Gatsby Computational Neuroscience Unit
Theoetical Neuroscience

Final Examination
26 Jan 2009

Part I

This part has 20 short questions arranged in 5 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. For a Hodgkin-Huxley neuron, the sodium current is given by

$$I_{Na} = g_{Na} m^3 h (V - \mathcal{E}_{Na})\,.$$

Consider a very simplified model of a spike,

$$\begin{aligned} t < 0 \text{ ms} : &\quad m = 0, h = 1 \\ t = 0 \text{ ms} : &\quad m \to 1 \\ t = 1 \text{ ms} : &\quad h \to 0\,. \end{aligned}$$

Assume that during the spike, $V(t)$ rises linearly from V_{rest} to $\mathcal{E}_{Na}$ in 1 ms.

Write down an expression for the total current that flows into the cell between the times 0 and 1 ms.

2. The propagation speed of an action potential in axons is linear in the radius if they are myelinated and proportional to the square root of the radius if they are unmyelinated. In what sense does this "explain" the fact that thin axons tend to be unmyelinated and thick ones tend to be myelinated?

3. For the following neurotransmitters, give the approximate reversal potential and time constant: AMPA, NMDA, GABA$_A$, GABA$_B$.

4. Provide a biophysical explanation for synaptic facilitation.

Coding. Answer any 3 questions.

1. A particular synapse invariably fails to release a vesicle for every **second** presynaptic spike. Assume that the presynaptic spike-train is drawn from a homogeneous Poisson process. What is the coefficient of variation (CV) of the interval between synaptic release events?
2. A self-exciting point process with counting process $N(t)$ and event times t_i has intensity

   $$\lambda(t) = \lambda_0(1 - e^{-(t-t_{N(t)})/\tau})$$

   Give a function $r(t)$ (in terms of λ_0, $t_{N(t)}$ and $r(t_{N(t)})$) such that the $r_i = r(t_i)$ appear Poisson distributed.
3. Suggest an experimental setting in which we might choose to fit a multidimensional LNP model by maximising mutual information rather than by diagonalising the spike-triggered covariance.
4. Does a population code with broad tuning curves always represent values less accurately than one with narrower tuning? Explain why (not).

Networks. Answer any 3 questions.

1. Suppose that h_i, the steady state synaptic drive to a cell, is given by

$$h_i = \frac{1}{K^{1/2}} \sum_{j=1}^{N} c_{ij} w_{ij} \nu_j \,.$$

Here ν_j is the firing rate of neuron j, w_{ij} is the connection strength from neuron j to neuron i, c_{ij} is a Bernoulli random variable,

$$c_{ij} = \begin{cases} 1 & \text{with probability } c \\ 0 & \text{with probability } 1 - c \,, \end{cases}$$

and $K \equiv cN$ is the average number of connections per neuron. Assume that c_{ij}, w_{ij} and ν_j are uncorrelated and drawn iid. Write down an expression for the mean and variance of h_i (with respect to index, i) in terms of: K, the first and second moments of the firing rates, and the mean and variance of w_{ij}. As usual, assume that $N \gg 1$.

2. Consider a standard Hopfield model with sparse connectivity,

$$x_i(t+1) = \text{sign}\left[\frac{1}{N} \sum_{j=1}^{N} \left(c_{ij} \sum_{\mu=1}^{p} \xi_i^\mu \xi_j^\mu \right) x_j(t) \right]$$

where c_{ij} is the same Bernoulli random variable as in the previous short question ("sparse" means $c < 1$; typically it's about 0.1) and the ξ_i^μ are, as usual, uncorrelated random variables that take on two values, $+1$ and -1, each with probability $1/2$. Use simple arguments about signal and noise to show that the capacity of the network (the maximum number of memories, p, that can be stored) is proportional to K ($\equiv cN$).

3. In our mean field analysis, we always assume that neurons operate in the asynchronous regime, meaning they fire with more or less Poisson statistics, and their firing times are uncorrelated. Why is this assumption important?

