Gatsby Computational Neuroscience Unit
Theoretical Neuroscience

Final examination, theoretical neuroscience
11 May 2020

Part I – short questions

There are four sections with four questions each. Please answer three out of each four, starting the answers for each section on a new page. Don't forget to write your name at the top of each block of answers.

You have a maximum of 7 hours for this exam.

Good luck!

1 Biophysics

1. Consider the usual Hodgkin Huxley equations,

$$\begin{aligned}
\tau \frac{dV}{dt} &= -(V - \mathcal{E}_L) - g_{Na} m^3 h (V - \mathcal{E}_{Na}) - g_K n^4 (V - \mathcal{E}_K) + V_{\text{ext}} \\
\tau_m \frac{dm}{dt} &= -(m - m_\infty(V)) \\
\tau_h \frac{dh}{dt} &= -(h - h_\infty(V)) \\
\tau_n \frac{dn}{dt} &= -(n - n_\infty(V)) \,.
\end{aligned}$$

   Show that in the limit $\tau_m, \tau_h, \tau_n \to 0$, the neuron can *not* exhibit repetitive spiking.

2. A simple model for synaptic adaptation has the form

$$\tau \frac{dx}{dt} = 1 - x - \alpha \sum_i \delta(t - t_i)$$

   where x is the normalized synaptic strength, the t_i are presynaptic spike times, δ is the Dirac δ-function, and α is a constant. Write down an exact expression for $x(t)$ in terms of the t_i. Why does this model break down when α is large? How could you modify the equations to fix that?

3. Consider an STDP model in which the weight change, Δw, at a synapse in response to a presynaptic spike at time t_i and a postsynaptic spike at time t_j is given by

$$\Delta w = \begin{cases} -\eta & t_i < t_j \text{ and } t_j - t_i < \tau \\ \eta & t_j < t_i \text{ and } t_i - t_j < \tau \,. \end{cases}$$

   In other words, it's STDP rule, but with square kernels rather than the usual exponentially decaying ones.

   Assume that the pre and postsynaptic neurons are firing with Poisson statistics at rates ν_{pre} and ν_{post}, respectively, chosen small enough that both $\tau\nu_{\text{pre}} \ll 1$ and $\tau\nu_{\text{post}} \ll 1$. What's the typical size of w after a (long) time T? Assume $w(t = 0) = 0$.

4. Suppose you have a Hodgkin Huxley neuron, to which you add a sodium current that slowly inactivates at high voltage,

$$\begin{aligned}
\tau \frac{dV}{dt} &= (\text{usual HH; see problem 1}) - x(V - \mathcal{E}_{Na}) \\
\tau_x \frac{dx}{dt} &= \frac{1}{1 + \exp((V + 50)/10)} - x
\end{aligned}$$

   where $\mathcal{E}_{Na} = 0$ mV and voltage, V, is measured in mV. Here "slowly" means τ_x is on the order of 100s of ms whereas τ is on the order of 10 ms.

   Can this additional sodium current lead to bursting? Explain your answer.

Solutions

1. In the limit that all three time constants go to 0, the left hand sides of the equations for m, h and n are zero. Consequently, the channel variables can be replaced by their asymptotic value ($m(t) = m_\infty(V)$, and the same for h and n). In that limit, the equation for V looks like

$$\tau \frac{dV}{dt} = f(V)$$

where $f(V)$ is a smooth function of V. That equation has fixed points, but it can't exhibit any interesting behavior. The reason may be obvious, but if not this should help: let's say $f(V)$ is positive for some V, the membrane potential will increase as long $f(V)$ is positive. Eventually $f(V)$ will go negative. Assuming $f(V_0) = 0$, we can linearize around V_0, leading to $\tau dV/dt = -\alpha(V - V_0)$ with $\alpha > 0$. The voltage will then relax exponentially fast to $V - 0$. For quadratic, cubic, etc. minima the analysis is a bit more complicated, but the conclusion is the same: V can only approach fixed points (or $\pm\infty$).

