TNI, 2010

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 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. Please mark any answers provided after 3 hours.

Good luck!

1 Biophysics

Consider a quadratic integrate and fire neuron whose voltage, V, evolves according to

$$\tau \frac{dV}{dt} = \frac{(V - V_t)(V - V_r)}{\Delta V} + \Delta V(I_0 + 1/4) + g(V - \Delta V(\mathcal{E} + \overline{V}))$$

where $\Delta V \equiv V_t - V_r$ and $\overline{V} \equiv (V_t + V_r)/2$. An action potential is emitted when V reaches $+\infty$, at which point the voltage is set to $-\infty$. Except for question 2, assume that $I_0 > 0$ and $\mathcal{E} > 0$.

Note: it should be $(\Delta V \mathcal{E} + \overline{V})$, *not* $\Delta V(\mathcal{E} + \overline{V})$. It was wrong on the exam!

1. Show that under a suitable change of variables from V to u, this equation can be written

$$\tau \frac{du}{dt} = u^2 + I_0 + g(u - \mathcal{E})$$

   (5 marks).

2. Solve this equation for $u(t)$ when $I_0 = g = 0$ and $u(t = 0) = u_0 > 0$, and use the solution to compute how long it takes u to reach $+\infty$ (3 marks).

3. Find the range of g that allows repetitive firing (5 marks).

4. Whenever the neuron spikes, there is calcium influx, and the increased calcium concentration causes g to increase. This can be captured in the set of equations

$$\begin{aligned} \tau_c \frac{dc}{dt} &= \tau_c \sum_i \delta(t - t_i) - c \\ \tau_g \frac{dg}{dt} &= c - g\,. \end{aligned}$$

   where t_i is the time of the i^{th} spike (the time at which the voltage reaches $+\infty$). Explain why this neuron can burst. Provide approximate conditions on the time constants and I_0 that will ensure bursting. Plot, qualitatively, $u(t)$, $c(t)$ and $g(t)$ when the neuron is bursting (15 marks).

5. Suppose we had skipped the calcium influx step, and g had evolved according to

$$\tau_g \frac{dg}{dt} = \tau_g \sum_i \delta(t - t_i) - g\,.$$

   Show that in this case the cell cannot exhibit bursting (12 marks).

2 Networks

Consider a 2-memory Hopfield network in which the activity of neuron i evolves according to

$$\frac{ds_i}{dt} = \tanh\left[\frac{\beta}{N}\sum_j \left(\xi_i^1\xi_j^1 + \xi_i^2\xi_j^2\right) s_j\right] - s_i$$

where the ξ_i^μ are *iid* random variables drawn from a distribution $p(\xi)$. For all questions, assume that $\beta > 1$ and N is large.

1. Let

$$x_\mu = \frac{1}{N}\sum_i \xi_i^\mu s_i\,,$$

   $\mu = 1, 2$. Write down the time evolution equations for x_1 and x_2 in terms of averages over $p(\xi)$ (10 marks).

2. Show that when the ξ_i^μ take on only the values $+1$ or -1, both with probability 1/2 (i.e., $p(\xi_i^\mu = +1) = p(\xi_i^\mu = -1) = 1/2$), the x_μ evolve according to

$$\frac{dx_\mu}{dt} = \frac{1}{2}\left[\tanh\beta(x_\mu + x_\nu) + \tanh\beta(x_\mu - x_\nu)\right] - x_\mu$$

   where $\mu \neq \nu$ (i.e., $\nu = 1$ if $\mu = 2$ and $\nu = 2$ if $\mu = 1$) (5 marks).

3. For the equations you derived in question 2, draw the nullclines (15 marks).

4. For each equilibrium, tell us whether it is stable or unstable and draw representitive trajectories near the equilibrium (10 marks).

3 Coding

Gain control

Consider a roughly linear cell with very simple contrast gain control, whose firing rate in response to a stimulus is given by

$$r(\mathbf{s}) = \alpha \, \frac{\mathbf{f}^\mathsf{T}\mathbf{s}}{\beta + \|\mathbf{s}\|} \, .$$

Here, $\mathbf{s}$ is the stimulus (say a vector of pixel contrasts), $\mathbf{f}$ is a unit vector in stimulus space indicating the preferred stimulus of the cell, $\|\mathbf{s}\| = \sqrt{\mathbf{s}^\mathsf{T}\mathbf{s}}$ and $\alpha, \beta > 0$. Assume that firing is Poisson, given this rate.

1. Compute the Fisher information about the average stimulus contrast $\|\mathbf{s}\|$ assuming that the direction of the stimulus in stimulus space is held fixed and known. Does this result suggest contrast invariance? (4 marks)
2. Why did we need the assumptions about the direction of the stimulus? How would the calculation need to change if those assumptions were untrue. [Don't redo the calculation – just explain how it would differ.] (3 marks)

You collect responses from the cell using a zero-mean Gaussian stimulus distribution with identity covariance matrix, in order to characterise it by the spike-triggered covariance.

3. Can this cell be described by an LNP model? (2 marks)
4. Argue geometrically that the spike-triggered average stimulus will tend to lie in the direction of $\mathbf{f}$. (3 marks)
5. Find an expression for the expected length of the STA in terms of an expected value of a function of $\|\mathbf{s}\|$. (5 marks)
6. Sketch the eigenspectrum you would expect to see when running STC, having projected out the STA, in the limit of infinite data. Explain your answer. (10 marks)

Suppose that the stimuli are presented independently and sequentially, switching every 50 ms (think of a 20Hz white noise movie). You fit a GLM model with exponential link function and spike-dependent feedback current to the responses.

7. Sketch the form of the spike-dependent feedback current you would expect to find. Why? (5 marks)
8. Would your answer change if the stimuli presented on successive frames were correlated (i.e. the noise were temporally pink instead of white)? Why? (7 marks)

Remember to justify all of your answers.

4 Learning

1. Define the actor-critic and Q-learning rules from reinforcement learning. What are the key differences between them?

2. A potential graduate student has suggested the following hybrid learning rule, which we will first study in the case of a single action. A Q value, and an action value S are updated as:

$$\delta_t = r_t - Q_t \tag{1}$$
$$Q_{t+1} = Q_t + \alpha\delta_t \tag{2}$$
$$S_{t+1} = S_t + \beta\delta_t \tag{3}$$

where β is not necessarily equal to α, but both are less than 1.

(a) If r_t is binomial with probability p of being 1, derive the asymptotic mean and variance of Q_t as $t \to \infty$. How does this depend on $Q_{-\infty}$? (5 points)

(b) Derive the asymptotic mean and variance of S_t as $t \to \infty$. How does this depend on $S_{-\infty}$? (10 points)

(c) Consider a paradigm in which $r_t = 1$ for $t = -\infty \ldots -1$, and $r_t = 0$ for $t \geq 0$. If $\beta = \alpha$ for $t < 0$ and $\beta = \alpha/2$ for $t \geq 0$, what are the values of Q_0 and S_0 and Q_t and S_t as $t \to \infty$ in terms of $Q_{-\infty}, S_{\infty}$? (10 points)

(d) The student interpreted this result as suggesting that setting $r_t = 1$ for $t < 0$ reduced the learning rate for decreases in the value of S_t when $\delta_t < 0$; thus trying to explain compulsions in addiction. Evaluate the computational credibility of this claim? (5 points)

(e) What would happen in the actor-critic in this case? Sketch the consequences of endowing the agent permanently with the additional choice of an action for which $r = 0.1$? (10 points)