

# **BRIGHTNESS DECODING IN THE FLY VISUAL SYSTEM: SPIKE TIMING AND LINEAR APPROXIMATION**

Jonathon Shlens

28 June 2002

## **I. Introduction**

The fly visual system is a popular system of study in neuroscience, because of its reduced scale and the robustness of the experimental preparation. These favorable conditions have made it possible to readily identify individual neurons in the fly's visual information processing system [1]. One prominent neuron of study is the H1 neuron, a neuron roughly analogous to a neuron in primary visual cortex in humans. This neuron responds preferentially to horizontal motion across the fly's visual field, obeying a tuning curve dependent on the angular velocity, contrast and average intensity of a visual stimulus.

Recording from the H1 neuron while manipulating the full-field visual stimuli presented to the fly, for protracted periods of time, provides a means of examining how the neuron, and presumably the fly, encode visual information in spike trains. Because the tuning curve is dependent on a multitude of parameters, it is necessary to judiciously select only a subset of parameters to examine experimentally while holding the rest constant. This set of experiments chooses to focus on how the H1 neuron encodes full-field luminance changes. Although probably only a peripheral function of H1, luminance changes do provide visual clues about the full field horizontal motion in the environment.

These set of experiments specifically examine two main questions (1) at what temporal precision does the H1 neuron reliably place spikes; (2) how well can the H1 neuron be characterized as a linear filter in response to luminance changes. The first question addresses an issue of time scale. Namely, understanding the temporal precision of the spike trains provides a rough upper bound on the temporal precision of neurons downstream of H1. Furthermore, a measure of the temporal precision provides a rough time scale for quantifying the information flow of the system. The second question addresses the issue of how well can the neuron be modeled as a linear system. Any variability not attributable to a linear reconstruction represents hysteresis effects due to

previous neural activity. In other words, testing the linearity of the H1 neuron response addresses how well the neuron resembles the idealized Poisson process [5].

## **II. Methodology**

The two main goals of this experiment are to (1) systematically manipulate the full field visual imagery of the fly and (2) concurrently record the extra-cellular spiking activity of the H1 neuron. The experimental set-up, including the fly preparation and visual stimulus, have been developed and discussed in detail already [3]. Therefore, a brief overview of these aspects will suffice and this section will instead focus on the specific experimental procedures used to address the two main questions of this project.

### **A. Fly Preparation**

All of the following experiments make use of the blowfly, *Sarcophaga bullata*. The preparation of the blowfly consists of anesthetizing it in a freezer, removing all its appendages (including the probiscus), and cementing it down to a metallic platform using melted dental wax. In addition, drops of dental wax are placed on the stubs of the wings and probiscus in order to minimize movement artifacts. Finally, the posterior portion of the head is exposed by bending down and affixing the head to the thorax.

The second step is to prepare the fly for the electrophysiology. This consists of cutting open the posterior portion of the head contralateral to the eye being examined, removing intervening air sacs and applying a special “fly saline” to the exposed area [4]. It is now possible to poke a ground electrode in the abdomen and place a 5 M-Ohm 0.005 inch diameter tungsten-coated electrode into the exposed area of the head as diagrammed previously [3]. Presenting the appropriate visual stimulus pattern functionally identifies the H1 neuron.

### **B. Visual Stimulation**

Figure 1 diagrams the set-up of the visual presentation to the fly. The fly platform resides in the center of an LED (light-emitting-electrode) arena that subtends 288 horizontal degrees of the fly visual system. The green LEDs are arranged in 96 vertical columns, each column individually encompassing 3 degrees of visual imagery. Each LED column can be switched on and off in sequential patterns via a digital I/O card connected to a

computer. Several scripts in LabView ® developed by Evren Tumer both record the electrode potential and synchronize this with the controls for the LED lights. These LabView scripts provide a means of manipulating the contrast, average intensity and angular velocity of the visual stimulus.