4. Consider the following differential equation, written in polar coordinates,

$$\begin{aligned} \frac{dr}{dt} &= r - 2r^2 + 1.01r^3 \\ \frac{d\theta}{dt} &= \omega \,. \end{aligned}$$

Sketch the trajectories in the x-y plane, with initial conditions $r \approx 0.1$ and $\theta = 0$. Note that r and θ are related to x and y via $x = r\cos\theta$ and $y = r\sin\theta$.

Learning. Answer any 3 questions.

1. What forms of learning rules are called covariance rules? Show why.
2. What is the baseline firing rate of dopamine neurons? Discuss two implications this has for the representation of negative prediction errors.
3. How does Oja's rule determine the norm of the weights? What learning rule would force the norm to be constant throughout learning?
4. How should uncertainty affect a Bayesian learner?

Systems. Answer any 3 questions.

1. A simple cell with a receptive field that is separable in space and time cannot report the direction of motion of a bar. What about a population of such cells? [Support your answer. A simple "yes" or "no" will not receive any marks.]
2. Describe the representation of visual space in V1. Why is it sometimes described as a "complex logarithm"?
3. What is whole cell recording? How might the sample of cells measured by whole cell methods be biased, and how would this bias differ from that of conventional extracellular recording?
4. List 4 different neuromodulatory systems, and a hypothesised computational role for each.

Gatsby Computational Neuroscience Unit
Theoetical Neuroscience

Final Examination
27 Jan 2009

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.

1. The Izhikevich Neuron.

Consider the now famous Izhikevich Neuron, which obeys the set of equations

$$\begin{aligned} \frac{dV}{dt} &= 0.04V^2 + 5V + 140 - U + W \\ \frac{dU}{dt} &= a(BV - U)\,. \end{aligned}$$

When V reaches V_0 (typically set to 30 mV), a spike is emitted, and V and U are reset as

$$\begin{aligned} V &\rightarrow C \\ U &\rightarrow U + D\,, \end{aligned}$$

where C and D are constants.

(a) Explain why this is a reasonable model of a real neuron. Provide an interpretation of the variable U. What are reasonable values of C and D? (15 marks)

(b) Find a change of variables (of the form $V = \alpha + \beta v$ and $U = \gamma + \delta u$, where α, β, γ, and δ are constant) that transforms this equation to

$$\begin{aligned} \frac{dv}{dt} &= v^2 - u + w \\ \frac{du}{dt} &= a(bv - u)\,. \end{aligned}$$

When $v = v_0$, v and u are reset according to

$$\begin{aligned} v &\rightarrow c \\ u &\rightarrow u + d\,. \end{aligned}$$

Provide explicit expressions for the new constants (w, b, c, d, and v_0) in terms of the old ones (W, B, C, D, and V_0). (10 marks)

(c) Draw nullclines for a set of parameters in which there are two equilibria. (5 marks)

(d) Show that if there are two equilibria and a is sufficiently large and positive, then one of the equilibria is stable and the other is unstable. (5 marks)

(e) This cell can burst, in the sense that initial settings of v and u can lead to several spikes, followed by relaxation to a stable equilibrium. Draw a set of nullclines that can support this type of bursting, and sketch the trajectories. (5 marks)

2. Locust collision detection.

The lobula giant motion detector (LGMD) is a neuron in the locust's lobula plate that responds to looming stimuli and seems to mediate an escape response. Hatsopolous et al (1995) and Gabbiani et al (1999) report that the firing rate of the neuron depends on both the angle subtended by a looming stimulus, and the angular velocity of its edges, according to:

$$f(t+\tau) = \dot{\theta}(t) e^{-\alpha\theta(t)} \tag{1a}$$

where $\theta(t)$ is **half** the angle subtended by the object at time t relative to impact (the half will be useful to avoid having to carry many factors of 2); $\dot{\theta}$ is then the angular velocity of an edge; and $\alpha > 0$ and the delay τ are parameters that vary from animal to animal. Here, we will set τ to 0 for convenience.

