla, James M. Rebesco, Eric Pohlmeier, and Lee E. Miller
Northwestern University

The activity of neurons in primary motor cortex provides the signals that control our ability to execute movements. One of the crucial questions, still unresolved, is that of identifying the code used by this neural ensemble. Do neurons in primary motor cortex use a representation of the external physical space, or an actuator representation of the state of the muscles? We address this question through the analysis of data obtained for awake behaving monkeys executing stereotyped reaches. The data includes both simultaneous recordings of the activity of about one hundred neurons in primary motor cortex, and of the activity of about ten muscles in the relevant limb. The analysis of this data involves a variety of techniques, from linear regression models to nonlinear methods for dimensionality reduction. I will review the current level of achievement in this active area of research and discuss its implications, both for understanding aspects of neural information processing that relate to natural behaviors and for extracting from these neural signals the information needed to guide prosthetic limbs and other types of external devices.

The presentation will focus on two projects: i) A natural representation for the ensemble activity is provided by a high-dimensional space in which each axis represents the activity of a single neuron as an independent degree of freedom. However, the observed correlations among neurons whose activity is detectably modulated by a task suggest that the population activity defines a low-dimensional space within the high-dimensional space of independent firing activities. We have used linear and nonlinear methods for dimensionality reduction to find the low-dimensional structure that captures the underlying

relationship between population neural activity and behavioral task. The use of multidimensional scaling in conjunction with an empirical measure of geodesic distances yields a low-dimensional manifold whose intrinsic coordinates capture the geometry of the task in the external physical space. This nonlinear mapping provides a highly informative representation for the prediction of reach directions. ii) Ensemble activity is also predictive of electromyographic (EMG) recordings of muscle activity during grasps. Appropriately regularized linear models easily achieve 60% to 80% prediction of dynamic EMG during palmar, lateral, and precision grasps. The ability to predict EMG signals from neural recordings can then be used to generate useful patterns of muscle activation to guide limb motion during nerve block. This paradigm provides biophysically generated signals for the real-time functional electrical stimulation (FES) of muscles that have lost nerve innervation.

Detecting the existence of higher-order correlations in massively parallel spike trains

Benjamin Staude¹, Stefan Rotter^{2,3} and Sonja Gruen^{1,4}

¹Computational Neuroscience Group, RIKEN Brain Science Institute, Waco-Shi, Japan ²Theory and Data Analysis, Institute for Frontier Areas of Psychology and Mental Health, Freiburg, Germany ³Bernstein Center for Computational Neuroscience, Freiburg, Germany ⁴Bernstein Center for Computational Neuroscience, Berlin, Germany

The cell assembly hypothesis [1] postulates dynamically interacting groups of neurons as building blocks of cortical information processing. Synchronized spiking across large neuronal groups was later suggested as a potential signature for active assemblies [2], predicting higher-order correlations (HOCs) among the spike trains of assembly members. However, the estimation of the necessary parameters of present analysis techniques for HOCs poses serious problems, mainly because their number grows exponentially with the number of recorded neurons [3,4]. As a consequence, most attempts to detect active cell assemblies resort to pairwise correlations. Such pairwise correlations, however, do not reflect potential HOCs and are insensitive for sparse synchronous events [5]. As massively parallel extracellular recordings are becoming more and more available, the limited experimental evidence in favor of the cell assembly hypothesis has to a large extent be assigned to a lack of suitable analysis tools [6].

Here we present a novel procedure to detect HOCs in massively parallel spike trains. Based on estimates of only a few low-order cumulants of the summed activity across all neurons (the 'population histogram') we devise a statistical test for the presence of HOCs among the recorded spike trains. The test exploits the fact that absence of HOCs in a neuronal population also imposes constraints on (population-average) correlations of lower order. The latter can, however, be estimated via the respective cumulants of the distribution of the entries in the population histogram. Under a compound Poisson assumption, where correlations of various orders are induced by 'inserting' appropriate patterns of near-synchronous spikes [7], the upper bounds for these lower order cumulants in the absence of HOCs can be derived analytically, together with the necessary confidence intervals of the respective k-statistics. This makes the test computationally very modest and hence applicable to large amounts of data without the need for time consuming bootstrap approaches. Furthermore, the inference of HOCs from cumulants of lower order circumvents the need to estimate large numbers of higher-order parameters, making the test less susceptible to the limited sample sizes typical for in vivo recordings than previous approaches [3,4]. We illustrate the test on data which was simulated using a compound Poisson model, and find that cumulants of third order are already surprisingly sensitive for present HOCs. Furthermore, the proposed test detects HOCs even if their effects on pairwise correlation coefficients c are very small (in the range of $c \sim 0.01$, compare [5]).

- [1] D.O. Hebb. Organization of behavior. Wiley (1949)
- [2] M. Abeles. Local cortical circuits. Springer (1982)
- [3] L. Martignon et al. *Biol Cyber.* **73**:69-81 (1995)
- [4] H. Nakahara & S.-I. Amari *Neural Comput.* **14**:2296-2316 (2002)
- [5] E. Schneidman et al. *Nature* **440**:1007-1012 (2006)
- [6] E.N. Brown, R.E. Kass, P.P. Mithra *Nat. Neurosci.* **7**:456-461 (2004)
- [7] A. Kuhn, A. Aertsen and S. Rotter. *Neural Comput.* **15**:67-101 (2003)

Mechanisms and neuropharmacology of visual attention in primate V1

Alex Thiele

University of Newcastle upon Tyne

Attention can selectively enhance neuronal responses and exclude external noise, but the neuronal computations underlying these effects and the neuropharmacology involved remain unknown. We hypothesized that the effects of attention are mediated by increasing the synaptic efficacy of feed-forward connections while simultaneously reducing the synaptic efficacy of lateral connections, possibly through contributions from cholinergic mechanisms. We tested this proposal by recording neuronal activity and length tuning in macaque V1 when attention was directed towards or away from stimuli presented in the neuron's classical receptive field. For cells with central/parafoveal receptive fields, attention indeed reduced spatial integration demonstrated by a reduction in preferred length and in the size of the spatial summation area. Conversely, in cells representing more peripheral locations, attention increased spatial integration by increasing the cell's summation area. This previously unknown dichotomy between central and peripheral vision could support accurate analysis of attended foveal objects and target selection for impending eye-movements to peripheral objects. We further tested whether cholinergic mechanisms are an integral part of attentional modulation in V1. Cholinergic antagonists reduced the attentional modulation in the attend RF condition by means of muscarinic receptor mechanisms. Nicotinic receptor mechanisms are necessary for adequate sensory transmission, but do not seem to contribute to attentional modulation. These data demonstrate that the cholinergic system plays an important role in mediating effects of attention in V1 of the macaque monkey. Supported by: Wellcome Trust, BBSRC, MRC.

Population codes, correlations and coding uncertainty

Andreas S. Tolia¹ and Alexander Ecker² and Georgios A. Keliris² and Fanis Panagiotaropoulos² and Stefano Panzeri³ and Nikos K. Logothetis²

¹Department of Neuroscience, Baylor College of Medicine, Biological Cybernetics Max-Planck-Institute,

²Biological Cybernetics Max-Planck-Institute, ³University of Manchester,

Despite progress in systems neuroscience the neural code still remains elusive. For instance, the responses of single neurons are both highly variable and ambiguous (similar responses can be elicited by different types of stimuli). This variability/ambiguity has to be resolved by considering the joint pattern of firing of multiple single units responding simultaneously to a stimulus. Therefore, in order to understand the underlying principles of the neural code it is imperative to characterize the correlations between neurons and the impact that these correlations have on the amount of information encoded by populations of neurons. We use chronically implanted tetrode arrays to record simultaneously from many neurons in the primary visual cortex (V1) of awake, behaving macaques. We find that the correlations in the trial-to-trial fluctuations of their firing rates between neurons under the same stimulation conditions (noise correlations) in V1 were very small (around 0.01 in 500 ms bin window) during passive viewing of sinusoidal grating stimuli. We are also measuring correlations in extrastriate visual areas and investigating the impact of correlations on encoding stimulus uncertainty by neuronal populations, under different stimulus and behavioral conditions.

A new class of neural encoding models for high dimensional stimuli and nonlinear contextual effects

Misha B. Ahrens¹, Liam Paninski² Jennifer F. Linden³ and Maneesh Sahani¹

¹Gatsby Computational Neuroscience Unit, UCL ²Dept. of Statistics and Center for Theoretical Neuroscience, Columbia University ³Ear Institute and Dept. of Anatomy and Developmental Biology, UCL

The fitting of meaningful models to the stimulus-response functions of neurons is often hampered by several factors. Compact models might lack the flexibility to adequately capture the nonlinear dynamics of the neural responses, but elaborate models may be hard to estimate due to a lack of good estimation algorithms and large numbers of model parameters. Here we describe a class of nonlinear neural encoding models based on multilinear (tensor) mathematics, which share many of the conveniences of linear models – such as robust estimation algorithms and low numbers of parameters – yet are able to capture nonlinear effects such as short-term stimulus-specific adaptation. They achieve this through an (interpretable) multiplicative factorization in an extended stimulus space. The effectiveness of the methods is illustrated on data from primary auditory cortex.

