Assignment 5

1. In class, we derived the maximum likelihood estimate θ_{ML} for the parameter θ of a model in which the probability of observing x spikes in response to the presentation of a stimulus s is given by a linear-Poisson model:

$$P(x) \sim \text{Pois}(\lambda),\ \lambda = \theta s$$

given a set of datapoints $\{x_i, s_i\}$. Use the same approach to derive an expression that θ_{ML} must satisfy for the linear-nonlinear-Poisson model:

$$P(x) \sim \text{Pois}(\lambda),\ \lambda = f(\theta s),$$

where f is some known function. The solution will be an implicit expression for θ_{ML} involving f, f', x_i, and s_i.

2. From Dayan & Abbott: Simulate the responses of four interneurons in the cercal system of the cricket and check the accuracy of a simple population decoding scheme. In this system, there are four interneurons whose firing rate responses are dependent on the wind direction θ.

For a true wind direction θ the average firing rates of the four interneurons should be generated as $\text{E}[r_i] = 50 \text{ Hz} \cdot f(\cos(\theta - \theta_i))$, where f is rectified-linear ($f(x) = x$ if $x > 0$, $f(x) = 0$ otherwise), and $\theta_i = \pi/4, 3\pi/4, 5\pi/4, 7\pi/4$ for $i = 1, 2, 3, 4$. The actual rates, r_i, are then obtained by adding to these mean rates a random number chosen from a Gaussian distribution with zero mean and a standard deviation of 5 Hz (set any rates the come out negative to zero). Assume the noise is independent across neurons.

From these rates, construct the x and y components of the population vector:

$$x = \sum_{i=1}^{4} r_i \cos(\theta_i),\ y = \sum_{i=1}^{4} r_i \sin(\theta_i) \tag{1}$$

and, from the direction of this vector, compute an estimate θ_{est} of the wind direction. Average the squared difference $(\theta - \theta_{\text{est}})^2$ over 1000 trials. The square root of of this quantity is the error. Plot the error as a function of θ over the range $[-90°, 90°]$.

3. Compute the Fisher information $I_F(\theta)$ of the above neural population. Use the Cramér-Rao bound to plot a lower bound on the error of the optimal unbiased estimator over the same range as above.

You can make the simplifying assumption that the evoked firing rates are $\text{E}[r_i]$+Gaussian noise, without thresholding. In other words, $p(r_i|\theta)$ is Gaussian distributed, centered at $\text{E}[r_i]$, and the distribution includes negative firing rates. This will lead to a slight overestimation of the variance because of an overestimation of noise.

To compute $I_F(\theta)$, first find an expression for the Fisher information of a single neuron with preferred stimulus θ_i as a function of θ. You will find that the resulting expression depends on the slope of the tuning curve and on the variance of the noise. Then plot numerically the sum of these functions for the four values of θ_i to find that the Fisher information is flat across values of θ.