(a) Assume an object with a linear size of $2l$ approaches the locust's eye at a constant linear speed v. Let $t = 0$ be the time of impact. Suppose that the object is centred with respect to the eye. What angle does the object subtend at some time $t < 0$, expressed as a function of l and v? (3 marks)

(b) Show that, if the LGMD neuron fires as in equation (1a), its firing rate will peak at a time

$$t_{\text{peak}} = -\frac{1}{2}\alpha\frac{l}{v}$$

with a value

$$f_{\text{peak}} = F_{\text{max}}(\alpha)\frac{v}{l}$$

where F_{max} is a function that depends on α alone. (12 marks)

Based on this result, the authors above suggested that the occurrence of the peak is an important cue to escape. This raises the question of how easy it might be for the downstream system to recognise the peak in the firing rate. One way to do so, would be to detect a negative rate of change in the LGMD firing rate. We will analyse the accuracy with which this might be possible, using a discretised simplification.

Suppose that LGMD fires according to an inhomogeneous Poisson process, and that the putative peak detector integrates spikes from the LGMD within two non-overlapping windows of the same size. Let the first have a mean spike count of μ (which we will later set to the value derived from f above), and the second a mean count of $\mu + \delta\mu$.

(c) Calculate the Fisher Information matrix for the parameter vector $(\mu, \delta\mu)$. (5 marks)

(d) Use this to derive the Cramér-Rao lower bound for the variance of an unbiased estimator of $\delta\mu$. (5 marks)

(e) Assume an unbiased efficient estimator with a Gaussian distribution is available (not a good assumption). How large must $|\delta\mu|$ be (in terms of μ) for a simple peak detector that just looks at the sign of such an estimate to experience a false negative rate of about 2.5% [recall that if $\Phi(z)$ is the standard normal CDF, $\Phi(2) \approx 0.975$]. (5 marks)

(f) Now, suppose the peak firing rate for a particular looming object is around 240 spikes s^{-1}, and that we can approximate equation (1a) by a linear fall to 0 during the following 150 ms. Assume the peak detector integrates over 50 ms windows, and consider an ideal setting where the first window is centred over the peak. (Ignore the edge effects, and assume incorrectly that the mean rate of 240 spikes s^{-1} is maintained throughout this first window). How much later must the other 50 ms window be to achieve the 2.5% false negative rate? (5 marks)

(g) How would this duration change if the looming object approached at the same speed, but was half the size (assume the time from peak to silence is the same)? If the locust had time to escape the first object, can we say if it would also have time to escape this one? (5 marks)

3. Networks

Consider a network of excitatory and inhibitory neurons whose firing rates obey the time-evolution equations

$$\begin{aligned}\tau_E \frac{d\nu_{E,i}}{dt} &= \phi_E \left(W_{EE}\nu_E - W_{EI}\nu_I + \frac{\beta}{Nf(1-f)} \sum_{j=1}^{N} \eta_i(\eta_j - f)\nu_{E,j} \right) - \nu_{E,i} \\ \tau_I \frac{d\nu_{I,i}}{dt} &= \phi_I \Big(W_{IE}\nu_E - W_{II}\nu_I \Big) - \nu_{I,i} \,.\end{aligned}$$

where τ_E and τ_I are time constants, ν_E and ν_I are the average excitatory and inhibitory firing rates,

$$\nu_E \equiv \frac{1}{N} \sum_{i=1}^{N} \nu_{E,k} \ , \qquad \nu_I \equiv \frac{1}{N} \sum_{i=1}^{N} \nu_{I,k} \ ,$$

N is the number of excitatory and inhibitory neurons (assumed to be the same), η_i is a Bernoulli random variable,

$$\eta_i = \begin{cases} 1 & \text{with probability } f \\ 0 & \text{with probability } 1 - f \,, \end{cases}$$

$0 < f < 1$, all weights are positive, and the two gain functions, ϕ_E and ϕ_I, are sigmoidal.

