

2 Spikes and Reliability

We now switch gears and consider just well animals do when their brains communicate with action potentials. In particular, we examine how well neurons in various mammals faithfully reproduce aspects of their sensory environment. We are clearly jumping over a large philosophical gap by ignoring the issue of whether the nervous system actually wants all the information in an environment, since what animals do best are to make simple behavioral decisions based on a complex sensory stream that reports features in the environment. But we can philosophize plenty when we get too old to measure things.

2.1 Coding of Angular Velocity by H1 in the Fly

The first system is the encoding of angular position in the fly. We consider the response of cell H1 in the lobular plate. These are neurons with very large receptive fields that are excited by back-to-front motions across the visual field and inhibited by front-to-back motions, i.e., yaw sensitive. They have essentially no response to vertical motion, i.e., pitch insensitive. There are two H1 cells, one for each side of the head, and these are arranged in antagonism so that rigid rotation of the fly in a static background excites one cell and inhibits the other.

FIGURE - chapt-2-fly-1.eps

FIGURE - chapt-2-fly-2.eps

These cells function to convey information about the rotation of the entire fly during flight. They are three synapses removed from the photoreceptors, i.e., optic ganglion \rightarrow lamina \rightarrow medulla \rightarrow lobula. There is a strong resemblance of this cell and its input structure to type 1 cells in the colliculus of vertebrates; just even the mammal-centric among us can love them. While we focus only on the sensory part, the output from H1 is ultimately combined with sensory input from wind hairs in the segmental ganglia in the cord of the fly. The error signal controls the amplitude of the wing beats and thus the steering behavior of the animal.

FIGURE - chapt-2-fly-h1.eps

We consider the experiments of Rob deRuyter, who rotated the visual world of a fly that was tethered, i.e., pasted to a stick, so that fly became a function! Rob measured the spike output from H1 for many epochs of the stimulus. Thus he had two pieces of information,

- $V_{app}^k(t)$ is the applied visual stimulus for the k-th trial.
- $S_{meas}^k(t) = \sum_{spike\ times} \delta(t - t_S^k)$ is the measured spike time for the k-th trial.

The key is to ask if we can use this measured information to construct a filter that allows us to predict the stimulus for an unknown trial from the measured spike train. In a sense, we ask, "How well can we reconstruct the stimulus from the spikes". We consider this via a linear transfer function - a.k.a. optimal filter - a.k.a. - linear kernel, an idea that is at least a century old, although it came into use only at the time of WW II when there was a big push at the MIT Radar Laboratory to formulate the mathematics of optimal filtering and prediction. The procedure is as follows:

- We define $T(t)$ as the sought after transfer function.
- We define $V_{pred}^k(t)$ as the predicted stimulus for the k-th trial, based on the measured spike train, where

$$\begin{aligned}
V_{pred}^k(t) &= \int_{-\infty}^t dt' T(t-t') S_{meas}^k(t') & (2.1) \\
&= \int_{-\infty}^t dt' T(t-t') \sum_s \delta(t' - t_S^k) \\
&= \sum \int_{-\infty}^t dt' T(t-t') \delta(t' - t_S^k) \\
&= \sum_{\text{spike times}} T(t - t_S^k)
\end{aligned}$$

is the predicted output.

FIGURE - chapt-2-convolution.eps

To get $T(t)$, we minimize the difference between the actual and the predicted stimulus, averaged over all trials and time, i.e.,

$$\begin{aligned}
Error &= \sum_k \int dt \left(V_{pred}^k(t) - V_{app}^k(t) \right)^2 & (2.2) \\
&= \sum_k \int dt \left(\int_{-\infty}^t dt' T(t-t') S_{meas}^k(t') - V_{app}^k(t) \right)^2
\end{aligned}$$

