Gatsby Computational Neuroscience Unit
Theoretical Neuroscience

Final Examination
26 Apr 2018

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 the simplest neuron in the world,

$$\tau \frac{dV}{dt} = -\frac{V[(V-V_r)^2 - \mu^2]}{2\mu(\mu - V_r)} \tag{1}$$

where $\mu = 10$ mV and $V_r =$ -60 mV. The denominator gives the right hand side units of voltage.

1. This "neuron" has three fixed points. Write down the linearized the dynamics around each of them, and show explicitly that two are stable and one is unstable.

   (10 marks)

   Solution

   The three fixed points occur at $V = 0$ and $V = V_r \pm \mu$. The resulting linearized equations are

$$\begin{aligned}
V \text{ near } V_r - \mu: \quad & \tau \frac{dV}{dt} = -\big(V - (V_r - \mu)\big) \\
V \text{ near } V_r + \mu: \quad & \tau \frac{dV}{dt} = \frac{V_r + \mu}{V_r - \mu}\big(V - (V_r + \mu)\big) = \frac{5}{7}\big(V - (V_r + \mu)\big) \\
V \text{ near } 0: \quad & \tau \frac{dV}{dt} = \frac{V_r + \mu}{2\mu} V = -\frac{5}{2} V\,.
\end{aligned}$$

   The equilibria at $V = 0$ and $V_r - \mu$ are stable; the one at $V_r + \mu$ is unstable.

2. Equation (1) has some of the flavor of a real neuron: it has a reasonable resting membrane potential, and when the voltage is above a threshold it it increases rapidly. The problem, however, is that the voltage never goes back to rest. To remedy this, you add a current, yielding the set of equations

$$\begin{aligned}
\tau \frac{dV}{dt} &= -\frac{V[(V-V_r)^2 - \mu^2]}{2\mu(\mu - V_r)} - g_0 x (V - \mathcal{E}) \\
\tau_x \frac{dx}{dt} &= -(x - x_\infty(V))\,.
\end{aligned}$$

   Choose $\mathcal{E}$, $x_\infty(V)$ and τ_x so that the neuron can still be at rest (it has a stable equilibrium near V_r), but a brief, sufficiently strong, current pulse can make it spike: the membrane potential can increase rapidly toward $V = 0$, then go back to rest. Nullclines are not necessary, but you do need to justify your answer.

   (15 marks)

   Solution

   The added current should have the following properties. First, x should turn on when the voltage is high, so that the current is active only during a spike; I used

$$x_\infty(V) = \frac{e^{(V+20)/5}}{1 + e^{(V+20)/5}}\,,$$

and set g_0 to 0.4. Second, $\mathcal{E}$ must be low (below V_r), so that the voltage is pulled down to rest; I chose $\mathcal{E} = -80$ mV. (Exact numbers aren't so important; I picked these so that my favorite draw program could make figures for me.) Interestingly, τ_x must be somewhat large compared to τ; otherwise, as the voltage decreases, the current will shut off. This means spikes will be slow – not exactly what you want in a neuron. A sketch of $x_\infty(V)$, along with $\tau dV/dt$ evaluated at $x = 0$ (the added conductance is off) and $x = 1$ (the added conductance is on), is shown here,

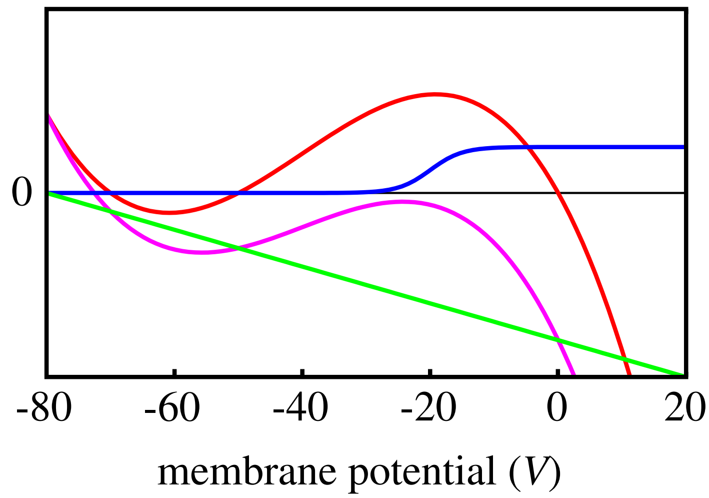

The red and magenta curves are dV/dt evaluated at $x = 0$ and $x = 1$, respectively; the blue curve is $x_\infty(V)$ and the green line is $-(V - \mathcal{E}) = -(V + 80)$. This has the desired properties: if the voltage can manage to get above the unstable fixed point associated with the red curve, it will quickly jump to $V = 0$; at that point, the added conductance will turn on (x will increase), the fixed point near $V = 0$ will disappear, and the voltage will drop down near rest.

