Problem Set 8, CNS 185 1999-2000

Handed out:	18 Nov 99
Due:	30 Nov 99
Points:	12 points out of 100 for the term

8.1 Decoding Spike Trains by Linear Reconstruction (6 points)

In the previous lecture (11/16/99, see class website for notes), Prof. Koch introduced concepts of Wiener filtering in the context of optimal linear filtering. Since the treatment in class was cursory, glossing over some interesting asides, we hope to consolidate and extend here some of the issues discussed in class.

Techniques from statistical estimation can be used to extract information from one time series about another related one. There has been a flurry of activity over the past decade in trying to apply these techniques to problems of theoretical neurobiology. One of the most exciting questions in theoretical neuroscience today is that of neural coding *i.e.,* how are signals from the outside world represented and coded in neural responses and what influence does the nature of this code (or even a class of codes) have on the possible computation that can be carried out by the nervous system. Towards this end, Bialek has popularized the "linear reconstruction" approach to understand the nature of neural codes. In order to understand how well neural responses can represent the plethora of sensory signals an animal receives from the outside world, one can ask the question 'how much can we say about the stimulus which causes a particular neural response if we are just given the response.' In absence of any other information, the answer is preciously little. If however, the organism knows something about the statistical distribution of the stimuli it is expected to receive and we have a model about how the neural responses are modulated by sensory inputs, we are in better shape.

Consider the following scenario: we are given the choice of a particular stimulus ensemble with which we can excite a neuron. We are also given partial information about the response of the neuron to each member of the stimulus ensemble we choose. Now given that we have the occasion to observe one particular response, how closely can we estimate the stimulus that could have caused this response. Since the best possible estimator of the stimulus given the response is highly complicated to analyze, we shall restrict ourselves to the best linear estimator (in the mean squared sense) instead. Note that it has been shown for several cases that the above restriction is not very stringent, as non-linear estimators do only marginally better on certain examples of real neural data. This setting allows us to use Wiener filtering to derive the best linear filter which reconstructs the stimuli from their responses. In this problem you will abstract a neuron by an leaky integrate and fire model, excite the neuron with stimuli from a particular ensemble, use the stimulus-response pairs to compute the optimal filter and characterize the quality of the optimization. For more details on this process, see section 3.2 in the above-cited lecture notes.

1. Use the MATLAB function `gensig.m` to generate several Gaussian noise stimuli with different bandwidths B_s and standard deviations σ_s (the parameter ranges are mentioned in `readme.ps8` supplied with the package). These stimuli are to be used as the input to a neuron, so generate the output of an integrate-and-fire neuron for each input (using `iafsim`).

   Use `reconst` (supplied in the MATLAB package) to get the best linear filter for each value of σ_s and B_s, and record the minimum mean squared error E in each case. On one graph, plot E as a function of stimulus variance (σ_s) for different values of stimulus bandwidth (B_s). Briefly explain the behavior of your plot. You might also want to plot the optimal reconstruction filters, h(t), etc.

2. In the programs supplied, find where computation of the optimal filter $h(t)$ occurs, and copy the corresponding line(s) of code as your answer. Be careful to copy *only* the relevant lines, in particular you should copy no more than 3 lines. We really mean it, if you copy a huge chunk of a program as your answer you will not get credit for this question. Just to be crystal clear: the goal of this problem is for you to *look* at the code and *understand it*, not for us to see reprinted versions of something we handed out.

3. For a particular value of σ_s and of B_s plot a certain segment (not more than 1000 points long) of the stimulus and the corresponding reconstruction on the same graph as well as the reconstruction error and the spike train. For plotting clarity, use different line types or colors in MATLAB for each of these three signals, and offset the error and the spike train to make them clearly visible – but be sure to indicate by how much the error has been offset.

4. Choose a particular value of σ_s and vary B_s. Obtain the optimal filters in each case. Plot these filters (by this we mean plot the kernels that you get) and see how the properties of the filter vary with B_s. Explain briefly what you see.

8.2 Signal Detection and ROC analysis (6 points)

This problem shall provide you with a flavor of signal detection theory used in the context of neurophysiology. As mentioned in class (see Prof. Koch's lecture notes for 11/18/99, found on the class website), a classic and highly successful application of signal detection theory and ROC analysis to neurophysiology is by Newsome and colleagues. We strongly encourage you to read the original papers (referenced in the lecture notes) for your edification and understanding. The simple problem we consider here should provide you with the insight to deal with the more complicated scenario discussed in the papers.

Assume that a single cell responds to two different stimuli a and b by firing independent Poisson-distributed action potentials at mean rates λ_a and λ_b respectively, with $\lambda_a > \lambda_b$.

1. Let's first consider a Poisson neuron with a mean firing rate λ. Show that the probability of observing exactly n action potentials in the time interval $[0;T]$ due is given by

$$p_{\lambda T}(n) = \frac{(\lambda T)^n e^{-\lambda T}}{n!} \quad n = 0, 1, ..., \tag{1}$$

given that the probability density of the inter-spike interval t for a Poisson process of rate λ is proportional to $\exp(-\lambda t)$. This should sound familiar, remember problem 7.1.1? (Hint: Using induction will allow you to eschew solving a lot of integrals.)

