CNeuro 2021

Estimating the information in spike trains

Vijay Balasubramanian

A Introduction

In this assignment, you will study the process of estimating entropy and mutual information from empirical data. First, you will derive a few of the basic properties of the entropy and mutual information for discrete variables. You will then simulate an idealized neural recording dataset, and then estimate entropy and mutual information from it.

We will consider the following scenario: an experimenter is recording from two neurons in visual cortex, while the animal is presented with one out of two possible sensory stimuli — drifting gratings moving in different directions. The experiment is repeated N times, and at each time both stimuli are shown, one after the other. In each experimental trial, the response of neuron n to stimulus k is well described by a Poisson distribution

$$r_n \sim \text{Poisson}(\lambda_k^n)$$

with the following mean:

	Neuron 1	Neuron 2
Stimulus 1	4 spikes	5 spikes
Stimulus 2	6 spikes	5 spikes

In other words, neuron 1 responds with 4 or 6 spikes on average, depending on the stimulus, while neuron 2 always responds with 5 spikes on average, regardless of the stimulus. As the response of neuron 2 does not depend on the stimulus, we will call neuron 2 the *uninformative* neuron, and neuron 1 the *informative* one.

B Properties of Entropy and Mutual Information

The *entropy* of a random variable measures the diversity found in samples of that variable. For a discrete random variable R it is defined as

$$H[R] = -\sum_r p(r) \log p(r) \tag{1}$$

where r are the various values that R can take. As described in class, the entropy quantifies the average number of bits of information that we gain by observing R, and the log is 2-based. Equivalently, the entropy is related to the logarithm of the number of different messages that can be sent in a sequence of x values that are distributed according to $p(r)$.

If S is another random variable, the *conditional entropy* of R given S is defined as the average entropy of the conditional distributions $p(r|s)$ over all values s that S can take:

$$H[R|S] = -\sum_s p(s) \sum_r p(r|s) \log p(r|s) \tag{2}$$

Suppose we think of S (stimulus) as the inputs to a neuron and R (response/rate) as the outputs of the neuron. Then, the Shannon's Mutual Information between R and S measures the amount of information that R provides about S, quantified as the entropy of R (i.e. number of bits of information transmitted on average by each neural response), minus the bits of information that were lost to noise because of the variable responses r that can happen for identical inputs s:

$$I[S:R] = H[R] - H[R|S] \tag{3}$$

As we discussed in class, we can use the formulae for entropy and the conditional entropy to write the mutual information as:

$$I[S:R] = \sum_{r,s} p(r,s) \log \frac{p(r|s)}{p(r)} \tag{4}$$

Problem 1:

(**a**) Show that $H[R] \geq 0$ for any discrete random variable R. That is, show that entropies are nonnegative for discrete random variables. For completeness, note that this does not hold for the *differential entropy* of a continuous random variable.

(**b**) Show that $I[R:S] = I[S:R]$ for any R and S. That is, show that the Mutual Information is symmetric between the input and the output. We can also say that the amount of information carried by R about S is the same as the amount of information carried by S about R. Hint: to show this you need to use Bayes' Rule in the expression for the mutual information that we gave in class.

(**c**) Show that $0 \leq I[R:S] \leq \min(\{H[R], H[S]\})$. In words, the amount of information that R can carry about S, or vice-versa, is nonnegative and is bounded from above by the total uncertainty/variability of R or S, whichever is smallest. One way to prove the first inequality is to use the fact that the logarithm is a convex function of its arguments. (This is a famous proof, so feel free to look it up if you want, understand it, and then repeat it in your own words.) The second inequality can be proven from the definition of the mutual information given above and the result in (b).

C Simulating The Trial Results

Next we will simulate the responses of the two neurons described in the Introduction responding to Stimuli 1 and 2. Then, we will use this "experimental" data to calculate the empirical distributions of the responses. Finally we will use these sampled distributions (instead of the true Poisson distributions, which we know in this case) in the formulae for entropy and mutual information.

Most scientific programming environments have built-in methods for sampling from standard distributions like the Poisson. Here is how to generate an N-by-1 array of values sampled from a Poisson distribution with mean `mean_n_spikes` in Python and Matlab/Octave. Python:

```
import numpy as np
spikes = np.random.poisson(lam=mean_n_spikes, size=(N,1))
```

Matlab/Octave:

```
spikes = poissrnd(mean_n_spikes, N, 1);
```

You can use the method above to sample the response of a give neuron to a specific stimulus in all trials. For instance, if `mean_n_spikes=6` then `spikes` will be an array of N integer numbers that represent the number of spikes fired by neuron 1 upon presentation of stimulus 2 across the N trials.

Problem 2:

(**a**) Simulate the spike counts in one run of the experiment with $N = 100$ trials. Note that the Poisson random variable here is the number of spikes in one trial, since we assume that the trials are Poisson processes with the same mean when the external parameters are fixed. Display a plot of the 100 spike counts for neuron 1, for stimulus 1 and stimulus 2.

(**b**) For each neuron, and for each value of the stimulus, build a histogram of the relative frequency of each possible response. These histograms estimate the conditional distributions $p(r^n|s = s_k)$. You should find that the conditional response histograms for neuron 1 are different, while the responses of neuron 2 to the two stimuli look more similar.

(**c**) Build the "marginal" histogram of the overall response of the neuron during the experiment. That is, estimate the distribution of responses $p(r^n)$ without conditioning on the stimulus. Note that the each stimulus is presented with equal probability $p(s_1) = p(s_2)$, because the number of trials N is fixed.

D Estimating Entropy and Mutual Information

Now that you have gathered the empirical distributions of the responses, you are well equipped to compute the entropy, conditional entropy and the mutual information between the neurons and the stimuli.

Problem 3: For problems (a)-(c) below, if your histogram contains exactly zero samples in some bin, you can set $p(x) \log p(x) = 0$ for that bin when you compute the entropy. Convince yourself that this is trues because $x \log x \to 0$ for $x \to 0^+$, as can be shown by application of l'Hôpital's rule.

(**a**) By plugging the marginal distribution you obtained into the definition of entropy in Sec. B, compute the entropy of the response $H[R]$ for each neuron.

(**b**) Similarly, plugging the conditional distribution you obtained to compute the conditional entropy given the identity of the stimulus $H[R|S]$ for each neuron.

(**c**) From (a) and (b), compute the mutual information $I[R : S]$ for each neuron. You should find that the mutual information is larger for neuron 1 than neuron 2. Explain why this is expected.

Now compare the empirical mutual information with the following analytic calculation:

(**d**) Calculate analytically the mutual Information $I[R : S]$ between the responses of neuron 2 and the stimulus. You can do this by employing the definition of the mutual information and the known Poisson distributions of the responses for stimuli 1 and 2. Compare the result you get to the empirical estimate you obtained in (c) from the simulated data. Why do you think the true analytical value differs from the estimate?

(**e**) Repeat the simulation 10 times, and report the values of the mutual information that you get in each case (you don't have to show the response histograms again; just report the values of the mutual information that you get in each run of the simulation). Report the results as a mean $\pm$ standard deviation. You will find that the difference between the empirical estimate and the true value of the mutual information for neuron 2 is nonzero even after averaging over multiple realizations of the experiment. In other words, the mutual information estimate is **biased** - what you are seeing here is an example of the so-called undersampling bias in the estimate of mutual information (or more generally, entropy-related quantities). In class we discussed a method for extrapolation from finite data to attempt to correct for this bias in estimates of entropy and mutual information from spike trains.