This experiment focused on solely modulating the contrast of the visual stimulus. All of the experiments used the full-field LED presentation pattern as diagrammed in figure 2. Every other LED column was permanently left off while the rest of the LED lights modulated their brightness level. This full-field pattern delivers an overall intensity of light that optimizes the H1 response. All of the experiments modulated the full-field brightness of the LED lights by sampling brightness values over time from the Lorenz system of differential equations:

$$\begin{aligned}\frac{dx}{dt} &= \sigma(y - x) & \sigma &= 16 \\ \frac{dy}{dt} &= rx - y - xz & \text{where } r &= 45.92 \\ \frac{dz}{dt} &= xy - bz & b &= 4\end{aligned}$$

This set of coefficients places the Lorenz system in a chaotic attractor with a positive Lyapunov exponent and a “1/f” power spectrum. Two important properties of this chaotic system are that each dynamical variable (e.g.  $x(t)$ ,  $y(t)$ ,  $z(t)$ ) has only short-term correlations and provides a good sampling of its phase space. Thus, all of these experiments use a normalized form of  $x(t)$  as the brightness value for the visual stimuli.

### **C. Procedures**

The first goal of the experiment is to examine the reliability and precision of the H1 neuron. The traditional method of approaching this problem is to compare the spike train responses of H1 to multiple trials of the exact same stimulus. The first set of experiments accomplished this by presenting the same 5 second sequence of brightness modulation using  $x(t)$  from the Lorenz system. Repeating a full-field, brightness-modulated stimulus generates multiple spike trains, all presumably responding to the same stimulus. Examining the variability in the ensuing spike train provides a measure of the reliability and precision of the spiking activity of the neuron.

The second goal of this paper is to test how well the brightness encoding of the H1 neuron can be modeled by an optimal linear filter. A protracted version of a roughly uncorrelated stimulus (e.g. the Lorenz system) provides a good sampling of response space. As discussed further in the analysis section, this data set can be used to characterize the linear impulse response (or equivalently, the frequency response) of a system. In this set of experiments, brightness was modulated according to the Lorenz system for 26 minutes, and the corresponding spiking activity was recorded.

### **III. Analysis**

The end result of the procedures outlined above is a set of two raw data files: (1) an input time series consisting of pattern brightness and (2) the recorded extra-cellular potential. Several steps of analysis are necessary to extract relevant information from the data sets and answer the questions posed in this paper. The first steps consist of spike identification and a preliminary characterization of the response with a firing rate. The main goal of this preliminary analysis is to identify the spikes associated with H1 neuron. Finally, the main thrust of the analysis is to calculate the linear filter associated with the H1 neuron. The derivation and calculation of the linear filter, also called the “spike triggered average,” is discussed in full below.

#### **A. Spike Identification and Sorting**

A raw time series of the extra-cellular potential appears much like figure 3. There exist sharp upward or downward deflections (depending on the polarity of the amplifier) corresponding to action potentials recorded in the extra-cellular medium. Hopefully, due to good electrode placement and functional identification, the prominent spikes in the time series are from the H1 neuron. Isolating the timing of neural spikes is a simple procedure. A common method replicated here is to take the derivative of the time series and select all points, which cross a particular threshold. Generally, this threshold is subjective and dependent on the noise level and how well isolated the neuron(s) are in the data set.

The set of isolated spike times and spike amplitudes consists of spiking activity from H1, spiking activity from other neurons, movement artifacts and even “double detections” of single spiking events. Movement artifacts as shown in figure 4 are readily identifiable in the data set. It is easiest to remove these artifacts manually. Double detections can be removed by eliminating all spikes that fall below a minimal interspike interval (ISI). In practice, a minimum ISI of 0.7 ms works well. Spike sorting - removing spikes due to other cells - is naively, if imperfectly, performed by identifying the appropriate range of spike amplitudes that appear to correspond to H1 activity. Again, this range is subjective and best determined for each individual experiment.

## **B. Calculating Firing Rates**

The true firing rate is an instantaneous measurement at best *estimated* from the data set. The estimated firing rate is not displayed in any of the figures; however, it is extensively used in order to functionally identify the H1 neuron. The standard method is to impose a binning resolution (in practice, ~100 ms) and calculate the number of spikes that occur on average a particular time after the onset of the stimulus [5]. This firing rate  $r(t; \Delta t)$  is a function of time and an implicit function of the binning resolution  $\Delta t$ . Assuming a constant stimulus, the mean and standard deviation characterize this distribution to provide an average firing rate and associated error bar for the duration of the stimulus presentation.