2. What is the mean number of action potentials and the variance around the mean?

As T becomes large, $p_{\lambda T}(n)$ resembles a bell-shaped curve centered around its mean and can be approximated by a Gaussian distribution with mean μ and variance σ^2,

$$p_{\mu,\sigma^2}(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}},$$

where $\mu = \lambda T$, and $\sigma^2 = \lambda T$.

3. Plot $p_{\lambda T}(n)$ and the corresponding Gaussian approximation $p_{\mu,\sigma^2}(n)$ for $\lambda = 80$ Hz, $T = 2$ seconds. You can use the built-in MATLAB functions `normpdf.m` and `poisspdf.m` to generate the probability distributions.

4. For which value of n within one standard deviation σ around the mean μ, is the difference in the two distributions maximal? Numerically compute and plot the mean square error between $p_{\lambda T}(n)$ and $p_{\mu,\sigma^2}(n)$ for n within one standard deviation σ around the mean μ.

Assume now that either stimuli a and b are presented randomly with probability 1/2 throughout a trial; that is, $p(a) = p(b) = 1/2$. Let n be the observed spike count (the number of action potentials fired by the cell) in the time interval $[0; T]$. Assume that we have an observer whose goal is to determine which stimulus occurred (a or b) and has access only to the spike count in the interval T to do so. Given the observed spike count n, how should the observer decide between the two cases optimally? The criterion of optimality we choose is to minimize the probability of error P_E. Notice that there are two possible types of error. We say that a *miss* error occurs when the stronger stimulus a is classified as the weaker stimulus b. Similarly, a *false alarm* error occurs when the weaker stimulus b is classified as the stronger stimulus a. These errors are also known as *false negatives* and *false positives* respectively. Let the probabilities of the two errors be denoted by P_M and P_F respectively. For our case, the probability of error P_E can be written as

$$P_E = \frac{1}{2}P_M + \frac{1}{2}P_F \tag{2}$$

5. Show that under the Poisson assumption stimulus a is more probable than stimulus b if the cell fires more than n_0 action potentials, while stimulus b is more probable than a if the cell fires fewer than n_0 action potentials, where

$$n_0 = \frac{\delta\lambda T}{\log\rho}, \quad \delta\lambda = \lambda_a - \lambda_b, \quad \rho = \frac{\lambda_a}{\lambda_b}. \tag{3}$$

Such a decision rule (one that always picks the most probable option) is called a *likelihood rule.*

6. Show that the likelihood rule minimizes P_E. (No math needed, just use your powers of logical and analytical reasoning)

Notice that the likelihood rule is of the form of comparison with a threshold. This is true for a lot of other rules which are optimized for criteria other than minimization of P_E.

8.2.1 The ROC curve

For a decision threshold n_v, the signal detection (*true positive*) probability $P_D(n_v) = 1 - P_M$ is equal to the probability that the spike count n for the stronger stimulus a exceeds n_v. Thus,

$$\begin{aligned} P_D(n_v) = P(n > n_v|a) &= \sum_{n=n_v+1}^{\infty} e^{-\lambda_a T}\frac{(\lambda_a T)^n}{n!} \\ &= 1 - e^{-\lambda_a T}\left(1 + \frac{(\lambda_a T)}{1!} + \ldots + \frac{(\lambda_a T)^{n_v}}{n_v!}\right), \end{aligned} \tag{4}$$

Similarly the probability of false alarm is equal to the probability that for the weaker stimulus b, n exceeds the threshold n_v. Thus, $P_F(n_v)$ is given by

$$\begin{aligned} P_F(n_v) = P(n > n_v|b) &= \sum_{n=n_v+1}^{\infty} e^{-\lambda_b T} \frac{(\lambda_b T)^n}{n!} \\ &= 1 - e^{-\lambda_b T}\left(1 + \frac{(\lambda_b T)}{1!} + \ldots + \frac{(\lambda_b T)^{n_v}}{n_v!}\right). \end{aligned} \quad (5)$$

As the threshold value n_v is varied in the interval $[0; \infty)$, we obtain P_D and P_F as functions of n_v. The *receiver-operating-characteristic* (ROC) curve is a parametric plot of the detection probability $P_D(n_v)$ against the false alarm probability $P_F(n_v)$.

1. Use the function `roc_poisson.m` to obtain P_D and P_F and plot the receiver-operating-characteristic (ROC) curves for different values of λ_a, λ_b, and T: $\lambda_b = 20$ Hz $\quad \lambda_a = 30, 50$ Hz, $\quad T = 250, 2000$ msec.

2. Use the function `roc_gaussian.m` to compare these ROC curves with those obtained from the Gaussian approximation of the Poisson distributions. Comment on the usefulness of the Gaussian assumption in these exercises (think back to question 8.2.4 as well).

3. Plot P_E as a function of n_v and verify that P_E is minimized when $n_v = n_0$.

4. How does P_E (optimized with respect to n_v) vary with T for fixed values of λ_a and λ_b ? Choose $\lambda_b = 20\text{Hz}$ and $\lambda_a = 30\text{Hz}, 50\text{Hz}$.