3. While you have improved your neuron, it still can't spike repetitively. Show that if you inject sufficiently large positive current, it can. However, if the current is too large, there will be only one fixed point at high voltage. Write expressions for upper and lower bound on the current. Do not try to solve them – there are no exact analytic solutions.

   (15 marks)

   Solution

   Here's the picture we have so far (without injected current): if the neuron exceeds threshold, it moves to a fixed point near $V = 0$. When that happen, x increases until the fixed point near $V = 0$ disappears, at which point the voltage drops again, to near V_r. When it drops to V_r, x decreases, but the fixed point at V_r remains.

   To remove that fixed point, the current, denoted I, must be high enough to remove the equilibrium near V_r when x decreases. This happens if

$$-\frac{V[(V - V_r)^2 - \mu^2]}{2\mu(\mu - V_r)} - g_0 x_\infty(V)(V - \mathcal{E}) + I$$

has a single zero near $V = 0$. It's OK to use $x_\infty(V)$ in this expression because we want to remove fixed points.

For the neuron to drop from near $V = 0$ to low voltage, the current can't be too high; otherwise, the fixed point near $V = 0$ won't disappear. We must must, therefore, choose I sufficiently small that

$$-\frac{V[(V-V_r)^2-\mu^2]}{2\mu(\mu-V_r)} - g_0 x_\infty(V)(V-\mathcal{E}) + I$$

has a single zero near $V = V_r$. Again, it's OK to use $x_\infty(V)$ because we want to remove fixed points.

All this should be explained in the following picture,

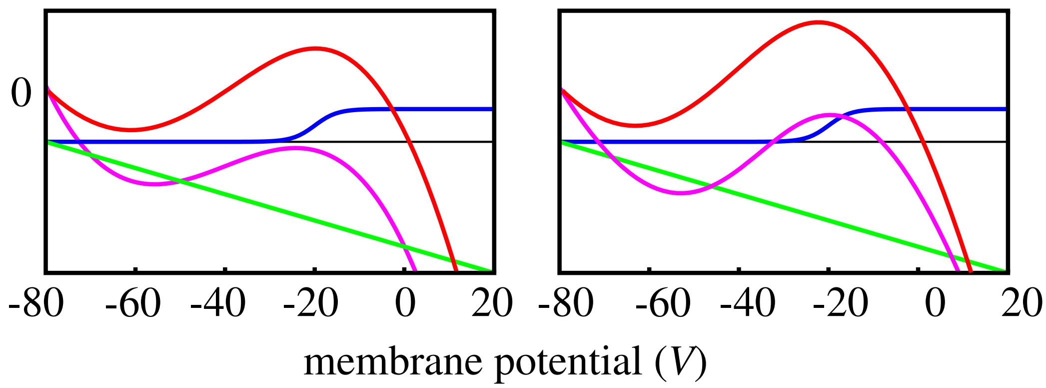

The color code is the same as it was above: the red and magenta curves are dV/dt evaluated at $x = 0$ and $x = 1$, respectively; the blue curve is $x_\infty(V)$ and the green line is $-(V - \mathcal{E}) = -(V + 80)$. These were drawn freehand to illustrate the ideas, so they don't correspond to any particular parameters.

In the left panel, the input current is just right: when $x = 0$ there's a fixed point only when V is near 0 and when $x = 1$ there's a fixed point only when V is near rest.

In the right panel, on the other hand, the input current is too high: even when $x = 1$, the fixed point near $V = 0$ doesn't disappear. Thus, the membrane potential will be stuck near $V = 0$. This is bad for a neuron: it will quickly die.

