CO6: Introduction to Computational Neuroscience

http://iec-lnc.ens.fr/group-for-neural-theory/teaching-260

Lecturer: Matthew Chalk
matthewjchalk@gmail.com
École Normale Supérieure
29 rue d'Ulm, 2nd floor

Solutions to the 5th exercise sheet

If you have any questions regarding these solutions or the exercises in general, feel free to contact me by e-mail (matthewjchalk@gmail.com).

Math reminder

Some of you seemed to be unsure about how to compute an integral over two variables. I believe that the solution for exercise 1e provided here might be sufficiently clear to help you solve the issues you encountered. If that's not the case, don't hesitate to contact me.

Attribution of points

For each exercise sheet, you can obtain up to 20 points, plus a variable number of bonus points for the exercises labeled as 'advanced'. Points were attributed as follows

exercise	points
1a	4
1b	6
1c	4
1d	3
1e	+ 3
2	3

Signal detection theory

Exercise 1a

An example of probability distributions of neural responses associated with two stimuli, and for two different neuronsare shown in Figure 1. Note that within a panel you can choose to either represent the distributions for one neuron and the two different stimuli (shown here), or for both neurons yet only one stimulus. The latter representation has the advantage that we can easily identify hit rates and false alarm rates for a given stimulus and threshold.

Exercise 1b

The function $\alpha(z)$ represents the *false alarm rate* for a decision threshold z. The false alarm rate is the probability that the stimulus underlying the neural activity is inferred to be $(+)$, when in fact it is $(-)$.

As was shown in the lecture, the false alarm rate can be visualized as the area under $p(r|-)$ for all values of r higher than the decision threshold z. This area corresponds to the following integral

$$\alpha(z) = \int_z^\infty p(r|-)\,\mathrm{d}r, \qquad (1)$$

which corresponds to the blue shaded area depicted in this diagram:

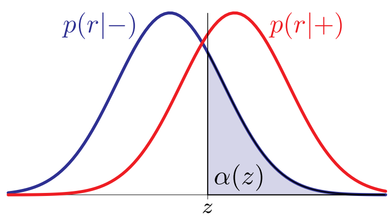

The function $\beta(z)$ represents the *hit rate* for a decision threshold z. The hit rate is the probability that the stimulus underlying the neural activity is correctly inferred to be $(+)$.

The hit rate corresponds to the area under $p(r|+)$ for all values of r higher than the decision threshold. This area is just

$$\beta(z) = \int_z^\infty p(r|+)\,\mathrm{d}r, \qquad (2)$$

and is illustrated in this diagram

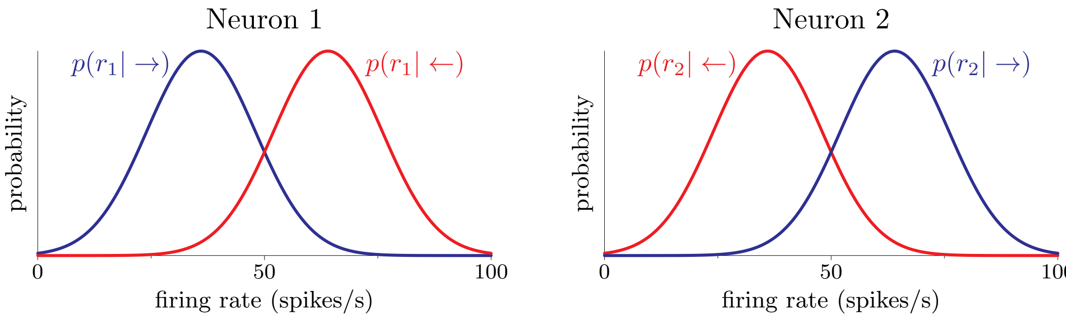

Figure 1. Example of gaussian probability distributions like those used in exercises 1a to 1e.

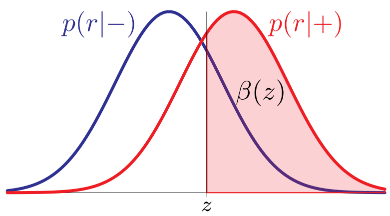

To plot the functions $\alpha(z)$ and $\beta(z)$ we simply sweep over the threshold z. For the particular distributions $p(r_1|+)$ and $p(r_1|-)$ shown in Figure 1,the false alarm rate $\alpha(z)$ and the hit rate $\beta(z)$ look like this

