Gatsby Computational Neuroscience Unit
Theoetical Neuroscience

Final Examination
25 Jan 2008

Part I

This part has 20 short questions arranged 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. Define electrotonic length.
2. If you increase the extracellular concentration of potassium, does the leak reversal potential go up or down?
3. The H-current is an inward sodium current that activates at low voltages (around -60 mV) and de-activates slowly at high voltages (around -50). Explain why this current can lead to bursting.
4. What happens to a cell's time constant when you double the thickness of the cell membrane? Assume the area, channel density, and channel conductance do not change.

Coding. Answer any 3 questions.

1. Describe two different types of point process that can be used to model refractory firing.
2. A colleague measures the mutual information between a dynamic stimulus and a recorded spike train in two ways. In the first case, the spike-train is represented by a series of 0's and 1's. In the second, each 1 is replaced by the angular phase of the LFP (filtered in the alpha band) at the time of the spike. It turns out that the mutual information in the second case is higher. Does this indicate synergistic encoding between the spike-train and the LFP? Justify your answer.
3. Explain how you would use an ROC curve to calculate the discriminability index d'.
4. A colleague has made an error in performing STC analysis. Instead of diagonalising the spike-triggered covariance, he has diagonalized the spike-triggered correlation matrix (that is, he did not subtract or project out the spike-triggered average vector). How might this affect his results?

Networks. Answer any 3 questions.

1. Consider a cultured neuronal network, assumed to be operating in the balanced regime. If you increase the extracellular concentration of potassium, can you say definitively what happens to the equilibrium firing rate of the excitatory population? The inhibitory population?

2. Suppose you want to build a feedforward network that propogates firing rate over multiple layers (i.e., for a range of firing rates, the firing rate in the input layer is the same as the firing rate in the output layer). Why does this get harder and harder as you add more layers?

3. In the brain, neurons receive both local (within a few hundred microns) and long distance (millimeters to centimeters) connections. In cortex, which one dominates, and by how much (approximately).

4. Consider the differential equations

$$\frac{dx}{dt} = ax + by$$
$$\frac{dy}{dt} = cx + dy\,.$$

   Find values of a, b, c and d for which the solution admits a limit cycle.

Learning. Answer any 3 questions.

1. Describe one conditioning paradigm that violates the Rescorla-Wagner learning rule and indicate how it fails.
2. Describe how both information maximization *and* information minimization could both have been suggested as principles of self-organization of cortical representations.
3. Discuss two ways that uncertainty should govern learning rates in conditioning. Should expected and unexpected uncertainty act in the same way in these regards?
4. What is the difference between indirect and direct policy (actor) learning?

Systems. Answer any 3 questions.

1. The populations codes for current eye position and intended next eye position in the parietal cortex are of different types. Suggest a theoretical reason why this might be so.
2. Which senses follow the pattern of thalamo-neocortical projection? Do they each have the same number of synapses between receptors and cortex?
3. Give two examples of critical periods for development. What purpose might a critical period serve?
4. What happens to the receptive fields of retinal ganglion cells as the overall illumination level changes. Why might this be?

Gatsby Computational Neuroscience Unit
Theoetical Neuroscience

Final Examination
25 Jan 2008

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 two gap-junction-coupled quadratic integrate and fire neurons,

$$\frac{dv_1}{dt} = \mu_1 + v_1^2 - g_1(v_1 - v_2) \tag{1a}$$
$$\frac{dv_2}{dt} = \mu_2 + v_2^2 - g_2(v_2 - v_1)\,. \tag{1b}$$

Here the v_i are normalized voltages. A spike is emitted from neuron i when v_i reaches $+\infty$, at which point voltage is reset to $-\infty$. Both g_1 and g_2 are positive.

(a) Explain why one would expect, in a real cell, that the coupling terms ($g_1(v_1 - v_2)$ and $g_2(v_2 - v_1)$) should depend only on the difference in voltage. Be explicit about your assumptions. In particular, how would this be modified if the ionic concentrations in the two cells were different? (5 marks)

(b) Draw the nullclines for Eq. (1). Show that there can be 0, 1, 2, 3 or 4 fixed points. Show which are stable, which are unstable, and which, if any, are saddle points. (10 marks)

(c) *For this question only,* rather than emitting a spike at $+\infty$ and resetting at $-\infty$, emit a spike at v_{thresh} and reset to v_{reset}. For this regime, find parameters in which one neuron fires repetitively and the other is silent. Sketch the resulting nullclines. (5 points)

(d) For this and the next question, let $\mu_1 = \mu_2 = \mu$ and $g_1 = g_2 = g$. Derive conditions on μ and g that ensure repetitive spiking. Draw the nullclines when the parameters are at the boundary between repetitive spiking and no spiking. (5 marks)

(e) Choose μ and g to allow repetitive spiking. Let μ obey the following dynamics:

$$\mu \rightarrow \mu - \delta\mu \quad \text{whenever either neuron spikes} \tag{2a}$$
$$\tau\frac{d\mu}{dt} = -(\mu - 1) \quad \text{between spikes}\,. \tag{2b}$$

Show that this can produce bursting. (5 marks)

(f) Show that an upper bound on the inter-burst interval is $\tau \log(1 + \delta\mu)$. (10 marks)

2. Fisher information and refractory firing.

   Consider a hypothetical cell, which responds to the presentation of a stimulus with a continuous feature s by firing at a homogeneous rate $f(s)$ in a (fixed) interval $[0, T]$. Assume that the firing rate is 0 outside this interval. We will be interested in the contributions made to the Fisher information by spike-timing, with and without a refractory period.

