Gatsby Computational Neuroscience Unit
Theoretical Neuroscience

Final Examination
30 Jan 2015

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 are allowed to use your own handwritten notes, but no other online or offline reference sources.

You have a maximum of 7 hours for this exam.

Good luck!

1 Biophysics

Comparison of LIF and QIF

There are several approximate models of single neuron dynamics. Two of the most popular are the leaky integrate and fire (LIF) and the quadratic integrate and fire (QIF). For the former,

$$\tau \frac{dV_{LIF}}{dt} = -V_{LIF} + V_{\text{in}}(t)\,,$$

and for the latter,

$$\tau \frac{dV_{QIF}}{dt} = \frac{V_{QIF}(V_{QIF} - V_{\text{thresh}})}{V_{\text{thresh}}} + V_{\text{in}}(t)\,.$$

Here $V_{\text{in}}(t)$ models synaptic drive. For the LIF, a spike is generated when $V_{LIF} = V_{\text{thresh}}$, at which point the voltage is reset to 0; for the QIF, a spike is generated when $V_{QIF} = +\infty$, at which point the voltage is reset to $-\infty$.

1. Which is a better model of real neurons, and why?

   (5 marks)

2. Show that when $V_{\text{in}}(t)$ is infinitesimally small (but time-varying), the two model neurons exhibit exactly the same behavior. That is, if they receive the same input, $V_{\text{in}}(t)$, and at $t = 0$ both are equal to zero, then $V_{LIF}(t) = V_{QIF}(t)$ for $t \geq 0$.

   (5 marks)

3. Plot the response of both model neurons to a δ-function input,

   $$V_{\text{in}}(t) = \tau V_0\, \delta(t)\,.$$

   Use two values of V_0: $V_0 = 0.9V_{\text{thresh}}$ and $V_0 = 1.1V_{\text{thresh}}$. Assume in both cases that the neurons were at rest for $t < 0$.

   (5 marks)

4. Show that if the input to the two model neurons is identical, and they have the same initial conditions, then, so long as the membrane potential remains below threshold (below V_{thresh}), $V_{QIF}(t) \geq V_{LIF}(t)$.

   (5 marks)

5. This does *not* imply that the QIF will always spike before the LIF. Sketch, qualitatively, an input, $V_{\text{in}}(t)$, for which the LIF fires but the QIF does not. Assume that $V_{LIF}(t = 0) = V_{QIF}(t = 0) = 0$.

   (5 marks)

6. But, of course, the QIF can have a higher firing rate. Find a fixed input, V_{in}, such that the QIF fires and the LIF does not.

   (5 marks)

7. Finally, we want to investigate spike timing jitter. We'll ask what happens when the two model neurons receive input that is barely enough to make them fire. For the LIF, let

$$V_{\text{in}}(t) = \begin{cases} 0 & t < 0 \\ V_{\text{thresh}} + \epsilon V_{\text{thresh}} & t \geq 0 \end{cases}$$

While for the QIF, let

$$V_{\text{in}}(t) = \begin{cases} 0 & t < 0 \\ V_{\text{thresh}}/4 + \epsilon V_{\text{thresh}} & t \geq 0\,. \end{cases}$$

Assume that the voltage was set were well in the past, so at $t = 0$, the membrane potential is at its fixed point.

Show that when ϵ is small (and positive), the time to spike, τ_{spike}, has different scaling for the LIF versus the QIF. For the LIF,

$$\frac{\tau_{spike}}{\tau} \approx -\log \epsilon\,,$$

while for the QIF,

$$\frac{\tau_{spike}}{\tau} \approx \frac{1}{\epsilon^{1/2}} \int_{-\infty}^{\infty} \frac{dz}{1+z^2}\,.$$

(5 marks)

8. Based on this result, which neuron exhibits more spike timing jitter in response to near threshold input, and why?

(5 marks)

2 Networks

For spike timing dependent plasticity (STDP), the change in the connection strength from pre-synaptic neuron i to post-synaptic neuron j is

$$\Delta W_{ij} = \eta K(t_i - t_j)$$

where t_i is the time of the post-synaptic spike and t_j the time of the pre-synaptic spike. The kernel, $K(\tau)$, is given approximately by

$$K(\tau) = \begin{cases} f_+ e^{-\tau/\tau_+} & \tau > 0 \\ -f_- e^{\tau/\tau_-} & \tau \leq 0\,. \end{cases}$$

This is, of course, the change in weight for a pair of spikes; if there are many spikes, distributed over some time T, then the time-averaged change in the weight is

$$\frac{\Delta W_{ij}}{T} = \frac{1}{T}\sum_{kl} K(t_i^k - t_j^l) \tag{1}$$ {dw}

where t_i^k is the time of the k^{th} spike on post-synaptic neuron i and, similarly, t_j^l is the time of the l^{th} spike on pre-synaptic neuron j.