(a) Derive a set of time-evolution equations for ν_E, ν_I, and the overlap, m, the latter defined to be

$$m \equiv \frac{1}{Nf(1-f)} \sum_{j=1}^{N} (\eta_j - f)\nu_{E,j} \,.$$

Assume that $N \gg 1$. (5 marks)

(b) Consider the limit $\tau_I \to 0$. Show that the time evolution equations can be written

$$\tau_E \frac{d\nu_E}{dt} = \phi_E\Big(-\gamma_E(\nu_E)\Big) + f\left[\phi_E\Big(-\gamma_E(\nu_E) + \beta m\Big) - \phi_E\Big(-\gamma_E(\nu_E)\Big)\right] - \nu_E \quad \text{(1b)}$$

$$\tau_E \frac{dm}{dt} = \left[\phi_E\Big(-\gamma_E(\nu_E) + \beta m\Big) - \phi_E\Big(-\gamma_E(\nu_E)\Big)\right] - m \quad \text{(1c)}$$

where $\gamma(\nu_E)$ may be a complicated function of ν_E. (5 marks)

(c) In the limit that W_{EE}, W_{EI}, W_{IE}, and W_{II} are very large, derive an explicit expression for $\gamma(\nu_E)$. (5 marks)

(d) Show that if the network exhibits stable operation in the balanced regime when $\beta = 0$, then $\gamma(\nu_E)$ is an increasing function of ν_E. (10 marks)

(e) Draw the nullclines for equation (1), assuming that $\gamma(\nu_E)$ is an increasing function of ν_E. Include arrows indicating flow. Put m on the x-axis and ν_E on the y-axis. Show that parameters exist for which there is bistability with respect to m; that is, a stable equilibrium exists with $m = 0$ and $m > 0$. (10 marks)

(f) Show that if $f = 0$, the equilibrium with $m > 0$ is at high firing rate. (5 marks)

4. The rewards of waiting.

   Barry Richmond has investigated a reward schedule task illustrated in the figure. Ostensively, it involves repeated performance of a simple vigilance task (releasing a touch-bar between 200-800ms after a visually presented dot turns from red to green). However, the monkey is not rewarded after each successful performance; rather reward comes at the end of schedules, *ie* after performing certain numbers of successful trials (up to three here). The schedules are labelled as `trial/schedule`, so `1/1` is the sole trial in a schedule involving just one single vigilance task, `1/2` is the first trial of a two schedule task, and so forth. Reward of value 1 is provided only at the final trial of each schedule. After state `0/0`, the next schedule is chosen uniformly from `1/1`, `1/2`, `1/3`. Assume optimal performance, so the only task is estimation, with each task taking one step.

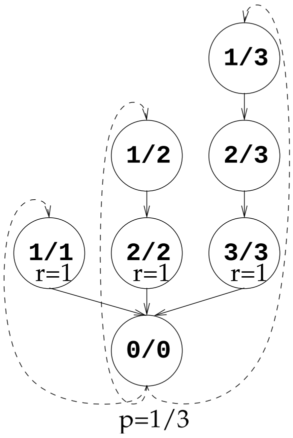

   (a) Assuming a discount factor of γ, what are the discounted values of each state (including `0/0`)? Sketch these values for each point in each schedule. (9 marks)

   (b) Average case reinforcement learning (RL) is an alternative to discounted RL. Here, the goal is to estimate

   $$\rho = \lim_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} r(s_t)$$

   where s_t is the state at time t. Write down a simple expression for ρ that does not involve an infinite sum. (13 marks)

   (c) Under average case, the *differential* values of a state s is defined as

   $$V(s) = E\left[\lim_{T \to \infty} \sum_{t=1}^{T} (r_t - \rho)\right]_{s_1 = s}$$

   Show that these satisfy

   $$V(s) + \rho = r(s) + E[V(s_2)]_{s_1 = s}$$

   What are these differential values in the schedule task? (10 marks)

   (d) One way to estimate these is to maintain a running estimate of ρ updated as

   $$\hat{\rho}_{t+1} = \hat{\rho}_t + \epsilon(r(s_t) - \hat{\rho}_t)$$

   and then to use

   $$\Delta V(s_t) \propto r(s_t) - \rho_t + V(s_{t+1})$$

   Sketch the long run average values of ρ_t at the various steps in the schedule task (for $\epsilon \simeq 0.2$), and comment on the resulting distortions to the values of $V(s)$. (8 marks)