The error is computed in terms of measured quantities, except for $T(t)$, which we find by the criteria that we want to choose $T(t)$ to minimize the error. This is much easier to solve in the frequency domain, where convolutions turn into products. As a mathematical aside, we consider the Fourier transformed variables:

$$V_{app}^k(t) \iff \tilde{V}_{app}^k(f) \quad (2.3)$$

$$S_{meas}^k(t) \iff \tilde{S}_{meas}^k(f) \quad (2.4)$$

$$V_{pred}^k(t) \iff \tilde{V}_{pred}^k(f) \quad (2.5)$$

$$T(t) \iff \tilde{T}(f) \quad (2.6)$$

$$(2.7)$$

where

$$\tilde{T}(f) = \int_{-\infty}^{\infty} dt e^{i2\pi ft} T(t) \quad (2.8)$$

$$T(t) = \int_{-\infty}^{\infty} df e^{-i2\pi ft} \tilde{T}(f) \quad (2.9)$$

so that (ignoring causality for the moment - hey, it's only biology) the convolution becomes

$$\int_{-\infty}^{\infty} dt' T(t-t') X(t') = \tilde{T}(f) \tilde{X}(f) \quad (2.10)$$

and the other relation we need is Parseval's theorem, effectively a conservation of energy, i.e.,

$$\int_{-\infty}^{\infty} dt |T(t)|^2 = \int_{-\infty}^{\infty} df |\tilde{T}(f)|^2 \quad (2.11)$$

where $|\tilde{T}(f)|^2 = \tilde{T}(f) \tilde{T}^*(f)$.

We put the above together to write:

$$\begin{aligned} Error &= \sum_k \int df |\tilde{V}_{pred}^k(f) - \tilde{V}_{app}^k(f)|^2 \quad (2.12) \\ &= \int df \sum_k \left(|\tilde{V}_{pred}^k(f) - \tilde{V}_{app}^k(f)|^2 \right) \\ &= \int df \sum_k \left(|\tilde{T}(f) \tilde{S}_{meas}^k(f) - \tilde{V}_{app}^k(f)|^2 \right) \\ &= \int df \sum_k \left(\tilde{T}(f) \tilde{S}_{meas}^k(f) - \tilde{V}_{app}^k(f) \right) \left(\tilde{T}^*(f) \tilde{S}_{meas}^{k*}(f) - \tilde{V}_{app}^{k*}(f) \right) \\ &= \int df \sum_k \left(\tilde{T}(f) \tilde{T}^*(f) |\tilde{S}_{meas}^k(f)|^2 - \tilde{T}(f) \tilde{S}_{meas}^k(f) \tilde{V}_{app}^{k*}(f) \right. \\ &\quad \left. - \tilde{V}_{app}^k(f) \tilde{T}^*(f) \tilde{S}_{meas}^{k*}(f) + |\tilde{V}_{app}^k(f)|^2 \right) \\ &= \int df \tilde{T}(f) \tilde{T}^*(f) \sum_k |\tilde{S}_{meas}^k(f)|^2 - \int df \tilde{T}(f) \sum_k \tilde{S}_{meas}^k \tilde{V}_{app}^{k*}(f) \\ &\quad - \int df \tilde{T}^*(f) \sum_k \tilde{V}_{app}^k \tilde{S}_{meas}^{k*}(f) + \int df \sum_k |\tilde{V}_{app}^k(f)|^2 \end{aligned}$$

The next step is to minimize the error with respect to the transfer function. We compute the function derivative

$$\frac{\partial(Error)}{\partial \tilde{T}^*(f)} = 0 \quad (2.13)$$

so that

$$\begin{aligned} 0 &= \int df \tilde{T}(f) \sum_k |\tilde{S}_{meas}^k(f)|^2 - \int df \sum_k \tilde{V}_{app}^k \tilde{S}_{meas}^{k*}(f) \quad (2.14) \\ &= \int df \left(\tilde{T}(f) \sum_k |\tilde{S}_{meas}^k(f)|^2 - \sum_k \tilde{V}_{app}^k \tilde{S}_{meas}^{k*}(f) \right) \end{aligned}$$