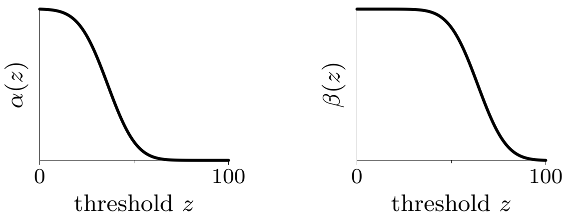

Although the homework didn't ask for it, it is worth checking that we obtain a sensible *receiver operating characteristic* (ROC) curve. Remember that this is the curve defined as the function of the false alarm rate versus the hit rate, that is, the function $\beta(\alpha)$. From the figures above, we don't have the explicit dependence $\beta(\alpha)$', but we do have the value of α and β at each threshold z. We can therefore vary continuously the value of z to trace out the ROC curve, which for the diagrams above has the shape:

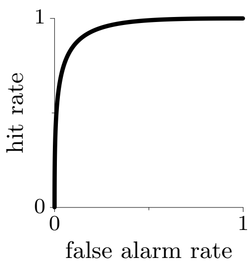

Exercise 1c

Equation (1) tells us that $\alpha(z)$ is related to the *primitive* (or *antiderivative*) of the function $p(r|-)$. Let's show this in more detail. If we denote by $P(r|-)$ an arbitrary primitive of $p(r|-)$, we have

$$\begin{aligned}
\frac{\mathrm{d}\alpha(z)}{\mathrm{d}z} &= \frac{\mathrm{d}}{\mathrm{d}z}\int_z^\infty p(r|-)\,\mathrm{d}r \\
&= \frac{\mathrm{d}}{\mathrm{d}z}\left(1-\int_{-\infty}^{z} p(r|-)\,\mathrm{d}r\right) \\
&= \frac{\mathrm{d}}{\mathrm{d}z}1-\frac{\mathrm{d}}{\mathrm{d}z}\left[P(r|-)\right]_{-\infty}^{z} \\
&= 0-\frac{\mathrm{d}}{\mathrm{d}z}\Big(P(z|-)-\lim_{r\to-\infty}P(r|-)\Big) \\
&= -\frac{\mathrm{d}}{\mathrm{d}z}\left(P(z|-)-0\right) \\
&= -p(z|-). \qquad (3)
\end{aligned}$$

Analogously to $\alpha(z)$, β can be regarded as the primitive of the function $p(r|+)$:

$$\frac{\mathrm{d}\beta(z)}{\mathrm{d}z} = \frac{\mathrm{d}}{\mathrm{d}z}\int_z^\infty p(r|+)\,\mathrm{d}r = -p(z|+). \qquad (4)$$

Exercise 1d

The probability of observing a particular combination of firing rates r_1 and r_2 as a result of a particular stimulus (for example, a movement to the left) is given by:

$$p\left(r_1, r_2|\,\mathrm{left}\right) = p\left(r_1|\,\mathrm{left}\right)\cdot p\left(r_2|\,\mathrm{left}\right) \qquad (5)$$

To obtain a more concise notation, we would like to replace the designation of the stimulus by the symbols $(+)$ and $(-)$ we used before.

For the probabilities on the right-hand side of the above equation, this is easy: A leftward movement is the preferred stimulus of neuron 1, and we can thus write $p\left(r_1|\,\mathrm{left}\right) = p\left(r_1|+\right)$. For neuron 2, this is the non-preferred stimulus, and we can write $p\left(r_2|\,\mathrm{left}\right) = p\left(r_2|-\right)$.

But what are we going to do about $p\left(r_1, r_2|\,\mathrm{left}\right)$? Well, we will simply have to define what we mean by

(+) and (−) in this context. For example, we may decide to use (+) when the stimulus is preferred by neuron 1. Equation (5) then becomes:

$$p\left(r_1, r_2|+\right) = p\left(r_1|+\right) p\left(r_2|-\right).$$

Arguably, this notation is tricky, but you will soon get used to it.

Advanced exercise: Exercise 1e

The most general way to solve this exercise is to assume that r_1 and r_2 are statistically dependent on each other, a case that we will consider on the next page. Before doing so, though, we will present the simpler case in which r_1 and r_2 are independent, which has a particularly intuitive solution. Both ways of addressing the problem were accepted in the homework.

The expression for the hit rate $\beta(z)$ in (2) assumes that we have a constant decision threshold z. What happens if we decide that the stimulus (+) was present whenever $r_1 > r_2$?

