# Homework 3
# Systems & Theoretical Neuroscience [SWC]

Due: Mon, 13th November

## 1 Signal detection theory

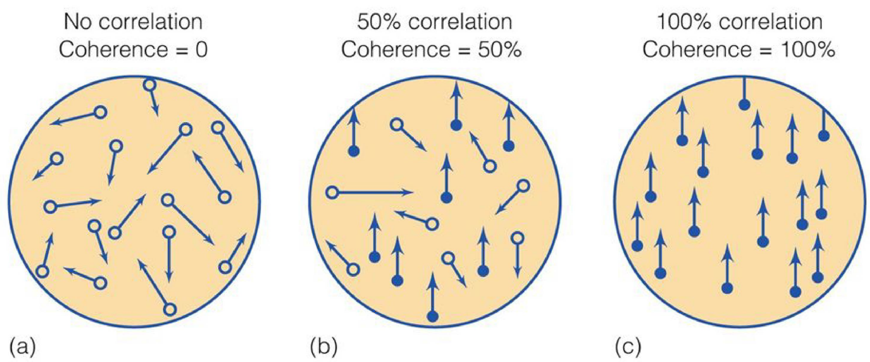


Figure 1: Schematic illustration of stimuli used.

In the experiment illustrated above, a visual stimulus consisting of moving dots is presented to a subject. The dots can either move coherently in a specified direction or in random directions. The subject is required to identify the direction of movement of the coherent dots, which can be either up (+) or down (-). The experimenter has control over the proportion of dots that move coherently, and can thus modify the difficulty of the task.

Consider a decoder that uses the firing rate of a single neuron to determine which of the two possible directions was presented as a stimulus. Assume the neuron has a Linear tuning curve and fires a higher firing rate in the plus direction than the minus. A simple decoding procedure is to determine the firing rate $r$ during a trial and compare it to a threshold $z$. If $r \geq z$ we report plus; otherwise we report minus.The success of our decoder in this case depends on two things, the separation of the mean firing rates in response to each stimulus, and the variance. This quantity, known as the discriminability is given by $d' = \frac{\mu_+ - \mu_-}{\sqrt{\frac{1}{2}(\sigma_+^2 + \sigma_-^2)}}$.

a) Simulate this random-dot discrimination experiment (using e.g. Jupyter Notebook). Denote the stimulus by plus or minus, corresponding to the two directions of motion. On each trial, choose the stimulus randomly with equal probability for the two cases. When the minus stimulus is chosen, generate the responses of the neuron as 20 Hz plus a random Gaussian term with a standard deviation of 10 Hz (set any rates that come out negative to zero). When the plus stimulus is chosen, generate the responses as 20 + 10d Hz plus a random Gaussian term with a standard deviation of 10 Hz, where d is the discriminability (again, set any rates that come out negative to zero). Generate results for 1000 trials for d=2, d=5 and d=10. Plot a histogram of + and - responses for each.

b) Choose a threshold $z = 20 + 5d$, which is half-way between the means of the two response distributions. Whenever $r \geq z$ guess "plus", otherwise guess "minus". Over a large number of trials (e.g. 1000) determine how often you get the right answer for different values of d. Plot the percent correct as a function of d over the range $0 \leq d \leq 10$.

c) The receiver operating characteristic curve (ROC) provides a way of evaluating how test performance depends on your choice of the threshold z. By allowing z to vary over a range, plot ROC curves for several values of d (starting with d = 2). To do this, determine how frequently the guess is "plus" when the stimulus is, in fact, plus (this value is known as the power, $\beta$), and how often the guess is "plus" when the real stimulus is minus (this is the false-positive rate, $\alpha$). Plot $\beta$ versus $\alpha$ for z over the range $0 \leq z \leq 140$.

