Gatsby Computational Neuroscience Unit
Theoretical Neuorscience 2007

Written Examination
12 Jan 2007

Part I

This part has 20 short questions arranges in blocks of 4. Answer any 3 questions in each block. Each is worth 5 marks. No reference materials are allowed.

This part should take 1 hour. You may continue to work for another 30 minutes once that time is up, but indicate clearly which answers (or parts of answers) were written afterward.

Biophysics. Answer any 3 questions.

1. Consider an infinitely long passive cable. Suppose you inject current at a single point, wait until the system reaches steady state, and then turn off the current. What is the subsequent time-evolution of the voltage? To answer, sketch voltage versus distance along the cable at a few well-chosen time points.

2. A cell has a membrane resistance of 100 MΩ and a time constant of 20 ms. What is its capacitance?

3. Consider a facilitating synapse with failures whose weight, W, evolves according to

$$\frac{dW}{dt} = -W + \sum_i \xi_i \delta(t - t_i)$$

   where t_i is the time of the i^{th} spike, the spikes arrive at rate ν, and ξ_i is 1 with probability p and 0 with probability $1 - p$. What is the mean value of W after the system has evolved for a long time? Assume that the ξ_i are uncorrelated with the spike times.

4. What are the approximate concentrations of Na^+, Cl^- and K^+ inside a neuron? What are the approximate concentrations outside it (in the extracellular medium)?

Coding. Answer any 3 questions.

1. Under what conditions on a channel does histogram equalisation maximise the transmitted information?

2. The Venn diagram representation of entropy and information is correct for 2 variables but fails for 3. Give an example of an inequality implied by the Venn diagram that is incorrect. Give an example of a putative neural code which violates it.

3. Two receptor neurons with inhomogeneous Poisson firing statistics step their firing rates from 0 to λ in response to an external stimulus. Cell 1 does so at the time of the stimulus ($t = 0$) and cell 2 follows at a delay Δ which depends on a stimulus parameter θ. Show that the probability that the *first* post-stimulus spike from cell 2 comes precisely τ after the first spike from cell 1 is
$$p(\tau) = \frac{\lambda}{2} e^{-\lambda|\tau - \Delta(\theta)|}$$
(Note the absolute value in the exponent.)

4. For the interval probability given above, calculate the Fisher information $J(\theta)$. [Hint: it's easiest to use the gradient definition and ignore the discontinuity in the derivative.]

Networks. Answer any 3 questions.

1. The neurons in a Hopfield network update according to

$$x_i(t+1) = \text{sign}\left[\sum_{j=1}^{N} J_{ij} x_j(t)\right].$$

Assume that sign(0)=1. The matrix J is given by (assume that N is even)

$$J = \text{sign}[(N+1)/2 - i]\,\text{sign}[(N+1)/2 - j]$$

(in other words, $J = 1$ if either both i and j are less than or equal to $N/2$ or both are greater than $N/2$, and -1 otherwise).

   (a) What are the possible equilibrium configurations?

   (b) Assuming synchronous update (all neurons are updated at once), how many iterations does it take to reach equilibrium if you start from an arbitrary initial condition?

2. Consider the standard set of nullclines for a randomly connected network in the balanced regime (shown below). What happens to the stability of the fixed point as the angle between the nullclines (labeled θ) goes to zero?

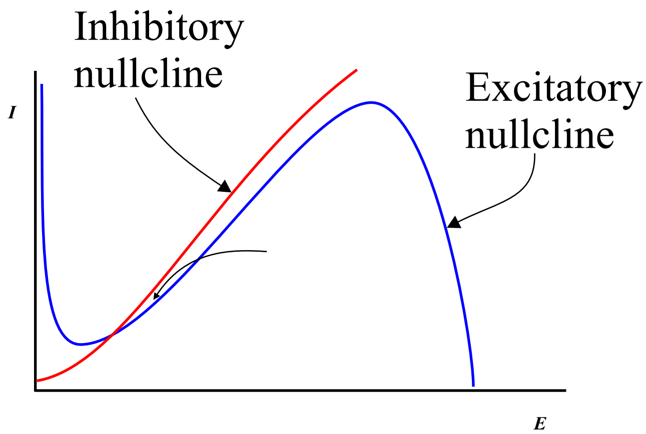

3. Explain why fast inhibition has a stabilizing effect on a network of randomly connected neurons operating in the balanced regime (i.e., operating at the equilibrium shown in the figure above).

4. Explain (using nullcines if you would like) why spike-frequency adaptation can cause bursting.

Learning. Answer any 3 questions.

1. Why can't the long-term depression inherent in a covariance version of a Hebbian learning rule act to stabilize learning?
2. Sketch out the BCM learning or sliding threshold learning rule. How might the threshold be set?
3. What shape of learning curves would be predicted from the Rescorla-Wagner learning rule?
4. In a square Independent Component Analysis model of visual cortex, what is the relationship between the receptive and projective field of the units?

Systems. Answer any 3 questions.