## **C. Constructing a Linear Filter: Spike Triggered Averages**

Modeling a system with a linear filter is common to many aspects of engineering and physics. Thus, each domain has renamed this tool: a first order Wiener or Volterra kernel in mathematics, a Green's function in physics, a reverse correlation function in signal processing or in the domain of this neuroscience application, the *spike triggered average* [2]. Intuitively, the spike triggered average  $K(\tau)$  represents “what the stimulus did, on average” just before any spike occurred. This is schematically diagrammed in figure 5. By definition, *positive* values of  $\tau$  represent the average stimulus *before* a spike occurs and *negative* values of  $\tau$  represent the average stimulus *after* a spike occurs. Assuming causality between the stimulus and spiking response,  $K(\tau)$  should be zero for *negative*  $\tau$  and the “interesting” activity should occur for *positive* values of  $\tau$ .

The goal of this analysis is to derive a linear filter that finds the average *impulse-response* associated with a spike. The “response” that we are searching for is actually the average stimulus activity that occurred *before* a spike. Although slightly counter intuitive, this is the reason for the funny definition of  $\tau$ . Hence, by this characterization we can treat our spike times as a set of impulses that occur at times  $t_i$ :

$$\rho(t) = \sum_i \delta(t - t_i) \quad \text{where } \delta(t) \text{ is the Dirac delta function}$$

Furthermore, in the guise of linear systems theory, we can reconstruct an approximation of the stimulus given our linear filter using a convolution.

$$s_{est}(t) = \int_0^{\infty} K(\tau)\rho(t - \tau)d\tau \quad \text{where } s_{est}(t) \text{ is the approximate stimulus}$$

Now we can formulate the spike triggered average  $K(\tau)$  as the filter which minimizes the mean squared error between the actual stimulus  $s(t)$  and the approximate stimulus.

$$E = \frac{1}{T} \int_0^T dt (s_{est}(t) - s(t))^2$$

To find the optimal filter we need to find the  $K(\tau)$  that minimizes  $E$ . Because we are examining  $E$  with respect to a function  $K(\tau)$  (and not a single variable), this requires calculus of variations. The math for this section is far too detailed to explain in full here. The main idea though is to discretize  $K(\tau)$ :

$$K(\tau) \rightarrow \{K(t_1), K(t_2), \dots, K(t_n)\}$$

This trick changes  $E$  into a function of a finite number of parameters  $n$ . Using standard calculus, we can take the derivative of  $E$  with respect to each of these parameters and set the derivative equal to zero. Finally, the limit of small time steps is taken  $\lim_{\Delta t \rightarrow 0} t_i + \Delta t = t_{i+1}$  to return the equation to the continuous case, resulting in the following equation [5]:

$$\int_{-\infty}^{\infty} K(\tau) Q_{ss}(t - \tau) d\tau = Q_{ps}(-\tau)$$

In this equation  $Q_{ss}$  is the auto-correlation of the stimulus<sup>†</sup> and  $Q_{rs}$  is the cross correlation between the stimulus and the response. Finally, we can recognize this equation as the convolution. Remembering that the convolution in time is equivalent to multiplication in frequency space, we can apply the Fourier transform to the above equation and arrive at our final form.

$$\tilde{K}(\omega) = \frac{\tilde{Q}_{rs}(-\omega)}{\tilde{Q}_{ss}(\omega)}$$

The inverse Fourier transform of  $K(\omega)$  then represents the final answer - the spike triggered average. This quantity is be easily calculated in Matlab ® or any numerical software application.

## IV. Results

### A. Spike Timing

In two separate sets of experiments the Lorenz attractor was presented as a repeating brightness input for 5 seconds per trial (plus one second of zero brightness in between). The input stimulus and the corresponding raster plots and histogram are plotted for both experiments in figure 6. These results demonstrate the reliability and precision of the H1 neuron to multiple presentations to the same input. Namely, for given a particular input sequence, the H1 neuron qualitatively reproduces a very similar complex spiking pattern.

A rough measure of the precision of the H1 neuron is to take the standard deviation of each peak of the histogram [8]. The standard deviation of a histogram peak quantifies how precisely a spike is placed at a particular portion of the H1 response. This measure is rough in the sense that it confounds precision with reliability since “missed” spikes and spikes that occur outside of each histogram peak do not contribute to the standard deviation. The standard deviation is also quite sensitive to the bin resolution of the histogram, requiring a judicious choice of bin size. In addition, the metric is biased by the algorithmic choice of what constitutes a “peak” in the histogram. In this paper we choose

---

<sup>†</sup> If the stimulus is perfect (or nearly perfect) white noise, then the auto-correlation of the stimulus becomes the Dirac delta function. In this special case the equation reduces the spike triggered average to be equal to the reverse cross correlation.

the simplest possible peak detector as any peak that goes above one standard deviation from the mean histogram count. In spite of these drawbacks, the set of standard deviations associated with the peaks in the histograms are plotted in figure 7. The average standard deviation of the peak widths for both experiments are  $\sim 20$  ms.. This result is on par with previously published results of  $\sim 10$  ms determined from stimuli which vary the angular velocity of the visual pattern [2].