Performance of different local measures of spike time variability and estimating internal state dynamics of cortical neurons.

Adrian Ponce Alvarez¹ Kosuke Hamaguchi^{2,3} Nicolas Brunel² Alexa Riehle¹

¹INCM-Universite Aix-Marseille II, France, ²LPN-Universite Rene Descartes, France, ³RIKEN BSI, Japan

Irregularity of spike timing can be measured by the coefficient of variation (CV), which is a global measure defined as the dispersion of the inter-spike intervals (ISI). However, the CV undergoes large changes by rate fluctuations. Several local and relatively rate-independent measures of this irregularity are found in the literature. Unfortunately, as far as we know, these new measurements have not been compared to each other. Here we present a comparison of four of them: the "local coefficient of variation" named CV2 [1], the local variation LV [2], the metric IR [3] and the metric SI [4]. The LV measure is thought to be useful for cell classification [2]. SI is based on the best estimating function (asymptotically) of the shape parameter in a gamma-distributed assumption of spiking activity without estimating the time-dependent rate, which is an unknown function. The question arises, which of these coefficients is the most efficient for analysing experimental data where the number of inter-spike intervals is limited. The classical CV of cortical neurons usually falls in the range between 0.2 and 2.

We first compared the performance of these measures by simulating point processes with gamma-distributed intervals and using time sliding windows as used in experimental data analysis. We tested the resistance of these measures to non-stationarities and discontinuities of the firing rate. By calculating the distributions of the individual local values we showed that the measure CV2 is the less sensitive to the shortage of samples: it has the lowest estimation bias. In addition, one can calculate analytically the value of CV2 for any gamma-process.

The irregularity of the spiking activity of neurons recorded by A. Riehle and collaborators was analysed using CV2. The neurons from the motor cortex of a Rhesus monkey were recorded during a Choice Reaction-Time task [5]. In this task the animal was trained to perform arm movements in opposite directions. On a vertical panel, three touch sensitive targets were mounted in a horizontal line. Both lateral targets could be lit either in red or green, whereas the central target was lit in yellow. The animal had to initialize the trial by holding the central target. After a delay, a preparatory signal (PS) was presented: both peripheral targets were presented simultaneously, one in red and the other in green. An auditory response signal (RS) then followed after either 600 or 1200 ms. The animal had been trained to associate to each colour one of these two possible delay durations.

Interestingly, the distribution of CV2 statistics of the whole population of neurons calculated from the delay period in the Choice-reaction task shows a clear bi-modal distribution.

Although the firing rate is highly modulated during the task, the CV2 is essentially constant in time. The fact that CV2 remains constant is likely to constrain the architecture of the network in which the recorded neurons are embedded. In fact, if the increase in rate is due to purely excitatory feedforward inputs, from the analysis of simple integrate-and-fire (LIF) type models one would expect the CV2 to decrease as the rate increases. Thus, the basic observation that the CV2 remains constant as the rate

increases seems to indicate that this rate increase is at least partly due to recurrent inputs.

Using a simple LIF model one can compute the mean, m , and fluctuations, s , of the inputs that produces a given firing rate and CV². Then, we can plot contour lines of both firing rate and CV on the m/s plane, and superimpose single neuron data on that contour plot. This tells us how the statistics of that particular neuron changes during the task, if we assume that the neuron can be well approximated by a LIF neuron. The next step is to take a recurrent network model with specific parameters, and compute how the firing rate/CV of neurons in such networks vary as a function of external inputs. This gives a trajectory in the m/s plane, that depends in a pronounced fashion on the excitation-inhibition balance. This analysis suggests that most of the neurons would be operating in an inhibition-dominated network (for many neurons m is far below the threshold although their firing rate is $10\sim 40$ Hz).

- [1] Holt G. R., Softky W. R., Koch C. and Douglas R. J. Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. *J. Neurophysiol.* **75**:1806-1814 (1996).
- [2] Shinomoto S., Shima K., Tanji K. Differences in spiking patterns among cortical neurons. *Neural Comput.* **15**:2823-2842 (2003).
- [3] Davies R. M., Gerstein G. L. and Baker N. Measurement of time-dependent changes in the irregularity of neural spiking. *J. Neurophysiol.* **96**:906-918 (2006).
- [4] Miura K., Okada M. and Amari M. Estimating spiking irregularities under changing environments. *Neural Comput.* **18**:2359-2386 (2006).
- [5] Roux S, Coulmance M, Riehle A. Context-related representation of timing processes in monkey motor cortex. *Eur J Neurosci.* **18**:1011-1016 (2003).

Efficient supervised learning in networks with binary synapses

Carlo Baldassi¹, Alfredo Braunstein^{1,2}, Nicolas Brunel^{1,3} and Riccardo Zecchina^{1,2}

¹Institut for Scientific Interchange Foundation, Torino, Italy, ²Politecnico di Torino, Italy, ³Laboratoire de Neurophysique et Physiologie, UMR 8119, Universite Rene Descartes, France

Recent experiments[1,2] have suggested single synapses could be similar to noisy binary switches. Binary synapses would have the advantage of robustness to noise and hence could preserve memory over longer time scales compared to analog systems. Learning in systems with discrete synapses is known to be a computationally hard problem though. We developed and studied a neurobiologically plausible on-line learning algorithm that is derived from Belief Propagation algorithms.

This algorithm performs remarkably well in a model neuron with N binary synapses, and a discrete number of ‘hidden’ states per synapse, that has to learn a random classification problem, or to learn a classification rule from a teacher. In the first case, such a system is able to learn a number of associations which is close to the information theoretic limit, in a time which is sub-linear in system size, corresponding to very few presentations of each pattern. In the second case, perfect generalization is shown, from experiments and analytical calculations, to be achieved in finite (and short) time.

The algorithm is similar to the standard ‘perceptron’ learning algorithm, but with an additional rule for synaptic transitions which occur only if a currently presented pattern is ‘barely correct’ (that is, a single synaptic flip would have caused an error). In this case, the synaptic changes are meta-plastic only (change in hidden states and not in actual synaptic state), and go towards stabilizing the synapse in its current state. This rule is crucial to the algorithm’s performance, in both learning protocols, and we suggest that it is sufficiently simple to be easily implemented by neurobiological systems.

- [1] C.C. Petersen, R.C. Malenka, R.A. Nicoll and J.J. Hopefield, *Proc. Natl. Acad. Sci. USA* **95**:4732-4737 (1998).
- [2] D.H. O’Connor, G.M. Wittenberg and S.S-H. Wang, *Proc. Natl. Acad. Sci. USA* **102**:9679-9684 (2005).

Irregular persistent activity induced by synaptic excitatory feedback

Francesca Barbieri¹ and Nicolas Brunel^{1,2}

¹I.S.I. Foundation, Torino, Italy, ²Laboratory of Neurophysics and Physiology, CNRS-Universite Paris Descartes, Paris, France

Neuro-physiological experiments on monkeys [1] have reported highly irregular persistent activity during the performance of an oculomotor delayed-response task. These experiments show that during the delay period the ISI's coefficient of variation (CV) of prefrontal neurons is above 1, on average, and larger than during the fixation period.

Previous models [2-3] of spontaneous and selective persistent activity in the cortex based on excitatory synaptic feedback do not reproduce this feature because the excitatory feedback during persistent activity brings neurons in a region of the f-I curve in which the firing is relatively independent from fluctuations and hence the CV is small.

We show that this feature of irregularity can be reproduced in a network in which persistent activity is induced by excitatory feedback, provided that (i) the post-spike reset is close enough to threshold, (ii) synaptic efficacies are a non-linear function of the pre-synaptic firing rate. Non-linearity between pre-synaptic rate and effective synaptic strength is implemented by a standard short-term depression mechanism (STD) [4].

First, we consider the simplest possible network with excitatory feedback: a fully connected homogeneous network of excitatory leaky integrate-and-fire neurons, both with numerical simulations and analytical techniques. The results are then confirmed in a network with selective excitatory neurons and inhibition. In both cases there is a large range of values of the synaptic efficacies for which the statistics of firing of single cells is similar to experimental data.

[1] Compte A, Constantinidis J T, Raghavachari S, Chafee M V, Goldman-Rakic P s, Wang X J, Temporally irregular mnemonic persistent activity in prefrontal neurons of monkeys during a delayed response task. *J. Neurophysiol.* **90**: 3441-3454 (2003).

[2] Amit D J and Brunel N, Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex* **7**: 237-252 (1997b).

[3] Brunel N, Persistent activity and the single-cell frequency-current curve in a cortical network model. *Comput. Neural. Syst.* **11**: 1-20 (2000).