2 Networks

Consider an all-inhibitory rate network with one Hopfield-like memory,

$$\frac{d\nu_i}{dt} = \phi\left(\sqrt{n}I - \frac{1}{\sqrt{n}}\sum_{j=1}^{n} w_{ij}\nu_j + \frac{\beta}{nf(1-f)}\sum_{j=1}^{n}\xi_i(\xi_j - f)\nu_j\right) - \nu_i \tag{4}$$

where ϕ is the gain function (typically sigmoidal, and definitely monotonic increasing), n is the number of neurons, I and β are constant, the elements of w_{ij} are drawn *i.i.d* from a distribution with mean 1 and variance σ^2, and ξ_i is a random binary vector,

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

It is convenient to define

$$m \equiv \frac{1}{nf(1-f)}\sum_{j=1}^{n}(\xi_j - f)\nu_j,$$

so that the equation for ν_i can be written

$$\frac{d\nu_i}{dt} = \phi\left(\sqrt{n}I - \frac{1}{\sqrt{n}}\sum_{j=1}^{n} w_{ij}\nu_j + \beta\xi_i m\right) - \nu_i\,.$$

1. Show that Eq. (4) can be written

   $$\frac{d\nu_i}{dt} = \phi\left(\sqrt{n}(I-\nu) + \beta\xi_i m + \eta_i\right) - \nu_i$$

   where ν is the population averaged firing rate,

   $$\nu \equiv \frac{1}{n}\sum_{j=1}^{n}\nu_j$$

   and η_i is a zero mean Gaussian random variable with variance given by

   $$\mathrm{Var}[\eta] = \frac{\sigma^2}{n}\sum_j \nu_j^2\,.$$

   Don't forget that the mean value of w_{ij} is 1!
   (5 marks)

   Solution

   This is a standard calculation: using $w_{ij} = 1 + \delta w_{ij}$, we have

   $$\frac{1}{\sqrt{n}}\sum_j w_{ij}\nu_j = \frac{1}{\sqrt{n}}\sum_j \nu_j + \frac{1}{\sqrt{n}}\sum_j \delta w_{ij}\nu_j\,,$$

The first term on the right hand side is just $\sqrt{n}\nu$. The second term has variance given by

$$\text{Var}\left[\frac{1}{\sqrt{n}}\sum_j \delta w_{ij}\nu_j\right] = \frac{1}{n}\sum_i \frac{1}{n}\sum_{jj'} \delta w_{ij}\delta w_{ij'}\nu_j\nu_{j'} = \frac{1}{n}\sum_{jj'} \nu_j\nu_{j'}\frac{1}{n}\sum_i \delta w_{ij}\delta w_{ij'}\,.$$

The sum over i is equal to $\sigma^2\delta_{jj}$; the sum over j and j' then collapses to a single sum over j, which gives us our desired result.

For the remainder of this question, use $\sigma^2 = 0$.

2. Show that ν and m obey the equations

$$\frac{d\nu}{dt} = \phi(\sqrt{n}(I-\nu)) + f[\phi(\sqrt{n}(I-\nu)+\beta m) - \phi(\sqrt{n}(I-\nu))] - \nu \tag{5a}$$
$$\frac{dm}{dt} = \phi(\sqrt{n}(I-\nu)+\beta m) - \phi(\sqrt{n}(I-\nu)) - m\,. \tag{5b}$$

(10 marks)

Solution

The equation for the mean firing rate, ν, is

$$\frac{d\nu}{dt} = \frac{1}{n}\sum_j \phi\left(\sqrt{n}(I-\nu)+\beta\xi_i m\right) - \nu\,.$$

The random variable is ξ_i is 1 with probability f and zero with probability $1-f$. This equation thus becomes, in the large n limit,

$$\frac{d\nu}{dt} = (1-f)\phi\left(\sqrt{n}(I-\nu)\right) + f\phi\left(\sqrt{n}(I-\nu)+\beta m\right) - \nu\,.$$

Rearranging terms produces Eq. (5a).

The equation for the overlap, m, is

$$\frac{dm}{dt} = \frac{1}{nf(1-f)}\sum_j (\xi_j - f)\phi\left(\sqrt{n}(I-\nu)+\beta\xi_i m\right) - m\,.$$

The random variable $\xi_i - f$ is $1-f$ with probability f and $-f$ with probability $1-f$; we thus have

$$\begin{aligned}\frac{dm}{dt} &= \frac{1}{f(1-f)}[f(1-f)\phi\left(\sqrt{n}(I-\nu)+\beta m\right) - (1-f)f\phi\left(\sqrt{n}(I-\nu)+\beta m\right) - m \\ &= \phi\left(\sqrt{n}(I-\nu)+\beta m\right) - \phi\left(\sqrt{n}(I-\nu)+\beta m\right) - m\,.\end{aligned}$$

This is the same as Eq. (5b).