## 2 Tuning curves

A tuning curve describes how the firing rate of a particular neuron varies according to the stimulus it is encoding. Typically a population of neurons will work together to encode some stimulus value, in which case the tuning curves of all neurons in the population usually share a basic shape, but with different parameters (for instance, shifted toward a different preferred stimulus value). The tuning curve for a single neuron gives its *average* firing rate for a given stimulus; the actual firing rate will be affected by noise.

Figure 2 gives four possible scenarios for the tuning curves of different neurons within a population of neurons that is working together to encode the value of a single 1D stimulus. The tuning curve for an individual example neuron is given in red; tuning curves for other neurons in the population are given in grey.

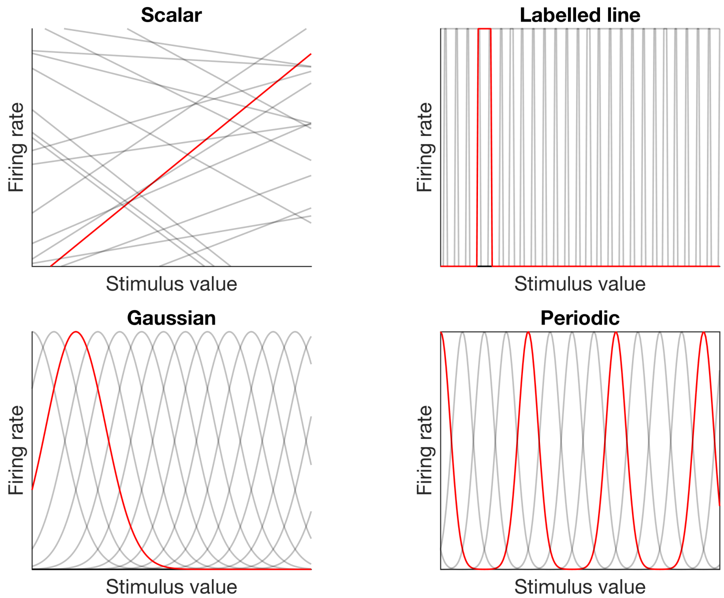


Figure 2: Different types of population tuning curves observed in the brain

a) Find examples from the literature of brain areas or functions that use each of the above types of encoding. In each case, describe the encoding and draw a cartoon tuning curve, making it clear what stimulus is being encoded.

b) Consider one or more neurons are enrolled by the brain to encode a single 1D stimulus using scalar coding, for instance the temperature of the environment. The brain may choose between the population tuning curves given in Figure 3. What are the advantages/disadvantages of each option?

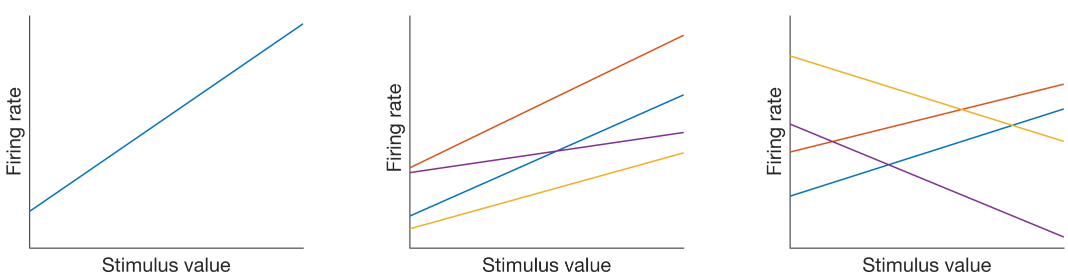


(a) Option A: 1 neuron

(b) Option B: 4 neurons

(c) Option C: 4 neurons

Figure 3: Possible population tuning curves using scalar coding

c) Consider a population of neurons that is encoding a 2 dimensional stimulus, for instance the location of berries within your field of view. Your brain may choose to use one of the tuning curve populations given in Figure 4. In each figure, an ellipse describes a 2-dimensional gaussian tuning curve.