1. What is the structure of a Reichardt motion detector?
2. What structure of horizontal interactions between orientation-tuned neurons in V1 could lead to contour integration and texture segregation?
3. How are the variables of orientation preference and ocular preference mapped on V1 of cats and old-world monkeys?
4. Describe, and sketch the proof of, the relationship between the Receiver Operating Characteristic curve and the probability of error in a two-alternative forced-choice psychophysical test.

Gatsby Computational Neuroscience Unit
Theoretical Neuorscience 2007

Written Examination
12 Jan 2007

Part II

This part contains 4 questions, of which you only need answer 3. You may consult your texts or notes, but **not** online resources.

This part should take 2 hours. You may continue to work for another 2 hours once this time is up, but indicate clearly which answers (or parts of answers) were written afterward.

1. **Biophysics.** Consider a simplified Hodgkin-Huxley type model,

$$\begin{aligned}
\tau \frac{dV}{dt} &= -(V - \mathcal{E}_L) - hm(V)V \\
\tau_h \frac{dh}{dt} &= h_\infty(V) - h \\
m(V) &= \frac{1}{1 + \exp(-(V - V_t)/\epsilon_m)} \\
h_\infty(V) &= \frac{1}{1 + \exp(+(V - V_h)/\epsilon_h)}
\end{aligned}$$

with parameters

$$\begin{aligned}
\mathcal{E}_L &= -65 \text{ mV} \\
V_t &= -50 \text{ mV} \\
\epsilon_h &= 10 \text{ mV} \\
\epsilon_m &\ll 1 \text{ mV} \,.
\end{aligned}$$

The remaining parameter, V_h, will be specified as needed (it will take on a range of values).

(a) Sketch the nullclines in V-h space for $V_h = -60, -50$ and -40 mV. Put voltage on the x-axis and h on the y-axis. Which of these values of V_h, if any allow spike generation? (10 marks)

Spike generation consists of the following: the neuron is instantaneously moved above V_t, under single neuron dynamics it rapidly moves to a higher voltage, and then returns to rest at a value *below* V_t.

(b) For a value of V_h that allows spike generation, sketch the trajectories starting at V slightly greater than V_t and $h = 1$. (10 marks)

(c) Show (graphically) that the amplitude of the spike is an increasing function of τ_h. (10 marks)

(d) As you probably noticed, this system cannot spike repetitively. Is it possible to fix this with a proper choice of ϵ_m, ϵ_h, V_h and τ_h? If so, how? (10 marks)

2. **Coding.** Consider two V1 cells whose responses to an instantaneous visual stimulus $s(\mathbf{x})$ are given by:

$$r_a(s) = \left(\int d\boldsymbol{\xi}\ D_1(\mathbf{x}_1 - \boldsymbol{\xi}) s(\boldsymbol{\xi}) \right) \left(\int d\boldsymbol{\xi}'\ D_2(\mathbf{x}_2 - \boldsymbol{\xi}') s(\boldsymbol{\xi}') \right) + \eta_a$$

$$r_b(s) = \int d\boldsymbol{\xi}\ D_3(\mathbf{x}_3 - \boldsymbol{\xi}) \left[s(\boldsymbol{\xi}) \left(1 + \int d\boldsymbol{\xi}'\ D_4(\boldsymbol{\xi}') s(\boldsymbol{\xi} - \boldsymbol{\xi}') \right) \right] + \eta_b$$

where $D_{\{1,2,3,4\}}(\mathbf{x})$ are zero-centered spatial filters contributing to the cell responses, $\mathbf{x}_{\{1,2,3\}}$ are constants, $\eta_{\{a,b\}}$ are Gaussian noise terms, and $r_a(s)$ and $r_b(s)$ are the **membrane potentials** for each cell in response to s, offset from rest. We will neglect all temporal filtering.

(a) Describe the encoding models defined by these equations, and give plausible network connections that might underlie them [the placement of parentheses might help]. Qualitatively, what stimulus (with constrained power) would make each cell fire most strongly? (8 marks)

A colleague obtains experimental data giving measured responses of the cells to a family of images $\{s_i(\mathbf{x})\}$ drawn from a distribution $P(s)$, such that $\langle s_i(\mathbf{x}) s_i(\mathbf{x}') \rangle \propto \delta(\mathbf{x} - \mathbf{x}')$. Assume that enough stimuli were presented so that all averages over presentations are equal to their expected values (both with respect to $P(s)$ and with respect to the Gaussian noise in the membrane potential).

(b) Write down (up to a constant of proportionality) expressions for the maximum likelihood estimators of **linear** spatial receptive fields for these cells, in terms of $D_{\{1,2,3,4\}}$ and the moments of $P(s)$. (8 marks)

To simplify calculations, suppose that the visual stimulus is a one-dimensional strip of pixels $x \in \{1, 2, 3, \ldots, N\}$, so that the integrals above can be replaced by finite sums. Your colleague used two classes of stimuli in the experiments:

- sparse noise, where each pixel $s_i(x_n)$ is chosen independently from some zero-mean distribution P_1.
- gratings, where $s_i(x_n) = \sin(\omega_i n \Delta x)$ and ω_i is chosen uniformly from $\{0, 2\pi/(N\Delta x), 4\pi/(N\Delta x), \ldots, 2N\pi/(N\Delta x)\}$

(c) Verify that both families satisfy the criterion $\langle s_i(x_n) s_i(x_{n'}) \rangle \propto \delta_{nn'}$. [Hint: for the gratings, it might help to think of $x_{n,n'}$ as frequency variables and ω_i as a time variable.] (4 marks)

Your colleague is excited, because when he carried out the linear analysis for each class of stimuli separately, he found different linear spatial filters in each case. He believes that this shows that V1 adapts to stimulus statistics in a sophisticated fashion.