[4] Romani S, Amit D J, Mongillo G, Mean-field analysis of selective persistent activity in presence of short-term synaptic depression. *J. Comp. Neurosci.* **20** : 201-217 (2006).

Near-maximum entropy models for binary neural representations of natural images

Philipp Berens and Matthias Bethge

Computational Vision and Neuroscience Group, Max-Planck-Institute for Biological Cybernetics, Tuebingen

Maximum entropy analysis of binary variables provides an elegant way for studying the role of pairwise correlations in neural populations [1,2]. Unfortunately, these approaches suffer from their poor scalability to high dimensions. In sensory coding, however, high-dimensional data is ubiquitous. Here, we introduce a new approach using a near-maximum entropy model, that makes this type of analysis feasible for very high-dimensional data - the model parameters can be derived in closed form and sampling is easy. We demonstrate its usefulness by studying a simple neural representation model of natural images. For the first time, we are able to directly compare predictions from a pairwise maximum entropy model not only in small groups of neurons, but also in larger populations of more than thousand units. Our results indicate that in such larger networks interactions exist that are not predicted by pairwise correlations, despite the fact that pairwise correlations explain the lower-dimensional marginal statistics extremely well up to the limit of dimensionality where estimation of the full joint distribution is feasible.

- [1] E Schneidman, MJ Berry, R Segev, and W Bialek. *Nature*, **440**:1007-1012 (2006).
 [2] J Shlens, JD Field, JL Gauthier, MI Grivich, D Petrusca, A Sher, AM Litke, and EJ Chichilnisky. *J Neurosci*, **26**:8254-8266 (2006).

Population coding of categories

Laurent Bonnasse-Gahot¹ and Jean-Pierre Nadal^{1,2}

¹ Centre d'Analyse et de Mathematique Sociales (CAMS, UMR 8557 CNRS-EHESS) Ecole des Hautes Etudes en Sciences Sociales, Paris, ² Laboratoire de Physique Statistique (LPS, UMR 8550 CNRS-ENS-Paris 6-Paris 7) Ecole Normale Supérieure, Paris

We analytically study the coding of a discrete set of categories (e.g. colors, phonemes) by a large assembly of neurons, each category being represented by input stimuli in a continuous space. One main outcome is to explain why and when one should expect more cells activated by stimuli close to the class boundaries, than for stimuli close to the prototypes of a class. For this analysis we consider population coding schemes, which can also be seen as instances of exemplar models proposed in the literature to account for phenomena in the psychophysics of categorization. We quantify the coding efficiency by the mutual information between the discrete categories and the neural code. From information theoretic bounds (Fano), we expect any learning mechanism that aims at minimizing the probability of error to maximize this mutual information. We characterize the properties of the most efficient codes in the limit of a large number of coding cells.

We show that in this limit, up to a constant this mutual information is given by (minus) the average over the input space of the ratio between two Fisher informations: at the denominator, the (usual) Fisher information $F_{\text{code}}(x)$, giving the sensibility of the population code to small changes in the input x ; and at the numerator, the category related Fisher information $F_{\text{cat}}(x)$, given in terms of the posterior probability of each class. This latter Fisher information can be written as the sum of the terms $P'(c|x)^2/P(c|x)$ over all the categories c . Typically the posterior probability $P(c|x)$ has a smooth S-shape, which entails that $|P'(c|x)|$ – and therefore $F_{\text{cat}}(x)$ – is the greatest in the regions near the boundary between categories. Given the limited resources of a large but finite number of neurons, maximizing the mutual information between the categories and the neural activity implies Fisher information F_{code} to be the greatest at the boundary between categories. The higher $F_{\text{code}}(x)$, the more discriminable two sensory inputs x and $x+dx$ in the perceptual space given by the output of the neuronal population. In other words, category learning implies better cross-category discrimination than within-category discrimination, a perceptual phenomenon traditionally called categorical perception. We consider the optimal configuration of a population of neurons in terms of the preferred stimuli of the cells and the widths of their receptive field. Our results predict that, if there is adaptation to a given categorization task, then (1) cells will be specifically tuned to the dimensions relevant for the classification task at hand; (2) more neurons will be allocated at the class boundaries, and (3) these boundary-specific cells will have a sharper tuning curve along each relevant dimension. All these predictions find support in recent neurophysiological experiments done in the inferotemporal cortex of the monkey brain, a cortex area shown to be specifically involved in classification tasks.

These results concern regimes of reasonably well-defined classes and sufficiently large signal-to-noise ratio. We have also studied ill-defined classes and noisy/short-time processing: in that case, the main result is that, contrary to the previous situation, the cells receptive fields will avoid the class boundaries.

References: L. Bonnasse-Gahot and J.-P. Nadal, "Neural Coding of Categories", preprint, May 2007[?]

Interplay between a phase response curve and an activity-dependent rewiring rule of neurons leads to wireless clustering

Hideyuki Cateau and Tomoki Fukai

RIKEN Brain Science Institute, Wako, Saitama 351-0198, Japan

Synchrony tendency of coupled oscillators or neurons is predicted by the phase response curve (PRC) of a neuron that describes an amount of advance or delay by synaptic input given at a specific phase

in an interspike interval. It is intriguing to know how this useful theory based on the fixed coupling strength between neurons generalizes to the cases where synaptic strength varies as observed in the real brain. A number of experiments have established that synaptic strength changes depending on pre- and postsynaptic spike times and theoretical implications of such spike-timing-dependent plasticity (STDP) have been extensively studied. Since the PRC and STDP both refer to the timings of spikes, a natural question is how these two properties of a neuronal network interact each other to carve a functional network in the brain.

To answer this question, we use a neuron model whose PRC can be systematically controlled (Izhikevich, 2004) unlike the simpler leaky integrate-and-fire (LIF) model. The model neurons favors either asynchronous (Model A) or synchronous (Model B) firing depending on the values of the model parameters. Our simulations show that STDP working on the network of Model A neurons converts an asynchronously firing neurons into three or more cyclically activated clusters of neurons. Interestingly, Model A neurons can synchronize within a cluster despite their preference to asynchrony because, as we see later, STDP selectively disrupts intra-cluster connections, nullifying the asynchrony preference. If STDP works on the network of Model B neurons, however, the neurons simply get synchronized globally, analogous to what was observed in (Karbowski and Ermentrout, 2002), and nothing peculiar happens.

Thus, the PRC influences the way STDP works. Importantly, STDP in turn changes the network structure and influences the way how PRC is readout to decide the network activity. Before the STDP learning begins, the initial slope of an effective PRC (defined later) determines the stability of the global synchrony. After the STDP learning forms cyclic activity of n clusters, the slope at $2\pi(1-1/n)$ determines its stability. In this way, the two key features of a network of spiking neurons, the PRC and STDP, work synergetically to organize functional networks in the brain.

Inferring Neural Firing Rates from Spike Trains Using Gaussian Processes

John P. Cunningham¹, Byron M. Yu¹, Maneesh Sahani², and Krishna V. Shenoy¹

¹Stanford University, Stanford, CA, ²Gatsby Computational Neuroscience Unit, University College London

Neural signals present challenges to analytical efforts due to their noisy, spiking nature. Many studies of neuroscientific and neural prosthetic importance rely on a smoothed, denoised neural signal considered to be the spike train's underlying firing rate. Current techniques to find time varying firing rates require ad hoc choices of parameters, offer no confidence intervals on their estimates, and can obscure potentially important single trial variability. We present a new method, based on a Gaussian Process prior, for inferring probabilistically optimal estimates of firing rate functions underlying single or multiple neural spike trains. We simulate spike trains to test the performance of the method and demonstrate significant average error improvement over standard smoothing techniques.

Studying the effects of noise correlations on population coding using a sampling method

Alexander S. Ecker¹, Philipp Berens¹, Matthias Bethge¹, Nikos K. Logothetis¹, and Andreas S. Tolias^{1,2}

¹ Max-Planck-Institute for Biological Cybernetics, Spemannstr. 38, 72076 Tübingen, Germany,

² Baylor College of Medicine, One Baylor Plaza, Houston, TX 77030, USA

Responses of single neurons to a fixed stimulus are usually both variable and highly ambiguous. Therefore, it is widely assumed that stimulus parameters are encoded by populations of neurons. An important aspect in population coding that has received much interest in the past is the effect of correlated noise on the accuracy of the neural code.

Theoretical studies have investigated the effects of different correlation structures on the amount of information that can be encoded by a population of neurons based on Fisher Information. Unfortunately, to be analytically tractable, these studies usually have to make certain simplifying assumptions such as high firing rates and Gaussian noise. Therefore, it remains open if these results also hold in the more

realistic scenario of low firing rates and discrete, Poisson-distributed spike counts.

