Gatsby Computational Neuroscience Unit
Theoretical Neuroscience

Final Examination
22 Jan 2016

Part II – long questions

There are four questions, one from each main section of the course. Please answer three out of the four, starting the answers for each new question on a new page. Don't forget to write your name at the top of the answer to each question.

You have a maximum of 7 hours for this exam.

Good luck!

1 Biophysics

Consider a simplified Hodgkin-Huxley type model with an m-current that turns on very rapidly,

$$\tau \frac{dV}{dt} = -(V - \mathcal{E}_L) - ghV\Theta(V - V_m) + X(t) \tag{1a}$$

$$\tau_h \frac{dh}{dt} = \frac{1}{1 + \exp\big((V - V_h)/\epsilon_h\big)} - h \tag{1b}$$

where $\theta(V - V_m)$ is the Heaviside step function: it's 1 if $V \geq V_m$ and 0 otherwise. The parameters are

$$\mathcal{E}_L = -70 \text{ mV} \tag{2a}$$

$$V_m = -50 \text{ mV} \tag{2b}$$

$$V_h = -60 \text{ mV} \tag{2c}$$

$$\epsilon_h = 10 \text{ mV} \tag{2d}$$

$$\tau = 10\text{ms} \tag{2e}$$

$$\tau_h = 1\text{ms} \tag{2f}$$

$$g = 100\,. \tag{2g}$$

1. Draw the nullclines in V-h space, assuming $X(t) = 0$. Put h on the y-axis and V on the x-axis. Draw arrows indicating the direction in which V changes at fixed h and h changes at fixed V.

   (15 marks)

2. Show that there are two stable fixed points. For both you can either demonstrate stability by considering linearized dynamics, or argue about stability based on trajectories in V-h space.

   (10 marks)

3. Show that if the membrane potential is driven above V_m mV with h sufficiently large, there will be an all-or-none spike. Indicate on your nullcline plot the minimum amplitude of the spike.

   (5 marks)

4. At $t < 0$, h and V are at the stable fixed point. The input $X(t)$ is given by

   $$X(t) = V_0 e^{-t/\tau}. \tag{3}$$

   What is the minimum value of V_0 that will make the neuron spike?

   (10 marks)

2 Networks

Computing correlations self-consistently. This is a kind of long question, but each step is relatively straightforward, and most of the steps don't depend on the previous ones, so they can be done more or less independently.

Consider a network of n LIF neurons,

$$\tau \frac{dV_i}{dt} + V_i = \frac{1}{\tau_s} \sum_{j=1}^{n} w_{ij} g_j(t) + x_i(t) \tag{4a}$$

$$\frac{dg_i}{dt} + \frac{g_i}{\tau_s} = \sum_k \delta(t - t_j^k), \tag{4b}$$

where t_j^k is the time of the k^{th} spike on neuron j, and $x_i(t)$, which is assumed known, consists of a constant term plus stationary noise. A spike is emitted whenever V exceeds threshold, at which point it is set to zero. The details of the single neuron won't be important

1. Assume the network is operating in steady state, so that the firing rates of all the neurons are constant. Show that

   $$\frac{1}{\tau_s} \sum_j w_{ij} g_j(t) = \sum_j w_{ij} \nu_j + h_i(t) \tag{5}$$

   where ν_j is the average firing rate of neuron j and $h_i(t)$ is given by

   $$h_i(t) = \sum_j w_{ij} s_j(t) \tag{6a}$$

   $$s_j(t) = \frac{g_j(t)}{\tau_s} - \nu_j. \tag{6b}$$

   Be sure to show that ν_j is the firing rate of neuron j.

   (5 marks)

2. The correlation between any two (stationary) functions of time, say $y(t)$ and $z(t)$, is defined to be

   $$\text{Cov}[y(t), z(t+\tau)] \equiv \lim_{T\to\infty} \frac{1}{T} \int_0^T dt\, y(t) z(t+\tau) - \left[\lim_{T\to\infty} \frac{1}{T} \int_0^T dt\, y(t)\right] \left[\lim_{T\to\infty} \frac{1}{T} \int_0^T dt\, z(t)\right]. \tag{7}$$