   First, assume that the firing is Poisson.

   (a) What is the probability of observing spikes at times $\{t_1 \dots t_n\} \subset [0, T]$? (1 mark)

   (b) What is the Fisher information $J_{t,Poiss}(s)$ associated with this probability density function, assuming that the relevant interval $[0, T]$ is known? How does it compare to the Fisher information $J_{n,Poiss}$ associated with the distribution of spike counts $P(n|s)$? (9 marks)

   Now consider refractory firing. Recall that one way to model a refractory period is to use a gamma-interval renewal process in place of a Poisson process. Thus, now assume that the cell's firing follows a gamma-interval process with the same mean rate $f(s)$ and with integral gamma order γ.

   (c) What is the probability of observing spikes at times $\{t_1 \dots t_n\} \subset [0, T]$ from this process? (5 marks)

   (d) What is the Fisher information $J_{t,Gamma}(s)$ associated with the new probability density function? You may assume that T is long enough to neglect contributions due to the first spike, and due to the silence after the last spike. (10 marks)

   Finally, we wish to see how much of this information gain is available in the spike count.

   (e) Which signal (count or spike-timing) do you expect to carry more information for this process? Why? (5 marks)

   (f) Find an expression for the the distribution of spike counts $P(n|s)$ under the gamma-interval model. (5 marks)

   (g) Write down the expression for the corresponding Fisher information $J_{n,Gamma}$, and thus for $J_{n,Gamma} - J_{t,Gamma}$. You need not necessarily evaluate the expectation. Identify the term(s) responsible for the difference between $J_{n,Gamma}$ and $J_{t,Gamma}$. (5 marks)

3. Network dynamics

Consider a network of N analog neurons that obey the time-evolution equations

$$\tau \frac{dx_i}{dt} = \phi\left(\sum_j W_j x_j - \theta_i\right) - x_i \,. \tag{1}$$

(a) Assume that $\theta_i = \theta\ \forall i$. Show that Eq. (1) can be effectively reduced to a one-variable model,

$$\tau \frac{dz}{dt} = \phi\,(Jz - \theta) - z \,. \tag{2}$$

Write down expressions for z and J in terms of W_i and x_i. (5 marks)

(b) Let's go back to Eq. (1), where θ_i depends on i. Show that Eq. (1) can still be reduced to a one-variable model,

$$\tau \frac{dz}{dt} = \tilde{\phi}\,(Jz) - z \,. \tag{3}$$

Write down an exact expression for $\tilde{\phi}(Jz)$ in terms of J, W_i and θ_i. (5 marks)

(c) Assume that both W_i and θ_i are correlated random variables with joint distribution $p(W,\theta)$. Assuming $N \to \infty$, write down an expression for $\tilde{\phi}(Jz)$ as an integral over this joint distribution. (7 marks)

(d) Write down the much simpler expression that results if $p(W,\theta)$ factorizes: $p(W,\theta) = p(W)p(\theta)$. (3 marks)

(e) Let's go back to the case in which $\theta_i = \theta$, so that z evolves according to Eq. (2). Let $\phi(y) = \tanh(y)$. To model spike frequency adaptation, let θ evolve according to

$$\tau_0 \dot{\theta} = -(\theta - \theta_0 z)\,, \tag{4}$$

with $\tau_0 \gg \tau$. Assume that $\theta_0 > J - 1 > 0$. Sketch the nullcines. (10 marks)

(f) Show that the system exhibits bursting, and sketch $z(t)$ and $\theta(t)$ versus time. (10 marks)

4. Temporal representations in conditioning

   (a) define an absorbing Markov chain. Under what circumstances will it have a well-defined long-run expected value function? (5 marks)

   (b) if an absorbing Markov chain has a transition matrix $\mathcal{P}$ and average immediate reward vector $\mathbf{r}$, what is an expression for the long run expected value vector $\mathbf{v}$ of all the states? What happens if future rewards are exponentially discounted (5 marks)

   (c) what is the temporal difference error, and what is its relationship to your answer to the previous question? (5 marks)

   Consider the case that a conditioned stimulus (a cue) comes on at time $t = 2$ in a trial, a single reward ($r = 1$) is delivered at time $t = 11 \ldots 15$, each timestep with probability 0.2, the subject has a noise-free tapped-delay line representation of the time since the cue came on, and knows when each trial begins ($t = 1$) and ends ($t = 30$), and there is no temporal discounting.

   (d) if the reward is preceded (by one time-step) on a trial by a distinct sensory event (like the click of a food dispenser), draw a Markov chain that describes the experience the subject would have (including transition probabilities), and find the expected value of each state in this chain (10 marks)

   (e) what would the TD error look like in this case – stimulus locked to the cue and to the dispenser click (10 marks)

   (f) if the monkey's interval timing is noisy, describe some possible consequences for the course of learning and the TD error signals. (5 marks)