In order to address this question we have developed a straightforward and efficient method to draw samples from a multivariate near-maximum entropy Poisson distribution with arbitrary mean and covariance matrix based on the dichotomized Gaussian distribution [1]. The ability to extensively sample data from this class of distributions enables us to study the effects of different types of correlation structures and tuning functions on the information encoded by populations of neurons under more realistic assumptions than analytically tractable methods.

Specifically, we studied how limited range correlations (neurons with similar tuning functions and low spatial distance are more correlated than others) affect the accuracy of a downstream decoder compared to uniform correlations (correlations between neurons are independent of their properties and locations). Using a set of neurons with equally spaced orientation tuning functions, we computed the error of an optimal linear estimator (OLE) reconstructing stimulus orientation from the neurons firing rates. We find supporting previous theoretical results that irrespective of tuning width and the number of neurons in the network, limited range correlations decrease decoding accuracy while uniform correlations facilitate accurate decoding. The optimal tuning width, however, did not change as a function of either the correlation structure or the number of neurons in the network. These results are particularly interesting since a number of experimental studies report limited range correlation structures (starting at around 0.1 to 0.2 for similar neurons) while experiments carried out in our own lab suggest that correlations are generally low (on the order of 0.01) and uniform.

[1] P. Berens and M. Bethge, *NCCD* (2007).

Attention improves object encoding in monkey area V4

D. Rotermund¹; K. Taylor²; K.R. Pawelzik¹; A.K. Kreiter²; U.A. Ernst³

¹Institute for Theoretical Physics, University of Bremen, Bremen, Germany. ²Brain Research Inst., University of Bremen, Bremen, Germany. ³Group for Neural Theory, Ecole Normale Supérieure and Collège de France, Paris, France.

Attention is thought to improve processing of selected stimuli. It has been demonstrated that firing rates can be modulated by attention and that their enhancement can improve the signal-to-noise ratio for discriminating stimulus features. However, from these often small effects observed at the single neuron level, it is not clear how to explain the massive enhancements in perceptual performance known from psychophysical studies.

Combining multielectrode recordings of field potentials with current methods from computational neuroscience, we investigate effects of selective attention on object representation using recordings of epidural field potentials (EFPs). We trained two monkeys to attend to one of two sequences of objects, which were simultaneously presented in both hemifields. The task required the monkeys to identify the re-occurrence of the initial object in the attended hemifield. LFPs were recorded with an array of 37 epidural electrodes, covering parts of area V4 and V1, while the monkeys were performing the task.

Analysis of object encoding was done with standard support vector machines using EFP wavelet power coefficients, allowing to estimate the probability of correct classification of the objects presented in dependence on the attentional state. Classifying on all electrodes, we found a performance of up to 94% correct (1200 ms window, chance level 17%), i.e. the EFP-data enabled nearly perfect object identification. Almost all stimulus-specific information was concentrated in the frequency range from 50 to 100 Hz. Classification on data from 5 electrodes covering V4 resulted in 52% performance for non-attended stimuli, and 63% for attended stimuli. We also classified the direction of attention from the EFP-signal, which reached a performance of up to 91% for all electrodes, and of up to 83% for the 5 electrodes from V4 (chance level 50%).

Our results clearly show that attention substantially improves the encoding of visual stimuli in gamma-band neuronal activity. Using artificial data sets derived from the recorded EFP activities, we investigated which specific features in the data were responsible for the attentional gain in performance. In particular, we first scaled independently either the mean values or the variances of the data in the non-attended condition to the corresponding numbers observed in the data for the attended condition. Then we quantified the resulting gain in performance by classifying with support vector machines trained on these new,

artificial data sets. The analysis revealed that the small improvements found in the signal-to-noise ratios under attention contribute only to a small extent to the large improvements in performance. The major part of the attentional enhancement of the neural representation can rather be explained by stimulus- and electrode-specific shifts in the mean wavelet amplitudes, which render the neural activity pattern more distinct from each other. This result uncovers a novel mechanism for attentional enhancement, which could only be discovered by using multichannel recordings combined with a thorough data analysis guided by specific hypotheses about principles of neural coding.

Support contributed by: DIP METACOMP, BMBF Bernstein Group Bremen, and EU Grant BIND MECT-CT-20095-024831 and BACS FP6-IST-027140.

A method for recording neural activity simultaneously from independently moveable, chronically implanted electrodes in cortical and sub-cortical areas of monkeys.

Joseph Feingold^{1*}, Naotaka Fujii^{2*}, Theresa M. Desrochers-Feledy³, Ray Harlan⁴, Ann M. Graybiel³
¹HST, MIBR, MIT, Cambridge, MA, USA; ²BSI, RIKEN, Saitama, Japan; ³BCS, MIBR, MIT, Cambridge, MA, USA; ⁴Specialty Machining, Wayland, MA USA

The ability to record neural signals simultaneously from many sites with chronically implanted electrodes represents a difficult technical challenge which, if surmounted, could greatly enhance our understanding of the activity of networks of neurons, and particularly of learning-related changes in them. Most multi-electrode implants for chronic neural recordings use arrays of electrodes rigidly fixed in place for the duration of the implant. These arrays are implanted surgically and have so far been used mainly in cortical areas. We have developed an alternative method using custom-designed screw micro-drives and recording chambers along with conventional microelectrodes to record from head-fixed or unrestrained monkeys. Using a custom-built grid with precisely machined holes (1mm center-to-center), we have been able to record single-unit and LFP activity from cortical and sub-cortical structures accessible through a large (~1600mm²) craniotomy over the frontal or parietal lobes. Recordings can be obtained from multiple structures with a minimal track spacing of 1-2mm. By adjusting the depth of each electrode (upward or downward) over the course of an implant, we can improve the isolation of units and record from new units at multiple depths along each track. Implants can be removed and reconfigured to target different regions or different tracks within the same region. We have used the method to implant repeatedly a total of 11 monkeys. We have recorded from as many as 126 electrodes implanted bilaterally, and for up to 9 months at a time in subsets of PFC, M1, SMA/pre-SMA, CMA, ACC, PM, FEF, SEF, parietal cortex, amygdala, nucleus accumbens, caudate nucleus, putamen, thalamus and globus pallidus. Following initial stabilization of the implants (< 3 weeks), we typically recorded well-isolated single units from 30-60% of electrodes in each session. Different subsets of electrodes had units in different sessions, but some electrodes had extraordinarily stable unit activity across weeks of recordings. In a typical implant, every electrode had single unit activity in at least one session. We are currently exploring new techniques to improve the daily yield. Our main effort is geared toward reducing the presumed gliosis associated with chronic implants. Our method has a number of advantages in this respect. First, the relative sparseness of our implants reduces the risk posed to neighboring electrodes from a localized inflammatory response. Second, once implanted, individual electrodes can be moved in a controlled and gradual manner, in order to improve stability and to minimize tissue damage around the tips. Third, if necessary, implants can be extracted and the monkey can be re-implanted with fresh electrodes in a matter of weeks. Fourth, the method can be adapted to work with many types of electrodes, enabling flexibility in battling gliosis. Ultimately, our method facilitates the search for changes within and among brain regions across multiple recording sessions, and promises to be a powerful tool for studying the neural basis of learning.

This work supported by: NEI, NINDS (Javits), NPF.

Repeating spatiotemporal spike patterns reflect functional network states in the visual cortex

Kai Gansel & Wolf Singer

Max-Planck-Institute for Brain Research, Frankfurt am Main

Do the neuromodulatory setting and the level of excitation inside a cortical network generate distinctive states of functionality, and can these states be characterized by the emerging spatiotemporal patterns of multineuronal suprathreshold activity?

To approach these issues, we analyzed single unit spiking activity from acute slices of rat visual cortex under two pharmacological conditions, namely without and with activation of cholinergic receptors using Carbachol. Recordings were performed with a 32-channel extracellular electrode array ("Utah probe"), yielding on average about 90 isolated single units in parallel. A second electrode matrix with flat electrodes was used to apply short current pulses to transiently raise the population's spike rate. Spatiotemporal firing sequences were detected using sliding windows with lengths of 5 to 50 ms and checked for significance by common Monte Carlo methods. In this way, we were able to collect thousands of statistically significant spatiotemporal spike patterns from ten slices, repeating with different timing precisions of up to 0.5 ms. The question now was, if the mere identities of the patterns - taken as indicators of the functional state of the network - as well as their qualities like complexity, spatial dimensions and variance change systematically depending on the experimental condition.

It turned out that both pattern complexities (in terms of participating units) and the average spatial distance between pattern-forming peers did not change significantly following pharmacological treatment or electrical stimulation of the neuronal tissue. The repeats of individual patterns were also constant across pharmacological conditions but increased two-fold during electrical stimulation, indicating a decreased variance among patterns. However, most importantly, about 77% of all patterns selectively occurred during a certain experimental condition (chance level: 51 ± 24 SD), showing an average specificity of about 90% (chance level: 78 ± 4.4 SD). About 23% of all patterns also selectively recurred and hence may be viewed as the "fingerprints" of that particular condition. Likewise, the functional connectivity maps, as derived from the patterns, display a clear dependence on the experimental condition, although unsystematic changes are present as well.