The expression for $T(f)$ must be valid at each frequency. Thus the frequency representation of the transfer function is

$$\tilde{T}(f) = \frac{\sum_k \tilde{V}_{app}^k(f) \tilde{S}_{meas}^{k*}(f)}{\sum_k |\tilde{S}_{meas}^k(f)|^2} \quad (2.15)$$

This is the central result. For the case of measured signal that is a spike train,

$$\tilde{T}(f) = \frac{\sum_k \tilde{V}_{app}^k(f) \sum_s e^{i2\pi f t_s^k}}{\sum_k \sum_{s,s'} e^{i2\pi f (t_s^k - t_{s'}^k)}} \quad (2.16)$$

In the time domain, this is just

$$T(t) = \int df e^{-i2\pi ft} \frac{\sum_k \tilde{V}_{app}^k(f) \sum_s e^{i2\pi f t_s^k}}{\sum_k \sum_{s,s'} e^{i2\pi f (t_s^k - t_{s'}^k)}} \quad (2.17)$$

Ugly! But we notice that this has a simple form when the spike arrival times may be taken to be a random, e.g., Poisson variable. This occurs if the spike rate is not too high, so that the refractory period plays no role. In this case the denominator is just

$$\sum_k \sum_{s,s'} e^{i2\pi f (t_s^k - t_{s'}^k)} \approx N \quad (2.18)$$

where N is the total number of spikes across all trials, and the denominator is just

$$\begin{aligned} \int df e^{-i2\pi ft} \sum_k \tilde{V}_{app}^k(f) \sum_s e^{i2\pi f t_s^k} &= \sum_k \sum_s \int df e^{-i2\pi f (t - t_s^k)} \tilde{V}_{app}^k(f) \\ &= \sum_k \sum_s V_{app}^k(t - t_s^k) \end{aligned} \quad (2.19)$$

Thus $T(t)$ is just the spike triggered average of the stimulus waveform, i.e.,

$$T(t) \approx \frac{1}{N} \sum_k \sum_s V_{app}^k(t - t_s^k) \quad (2.20)$$

and finally we see that all that happens is that the transfer function reports the waveform of the stimulus that is most likely to cause the neuron to fire.

FIGURE - chapt-2-fly-h1.eps (again!)

Rob and Bill Bialek calculated for transfer function for Rob's data, It looks largely like a derivative in time; this is not unreasonable for a system that is part of looking for deviations in the visual field. The fit of the prediction to the measured data is good. It is better for low frequencies than high frequencies. In fact, we expect that the filter should attenuate at time scales shorter than it's width of ≈ 40 ms, or in other words, at frequencies above the inverse of the width, or 2 Hz. Bialek did estimate. He quantified, as an average over trials in the frequency domain, the

error between the spectral response and the observed behavior. Not surprisingly, the error exceeds the signal at high frequencies, above 25 Hz in this case.

FIGURE - chapt-2-fly-h1-error.eps

In general, the fly can initiate visually guided course corrections in 30 ms, so the ≈ 20 ms width of $T(t)$ is consistent with this. It is also interesting to point out that the maximum rate that this cell fires is about 100 - 200 Hz. This gives 2 to 4 spikes per width, which is about the same as the correlation time. More typically, we are down to one spike per correlation time.

A final issue is that of causality. The filter here has been shifted away from the origin. so that $T(t < 0) = 0$. This is a potentially dirty trick. The critical thing is that there is a delay until the information gets out. It is likely that this delay is the time required to initiate a visually guided flight correction. A side remark is that a related technique, called reverse correlation, gives a causal filter directly.

2.2 Coding of Vibrissa Motion in S1 of Rat

The above example shows that a single neuron may have a high degree of reliability of relaying the status of the sensory world. One obvious possibility is that H1 in the fly is a very special cell, especially given that it is a singular neuron. We now consider an alternate example from the mammalian cortex, for which thousands of cells may respond to the same stimulus. In particular, we consider the case of neurons in the primary somatosensory (S1) cortex of rat that respond to motion of the vibrissae during a task in which the animal whisks in air. It was shown that the probability of spiking by these cells is proportional to the position of the vibrissae. Here, in analogy with the case for H1, we ask how well the spike trains from such cells can be used to predict the position of the vibrissae.

