# Assignment 5
# Theoretical Neuroscience

Maneesh Sahani

11 November 2008

1. **Encoding Models**

   (a) Prove Bussgang's Theorem. That is, show that if we have samples $\{\mathbf{x}_i, y_i\}$, where $y_i$ is a random variable whose expectation is given by $E[y_i|\mathbf{x}_i] = f(\mathbf{w}\cdot\mathbf{x}_i)$, then the cross-correlation $\sum_i y_i\mathbf{x}_i$ (i.e. the "spike-triggered average" if $y_i$ is binary) provides an unbiased estimate of $\alpha\mathbf{w}$ (i.e. $\mathbf{w}$ times an unknown constant $\alpha$) if:

      i. $P(\mathbf{x})$ is spherically symmetric, where we define spherical symmetry to mean that,

      $$\forall\, \mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^n, \|\mathbf{x}_1\| = \|\mathbf{x}_2\| \Rightarrow P(\mathbf{x}_1) = P(\mathbf{x}_2).$$

      ii. $E[y\mathbf{x}] \neq \mathbf{0}$ (i.e. the expected spike-triggered average is not zero).

   (b) Simulate the response of an LNP (Linear-Nonlinear-Poisson) model to a temporal stimulus. Let $\mathbf{w}$ be a 20-tap filter (sampled in 10-msec bins) with biphasic temporal structure (i.e. a short, large-amplitude peak and a longer, smaller-amplitude trough). Choose the nonlinear response function $f$ to be a sigmoid that saturates at 200 spikes/sec. Recall that the instantaneous rate for an LNP neuron is given by

      $$r_i = f(\mathbf{w}\cdot\mathbf{x}_i)$$

      and that Poisson spikes can be generated by flipping a biased coin in each (suitably small) time bin with the probability of "heads" equal to $(\Delta t)r_i$.

      i. Simulate the neuron with a 1-sec Gaussian white noise stimulus sampled at a framerate of 100-Hz. Generate 200 responses of the neuron to this stimulus. Compute the PSTH of these responses, and show that it matches the rate prediction given by convolving the stimulus with $\mathbf{w}$ and passing the output through $f$.

      ii. Simulate the response to a long Gaussian white noise stimulus, and compute the STA (spike-triggered average). Plot the STA rescaled as a unit vector, and show that it provides a reasonable match to $\mathbf{w}/\|\mathbf{w}\|$.

      iii. Reconstruct the nonlinearity of the cell: begin by filtering the raw stimulus with the STA. Make a histogram of the filtered stimulus values, and make another histogram of the spike-triggered filtered stimulus values, using in the same binning. Divide the latter histogram by the former and multiply by the inverse of the bin size. Plot this estimate of the nonlinearity against the true $f$.

      iv. Stimulate the model cell with correlated Gaussian white noise: take the original (Gaussian white noise) stimulus and filter it with a Gaussian whose standard deviation is 20ms). Rescale if necessary to ensure that the standard deviation of the new stimulus is the same as the old. Now simulate the neuron and compute the STA, and compare it to $\mathbf{w}$. Compute the decorrelated STA (obtained by "whitening" with the inverse of the stimulus covariance matrix), and compare with $\mathbf{w}$. If necessary, regularize by adding a small constant to the diagonal of the stimulus covariance matrix (this corresponds to doing "ridge regression"), and examine how this affects the estimate.

v. Change $f$ to be a symmetric function, such as $f(\xi) = \alpha\xi^2$. Simulate the new model neuron with Gaussian white noise, and compute the STA and largest eigenvector of the spike-triggered covariance (STC) matrix. Compare with $\mathbf{w}$. Reconstruct the nonlinearity using both filters, and compare with the true $f$.

vi. Decoding: estimate a linear decoder for the first LNP model described above (i.e. with sigmoidal nonlinearity). You should find a decoding filter $\mathbf{w}_d$ and a scalar $b$ such that we can estimate the stimulus via:

$$\hat{x} = y \circ \mathbf{w}_d + b$$

(where '$\circ$' indicates linear convolution). Once you have estimated the filter and constant, apply them to responses to a novel stimulus segment, and plot the true stimulus $x$ alongside the reconstructed stimulus $\hat{x}$. Create a temporally correlated stimulus whose temporal frequency structure is better matched to that of the encoding filter $\mathbf{w}$. Re-estimate the decoding filter and constant, and show reconstruction of a novel stimulus segment with the same temporal correlation.

2. **Estimation Theory**

(a) We derived the Fisher information $J(\theta)$ as the expected value of the second derivate (curvature) of the log-likeihood in the lecture.

i. Repeat the derivation for a *vector* parameter (or stimulus in our setting) $\boldsymbol{\theta}$, showing that the Fisher information in this case is given by a matrix.