Special case: Mutually independent firing rates

In this case, we have a different threshold in each trial, and this threshold is $z = r_2$. So for all trials, we have to sum up their contribution to the probability of making a correct inference:

$$p_{\text{correct}} = \frac{1}{N} \sum_{t=1}^{N} \int_{r_2(t)}^{\infty} p(r_1|+) \, \mathrm{d}r_1$$

To obtain the exact value of β, we have to take the limit of an infinite number of trials. In other words, we have to integrate over all possible values of r_2, weighted by their probabilities:

$$\begin{aligned} p_{\text{correct}} &= \lim_{N\to\infty} \left(\frac{1}{N} \sum_{t=1}^{N} \int_{r_2(t)}^{\infty} p(r|+) \, \mathrm{d}r_1 \right) \\ &= \int_0^{\infty} \left(p(r_2|-) \int_{r_2}^{\infty} p(r_1|+) \, \mathrm{d}r_1 \right) \mathrm{d}r_2. \end{aligned}$$

Inside the integral over r_2, we need to weigh the contribution of each particular value of r_2 with the probability $p(r_2|-)$ of that particular firing rate, because some firing rates will appear more often than others in the sum over all trials.

The integral above can be simplified using the definition (4) of the hit rate β:

$$p_{\text{correct}} = \int_0^{\infty} p(r_2|-)\beta(r_2) \, \mathrm{d}r_2. \tag{6}$$

To demonstrate that this is the same as the surface under the ROC curve, we will have to find a mathematical expression for that surface. The ROC (or receiver operating characteristic) curve is the curve describing the hit rate as a function of the false alarm rate, $\beta(\alpha)$. The area under it is

$$\int_0^1 \beta(\alpha) \, \mathrm{d}\alpha. \tag{7}$$

To compare equations (6) and (7), we will have to transform the integral in (6) to make it an integral over α. According to equation (3):

$$p(r_2|-) = -\frac{\mathrm{d}\alpha}{\mathrm{d}r_2}$$

and thus, using the integration-by-substitution rule, we can rewrite equation (6):

$$\begin{aligned} p_{\text{correct}} &= \int_0^{\infty} \left(-\frac{\mathrm{d}\alpha}{\mathrm{d}r_2} \beta(r_2) \right) \mathrm{d}r_2 \\ &= -\int_{\alpha(r_2=0)}^{\alpha(r_2=\infty)} \beta(r_2) \, \mathrm{d}\alpha \\ &= \int_0^1 \beta(\alpha) \, \mathrm{d}\alpha \end{aligned}$$

which is the same as the area under the ROC curve described by equation (7).

General case: Mutually dependent firing rates

In the general case, we must renounce the above arguments concerning the *summation* of different contributions. But we can obtain the same result by integrating $p(r_1, r_2|+)$ over half of the plane spanned by r_1 and r_2, where $r_1 > r_2$. We then have:

$$\begin{aligned} p_{\text{correct}} &= \int_0^{\infty} \left(\int_{r_2}^{\infty} p(r_1, r_2|+) \, \mathrm{d}r_1 \right) \mathrm{d}r_2 \\ &= \int_0^{\infty} \left(\int_{r_2}^{\infty} p(r_1|+)p(r_2|-) \, \mathrm{d}r_1 \right) \mathrm{d}r_2 \\ &= \int_0^{\infty} \left(p(r_2|-) \int_{r_2}^{\infty} p(r_1|+) \, \mathrm{d}r_1 \right) \mathrm{d}r_2 \\ &= \int_0^{\infty} p(r_2|-)\beta(r_2) \, \mathrm{d}r_2 \end{aligned}$$

which is the same as equation (6). From here on, the solution of the general case is the same as the one of the special case presented above.

Reinforcement learning theory - Exercise 2

The reinforcement learning framework laid out in the lecture is powerful, but there are certain environmental paradigms with which it cannot deal. Remember that we developed this theory based on the assumption that the processes we are dealing with can be described as *Markov chains*. In other words: the future states of our system, and their values, only depend on the present state. They do not depend on the past —the *history*— of the system.

Imagine a rat exploring a maze (= a labyrinth) in which the rewards at the different exits vary randomly. This is an example for a Markov process, because the same random variations could occur on a Monday or a Tuesday.

Now imagine a maze in which the rewards are varied deterministically as a function of the day of the week. On Mondays, the biggest reward is hidden behind exit 1, on Tuesdays, it is behind exit 2, and so on. In order to learn this paradigm, the rat has to establish a memory of the previous seven days. In other words, it will not be able to infer the rewards of tomorrow from today's events alone—it has to take into account the recent history of events.

Although we can easily imagine the rat to achieve this, we cannot describe this kind of learning within the framework of reinforcement learning we have developed. The only way to work around this issue would be to expand our state space to include an explicit representation of time. This way, the state "I'm standing at the first intersection to the left" would *not* be the same state on Mondays and Tuesdays anymore, and reinforcement learning could take over.