1. Assume we're in the stationary regime, and the neurons in the network are firing asynchronously, so the pre- and post-synaptic neurons are firing at constant rates ν_j and ν_i. As you know, the covariance between the two spike trains is given by

   $$C_{ij}(\tau) = \lim_{T\to\infty} \frac{1}{T}\int dt\, \sum_{kl}(\delta(t+\tau-t_i^k)-\nu_i)(\delta(t-t_j^l)-\nu_j)\,.$$

   Show that Eq. (1) becomes

   $$\lim_{T\to\infty}\frac{\Delta W_{ij}}{T} = \int_{-\infty}^{\infty} d\tau\, K(\tau)\big(\nu_i\nu_j + C_{ij}(\tau)\big)\,.$$

   (10 marks)

2. If the weights change slowly enough compared to the interspike interval, we can let $\Delta W_{ij}/T \to dW_{ij}/dt$. Let's make another approximation, which is that $C_{ij}(\tau)$ is small compared to $\nu_i\nu_j$. In that case,

   $$\frac{dW_{ij}}{dt} = \int_{-\infty}^{\infty} d\tau\, K(\tau)\nu_i\nu_j \equiv \kappa\nu_i\nu_j\,.$$

   We would like to derive a set of equations for the dynamics of the average connection strengths. We're going to consider a randomly connected network of excitatory and inhibitory neurons, so there are four types of weights,

   $$W_{QR} = \frac{1}{N_Q N_R}\sum_{i\in Q, j\in R} W_{ij}$$

where $Q, R \in E, I$, N_Q is the number of neurons of type Q, and the notation $i \in Q$ under the sum means sum over only neurons of type Q. Show that

$$\frac{dW_{QR}}{dt} = \kappa \nu_Q \nu_R \tag{2}$$ {dwdt}

where

$$\nu_Q \equiv \frac{1}{N_Q} \sum_{i \in Q} \nu_i \,.$$

(3 marks)

3. We now want to express the firing rates in terms of the weights. Assume the firing rates evolve according to Wilson and Cowan equations,

$$\begin{aligned} \tau_E \frac{d\nu_E}{dt} &= \phi_E\big(K(W_{EE}\nu_E - W_{EI}\nu_I + I_E)\big) - \nu_E \\ \tau_I \frac{d\nu_E}{dt} &= \phi_I\big(K(W_{IE}\nu_E - W_{II}\nu_I + I_I)\big) - \nu_I \end{aligned}$$

where K is the number of connections per neuron and ϕ_E and ϕ_I are approximately sigmoidal functions of their input (the main properties we'll need for ϕ_E and ϕ_I are the usual ones: they are non-negative, monotonic increasing, and they saturate at 0 for large negative input and at a few 10s of Hz for large positive input). Show that in the limit $K \to \infty$, the equilibrium firing rates are given by {rates}

$$\nu_E = \frac{W_{II} I_E - W_{EI} I_I}{W_{EI} W_{IE} - W_{EE} W_{II}} \tag{3a}$$

$$\nu_I = \frac{W_{IE} I_E - W_{EE} I_I}{W_{EI} W_{IE} - W_{EE} W_{II}} \,. \tag{3b}$$

(8 marks)

4. Inserting this into Eq. (2) gives us a set of self-consistent update equations for the weights. We will also add a weight decay term, yielding the equations

$$\frac{dW_{QR}}{dt} = \kappa \nu_Q \nu_R - \eta W_{QR}$$

where the firing rates, ν_Q and ν_R, are given by Eq. (3).

This gives us a set of four equations – too many to analyze easily. So we'll assume that only W_{EE} and W_{II} are plastic; W_{EI} and W_{IE} are fixed. We'll also let

$$\begin{aligned} I_E &= I_I = 1 \\ W_{EI} &= y_0 \\ W_{IE} &= x_0 \,. \end{aligned}$$

Also, to reduce writer's cramp, we'll let

$$W_{EE} \to x_0 \times x$$
$$W_{II} \to y_0 \times y\,.$$

Show that x and y evolve according to

$$\frac{1}{\eta}\frac{dx}{dt} = \epsilon_x^2 \left(\frac{y-1}{1-xy}\right)^2 - x$$
$$\frac{1}{\eta}\frac{dy}{dt} = \epsilon_y^2 \left(\frac{1-x}{1-xy}\right)^2 - y$$

where

$$\epsilon_x^2 \equiv \frac{\kappa}{\eta x_0^3}$$
$$\epsilon_y^2 \equiv \frac{\kappa}{\eta y_0^3}\,.$$

(3 marks)

5. Assume $\kappa > 0$. Sketch the nullclines in two regimes: $\epsilon_y < 1$ and $\epsilon_y > 1$. Be sure to put x on the x-axis and y on the y-axis. Include arrows indicating the sign of dx/dt and dy/dt. Take into account that fact that for the firing rates to be positive, we must have $y > 1$ and $x < 1$, and for the firing rate fixed point to be stable, we must have $xy < 1$.