Figure 4: Two possible sets of 2D tuning curves for a population of neurons

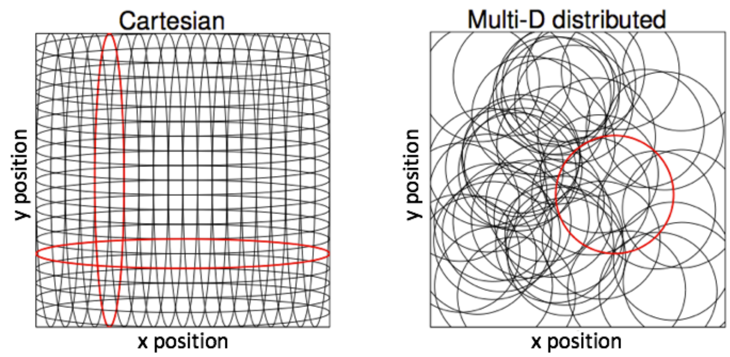


What are the advantages/disadvantages of each option? Which would you prefer? Does your answer change if you consider the scenario where there are multiple berries at different locations in your visual field? Think about how you would decode the stimulus.

d) In many animals, the current eye position appears to be encoded by a cartesian distributed "scalar" code in which the tuning curves are linear ramp functions of either the horizontal or vertical eye-position; while the intended target of a saccadic eye movement is more often represented with radial "bump"-shaped tuning curves distributed in 2D. Draw a picture of the tuning curves described above. Suggest why these different approaches might be chosen.

e) Zhang and Sejnowski [ZS99] analyse the relationship between Fisher Information and tuning curve width. Read the paper, and write a short summary:

   i) What assumptions do they make?

   ii) Under these assumptions, what result do they find?

   iii) What are the limitations of this analysis? How could you modify it to be more general?

# References

[ZS99] Kechen Zhang and Terrence J Sejnowski. "Neuronal tuning: To sharpen or broaden?" In: *Neural Computation* 11.1 (1999), pp. 75–84.

# 3 Optimal decoding and the Fisher information

## 3.1

After mapping out Jamie the barn owl's binaural sound localisation system in the previous assignment, you start wondering how Jamie's brain represents the loudness of particular tones in a population of auditory neurons. In an extremely cutting edge experiment, you manage to record from the full population of neurons that respond to a specific tone, while presenting that tone to Jamie at different loudness levels. You are interested in decoding the value of the continuous stimulus parameter $s$ (loudness) from the vector of firing rates $\vec{r} = \{r_1, r_2, ...r_N\}$. You determine the tuning curves of all measured neurons by systematically varying the stimulus and recording the (mean) firing rates, and you find something surprising. Across the population of N neurons, all neurons have a Gaussian tuning curve:

$$f_a(s) = r_{max}\, exp\Bigg( -\frac{1}{2}\bigg(\frac{s - s_a}{\sigma_a}\bigg)^2\Bigg), \tag{1}$$

and their means are evenly distributed over all possible stimulus values.

a) What does this tell you about the sum of all responses at any value of your stimulus. How does it depend on your stimulus value?

   The tuning curves you found give the average firing rates of neurons over multiple trials. In a single trial, firing rates will vary from that mean value. To decode the stimulus from these responses, you need to model this variability, which can be described as a probability density of responses given your stimulus, $p(\vec{r}|s)$. Let us assume that the firing rate $r_a$ of neuron $a$ is determined by counting $n_a$ spikes over a time interval $\Delta t$, so that $r_a = n_a/\Delta t$, and that the variability follows a Poisson distribution. This means the likelihood of a stimulus $s$ evoking $n_a = r_a \Delta t$ spikes when the average firing rate $r_a = f_a(s)$ is:

$$p(r_a|s) = \frac{1}{(r_a \Delta t)!}(f_a(s)\Delta t)^{r_a \Delta t} exp(-f_a(s)\Delta t) \tag{2}$$

b) Write down the log likelihood for the population, in which each neuron fires independently given the stimulus.

c) Differentiate the log likelihood. Hint: you can ignore **all** the terms that do not depend on $s$.

d) Set the derivative to zero to find the maximum likelihood estimate $S_{ML}$ for your stimulus given the fact that the tuning curves are Gaussian.