3. Consider first the limit $f \to 0$. In that limit, ν decouples from m, and we can consider only the equation for m. Assume that $\phi(x)$ is convex ($\phi''(x) > 0$) for $x < x_0$ and concave ($\phi''(x) < 0$) for $x > x_0$. Show that if $\sqrt{n}(I-\nu) < x_0$, then:

- For β sufficiently small, Eq. (5b) will have only one equilibrium, at $m = 0$, and it will be stable.
- As β increases, two new equilibria will appear at positive m **before** the equilibrium at $m = 0$ becomes unstable.
- Show that the stable equilibrium at positive m must occur at $\beta m > x_0 - \sqrt{n}(I - \nu)$. You can use a plot to demonstrate this.

(10 marks)

Solution

This figure

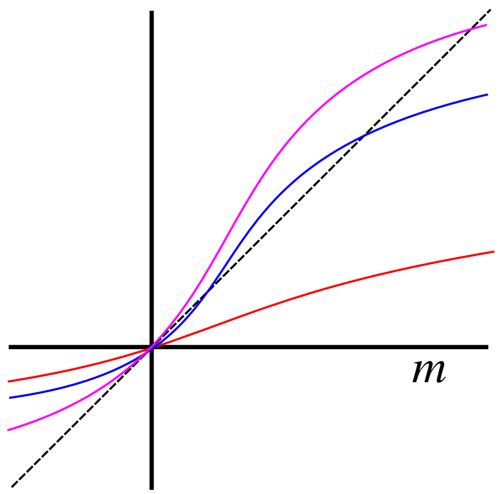

shows a plot of $\phi(\sqrt{n}(I - \nu) + \beta m) - \phi(\sqrt{n}(I - \nu))$ (the y-axis, which is unlabeled) versus m (the x-axis) for various values of β. The red curve has the smallest β; the blue curve the next smallest, and the magenta curve the largest. Equilibria occur at the intersection of the curves with the 45° line (the dashed black line), and the slope determines stability: if it's greater than 1 the fixed point is unstable; if it's less than 1 the fixed point is stable.

It's clear that when β is sufficiently small, the only fixed point is at $m = 0$, and it's stable (as indicated by the red curve). That fixed point becomes unstable when $\beta\phi'(\sqrt{n}(I - \nu)) = 1$ (magenta curve). However, because ϕ is convex here (remember, $x_0 > \sqrt{n}(I - \nu)$, which implies that ϕ is convex when evaluated at $\sqrt{n}(I - \nu)$), it's clear that there must have been another pair of equilibria occurring at positive m at a lower value of β (such as the blue curve).

Because the curves don't become concave until the argument of ϕ is greater than x_0, and the highest fixed point occurs when ϕ is convex, we must have $\beta m + \sqrt{n}(I - \nu) > 0$ at the highest fixed points.

4. In the large n limit, ν doesn't change much – it must be within $1/\sqrt{n}$ of I. This does not, however, imply that ν and m decouple. Explain why.

   (15 marks)

   Solution

   The key idea is that $\sqrt{n}(I-\nu)$ can change by an order 1 amount. This is kind of obvious: changes in $I-\nu$ are $\mathcal{O}(1/\sqrt{n})$; multiply that by $\sqrt{n}$ and you get back to $\mathcal{O}(1)$ changes. Thus, instead of considering activity in ν-m space, we consider activity in h-m space, where

   $$h \equiv \sqrt{n}(I-\nu). \tag{6}$$

   To find the behavior of h, use Eq. (5a)

   $$\frac{dh}{dt} = -h - \sqrt{n}[(1-f)\phi(h+\beta m) + f\phi(h) - I]\,. \tag{7}$$

   Because of the factor of $\sqrt{n}$, we can approximate the h-nullcline by

   $$I = (1-f)\phi(h+\beta m) + f\phi(h)\,. \tag{8}$$

   This is coupled to the m-nullcline,

   $$m = \phi(h+\beta m) - \phi(h)\,. \tag{9}$$

   These two equations can be solved simultaneously to determine the behavior when $f \neq 0$.

3 Coding

Consider a cell whose firing (viewed as a point process) in the absence of stimulation can be modelled by a conditional intensity function:

$$\lambda(t|t_n, t_{n-1}, ..., t_1) = \lambda(t|t_n) = \nu^2(t-t_n).$$

That is, the intensity depends only on the time of the most recent spike, and increases linearly from that time.