As a bottom line, changing the neuronal dynamics in a cortical network by cholinergic modulation or electrically evoked activity produces systematic changes in the network's functional connectivity, leading to the emergence of selectively recurring spatiotemporal spike patterns. Concomitant unsystematic changes, provoking non-selective or non-recurring spike sequences, may well be interpreted as the result of ongoing plasticity.

Evaluating information transfer between neurons: a different view on causality and integration within neuron assemblies

Boris Gourevitch and Jos J. Eggermont

Departments of Physiology and Biophysics, and Psychology, University of Calgary, Alberta, Canada

Neurons assemblies are often considered as simultaneous activations of a set of neurons leading to a certain representation or even a "coding" of the incoming information. Thus, cross-correlation between pairs or neurons has remained the most widely used tool to investigate such assemblies for decades. However, while the number of simultaneous recordings tends to increase, there is obviously no chance that multiple spike trains fully represent even one assembly. We only have access to a very subsampled representation of the brain states, with no clear knowledge of the physiological links between the neurons we're recording. In this case, notions of causality, through neurons driving other neurons, or analysis of synchrony or integration between neurons are not well defined.

With this in prospect, we have developed the transfer entropy, presented as a new tool for investigating neural assemblies. The transfer entropy quantifies the fraction of information (in a Shannon's sense) in a neuron found in the past history of another neuron. The maximization of the transfer entropy over several intervals defining the past history of a neuron gives insights about the memory span involved

in information transmission between neurons. Consequently, this tool has potential applications in investigating windows of temporal integration and stimulus-induced modulation of firing rate. The asymmetry of the measure allows feedback evaluations. Transfer entropy is also able to eliminate some effects of common history in spike trains and obtains results that are different from cross-correlation. Moreover, the measure of the information transferred from one neuron to another avoids the difficulty of giving a precise definition of what is causality between neurons.

The basic transfer entropy properties are illustrated with simulations. The information transfer through a network of 16 simultaneous multiunit recordings in cat's auditory cortex was examined for a large number of acoustic stimulus types. Application of the transfer entropy to a large database of multiple single-unit activity in cat's primary auditory cortex revealed that most windows of temporal integration found during spontaneous activity range between 2 and 15 ms. The normalized transfer entropy shows similarities and differences with the strength of cross-correlation; these form the basis for revisiting the neural assembly concept.

Decoding responses from populations of neurons in macaque primary visual cortex

Arnulf B. A. Graf, Adam Kohn, Mehrdad Jazayeri and J. Anthony Movshon
Center for Neural Science, New York University

One of the central questions in computational and systems neuroscience is how the responses from a population of neurons are combined to decode sensory signals. To address this question, we recorded from the superficial layers of primary visual cortex of anesthetized, paralyzed macaque monkeys using a 10 x 10 electrode array with 400 microns spacing between adjacent electrodes. We sorted the recorded waveforms offline to isolate between 44 and 86 simultaneously recorded orientation-selective units. We presented sinusoidal gratings large enough to cover the receptive fields of all recorded units, drifting in 72 directions 5 deg apart. Each grating was presented 50 times for 1280 ms in randomly-sequenced blocks, and was followed by a 1280 ms blank screen.

We considered linear decoders that discriminated between the population activity evoked by pairs of stimuli by computing the weighted sum of the neuronal responses (spike counts). First, we used a likelihood ratio (LR) decoder which compared the likelihood of two stimuli from the underlying response distributions. Our LR decoder assumed that these responses were statistically independent, and that Poisson statistics adequately described the spike count distributions. The LR decoder was linear, and its weights were given by the difference of the logarithms of the tuning curves. Second, we used a linear Support Vector Machine (SVM) which estimated a separating hyperplane in a space whose dimensions corresponded to the neural responses. This hyperplane maximized the distance between classes of neural response and minimized misclassified responses. Unlike the LR decoder, the SVM made no *a priori* assumptions about the data, and could take advantage of all the structure present in the empirical data.

Using the discrimination accuracy of the decoders as a measure of decoding performance, we constructed population neurometric functions to measure how well pools of neurons could discriminate between different stimulus directions. We explored the impact of correlations on decoding by comparing raw with shuffled population responses (interneuronal correlations removed by trial shuffling). The LR decoder performed better on the shuffled data than on the raw, because it assumed independence between neurons. The SVM decoder, however, performed better on the raw data, showing that it could make use of correlations to improve its decoding performance. This effect was present for different population subsamples and integration times. The linearity of the two decoders allowed us to get a mechanistic insight into decoding. For the SVM decoder, the neurons responding most vigorously across all directions were weighted more heavily than the least responsive neurons. For fine discrimination both decoders emphasized units with preferred directions that were further apart than the stimulus direction difference, a behavior predicted from theory and inferred previously in psychophysics. Furthermore, for the SVM decoder, the increase of decoding accuracy in the presence of correlations was accompanied by a sharpening of the shape of the weight vector.

The LR decoder can be constructed directly from neuronal tuning curves and is appealing because of its simplicity and biological feasibility. However, our results provide evidence that more sophisticated

linear classifiers such as the SVM, while retaining much of the simplicity of the LR decoder, can take advantage of neural correlations to decode the activity of neuronal populations with greater precision.

Extracting synfire correlation structure from massively parallel recordings

Sven Schrader³, Sonja Gruen^{1,2,*}, Markus Diesmann^{1,3} and George Gerstein⁴

¹Computational Neuroscience Group, RIKEN Brain Science Institute, Wako, Japan, ²Bernstein Center for Computational Neuroscience, Berlin, Germany, ³Bernstein Center for Computational Neuroscience, Albert-Ludwigs-University, Freiburg, Germany, ⁴Dept Neuroscience, University of Pennsylvania, Philadelphia, USA

The synfire chain was proposed to constitute the elementary building block of the substrate of information processing in the cortex [1,2]. If such a subnetwork is stimulated appropriately, synchronized spiking activity propagates through consecutive neuron groups. Despite intense research over the past 20 years or so, statistical tools to reliably and unambiguously detect synfire activity are still not available. Existing methods cannot cope with massively parallel spike trains (say 100 neurons or more) or the analysis is carried out for individual spatial or spatio-temporal spike patterns [e.g. 3,4] not respecting the variability expected in synfire activity. Here we present a new approach that overcomes both problems and allows to directly detect and visualize the signature of synfire activity in large but finite data sets. The analysis of random parallel recordings from large-scale neural network simulations with embedded synfire chains demonstrates sensitivity and specificity.

The first step in the analysis is to discretize time with a bin size h appropriate for the expected precision of spike synchronization, e.g. $h=2ms$. Each bin is identified by an integer i corresponding to the left border of the time interval. Let $S(i)$ be the subset of all recorded neurons active in bin i . We now compare network activity at discretized times i and j by computing the number of elements in the intersection of $S(i)$ and $S(j)$, i.e. the number of neurons active at both instances of time. This value is normalized to the range between 0 and 1 by dividing by the number of elements in the set with fewer elements. The two time axes i and j span a matrix $M(i,j)$ of the new measure. Clearly $M(i,i) = 1$ and $M(j,i) = M(i,j)$, therefore it suffices to study only the lower triangular matrix $i \geq j$.

The signature of a synfire chain that fires twice is as a short diagonal stripe of large entries. The stripe starts at a particular matrix element $M(i,j)$ corresponding to the start times of the synfire chain. The overlap $M(i,j)$ is large because the set $S(i)$ includes all active members of the first neuron group in the chain, and set $S(j)$ includes the active neurons of the group in the second run. The two sets corresponding to $M(i+d,j+d)$ include the active neurons of the second group and so on, where d depends on the synaptic delay. If a particular chain is activated at discretized times k_1 to k_n , $n(n-1)/2$ stripes are found originating at matrix elements $(k_1,k_2), (k_1,k_3), \dots, (k_1,k_n); (k_2,k_3), \dots, (k_2,k_n); \dots, (k_{n-1},k_n)$. Finally, if several chains are active, each produces its set of diagonal stripes. Synfire activity is detectable if the overlap activity is differentiable from chance overlap. The temporal aspect of the synfire signature is taken into account by filtering the matrix with an appropriate diagonal filter. This improves visualization and considerably enhances sensitivity. The amplitude distribution of the matrix elements allows to define a threshold separating signal from noise.

We calibrate the method by recordings from a balanced neural network model with (50.000 I&F neurons, 80% excitatory, 20% inhibitory, connectivity: 0.1). In contrast to the standard wiring scheme all excitatory to excitatory synapses are formed by synfire chains. The network embeds 50 chains composed of 20 consecutive groups each containing 100 neurons. The chains are activated by pulse packets, at random times and low rate. Including all 50.000 neurons into the analysis reveals clear signatures of the synfire runs. However, even when the number of recorded neurons is reduced to 200, synfire activity is reliably detected.