2. The equation has the general solution

$$x(t) = x(t_0)e^{-(t-t_0)/\tau} + \int_{t_0}^{t} \frac{dt'}{\tau}\, e^{-(t-t')/\tau} \left(1 - \alpha \sum_i \delta(t' - t_i)\right) .$$

It is convenient to let $t_0 \to -\infty$. In that case the first term vanishes and the first integral can be done analytically. And so can the second,

$$x(t) = 1 - \frac{\alpha}{\tau} \sum_{t_i < t} e^{-(t-t_i)/\tau} .$$

If α is too large, x can become negative. This would mean the synaptic strength changes sign, which it doesn't do.

An easy – and not unrealistic – way to fix this is to reduce the jump in x when x is small. For instance,

$$\tau \frac{dx}{dt} = 1 - x - x\alpha \sum_i \delta(t - t_i)$$

with the convention that when a spike arrives, x on the right hand side is evaluated immediately before the spike. Now when there's a spike, x drops by $x\alpha/\tau$. Put another way, when there's a spike

$$x \to x\left(1 - \frac{\alpha}{\tau}\right) .$$

So long as $\alpha < \tau$, x can never become negative.

3. By symmetry, the average increase in weight is equal to the average decrease, so the average weight change is zero. What we need to do, therefore, is estimate the variance, of w, as follows. That can be cone as follows:

Assume there are N postsynaptic spikes. N is, of course, a random number; later we're going to average over it. That means there are N windows of length τ before the postsynaptic spikes and N windows after the spikes. Thus the total length of time for both potentiation and depression is $N\tau$. Because τ is small, the potentiation and depression windows rarely overlap, so we can treat them as independent random variables. Because the spikes are Poisson distributed, the variance for both potentiation and depression is $N\tau\nu_{\rm pre}$, leading to a total variance of $2N\tau\nu_{\rm pre}$. To get the variance of w we need to average over N; that average is $T\nu_{\rm post}$. Thus, the variance of w is $2\tau T\nu_{\rm pre}\nu_{\rm post}$. The typical size of w is the square root of that.

4. For a cell to burst, it needs to slowly become more excitable when it's not firing and slowly become less excitable when it is firing. The x-channel clearly satisfies the former: at rest, around -65 mV, it turns on, and raises the membrane potential. When the neuron starts spiking, it sees a higher membrane potential on average, and starts to turn off. The one worry is that it turns off after one spike. However, because the x-channel activates slowly, by the time the neuron spikes both the h and n currents will be partially activated, and the neuron will be well above threshold. It will take a few spikes for these currents to inactivate. Thus, spiking should start when the neuron is sufficiently excitable that it will generate sustained firing until the x channel inactivates.

2 Networks

1. Consider a network of excitatory and inhibitory neurons in which the average firing rates evolve according to

$$\tau \frac{d\nu_E}{dt} = \phi_E\big(J_{EE}\,\nu_E - J_{EI}\nu_I + h_E\big) - \nu_E$$
$$\tau \frac{d\nu_I}{dt} = \phi_I\big(J_{IE}\,\nu_E - J_{II}\nu_I + h_I\big) - \nu_I$$

   Both gain functions, ϕ_E and ϕ_I are monotonic increasing. Assume the network is operating in the balanced regime, meaning its equilibrium is on the unstable branch of the excitatory nullclines. Suppose the average connection strength among the inhibitory neurons, J_{II}, increases by a small amount.

   (a) Does the average inhibitory firing rate go up or down?

   (b) Does the average excitatory firing rate go up or down?

   Justify both answers.

2. Consider a large network of excitatory and inhibitory neurons operating in the balanced regime (i.e., on the unstable branch of the excitatory nullcline). Is it possible to have a stable fixed point if the time constant (rise and decay times) of the inhibitory synapses are very long? Justify your answer, either qualitatively or quantitatively.