FIGURE - chapt-2-rat-1.eps

FIGURE - chapt-2-rat-2.eps

FIGURE - chapt-2-rat-3.eps

The first result is that the transfer function appears tuned, i.e., rather than a flat function or a bump, the transfer function appears as a dampened oscillation. The damping, or correlation time, is about 150 ms, slightly longer than the period of whisking. The second result is a single spike train can be used to predict the position of the vibrissae with excellent fidelity. In fact, if one looks only at the frequency for whisking, the magnitude of the coherence between the predicted and the measured signal is greater than 0.9. From a functional point of view, this results shows that some cells in S1 (about 1 out of 20) have high fidelity information about the position of the vibrissae. Further, it speaks against the necessity of a "distributed code" as a means to represent the sensory stream.

FIGURE - chapt-2-rat-4.eps

As in the case of H1, the high degree of reliability can be understood in terms of a sufficiently high spike rate, i.e., about one spike per correlation time, and a sufficiently high correlation between spiking and the EMG for whisking. In fact, an additional measurement, in which rats were trained to use their vibrissae to detect the presence of an edge, demonstrates that neurons in S1 produce about 1 spike within the time it takes the rat to make a decision. In this case, we plot the integrated number of spikes as the rat probes the edge, and note that she turns to make her decision after close to a single spike has accumulated.

FIGURE - chapt-2-rat-contact.eps

2.3 Comparison of Spike Coding and Behavioral Choice for Random Dot Motion

We now turn to a final example and ask how a sensory spike signal may be related to the behavior, or psychometric judgment, of an animal. The importance of this is that the animal is likely to pool the response of many sensory input neurons if the behavioral response is much better than the response of a given cell. This is a vote in favor of a "distributed code" across primary sensory cells. On the other hand, it is certainly possible that the output of a single neuron further down the sensory stream is an excellent predictor of behavior. This was hinted at in our discussion of the rat experiments, but there was no measure of correct or incorrect behavior in that case.

We consider Newsome's experiments on directional tuning in the visual cortex of monkey, and focus on the response of neurons in area MT. These cells receive input via two pathways, one is the retina \rightarrow thalamus \rightarrow primary visual (V1) cortex \rightarrow \dots \rightarrow medial temporal (MT) cortex pathway. The other is the retina \rightarrow colliculus \rightarrow pulvinar \rightarrow MT pathway.

FIGURE - chapt-2-monkey-1.eps

FIGURE - chapt-2-monkey-2.eps

The monkey is trained to fixate on a spot and then is presented with a random dot pattern. The motion of the dots has a large random component and a small biased component in one of two directions. The monkey must decide which direction the bias is in, and at the same time the spiking output from neurons in MT is recorded and used to calculate the preference.

FIGURE - chapt-2-monkey-3.eps

In this experiment, only cells that had a strong directional preference were selected. From about 50 to 100 trials of data, both psychometric and physiometric decision curves were constructed. For the psychometric case, the curve was simply the probability of a correct choice. For the physiometric case, it was (very) roughly the probability of the number of spikes in a 2 second interval coming from the null or unresponsive direction versus the preferred direction. Similar in spirit to the

above cases, the data from all trials was used to construct the distribution of response. Then, a maximum likelihood measure was used to determine if the response on a given trial best fit the distribution for the null direction versus the preferred direction. This was matched against the psychometric response.

FIGURE - chapt-2-monkey-4.eps

The essential result is that the average output from a single neuron could be used to predict the direction of the dots with a reliability that essentially equaled the psychometric choice of the animal.

2.4 Addendum on Transfer Functions

The error between the predicted performance and the measured performance may be quantified in terms of the variance of the transfer function. In particular, only a limited spectral range of stimuli may be used. One then calculates the error band for the transfer function, something one can do by a brute force method, like jackknife statistics, or in terms of asymptotic estimates of the variance.