   Define $H_{ij}(\tau)$, $S_{ij}(\tau)$ and $X_{ij}(\tau)$ to be the covariance between the h's, s's and x's, respectively,

   $$H_{ij}(\tau) = \text{Cov}[h_i(t), h_j(t+\tau)] \tag{8a}$$

   $$S_{ij}(\tau) = \text{Cov}[s_i(t), s_j(t+\tau)] \tag{8b}$$

   $$X_{ij}(\tau) = \text{Cov}[x_i(t), x_j(t+\tau)]. \tag{8c}$$

Note that because $x_i(t)$ is known, $X_{ij}(\tau)$ is also known. Assume that X_{ij} is diagonal,

$$X_{ij} = X_{ii}\delta_{ij} \, . \tag{9}$$

Our goal is to find a self-consistent relationship between H and C. Start by showing that

$$H_{ij}(\tau) = \sum_{lm} w_{il} S_{lm}(\tau) w^T_{mj} \tag{10}$$

where superscript T denotes transpose.

(3 marks)

3. Assume that $S_{ij}(\tau)$ is linear in $H_{ij}(\tau) + X_{ij}(\tau)$. Argue, *very* qualitatively, that the linearity includes a convolution,

$$S_{ij}(\tau) = \int d\tau' \, K(\tau - \tau')[H_{ij}(\tau') + X_{ij}(\tau')] \, . \tag{11}$$

Hint: $S_{ij}(\tau)$ should depend on the membrane time constant, τ, that appears in Eq. (4a).

(5 marks)

4. Show that Fourier transforming both sides diagonalizes the convolution operator,

$$S_{ij}(\omega) = K(\omega)[H_{ij}(\omega) + X_{ij}(\omega)] \tag{12}$$

where

$$S_{ij}(\omega) \equiv \int d\tau \, e^{i\omega\tau} S_{ij}(\tau), \tag{13}$$

and similarly for $K(\tau)$, $H_{ij}(\omega)$ and $X_{ij}(\omega)$. We are using bad physicist's notation, in which the argument of the function (τ versus ω) determines what function it is.

(5 marks)

5. Let's drop the explicit dependence on ω (it's still there; we'll just ignore it), and use Eq. (10) to rewrite Eq. (14) as

$$S_{ij} = K \left[\sum_{lm} w_{il} S_{lm} w^T_{mj} + X_{ij} \right] . \tag{14}$$

Note that K is just a number, not a function. This is a linear equation for S. However, it's an n^2-dimensional linear equation, which makes it hard to solve. We will, therefore, use a mean field approach. Our strategy is to consider separately the diagonal and off-diagonal elements of S_{ij}. The diagonal elements must be $\mathcal{O}(1)$, as they correspond to the variance, so what we're after is the size of the off-diagonal elements. We thus separate the diagonal and off-diagonal terms in Eq. (14), leading to

$$S_{ij} = K \left[\sum_{l} w_{il} w_{jl} S_{ll} + \sum_{l \neq m} w_{il} w_{jm} S_{lm} + X_{ii}\delta_{ij} \right] \tag{15}$$

where we used the fact that X_{ij} is diagonal.

Assume that the weights are drawn *iid* (so the network doesn't obey Dale's law), with mean zero and variance σ_w^2/n,

$$\text{Mean}[w_{ij}] = 0 \tag{16a}$$

$$\text{Var}[w_{ij}] = \frac{\sigma_w^2}{n}\,. \tag{16b}$$

Assuming that the weights and S_{lm} are uncorrelated, and n is large, derive the following relationships. (Note that equality holds only as $n \to \infty$.)

$$\sum_l w_{il}^2 S_{ll} = \sigma_w^2 \text{Mean}[S_{ll}] \tag{17a}$$

$$\sum_l w_{il} w_{jl} S_{ll} = \frac{\sigma_w^2}{n^{1/2}} \text{Mean}[S_{ll}^2]^{1/2} \xi_{ij} \tag{17b}$$

$$\sum_{l \neq m} w_{il} w_{jm} S_{lm} = \sigma_w^2 \text{Mean}[S_{lm}^2]^{1/2} \eta_{ij} \tag{17c}$$

where the ξ_{ij} and η_{ij} are a set of uncorrelated zero mean, unit variance Gaussian random variables, $i \neq j$ in Eq. (17b), and the moments are given by

$$\text{Mean}[S_{ll}^k] \equiv \frac{1}{n} \sum_l S_{ll}^k \tag{18a}$$

$$\text{Mean}[S_{lm}^k] \equiv \frac{1}{n(n-1)} \sum_{l \neq m} S_{lm}^k. \tag{18b}$$

Here k is a power, not a superscript.