3. Consider a 1-memory Hopfield network with analog rates,

$$x_i(t+1) = \phi\left(\frac{1}{f(1-f)}\frac{1}{N}\sum_{j=1}^{N}\eta_i(\eta_j - f)x_j(t)\right)$$

   where ϕ is sigmoidal and η_i is a random binary vector,

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

   Show that in the large N limit, the steady state solution to this equation is $x_i = \phi(0) + b\eta_i$. Write down an equation that can be solved (usually numerically) for b. Show that b can be zero, but that solution is unstable if $|\phi'(0)| > 1$.

4. Consider a recurrent neural network,

$$\frac{d\mathbf{x}}{dt} = \phi\left(\sum_{k=1}^{K}\mathbf{u}_k\mathbf{v}_k\cdot\mathbf{x}\right) - \mathbf{x}$$

   where $\mathbf{x}$ is an n-dimensional vector, $K < n$, and ϕ is a monotonic increasing pointwise nonlinearity. Show that if the dynamics is periodic, the trajectories relax to a manifold of dimension K. Argue that if the dynamics is chaotic, the trajectories fill the whole space.

Solutions

1. Increasing J_{II} lowers the inhibitory nullcline. Intuitively, that's because at fixed excitatory firing rate, more inhibitory to inhibitory coupling should lower inhibitory firing rate. It's also easy to show: let $J_{II} \to J_{II} + \delta J_{II}$, and fix ν_E. That induces a change in firing rate

$$\delta\nu_I = -\frac{\phi'\big(J_{IE}\,\nu_E - J_{II}\nu_I + h_I\big)\nu_I}{1 + \phi'\big(J_{IE}\,\nu_E - J_{II}\nu_I + h_I\big)}\,\delta J_{II}\,.$$

   Because ϕ is monotonic increasing, the right hand side is negative if δJ_{II} is positive.

   However, because the inhibitory nullcline intersects the excitatory one from below, with positive slope, lowering the inhibitory nullcline leads to an increase in both excitatory and inhibitory firing rates.

2. The answer is no: the system will oscillate. That's because when the excitatory firing rate rises, inhibition will also rise, but with a long lag. Eventually it will cause the excitatory firing rate to fall, which will drive inhibition down, but again with a long lag. This cycle will repeat forever. More formally, the linearized dynamics looks something like

$$\frac{d\nu_E}{dt} = \frac{W_{EE}\nu_E - W_{EI}\nu_I}{\tau_E}$$
$$\frac{d\nu_I}{dt} = \frac{W_{IE}\nu_E - W_{II}\nu_I}{\tau_I}$$

   for some coefficients $\mathbf{W}$. Because we're in the balanced regime, the determinant of the matrix on the right hand side is positive, independent of the time constants. However, the trace, which must be negative for stability, depends on the time constants,

$$\text{Trace} = \frac{W_{EE}}{\tau_E} - \frac{W_{II}}{\tau_I}\,.$$

   If τ_I is too large, the trace will become positive, and the system will go unstable via a Hopf bifurcation.

3. Because we're in steady state, we can drop the time dependence. The key observation is that x_i can take on only two values, depending on the value of η_i: if $\eta_i = 0$, $x_i = \phi(0)$ and if $\eta_i = 1$, x_i is a different number, which we'll call (for later convenience) $b + \phi(0)$. To complete the calculation, we just need to find b. We start by summarizing the above by writing

$$x_i = (1 - \eta_i)\phi(0) + \eta_i(b + \phi(0))\,.$$

   We chose this form so that $x_i = \phi(0) + \eta_i b$. To find b, we note first that

$$\frac{1}{f(1-f)}\frac{1}{N}\sum_{j=1}^{N}(\eta_j - f)x_j = \frac{1}{f(1-f)}\frac{1}{N}\sum_{j=1}^{N}(\eta_j - f)(\phi(0) + \eta_j b) \to b\,,$$

   valid in the large N limit. We can use this to find b self-consistently; when $\eta_i = 1$,

$$b + \phi(0) = \phi\left(\frac{1}{f(1-f)}\frac{1}{N}\sum_{j=1}^{N}(\eta_j - f)x_j\right) = \phi(b) \quad\Longrightarrow\quad b = \phi(b) - \phi(0)\,.$$

   One solution is $b = 0$. However, if $\phi'(0) > 0$, that solution is unstable: small increases in b will cause the mean value of x_i to rise, leading to an instability.