## 3.2

In the last question, you derived a maximum likelihood estimator for a continuous stimulus value given the firing rates of a population of neurons. In this question, we will explore further what estimators really are, and what constitutes a good estimator.

The accuracy of an estimate $s_{est}$ for a stimulus $s$ is described by two terms. The difference between the average of the estimate and the actual value is called the bias: $b_{est}(s) = \langle s_{est} \rangle - s$. The definition of the variance of an estimate is $\sigma^2_{est}(s) = \langle (s_{est} - \langle s_{est} \rangle)^2 \rangle$, and it is a measure of how much the estimate varies around its mean value. The angled brackets in these equations indicate the average.

a) Given the definitions of bias and variance, show that the mean squared error $\langle (s_{est} - s)^2 \rangle$ is equal to the sum of the variance and the squared bias $\sigma^2_{est}(s) + b^2_{est}$. Explain what this means for unbiased estimates, i.e. estimates for which $b_{est}(s) = 0$.

In the limit of a very large N, the maximum likelihood estimate (question 3.1) is unbiased. It turns out that, for unbiased estimates, we can use the estimates from our decoding method to find a limit to the accuracy with which a neural population can encode a stimulus. Through an inequality known as the Cramér-Rao bound, the variance of an estimator is limited by a quantity known as Fisher information:

$$\sigma^2_{est}(s) \geq \frac{1}{J(s)} \tag{3}$$

where $J(s)$ is the Fisher information of the firing rate distribution. This can be found by taking the negative of the expected value of the second derivative of the log likelihood:

$$J(s) = \left\langle -\frac{\partial^2 \log p[\vec{r}|s]}{\partial s^2} \right\rangle_{p[\vec{r}|s]} = \left\langle -\left( \frac{\partial \log p[\vec{r}|s]}{\partial s} \right)^2 \right\rangle_{p[\vec{r}|s]} \tag{4}$$

where $\langle x \rangle_{p[\vec{r}|s]}$ is the expectation of $x$ with respect to $p[\vec{r}|s]$, which can also be written as $\int d\vec{r}\, p[\vec{r}|s]\, \langle x \rangle$.

b) Show that the Fisher information for the population of neurons with uniformly arrayed tuning curves and Poisson statistics from question 3.1 can be written as:

$$J(s) = \Delta t \sum_{a=1}^{N} \frac{(f'_a(s))^2}{f_a(s)}, \tag{5}$$

where $f'(x)$ is the derivative of $f(x)$. Hint: you should remember what variable you take the expectation with respect to, and use the fact that the mean firing rate can be read directly off the tuning curve: $\langle r_a \rangle = f_a(s)$.

In the answer you found in question (b), you can see that each neuron contributes to the Fisher information of the stimulus to a degree proportional to the square of the slope of the tuning curve, and inversely proportional to the average firing rate for the stimulus being estimated.

c) For any particular neuron, how does the Fisher information of the stimulus relate to the tuning curve? Draw a tuning curve from a neuron of the population in question 3.1, and plot the shape of the Fisher information in the same graph.

d) Judging by your answer in question (b), what shape do you think the optimal Gaussian tuning curve shape should be in order to maximise the Fisher information for a single neuron? Should they be narrow or wide? Does your answer agree with that from Zhang and Sejnowski [ZS99]?

For biased estimators, the Cramér-Rao bound (equation 3) should be written as follows:

$$\sigma^2_{est}(s) \geq \frac{(1 + b'_{est}(s))^2}{J(s)} \tag{6}$$

e) Given the same variance, what happens to the Fisher information of a stimulus if a bias is induced in the estimator?

f) Could you imagine a situation where you would actually prefer a biased estimator over an unbiased one? Hint: think about your answer in question (a).