As mentioned in the lecture, there is an alternate definition in terms of the first derivative. For vector parameters this is:

$$J(\boldsymbol{\theta}_0) = \mathrm{Cov}_{\boldsymbol{\theta}_0}\left( \nabla \log p(n|\boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}_0} \right).$$

where $\mathrm{Cov}_{\boldsymbol{\theta}_0}$ means the covariance evaluated under $p(n|\boldsymbol{\theta}_0)$.

ii. Demonstrate that these two definitions are the same (or more precisely, give conditions under which these two definitions are the same).

(b) Consider an LNP model:

$$p(n|\mathbf{x}) = \mathsf{Poiss}(g(\mathbf{w} \cdot \mathbf{x}))$$

i. What is $J(\mathbf{x})$ (the Fisher Information about the stimulus value available to the rest of the brain)? How does it depend on $\mathbf{w}$?

ii. What is $J(\mathbf{w})$ (the Fisher Information about the weight vector available to an experimenter)? How does it depend on $\mathbf{x}$?

3. **Fisher information and refractory firing**

Consider a hypothetical cell, which responds to the presentation of a stimulus with a continuous feature $s$ by firing at a homogeneous rate $f(s)$ in a (fixed) interval $[0, T]$. Assume that the firing rate is 0 outside this interval. We will be interested in the contributions made to the Fisher information by spike-timing, with and without a refractory period.

First, assume that the firing is Poisson.

(a) What is the probability of observing spikes at times $\{t_1 \dots t_n\} \subset [0, T]$?

(b) What is the Fisher information $J_{t,\text{Poiss}}(s)$ associated with this probability density function, assuming that the relevant interval $[0, T]$ is known? How does it compare to the Fisher information $J_{n,\text{Poiss}}$ associated with the distribution of spike counts $P(n|s)$?

Now consider refractory firing. Recall that one way to model a refractory period is to use a gamma-interval renewal process in place of a Poisson process. Thus, now assume that the cell's firing follows a gamma-interval process with the same mean rate $f(s)$ and with integral gamma order $\gamma$.

(c) What is the probability of observing spikes at times $\{t_1 \dots t_n\} \subset [0, T]$ from this process?

(d) What is the Fisher information $J_{t,\text{Gamma}}(s)$ associated with the new probability density function? You may assume that $T$ is long enough to neglect contributions due to the first spike, and due to the silence after the last spike.

Finally, we wish to see how much of this information gain is available in the spike count.

(e) Which signal (count or spike-timing) do you expect to carry more information for this process? Why?

(f) Find an expression for the the distribution of spike counts $P(n|s)$ under the gamma-interval model.

(g) Write down the expression for the corresponding Fisher information $J_{n,\text{Gamma}}$, and thus for $J_{n,\text{Gamma}} - J_{t,\text{Gamma}}$. You need not necessarily evaluate the expectation. Identify the term(s) responsible for the difference between $J_{n,\text{Gamma}}$ and $J_{t,\text{Gamma}}$.

4. **Population Coding**

Shadlen and collaborators have claimed that if the activities of neurons in population codes are corrupted by *correlated* noise, then there is a sharp limit to the useful number of neurons in the population. *Prima facie* this is wrong – the stronger the correlations, the lower the entropy of the noise, and therefore the stronger the signal.

Resolve this issue for the case of additive and multiplicative noise by considering the following three models for the noisy activities $r_1$ and $r_2$ of two neurons which form a population code for a real-valued quantity $x$:

$$\text{a)} \quad \begin{cases} r_1^a = x + \epsilon_1 \\ r_2^a = x + \epsilon_2 \end{cases} \tag{1}$$

$$\text{b)} \quad \begin{cases} r_1^b = x(1-\delta) + \epsilon_1 \\ r_2^b = x(1+\delta) + \epsilon_2 \end{cases} \tag{2}$$

$$\text{c)} \quad \begin{cases} r_1^c = x(1-\delta)(1+\eta_1) \\ r_2^c = x(1+\delta)(1+\eta_2) \end{cases} \tag{3}$$

where $\delta \neq 0$ is known, and, $\epsilon$ and $\eta$ are Gaussian, with mean 0 and covariance matrices:

$$\Sigma = \begin{pmatrix} 1 & c \\ c & 1 \end{pmatrix}$$

(a) What is the maximum likelihood estimator (MLE) for $x$ on the basis of $r_1$ and $r_2$ in each case?

(b) What is the appropriate measure of the expected accuracy of the MLE, and why?

(c) How does the expected accuracy in each case depend on the degree of correlation $c$?

(d) What conclusions would you draw about the clash between Shadlen and common sense?