This level of sensitivity should allow experimental testing (e.g. by optically recorded data) in the fairly near future and enable either verification or rejection of the synfire idea.

[1] M Abeles (1991). *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press.

[2] E Bienenstock (1995) *Network: Comput Neural Systems* **6**, 179-224.

[3] M Abeles and G L Gerstein (1988) *J Neurophysiol* **60(3)**: 909-924.

[4] S Grun, M Diesmann, and A Aertsen (2002) *Neural Comput* **14**(1), 81-119.

Stable, highly irregular persistent activity induced by colored noise

Kosuke Hamaguchi^{1,2} and Nicolas Brunel²

¹RIKEN BSI, Japan, ²UMR 8119 CNRS-Universite Rene Descartes, France

The spike patterns of prefrontal cortex neurons are highly irregular (CV close to 1) during the delay period of a spatial working memory task[1]. During this task, some neurons switch from a background state to an active firing state when spatial cue is presented and maintain their active states during delay period [1]. Thus, the prefrontal neuron network has multi-stable states with high spike irregularity.

One robust mechanism to obtain high spike irregularity is to balance excitation and inhibition in recurrent connections [2]. However, self-sustained activity typically relies on strong recurrent excitation. In an excitation-dominated network, CV of spike trains typically decreases as the firing rate increases. Therefore, it is unclear how to obtain both high CVs and bistability.

Here we propose a model that can reconcile this apparent contradictions. Recent analysis of sparsely connected recurrent network with spiking neurons [3] indicates the existence of fluctuation-driven bistability. More recent study shows that it requires relatively wide distributed delay of synapses to stabilize the fluctuation-driven bistable states [4] if the synapses are instantaneous (delta-function). We found that the instability occurs due to a Hopf-bifurcation. Here we use exponential-decay synapses with realistic time constants to alleviate the oscillatory instability. The system is now driven by colored noise. We calculate the firing rate of colored noise driven Leaky Integrate-and-Fire (LIF) neuron by using numerical calculations of 2D Fokker-Planck equation and analytically calculated firing rate with first order correction in small synaptic time constant limit [5]. We confirm that the fluctuation-driven bistability also exists in colored noise driven LIF neuron network both from the self-consistent analysis and simulations. These results indicate that prefrontal cortex would be operating in the regime of balanced/inhibition-dominated network.

Acknowledgments: K.H. is supported by JSPS Research Fellowship.

[1] Compte A, Constantinidis C, Raghavachari S Tegner J, Chafee MV, Goldman-Rakic PS, and Wang X-J., Temporally irregular mnemonic persistent activity in prefrontal neurons of monkeys during a delayed response task. *J. Neurophysiol.* **90**:3441-3454 (2003).

[2] C van Vreeswijk and H Sompolinsky., Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **74**(5293):1724-1726 (1996).

[3] Brunel N., Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J. Comp. Neurosci.* **8**(3):183-208 (2000).

[4] Renart A, Moreno-Bote R, Wang X-J, and Parga N., Mean-Driven and fluctuation-driven persistent activity in recurrent networks, *Neural Computation* **19**: 1-46, (2007).

[5] Brunel N and Sergi S. Firing frequency of Leaky Integrate-and-Fire neurons with synaptic current dynamics, *J. Theor. Biol.* **195**:87-95 (1998).

Laplace's method in neural decoding

Shinsuke Koyama^{1,2}, Lucia Castellanos Perez-Bolde³, Cosma Rohilla Shalizi¹ and Robert E. Kass^{1,2}

¹Department of Statistics, ²Center for the Neural Basis of Cognition, ³Department of Machine Learning, Carnegie Mellon University, USA

State-space models are a promising technique for neural decoding, especially in domains like neural prostheses where the signal to be reconstructed has significant temporal structure. The optimal estimate of the state is its conditional expectation given the observed spike-train histories, but taking this expectation is computationally hard, especially when nonlinearities are present. Existing filtering methods, including sequential Monte Carlo, tend to be either inaccurate or slow. In this paper, we propose a new nonlinear filter which uses Laplace's method, an asymptotic series expansion, to approximate the

conditional mean and variance, and a Gaussian approximation to the conditional distribution of the state. This "Laplace-Gaussian filter (LGF)" gives fast, recursive, deterministic state estimates, with an error which is set by the stochastic characteristics of the model and is, we show, stable over time. We illustrate the decoding ability of the LGF by applying it to a simulation of the cortical control of hand motion, where it delivers superior results to sequential Monte Carlo in a fraction of the time.

High frequency organization and synchrony of activity in the Purkinje cell layer of the cerebellum

Clement Lena¹, Camille de Solages¹, German Szapiro¹, Nicolas Brunel², Vincent Hakim³, Philippe Isope¹, Charly Rousseau¹ and Boris Barbour¹

¹Laboratoire de Neurobiologie UMR 8544, Ecole Normale Supérieure, France, ²Laboratoire de Neurophysique et Physiologie, UMR 8119, Université René Descartes, France, ³Laboratoire de Physique Statistique UMR 8550, Ecole Normale Supérieure, France

The cerebellum controls complex, coordinated and rapid movements, a function requiring precise timing abilities. However, the network mechanisms that underlie the temporal organization of activity in the cerebellum are largely unexplored, because in vivo recordings have usually targeted single units. Here, we use tetrode and multi-site recordings to demonstrate that Purkinje cell activity is synchronized by a high-frequency (~200 Hz) population oscillation. We combine pharmacological experiments and modeling to show how the recurrent inhibitory connections between Purkinje cells are sufficient to generate these oscillations. A key feature of these oscillations is a fixed population frequency that is independent of the firing rates of the individual cells. Convergence in the deep cerebellar nuclei of Purkinje cell activity, synchronized by these oscillations, likely organizes temporally the cerebellar output.

Back to the future: episodic memories for control

Máté Lengyel¹ and Peter Dayan²

¹Collegium Budapest Institute for Advanced Study, ²Gatsby Computational Neuroscience Unit, UCL

Episodic (or episodic-like) memories have traditionally been discussed in terms of remembering specific events in one's personal *past* [1]. However, there is an increasing realisation that remembering the past is mainly important for planning for the *future* [2]. Indeed, for example, experimental evidence shows that amnesic patients have difficulties with imagining future events [3]. However, from a computational standpoint, *semantic*, rather than *episodic*, memories seem to be a much better substrate of future planning: why should it be better to act on the basis of the apparently statistically inefficient recollection of single happenings (episodic, memory-based control), rather than the seemingly normative use of accumulated experience from multiple events (semantic, model-based control)?

We suggest that the key issue is computational noise. Particularly in interestingly complex environments, semantic control, which depends on noisy recursive operations to calculate the values of actions, can readily be bested by episodic control, which does not, particularly in the face of very limited experience. Processing noise overwhelms the statistical inefficiency of basing decisions on individual episodes rather than their integrated statistics. We show this using detailed mathematical analyses and numerical simulations in the normative framework of reinforcement learning [4]. Our analysis parallels earlier work [5], which showed that, again because of computational noise, semantic, model-based control, is less accurate than habitual control after more substantial amounts of experience.

Besides giving a normative benefit for episodic memories, our analysis also allows us to interpret experimental data on systems-level consolidation (e.g. [6]) as the hallmark of the transfer of *control*, rather than the transfer of *memories*, from the hippocampus to the neocortex, or the striatum [7]. Furthermore, in the light of these results, the fact that neural networks tend to forget past episodes in favour of recent ones may seem less catastrophic than often mooted [8].

[1] E. Tulving. Episodic and semantic memory. In E. Tulving and W. Donaldson, eds, *Organisation of Memory*, pp. 381-403. Academic Press, 1972.

[2] Y. Dudai and M. Carruthers. The Janus face of Mnemosyne. *Nature* **434**:567, 2005.

- [3] D. Hassabis, D. Kumaran, S.D. Vann, and E.A. Maguire. Patients with hippocampal amnesia cannot imagine new experiences. *Proc. Natl. Acad. Sci U.S.A.* **104**:1726-1731, 2007.
- [4] R.S. Sutton and A.G. Barto. *Reinforcement Learning*. MIT Press, 1998.
- [5] N.D. Daw, Y. Niv, and P. Dayan. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* **8**:1704-1711, 2005.
- [6] D. Tse, R.F. Langston, M. Kakeyama, I. Bethus, P.A. Spooner, E.R. Wood, M.P. Witter, and R.G.M. Morris. Schemas and memory consolidation. *Science* **316**:76-82, 2007.
- [7] M.G. Packard and J.L. McGaugh. Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: further evidence for multiple memory systems. *Behav. Neurosci.* **106**:439-446, 1992.
- [8] S. Fusi and L.F. Abbott. Limits on the memory storage capacity of bounded synapses. *Nat. Neurosci.* **10**:485-493, 2007.

Decomposing neural circuits using model-based analysis: Applying the sSoTS model to fMRI data.

