Neural Information Processing: 2015-2016
Assignment 1

School of Informatics, University of Edinburgh

Instructor: Mark van Rossum

6th March 2016

Remember that plagiarism is a university offence. Please read the policy at:

`http://www.inf.ed.ac.uk/admin/ITO/DivisionalGuidelinesPlagiarism.html`

Practical information

You should produce a digital document for your assignment answers (e.g. with lyx) and submit this electronically using the `submit` command on a DICE machine. The format is e.g.

```
submit
msc nip 1 nipasst1.pdf
```

You can check the status of your submissions with the `show_submissions` command. NOTE: postscript or pdf formats are acceptable, other formats are not. Make sure that the file you submit prints ok on the DICE system, in particular when you produced it on a non-Unix machine.

Late submissions

Late submissions will receive a zero mark. Only evidence for illness or other serious reasons can prevent this at the discretion of the instructor. See

`http://www.inf.ed.ac.uk/teaching/years/msc/courseguide09.html#exam`

Population codes with multiple stimuli

In this assignment we analyse the Fisher Information when multiple stimuli are simultaneously encoded in the same population, e.g. two overlapping transparent gratings. Suppose we have N neurons. At each trial the firing rate of neuron a, denoted r_a, is modelled as

$$r_a = g_a(\theta_1, \theta_2) + \sigma\eta_a$$

where tuning curve $g_a(\theta_1, \theta_2)$ is the mean response of the neuron to stimuli θ_1 and θ_2. Furthermore, η_a are independent Gaussian random variables with standard deviation 1, and σ sets the noise level in the neurons. The response probability distribution is thus

$$P(r_a|\theta_1, \theta_2) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(r_a - g_a)^2}{2\sigma^2}\right]$$

The tuning curves are sums of Gaussians

$$\begin{aligned} g_a(\theta_1, \theta_2) &= \sum_{k=1,2} f_a(\theta_k) \\ f_a(\theta) &= A \exp\left[-\frac{(\theta - \phi_a)^2}{2w^2}\right] \end{aligned}$$

where A is the amplitude of the response, w is the width of the tuning curve, and ϕ_a is the preferred stimulus of the neuron. For numerical evaluation you can set $A = 1$, $w = 1$.

When considering the Fisher Information for multiple variables, the Fisher information becomes a symmetric matrix with entries i, j (i and j can both take values 1 and 2)

$$I_{i,j}(\theta_1, \theta_2) = -\int d\mathbf{r} P(\mathbf{r}|\theta_1, \theta_2) \frac{\partial^2 \log P(\mathbf{r}|\theta_1, \theta_2)}{\partial\theta_i \partial\theta_j}$$

The Cramer-Rao bound now bounds the covariance matrix C of any unbiased estimator of θ_1 and θ_2 as $C_{ij} \geq (I^{-1})_{ij}$ for all i, j.

Question 1: Calculate the Fisher information matrix. First express it in terms of $f()$. Next, assume that the neurons are densely and homogeneously spaced so that the sum over neurons can be replaced by an integral $\sum_a \Rightarrow \rho \int da$, where ρ is the coding density. Plot matrix entries as a function of $(\theta_1 - \theta_2)$.

We next compare the Fisher Information to a maximum likelihood estimator. Implement a maximum likelihood estimator of θ_1 and θ_2. Because the noise is Gaussian, the log-likelihood is $\log \mathcal{L} = -\sum_a [r_a - g_a(\theta_1, \theta_2)]^2$. Thus the maximal likelihood minimizes the mean squared error. Draw a number ($\sim$100) samples of the noisy response using N =50 neurons with preferred stimulus values equally spaced between -5 and +5, and find the stimulus values that maximize the likelihood. In Matlab you can use functions like `fminsearch` or a nonlinear curve fitter. Other parameters: $A = 1$, $w = 1$, $\sigma = 0.1$.

Question 2: Compare the covariance from the simulation to the minimal covariance as given by the Cramer Rao bound as a function of the difference in the stimuli (i.e. $\theta_1 - \theta_2$).