4. Define $z_k \equiv \mathbf{v}_k \cdot \mathbf{x}$. The differential equation the becomes

$$\frac{dz_k}{dt} = \mathbf{v}_k \cdot \phi\left(\sum_{l=1}^{K}\mathbf{u}_l z_l\right) - z_k\,.$$

   Now define another $n - K$ variables, $y_\mu = \mathbf{v}_\mu \cdot \mathbf{x}$ where the $\mathbf{v}_\mu$ span the space orthogonal to of the $\mathbf{v}_k$. The equation for the y_μ is

$$\frac{dy_\mu}{dt} = \mathbf{v}_\mu \cdot \phi\left(\sum_{l=1}^{K}\mathbf{u}_l z_l\right) - y_\mu\,.$$

   The y_μ are, therefore, driven by the z_k, which lie on a k-dimensional manifold. If the dynamics of the z_k is periodic, then the y_μ will also be periodic, and will be functions of the z_k. Thus, the full dynamics lives on a K-dimensional manifold. If, on the other hand, the dynamics of the z_k is chaotic, the y_μ never see exactly the same input, and so can – and generally do – live on an $(n - K)$-dimensional manifold.

3 Coding

1. Show that the coefficient of variation (CV) of the inter-spike intervals (ISIs) generated by a homogeneous Poisson process is 1. Is the converse (any renewal process with an ISI CV of 1 is Poisson) true?

2. A neuron has an absolute refractory period of Δ, but its firing is otherwise unconstrained. Assume it fires so that its ISI distribution has maximum entropy. What form will its ISI distribution have? Give an expression linking its mean firing rate and Δ. [HINT: first assume the mean ISI is also constrained, find the distribution in that case, then maximise over mean firing rate.]

3. An experiment on the visual system uses stimuli formed by a superposition of a family of Gabor wavelets $\{\mathbf{g}_i\}$, each with a coefficient drawn from a normal distribution:

$$\mathbf{s} = \sum_i \alpha_i \mathbf{g}_i; \qquad \alpha_i \sim \mathcal{N}\left(0, \sigma_i^2\right)$$

   with variances σ_i^2 chosen to emphasise low spatial frequency wavelets. Given a set of such stimuli (with α_i drawn independently) and corresponding neural responses $\{\mathbf{s}_n, r_n\}$, say (giving all the relevant equations) how would you obtain an unbiased estimate of the image subspace to which the recorded neuron is sensitive?

4. Give three arguments *against* the idea that population codes have evolved in part to maximize the average Fisher information in neuronal responses to single (scalar-valued) sensory stimuli.

4 Learning

1. The STDP rule for a single weight, w, between a pre and postsynaptic neuron can we written

$$\Delta w = \eta \sum_j \left[K(t^i_{\text{post}} - t^j_{\text{pre}}) - wK(t^j_{\text{pre}} - t^i_{\text{post}}) \right]$$

where t^i_{post} is the time of the i^{th} post-synaptic spike and t^j_{pre} is the time of the j^{th} pre-synaptic spike, a is a constant, and K is a temporal kernel,

$$K(t) = \Theta(t) e^{-t/\tau}$$

where Θ is the Heaviside step function.

Both the pre and postsynaptic neurons follow Poisson processes. The presynaptic firing rate is constant, and set to ν_{pre}. The postsynaptic firing rate, denoted ν, depends on incoming spikes,

$$\nu(t) = \nu_{\text{post}} \left(1 + \epsilon w \sum_j K(t - t^j_{\text{pre}}) \right)$$

where ν_{post} and ϵ are constant. Thus, whenever the presynaptic neuron fires, the postsynaptic firing rate increases for a short time.

Explain, qualitatively, why this rule is stable (in the sense that the weight fluctuates around an intermediate value), even when ϵ is large. Make a rough estimate of the average weight in the limit $\epsilon \ll 1$ and $\tau \nu_{\text{pre}} \ll 1$.