(10 marks)

6. Combine Eqs. (15) and (17) to show that, in the large n limit,

$$\text{Mean}[S_{ll}] = \frac{K \text{Mean}[X_{ll}]}{1 - K\sigma_w^2} \tag{19a}$$

$$\text{Mean}[S_{ll}^2] = K^2 \left[\frac{\text{Mean}[X_{ll}]^2}{(1 - K\sigma_w^2)^2} + \text{Var}\big[X_{ll}\big] \right] \tag{19b}$$

$$\text{Mean}[S_{lm}^2] = \frac{1}{n} \frac{K^2 \sigma_w^4 \text{Mean}[S_{ll}^2]}{1 - K^2 \sigma_w^4}. \tag{19c}$$

where $\text{Var}\big[X_{ll}\big]$ is the variance of X_{ll}, $\text{Var}\big[X_{ll}\big] = \text{Mean}\big[X_{ll}^2\big] - \text{Mean}\big[X_{ll}\big]^2$. This result implies that correlations among neurons are very weak in large networks that don't obey Dale's law.

(12 marks)

3 Coding

Consider an experiment in which a trial starts at time 0, and information arrives in the form of the time, t, at which an event (e.g., a spike) occurs.

First, consider the case in which this time follows gamma distributions with shape k and rate λ

$$p(t;k,\lambda) = \frac{\lambda^k}{\Gamma(k)} t^{k-1} \exp(-\lambda t)$$

Consider the case that there are two possible distributions, one with parameters (k_1, λ_1) and the other with parameters (k_2, λ_2). They are equally likely, *a priori*.

1. What is the form of a Bayesian test between the two distributions based on an observation of the time, t? (5 marks)
2. A lazy observer has decided upon a decision strategy of waiting until time T, and reporting one distribution if the spike occurred before that time, and the other distribution otherwise. Are there values of (k_1, λ_1) and (k_2, λ_2) for which this is optimal? Be as general as you can. (10 marks)
3. Consider exponential distributions, $k_1 = k_2 = 1$, and, for definiteness, let $\lambda_2 = 2\lambda_1$. If waiting costs £k/s and being incorrect costs £e (being correct costs and gains nothing), what time, T, would the lazy observer choose to minimize her expected cost? (10 marks)

Second, consider the slightly different case in which there is just one single Gamma distribution, but now information about an input x is conveyed in the *rate* parameter $\lambda(x)$ (the shape k is fixed).

4. What is the Fisher information in t about x (give your answer in terms of $\lambda(x)$ and its derivatives)? (10 marks)
5. Why does this depend on k in the way it does? (5 marks)

4 Decision-Making

A rat in Montreal has to decide how to allocate its time between working for brain stimulation reward and taking leisure. Say that it gets a reward r if it works for time P.

The labour supply theory view suggests a utility function such as

$$U(w, l) = (r(P/w)^s + k_l l^s)^{\frac{1}{s}}$$

where w is the time it works (thus getting P/w rewards), l the time it lazes, k_l the scaling of the benefit of leisure, and s is involved in substitutability.

1. Sketch a contour map of U (i.e., lines of equal utility) for $s = 1$, $s = 0.5$ and $s = 2$. (5 marks)
2. The animal has a fixed budget of time T to be allocated between work and sloth. Show the budget line on the contour maps from (1) - what relationship will there be between this line and the contours if the animal optimises? (5 marks)
3. What is the optimal time allocation in each of these three cases? Provide formulæ. (5 marks)
4. Take the case that $s = 0.5$. The experimenter increases the price of each reward to $2P$, but compensates the animal by providing some free rewards (for which it doesn't have to work) such that the solution from (3) remains possible in the budget. Will the animal choose to allocate its time in just the same way as this solution? If not, how will it change? (5 marks)

A microscopic view of this works differently. Imagine the rat working for time P for reward r, then choosing a leisure time l from some distribution $\pi_l(l)$, getting reward $k_l \times l$ for the choice, then working for P again, and so forth forever.

5. Write down an expression for the long-run average reward rate ρ^π. (8 marks)
6. If the animal optimizes
$$\rho^\pi + \frac{1}{\beta}\mathcal{H}[\pi]$$
with respect to the distribution π, where β is an inverse temperature, and $\mathcal{H}[\pi]$ is the entropy of the distribution π, what is its optimal policy? (12 marks)