   (10 marks)

6. Show that there can be a stable fixed point only if $1 < \epsilon_y < \epsilon_{cutoff}$. You don't have to compute ϵ_{cutoff}; you just need to show that it exists.

   (6 marks)

3 Coding

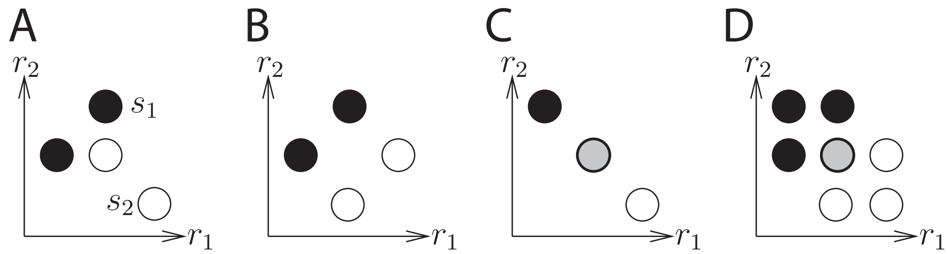

Consider the four distributions shown in the figure. These show the probabilities of the responses of two neurons r_1, r_2 to two stimulus values ($s = \{\text{black; white}\}$). Gray circles in (C;D) are overlapping – so the central circle has probability 0.5 under either white or black in (C); and probability 0.25 under either white or black in (D). The prior probability of each stimulus is 1/2, and, conditioned on a stimulus, all possible responses are equally likely (e.g., if stimulus s_2 is shown, then each of the white circles in panel A has probability 1/2).

1. Define signal and noise correlations, and evaluate them for the distributions in the figure. You may approximate the circles as point masses. (8 marks)
2. What are the mutual information values $I(\mathbf{r}; s)$ for all the distributions? (5 marks)
3. The synergy associated with $\mathbf{r}$ has been formalized as $I_{\text{syn}} = I(\mathbf{r}; s) - \sum_a I(r_a; s)$. Calculate this for the cases shown. (5 marks)
4. One construct often used to assess the importance of correlations is the distribution $p_{\text{ind}}(s, \mathbf{r})$ that would arise if all the trials for a single stimulus were shuffled. Write down a formal expression for this distribution and sketch it for each of the distributions in the figure. (8 marks)
5. What is the mutual information between s and $\mathbf{r}$ based on $p_{\text{ind}}(s, \mathbf{r})$ for each case. Why does it differ from the true mutual information in the way that it does? (7 marks)
6. What do these four distributions teach us about the role of correlations? (7 marks)

4 Learning

1. Consider a monotonic, non-linear, mapping (subject to minimal noise) from scalar x (with probability distribution $p(x)$) to $z = f(x)$. If z is confined to the interval $[0, z_{\text{max}}]$, and you seek to maximize the mutual information $\mathcal{I}(x, z)$, what is the best function $f(\cdot)$? (10 marks)

2. Now, consider the case of vector $\mathbf{x}$, a linear, square, mapping to $\mathbf{y} = R \cdot \mathbf{x}$, and then a point-wise logistic sigmoidal non-linearity $z_i = \sigma(y_i)$. Write down an expression for the mutual information between $\mathbf{x}$ and $\mathbf{z}$ in terms including the externally-determined distribution over $\mathbf{x}$ and the Jacobian of the transformation. Write down a learning rule that does gradient ascent in the mutual information (you may, if you wish, use the fact that $\nabla_W \ln|\det(W)| = (W^T)^{-1}$ for a square invertible matrix W). (15 marks)

3. An alternative derivation starts from a generative model with $p(\mathbf{y}) = \prod e^{f_i(y_i)}$ and a deterministic, invertible, transform $\mathbf{x} = G \cdot \mathbf{y}$. Write down the likelihood of an input $\mathbf{x}$, and derive a gradient ascent rule on this likelihood. Is there a form of $f_i(\cdot)$ for which this is the same as the information maximization version? (15 marks)