2. The BCM rule, written in terms of firing rates, looks something like

$$\tau \frac{d\mathbf{w}}{dt} = -(\lambda - \epsilon\lambda^2)\boldsymbol{\nu}$$
$$\tau_\epsilon \frac{d\epsilon}{dt} = f(\lambda) - \epsilon$$

where $\boldsymbol{\nu}$ is a vector of (constant) presynaptic firing rates, λ is the postsynaptic firing rate, and $f(\lambda)$ is a monotonic increasing function of λ. Not sure what I was thinking: if we want this rule to be stable, $f(\lambda)$ cannot be a monotonic increasing function of λ. See solution. The postsynaptic firing rate is given in terms of the presynaptic rates as

$$\lambda = \mathbf{w} \cdot \boldsymbol{\nu}$$

Assume that $\tau_\epsilon \ll \tau$. Find a function $f(\lambda)$ that will make the rule stable, in the sense that the weights don't converge to zero or ∞.

3. It is not likely that the brain uses backprop, at least not directly. An alternative is so-called node perturbation, in which noise is added, and the change in loss is used to adjust the weights. Consider, for instance, linear regression,

$$y = \mathbf{w} \cdot \mathbf{x}\,,$$

for which the loss is

$$L(\mathbf{w}) = \frac{1}{2} \langle (y - \mathbf{w} \cdot \mathbf{x})^2 \rangle$$

where the average is over the input. In the case of node perturbation, we add noise to $\mathbf{w}\cdot\mathbf{x}$: $\mathbf{w}\cdot\mathbf{x} \to \mathbf{w}\cdot\mathbf{x} + \sigma\xi$ where ξ is a zero mean, unit variance Gaussian random variable and σ is small. The change in loss (for a single $\mathbf{x}$ rather than an average) is, to first order in σ,

$$\Delta L = -\sigma(y - \mathbf{w} \cdot \mathbf{x})\xi\,.$$

If the loss decreases, it makes sense to move the weights such that $\mathbf{w} \cdot \mathbf{x}$ moves in the same direction as ξ; if the loss increases, it makes sense to move it in the opposite direction. In other words, we want to choose $\Delta\mathbf{w}$ so that $\Delta\mathbf{w} \cdot \mathbf{x}$ has the same sign as ξ.

Given that argument, write down a learning rule in terms of the noise and ΔL. One way to tell if you're right is to note that when averaged over the noise, your learning rule should reduce to SGD (with, possibly, a scale factor).

4. NMDA receptors provide a mechanism for coincidence detection: a postsynaptic spike causes an increase in voltage at a synapse; that removes the magnesium block; glutamate released by a presynaptic spike can then open NMDA channels; and the resulting influx of calcium can cause long term changes in synaptic strength.

   Based on this picture, sketch the amount of potentiation or depression that you expect as a function of the difference between the presynaptic and postsynaptic spike times. Does your curve look anything like the classic STDP curve?

Solutions

1. The average increase (or decrease) associated with a kernel is given approximately by the number of presynaptic and postsynaptic spikes that occur within the kernel. That number scales with the product of the pre and postsynaptic firing rates. The hard part is estimating the postsynaptic firing rate. For the kernel in which the presynaptic spike precedes the postsynaptic one, the term $\sum_j K(t - t^j_{\rm pre})$ is elevated in the post-spike window, and so averages to something that's about 1. We'll call it $k_{\rm post}$. For the kernel in which the presynaptic spike follows the postsynaptic one, the term $\sum_j K(t - t^j_{\rm pre})$ still makes a contribution, but now it needs to be averaged over all presynaptic spikes. This average is about $\tau\nu_{\rm pre}$; we'll call it $\nu_{\rm pre}\tau k_{\rm pre}$.