(d) Do you agree? Explain. (4 marks)

(e) If each filter $D_{\{1,2,3,4\}}(x)$ is a local bump, describe qualitatively (or sketch) the forms of the linear receptive field you expect to see for each cell probed with each stimulus class. Explain. (12 marks)

(f) Finally, for which (if any) of these cells might you use spike-triggered covariance methods to obtain estimates of the actual filters $D_{()}$. (4 marks)

3. **Network dynamics.** Consider a network of quadratic integrate-and-fire neurons whose voltages, V_i, evolve according to

$$\begin{aligned}
\tau \frac{dV_i}{dt} &= \frac{(V_i - V_r)(V_i - V_t)}{V_t - V_r} + I_0 - I_i \\
I_i &= \sum_j W_{ij} g_j(t)(V_i - \mathcal{E}_j) \\
g_j(t) &= \sum_k f(t - t_j^k) \\
f(t) &= \Theta(t) \frac{\exp(-t/\gamma)}{\gamma}
\end{aligned}$$

where V_r, V_t, $\mathcal{E}_j$, W_{ij}, γ and I_0 are constants, t_j^k is the time of the k^{th} spike on neuron j, and Θ is the Heaviside step function: $\Theta(t) = \max(0, t)$. Note that all the W_{ij} are non-negative.

Assume that $V_t > V_r$. When $V_i(t)$ reaches ∞, a spike is emitted and V_i is reset to $-\infty$. In case you don't remember from the homework, V_i both goes to ∞ and returns from $-\infty$ in finite time.

(a) Assume $W_{ij} = 0\ \forall i, j$ and $I_0 = 0$. What is the resting membrane potential? What is the threshold for the generation of an action potential? (3 marks)

(b) Assume $W_{ij} = 0\ \forall i, j$. At what value of I_0 does the neuron start to fire repetitively? (4 marks)

(c) The synaptic current, I_i, can be written $G_i(t)V_i - H_i(t)$ where

$$\begin{aligned}
G_i(t) &= \sum_j W_{ij} g_j(t) \\
H_i(t) &= \sum_j W_{ij} g_j(t) \mathcal{E}_j \,.
\end{aligned}$$

Assume that the i^{th} neuron fires at constant rate ν_i. Write down expressions for $\langle G_i \rangle$ and $\langle H_i \rangle$, the time averages of $G_i(t)$ and $H_i(t)$, respectively, in terms of the W_{ij}, $\mathcal{E}_j$ and ν_j. Hint: $f(t)$ integrates to 1. (8 marks)

(d) Assume that W_{ij} is drawn *i.i.d.* from a distribution that is independent of the number of neurons, N, and that V_i is independent of G_i. Show that as $N \to \infty$, there is no equilibrium in which all neurons fire at constant (and nonzero) rate. (10 marks)

(e) Let us go back to the finite N limit, so that all the neurons can fire steadily, at rate ν_i for neuron i. Derive an expression for the temporal fluctuations of H_i, denoted $C_{ij}^H(t - t')$. This quantity is defined to be

$$C_{ij}^H(t - t') = \Big\langle \Big(H_i(t) - \langle H_i \rangle\Big)\Big(H_j(t') - \langle H_j \rangle\Big)\Big\rangle$$

where, as above, the angle brackets denote a time average. Assume that

$$\Big\langle \Big(g_i(t) - \langle g_i \rangle\Big)\Big(g_j(t') - \langle g_j \rangle\Big)\Big\rangle = \nu_i c(t - t')\delta_{ij} \,,$$

the W_{ij} are drawn *i.i.d.* from a distribution with mean W_0 and variance σ^2, W_{ij} and ν_j are uncorrelated, and there are N neurons. (15 marks)

4. **Reinforcement learning (Niv).** Consider a rat pressing a lever to gain delivery of electrical stimulation to an appetitive location. Under a *variable ratio schedule,* the experimenter will provide a reward worth R units for every n lever presses on average. If a lever press that takes time τ costs K/τ units:

   (a) if he presses at regular intervals, what is the long run average rate of net reward (*ie* taking account of costs)? (4 marks)

   (b) what value of τ optimises this average rate? (4 marks)

   (c) what is the resulting optimal reward rate, and how does this depend on n? (6 marks)

   (d) in a *random interval schedule*, animals are rewarded for the first press after a time drawn from an exponential distribution with mean γ seconds has passed, with the next time being drawn at the time of the successful lever press. If the animal presses every τ seconds, what is the long run average number of lever presses per reward? The long run average time between rewards? The average reward rate? (13 marks)

   (e) Based on your expression for the average reward rate, indicate the qualitative dependence of the optimal τ on γ. (13 marks)