1. What are the dimensions of the constant ν? (2 marks)

   ANSWER: $(\text{time})^{-1}$

2. Assuming the model is correct, what is the cell's ISI distribution? (8 marks)

   ANSWER:

   $$\begin{aligned} p(\tau) &= \lambda(t_n+\tau)\exp\left(-\int_{t_n}^{t_n+\tau} dt\,\lambda(t)\right) \\ &= \nu^2\tau\exp\left(-\int_0^{\tau} d\tau'\,\nu^2\tau'\right) \\ &= \nu^2\tau e^{-\nu^2\tau^2/2} \end{aligned}$$

3. Show that the mean firing rate under this model is $\overline{\lambda} = \sqrt{\frac{2}{\pi}}\nu$ (assuming $\nu \geq 0$). (5 marks)

   ANSWER: Mean ISI:

$$\begin{aligned}
\int_0^\infty d\tau\ \tau p(\tau) &= \int_0^\infty d\tau\ \nu^2\tau^2 e^{-\nu^2\tau^2/2} \\
&= \tfrac{1}{2}\sqrt{2\pi}\nu \int_{-\infty}^{\infty} d\tau\ \tau^2 \frac{\nu}{\sqrt{2\pi}} e^{-\nu^2\tau^2/2} \\
&= \tfrac{1}{2}\sqrt{2\pi}\nu \frac{1}{\nu^2} \\
&= \sqrt{\frac{\pi}{2}}\frac{1}{\nu}
\end{aligned}$$

   So the mean firing rate is

$$\overline{\lambda} = \sqrt{\frac{2}{\pi}}\nu$$

4. Assume that we started recording from the cell (i.e. $t = 0$) a random time after the stimulus was turned off. Consider a small interval of time $[t, t + dt)$. What is the entropy H_t of the spiking process in this interval? (5 marks)

   ANSWER: A spike appears in the interval with probability $\overline{\lambda}dt$. Thus the entropy is

$$\begin{aligned}
H_t &= -\overline{\lambda}dt \log(\overline{\lambda}dt) - (1 - \overline{\lambda}dt)\log(1 - \overline{\lambda}dt) \\
&= \overline{\lambda}(1 - \log(\overline{\lambda}))dt
\end{aligned}$$

5. What is the entropy rate $\mathcal{H}$ of the process? [Treat it as a binary discrete-time stochastic process in discretised bins of size dt. You may leave the answer in integral form.] (5 marks)

Now, suppose that a stimulus $s(t)$ modulates the cell's intensity, so that

$$\lambda(t|s(t), t_n, t_{n-1}, ..., t_1) = \lambda(t|s(t), t_n) = (\nu s(t))^2(t - t_n)$$

and that $s(t)$ is white noise with power $\left\langle s(t)^2 \right\rangle = \rho^2$.

6. What is the new ISI distribution (collected over an infinitely long period of stimulation)? (8 marks)

7. What is the spike-triggered-average stimulus (again, based on an infinitely long stimulus sequence)? (7 marks)

4 Learning

A rat runs left to right on a linear track with constant velocity a. The i^{th} neuron in the hippocampus acts as an asymmetric exponential place cell with Poisson activity r_i of mean

$$\mu_i = \begin{cases} e^{-\lambda(x-\xi_i)} & \text{if } x > \xi_i \\ 0 & \text{otherwise} \end{cases}$$

when it is at location x. Here ξ_i are the place field modes, which you may assume to be suitably dense that each x is evenly covered, and λ is a space constant. The neurons are independent.

1. given a uniform prior, what is the posterior density $p[x|\boldsymbol{r}]$ and what estimate of x would minimize the squared prediction error? Hint: think about the allowable values of x given some collection of active neurons. (10 marks)

2. these neurons are connected with a temporally-inverted anti-Hebbian time-dependent learning rule such that the connection from neuron j to neuron i is

   $$\Delta W_{ij} = -\epsilon W_{ij} - \epsilon \int_t dt\; r_i(x(t-\delta t)) r_j(x(t)) dt$$

   where ϵ is a learning rate and $\delta t > 0$. Assuming that the rat learns from many identical left-right-sweeps (at velocity a) and you don't have to capture the effect that its activity is changing during learning, derive an expression for the asymptotic weight from neuron j to neuron i. (10 marks)

3. if the neurons are weakly recurrently connected:

   $$\tau \frac{d}{dt} \nu_i = -\nu_i + \mu_i + \eta \sum_j W_{ij} \nu_j$$

   where $0 < \eta \ll 1$, sketch the effect that these inputs will have on the population activity. (8 marks)

4. Discuss the effect that these inputs will have on the true posterior density, and on decoding if the rat performs decoding algorithmically as in equation 1 (without taking account of the change in activity). An approximate answer is fine. (8 marks)

5. How would inference change if the connections were forward and Hebbian instead of backwards and anti-Hebbian? (4 marks)