   Thus, very approximately,

$$\Delta w \propto \nu_{\rm pre}\nu_{\rm post}[(1 + k_{\rm post}\epsilon w) - (1 + k_{\rm pre}\nu_{\rm pre}\tau\epsilon w)w]\,.$$

   where both $k_{\rm post}$ and $k_{\rm pre}$ are about 1. Because of the w^2 term, this always has a solution. But we get to consider the limit $\epsilon \ll 1$ and $\tau\nu_{\rm pre} \ll 1$. In that limit, $w \approx 1 + k_{\rm post}\epsilon w$, or

$$w \approx \frac{1}{1 - k_{\rm post}\epsilon}\,.$$

2. Dotting both sides of the equation for $\mathbf{w}$ with $\boldsymbol{\nu}$, and using $\epsilon = f(\lambda)$ (valid because τ_ϵ is small), we can write a differential equation for λ,

$$\tau\,\frac{d\lambda}{dt} = -(\lambda - f(\lambda)\lambda^2)\lambda\,.$$

   For this rule to be stable, $d\lambda/dt$ needs to be positive as λ approaches 0 and negative as λ approaches ∞. There are lots of choices, but an easy one is $f(\lambda) = a/\lambda + 1/\lambda^3$ with $a > 1$.

3. Given the argument in the problem statement, it makes sense for the change in weight, $\Delta\mathbf{w}$, to be chosen so that $\Delta\mathbf{w}\cdot\mathbf{x}$ is proportional to $-\Delta L\xi$. Because the only direction in weight space for which we have information is $\mathbf{x}$, we'll let

$$\Delta\mathbf{w} \propto -\Delta L\xi\mathbf{x}\,.$$

   Using the expression for ΔL from the problem statement, this can be written

$$\Delta\mathbf{w} \propto (y - \mathbf{w}\cdot\mathbf{x})\xi^2\mathbf{x}\,.$$

   This does indeed reduce to SGD. Note also that we could have used $\Delta\mathbf{w} \propto -\Delta L f(\xi)\mathbf{x}$. So long as $f(\xi)$ is a symmetric function of ξ; this would also reduce to SGD. It would produce a slightly different prefactor, but that could be absorbed into the learning rate.

4. The key quantity that determines NMDA-dependent plasticity is the amount of calcium that flows through the NMDA channels. That's determined by the magnesium block, which is determined by the postsynaptic membrane potential, which is regulated by the backpropagating action potential. Kind of a long list, but it's all illustrated in Fig. 1, which plots various quantities relative the the time of a postsynaptic spike. Qualitatively, these are as follows:

   Red: postsynaptic membrane potential at a synapse. This rises shortly after a postsynaptic spike, and stays up for 20-40 ms.

   Blue: the probability that the magnesium block has been removed from the NMDA channel. Alternatively, we can think of this as the NMDA conductance should a pre-synaptic spike arrive.

   Solid green: calcium influx through the NMDA channels given the time of the presynaptic spike relative to the postsynaptic one. There's a baseline level of calcium influx because some of the NMDA channels are open even at rest; that's the green solid line before (and well after) zero. However, if the presynaptic

spike comes when the NMDA channels are open (when the blue curve is high), there's a lot more calcium influx.

Dashed green: threshold for induction of LTP. Above that line the weights increase; below it they decrease.

Magenta: change in weight.

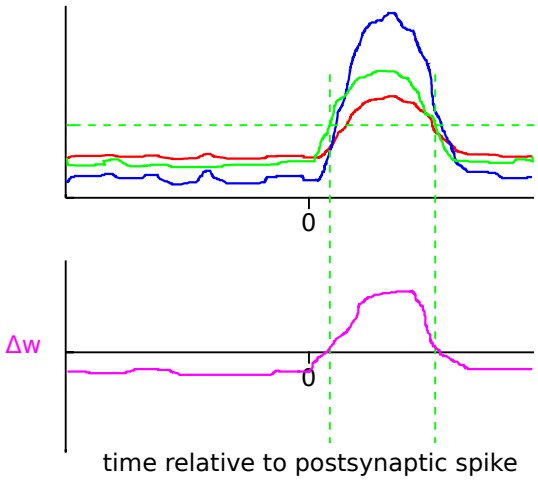

Figure 1: Various quantities relative to the time of the postsynaptic spike. Colors explained in the answer.