Eirini Mavritsaki, Harriet Allen and Glyn Humphreys

Behavioural Brain Sciences Centre, School of Psychology, University of Birmingham, B15 2TT

We show how a biological plausible model can be used to decompose complex neural circuits found in fMRI studies of human attention, separating circuits concerned with inhibition attention function from those concerned with top-down enhancement. The studies examined preview search as presented by Watson and Humphreys [1].

The model for visual search over time and space (SSoTS) used incorporates different synaptic components (NMDA, AMPA, GABA) and a frequency adaptation mechanism based on [Ca²⁺] sensitive K⁺ current. This frequency adaptation current can act as a mechanism that suppresses the previously attended items. It has been shown [2] that when the passive process (frequency adaptation) is coupled with a process of active inhibition, new items can be successfully prioritised over time periods matching those found in psychological studies.

We use the model to examine the time course of the fMRI BOLD signal from studies of preview search. Activity in the model is related to different brain regions (i) by convolving the synaptic activation with the haemodynamic response function as formulated by Glover [3] and comparing the model's BOLD response with fMRI data activation from predefined regions of interest (i.e. inferior parietal lobe and precuneus) and also (ii) we use the synaptic activation from the sSoTS's maps as regressors for brain activity using standard imaging analysis techniques (FSL). We find activation related to SsoTS processes in discrete brain areas including a network of areas relating to active inhibition

[1] D. Watson and G.W. Humphreys *Psychological Review*. **104**:90-122 (1997).

[2] E. Mavritsaki, D. Heinke, G.W. Humphreys and G. Deco, *Journal of Physiology*. **100**:110-124 (2006).

[3] G.H. Glover, *NeuroImage*. **9**:416-429 (1999).

Dynamics of spontaneous activity in neuronal networks

Alberto Mazzoni¹, Frederic D. Broccard², Elizabeth Garcia-Perez³, Paolo Bonifazi⁴, Maria Elisabetta Ruaro², and Vincent Torre²

¹Institute for Scientific Interchange, Turin, Italy, ²International School for Advanced Studies, Trieste, Italy, ³Centro de Investigacion Medica Aplicada, Pamplona, Spain, ⁴Institut de Neurobiologie de la Mediterranee, Marseille, France

Isolated neuronal networks still display spontaneous activity. This activity is not uncorrelated noise but is structured into bursts of concerted firing alternated with silent intervals. We analyzed the dynamics of spontaneous bursts in two very different neuronal networks: intact leech ganglia and dissociated cultures of rat hippocampal neurons. The following features were common to the two networks:

- i) single neurons displayed three classes of spontaneous firing: Poisson, periodic and bursting.
- ii) the power spectrum of the firing rate was 1/f at low frequencies, indicating the presence of long-term

correlations.

iii) bursts size and duration had power law distributions with no characteristic burst scale. Power laws were followed by an exponential breakdown depending on the network size.

Furthermore, we could switch the two networks from this "critical" state to states dominated by very small or very large events with the same pharmaceutical modulations. We concluded that diverse neuronal networks display similar spontaneous activity features independently from their architecture. [1]

In the leech, we analyzed then the spontaneous activity of the ganglion when it was connected to the rest of the nervous system. Stimuli from neighboring ganglia and from the two "brains" increased the correlation among the neurons and drove the ganglion to a state similar to the one associated to decreased inhibition, in which bursts have a characteristic size. Finally, we studied how the spontaneous bursts determine the onset of the spontaneous movements in the semi-intact preparation. From the previous results it was possible to determine the relation between the threshold duration and firing rate necessary for a burst to trigger a movement.

[1] Mazzoni A, Broccard FD, Garcia-Perez E, Bonifazi P, Ruaro ME, et al. (2007) On the Dynamics of the Spontaneous Activity in Neuronal Networks. PLoS ONE 2(5): e439. doi:10.1371/journal.pone.0000439

STDP results from neural dynamics, and not from a specific optimality criterion

Jeffrey M. Beck¹ and Lucas C. Parra²

¹University of Rochester, ²City University of New York

Since the discovery of spike-time dependent plasticity (STDP) a variety of descriptive and normative explanations have been offered for its origin. Here, we show that STDP learning rules result from the stochastic dynamics of individual neurons and not from any specific optimality criterion. Specifically, we show that each iteration of a gradient descent learning rule can be written as a local update rule which takes the form of traditional STDP but with a magnitude given by reward, i.e. network weights are updated according to the covariance between reward and the STDP function. As a result, the overall shape of the STDP curve is not affected by the specific choice of reward function. Moreover, a global reward signal which modulates learning rate is sufficient to optimize a wide class of objective functions.

This result is obtained for the standard linear-nonlinear Poisson model, wherein presynaptic spikes are linearly filtered in space and time to generate a "membrane potential". The likelihood of firing at any moment in time is then a non-linear function of the instantaneous membrane potential, i.e. spiking follows an inhomogeneous Poisson process. This model readily implements refractory behavior by treating after-hyperpolarization as a linear response to the neuron's own spikes and is an excellent generalization of integrate-and-fire dynamics. Moreover, this model is analytically tractable. Indeed, it has been shown (Paninski) that for a given input/output spike pattern a maximum likelihood learning rule will converge to a unique set of network parameters (under certain conditions on the non-linearity). In our formalism this corresponds to supervised learning with the log likelihood as reward function. Thus, STDP can learn to maximize the likelihood of observing a specific input/output spike pattern relationship.

We then explore a number of possible reward functions by means of numerical simulations. These reward functions are motivated by concepts such as predictive encoding, independent component analysis, minimum quantization error, optimal linear decoding, mutual information maximization, description length minimization, and energy efficiency.

The effect of single neuron properties and input diversity on the contrast response function in V1

E. Persi, D. Hansel, L. Nowak, C. van Vreeswijk

Laboratoire de Neurophysique et Physiologie, UMR 8119 CNRS, Paris 75270, FRANCE

The response (R) of V1 neurons to elongated stimuli at the preferred orientation varies sigmoidally as a function of contrast (C). This contrast response function (CRF) accelerates at low C and saturates

at high C . It is characterized by the maximum firing rate (R_{\max}), the steepness and the contrast at which R equals half R_{\max} ($C50$). Experiments reveal a large diversity in these parameters. What are the respective contributions of the feedforward input, the recurrent intra-cortical input and the intrinsic properties of the neurons in shaping the CRF? To answer this question we have developed a rate model of a hypercolumn in V1, in which we assumed power-law input-output transfer-function (TF) for the neurons. Assuming that the exponent of the TF is larger than 1, the width of the orientation tuning curve (OT) of the output is narrower than the OT width of the input, and it is approximately contrast-invariance (CI). The response varies sigmoidally with C , the acceleration at low C reflects the TF non-linearity, whereas the saturation is strongly affected by the intra-cortical inhibition and by the saturating properties of the input. We have found similar behavior while testing the model with conductance-based neurons, in which power-law TFs emerge naturally as a result of the presence of noise in the inputs. A great deal of the diversity observed in the shape of the CRF can be accounted for in our model by assuming a large diversity in the input and cellular properties. Based on these findings we predict a correlation between the steepness and $C50$ of the CRF of neurons in V1.

Time course of reactivation of memory-related cell ensembles in the rat medial prefrontal cortex during sleep.

A. Peyrache¹, K. Benchenane¹, M. Khamassi^{1,2}, V. Douchamps¹, P. L. Tierney³, S. I. Wiener¹, F. P. Battaglia⁴,

¹LPPA, CNRS College de France, Paris, France, ²Inst. des Systemes Intelligents et Robotiques, Univ. Pierre et Marie Curie, Paris, France, ³Inserm U667, College de France, Paris, France, ⁴Univ. van Amsterdam, Amsterdam, The Netherlands

The prefrontal cortex is implicated in the flexible learning of stimulus-outcome associations, which are consolidated in memory during offline periods. Reactivation of memory traces, in the form of the reinstatement of experience-related activity in prefrontal cell assemblies during sleep could be the basis for such a consolidation process. To study this, we developed a novel analysis which allows to follow the time course of task-related reactivation in simultaneously recorded cell ensembles. The correlation matrix of binned spike trains from multiple cells is decomposed in its principal components, the largest of which represents groups of cells whose activity was highly correlated during the reference recording period. The instantaneous cell pair co-activation matrix during sleep, weighted by the coefficients in a given principal component, and averaged over all cell pairs, can then be taken as a measure of the reactivation of the cell assembly corresponding to that principal component at a given time.

We analyzed medial prefrontal ensembles from five rats while learning a set-switching task on a Y-maze, and in rest sessions preceding and following the task. In 62 out of 86 sessions, cell assembly reactivation was significantly greater ($p < 0.05$) during slow wave sleep (SWS) after the session than in SWS before. There was a significant correlation (Pearson's correlation test, $p < 0.05$) between the eigenvalues associated with the principal components during task performance (indicating the strength of the encoding) and the increased re-activation in post-task SWS (compared to pre-task SWS). Moreover, in 67 out of 86 sessions, co-activation was correlated ($p < 0.05$) with the power of both delta and spindle cortical oscillations, and it was much weaker during rest periods that were classified as non-sleep. The increased co-activation in the post-experience sleep was attributable to discrete bouts of activation, typically 2-5 seconds in duration. This new technique permits to precisely follow the time course of neural ensemble re-activation. These data demonstrate that theta, ripple-sharp waves and spindles are important for prefrontal post-task SWS reactivation, a possible neural ensemble basis for memory consolidation.

Sequential non-parametric change-point detection for time series models: assessing the functional dynamics of neuronal networks

Fabio Rigat

CRiSM, Department of Statistics, University of Warwick, UK

This work illustrates the application of a novel sequential non-parametric method to detect significant changes of the functional connectivity of neuronal networks using in-vivo multiple spike trains recordings

obtained via multi-electrode arrays. Rather than relying on a parametric specification of the parameters' evolution, their dynamics are assessed as a change-point problem in discrete time within a hypothesis testing framework. The Kullback-Leibler divergence between the posterior distributions of different sets of data under the same model is proposed as a test statistic. Markov chain Monte Carlo posterior simulation is used in general to approximate the value of the Kullback-Leibler statistic and its critical region under the null hypothesis of no change. The changes in the network parameters detected by the Kullback-Leibler statistic explain variations of the baseline neuronal spiking rates as well as of the functional connections between neurons across different experiments.

Spike train decoding without spike sorting

Valerie Ventura

Statistics Department, Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh

We propose a novel paradigm for spike train decoding, which avoids entirely spike sorting based on waveform measurements. This paradigm directly uses the spike train collected at recording electrodes from thresholding the bandpassed voltage signal. Our approach is a paradigm, not an algorithm, since it can be used with any of the current decoding algorithms, such as population vector or likelihood based algorithms. Based on analytical results and an extensive simulation study, we show that our paradigm is comparable to, and sometimes more efficient than, the traditional approach based on well isolated neurons, and that it remains efficient even when all electrodes are severely corrupted by noise, a situation that would render spike sorting particularly difficult. Our paradigm will also save time and computational effort, both of which are crucially important for successful operation of real-time brain-machine interfaces. Indeed, in place of the lengthy spike sorting task of the traditional approach, it involves an exact expectation EM algorithm that is fast enough that it could also be left to run during decoding to capture potential slow changes in the states of the neurons.

V. Ventura *Neural Computation* To appear (2007).

Neuronal synchronization underlies effective inter-areal interactions in visual cortex

Thilo Womelsdorf¹, Jan-Mathijs Schoffelen¹, Robert Oostenveld¹, Andreas K. Engel², Wolf Singer³, and Pascal Fries¹,

¹F.C. Donders Centre for Cognitive Neuroimaging, Radboud University Nijmegen, The Netherlands, ²Dept. of Neurophys. and Pathophys., UMC Hamburg-Eppendorf, Germany, ³Max-Planck-Institute for Brain Research, Frankfurt, Germany

Our cognitive abilities rely on dynamic interactions of neuronal groups distributed across distant cortical areas. How the efficiency of these interactions is modulated flexibly is an unresolved question in neuroscience. Previous studies suggest that the efficacy of mutual inputs between two neuronal groups is shaped by their rhythmic activities. In particular, neuronal groups that synchronize their rhythmic activity could establish a communication structure that enhances their interaction dynamically, while reducing the influence from non-synchronized neuronal groups.

Here, we tested this hypothesis of communication through coherence directly by testing the strength of intra- and inter-areal interactions as a function of the precision of their synchronization. We specifically hypothesized that interactions between distant cortical neuronal groups should be enhanced during periods of precise synchronization.

We tested this by recording spike- and local-field potential (LFP) activity simultaneously from electrodes in areas 17, 18 and 21a of awake cats, stimulated with moving gratings. Each electrode assessed the activation in a local group of neurons. While stimulation was constant, the spontaneous response variability allowed our analysis. We investigated whether the precision of synchronization (within a given trial) between neuronal groups predicted the strength of their interaction. We estimated interaction strength by determining the Spearman rank correlation coefficient (across trials) between their firing rates, or between the spectral power of their spike trains and their LFPs.

We found that neuronal groups interacted stronger when they were precisely synchronized. Both the precision of synchronization and the strength of interaction fluctuated rapidly in time and precise synchronization preceded strong interaction by a few milliseconds, suggesting a causal role. Synchronization-dependent interactions were present between neuronal groups both within the same area and distributed in separate areas. Within an area, the effect was present for distances between 2 and 15 mm.

These findings suggest that intra-areal and inter-areal neuronal interactions are mechanistically subserved by precise neuronal synchronization. We hypothesize that neuronal synchronization establishes the dynamic and specific patterns of neuronal interactions that are at the core of our cognitive abilities.

Author Index

- Afshar A., 12
 Ahrens M. B., 17
 Allen H., 29
 Alvarez A. P., 17
 Averbek B. B., 5

 Baldassi C., 18
 Barbieri F., 19
 Barbour B., 28
 Bartho P., 7
 Battaglia F. P., 31
 Beck J., 6, 30
 Benchenane K., 31
 Berens P., 19, 21
 Bethge M., 19, 21
 Bonifazi P., 29
 Bonnasse-Gahot L., 20
 Braunstein A., 18
 Broccard F. D., 29
 Brody C., 9
 Brunel N., 11, 17–19, 27, 28

 Cateau H., 20
 Chen S., 6
 Chichilnisky E. J., 13
 Churchland M. M., 12
 Cunningham J. P., 12, 21
 Curto C., 7

 Dayan P., 28
 de Solages C., 28
 Desrochers-Feledy T.M., 23
 Diesmann M., 26
 Douchamps V., 31

 Ecker A. S., 16, 21
 Eggermont J. J., 24
 Engel A. K., 32
 Ernst U. A., 22

 Feingold J., 23
 Field G. D., 13
 Frank L. M., 6
 Fries P., 32
 Fujii N., 23
 Fukai T., 20

 Ganguli S., 7
 Gansel K., 24
 Garcia-Perez E., 29

 Gauthier J. L., 13
 Gerstein G., 26
 Gourevitch B., 24
 Graf A. B. A., 25
 Graybiel A. M., 23
 Greschner M., 9, 13
 Grivich M. I., 13
 Gruen S., 15, 26

 Häusler S., 8
 Hakim V., 11, 28
 Hamaguchi K., 17, 27
 Hansel D., 30
 Harlan R., 23
 Harris K. D., 7
 Humphreys G., 29

 Isope P., 28

 Jazayeri M., 25
 Juarez Paz L. M., 9

 Kass R., 8, 27
 Keliris G. A., 16
 Kelly R. C., 13
 Khamassi M., 31
 Kohn A., 13, 25
 Koyama S., 27
 Kreiter A. K., 22
 Kretzberg J., 9

 Lee T. S., 13
 Lena C., 28
 Lengyel M., 28
 Linden J. F., 17
 Litke A. M., 13
 Logothetis N. K., 16, 21
 Lubenov E., 13
 Luczak A., 7

 Maass W., 8
 Machens C., 9
 Marguet S., 7
 Mavritsaki E., 29
 Mazzoni A., 29
 Miller L. E., 14
 Movshon J. A., 25

 Nadal J.-P., 20
 Nikolic D., 8
 Nirenberg S., 10

- Nowak L., 30
- O’Keefe, J., 10
- Oostenveld R., 32
- Panagiotaropolulos F., 16
- Paninski L., 17
- Panzeri S., 16
- Parra L., 30
- Pawelzik K. R., 22
- Perez-Bolde L. C., 27
- Persi E., 30
- Petrusca D., 13
- Peyrache A., 31
- Pohlmeyer E., 14
- Ratcliff R., 11
- Rebesco J. M., 14
- Reimer I. C. G., 9
- Riehle A., 17
- Rigat F., 31
- Romo R., 9
- Rotermund D., 22
- Rotter S., 15
- Rousseau C., 28
- Roxin A., 11
- Ruaro M. E., 29
- Sahani M., 12, 17, 21
- Sakata S., 7
- Schoffelen J.-M., 32
- Schrader S., 26
- Shalizi C. R., 27
- Shenoy K. V., 12, 21
- Sher A., 13
- Shlens J., 13
- Siapas T., 13
- Singer W., 8, 24, 32
- Smith M. A., 13
- Solla S. A., 14
- Stauder B., 15
- Szapiro G., 28
- Taylor K., 22
- Thiel A., 9
- Thiele A., 16
- Tierney P. L., 31
- Tolias A. S., 16, 21
- Torre V., 29
- van Vreeswijk C., 30
- Ventura V., 32
- Wiener S. I., 31
- Winzenborg I., 9
- Womelsdorf T., 32
- Yu B. M., 12, 21
- Zecchina R., 18