## B. Brightness Kernel

The calculated spike triggered average is plotted in figure 8. This kernel is plotted in negative  $\tau$  time relative to spiking at  $t=0$ . In other words, the left side of the figure before  $t=0$  represents the spike triggered average. It is interesting to note that the width of the kernel implies that the H1 neuron can integrate brightness changes over several hundred milliseconds. Furthermore, there appear to be correlations beyond  $t=0$ . In other words, *after* a spike occurs, the stimulus on average follows a particular trend. As stated in the Methodology, the kernel beyond  $t=0$  should be entirely zero. The correlation beyond  $t=0$  indicates that the presented stimulus did not sample all of the stimulus space. In other words, this artifact indicates that we should have presented the stimulus for an even longer period of time or that we should have used a more uncorrelated, white noise stimulus.

To prevent cheating and ensure causality, the kernel was zeroed beyond  $t=0$  and the stimulus was reconstructed via convolving the spiking activity with the spike triggered average  $K(\tau)$ . Portions of these results are shown in figure 9. This figure shows a portion of the stimulus not used for constructing the kernel  $s(t)$  as well as its reconstruction  $s_{est}(t)$ . Qualitatively, the reconstruction captures most of the low frequency variation in the stimulus. A distribution of the squared errors is shown in Figure 10. The mean of the distribution of errors is  $-1.6e-3$  and the standard deviation is 1.5. Hence, the reconstructed stimulus is on average quite close to the actual stimulus.

## V. Conclusion

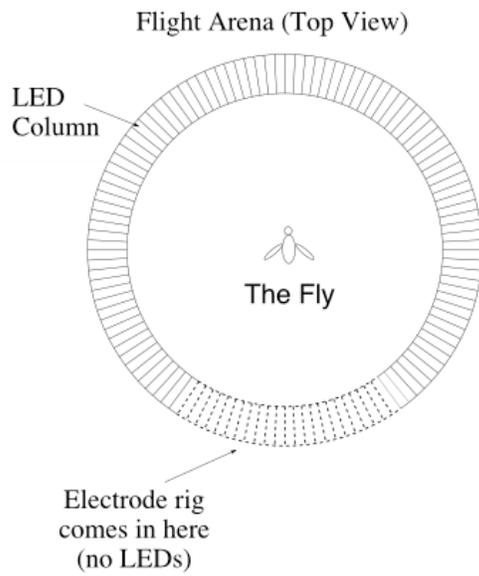
This set of experiments successfully recorded from the H1 neuron in the fly in response to controlled, brightness-modulated visual stimuli. The first set of experiments indicates that the fly H1 neuron places spikes in reproducible complex patterns with a precision of roughly  $\sim 20$  ms. The second set of experiments indicates that treating the spiking behavior of the H1 neuron as a linear filter is a reasonable approximation that captures a large amount of the variability of the neuron's spiking activity.

This set of experiments suggests future lines of inquiry. As discussed previously, the standard deviation of the peaks in the histogram provides a rough and flawed measure of the spike timing precision of H1. Furthermore, characterizing "how well" the linear filter approximates the activity of the H1 neuron by using the mean squared error fails to quantify how well a linear filter captures the variability of H1. These two difficulties can be improved by attempting an information theory form of analysis. Previous papers have suggested methods based on quantifying how much variability in the stimulus is *statistically correlated* with the variability in the spiking activity at different temporal resolutions [6]. Quantifying the amount of correlated activity that the linear model can account for provides a novel measure of understanding how well a linear model describes the activity of H1 [7]. Furthermore, quantifying and comparing the amount of correlated activity at varying temporal resolutions provides another novel measure for understanding the actual temporal precision used by the cell [8]. Although more involved, applying these forms of analysis on H1 data might provide a more rigorous means of understanding the neural coding of brightness modulation in the fly's visual system.

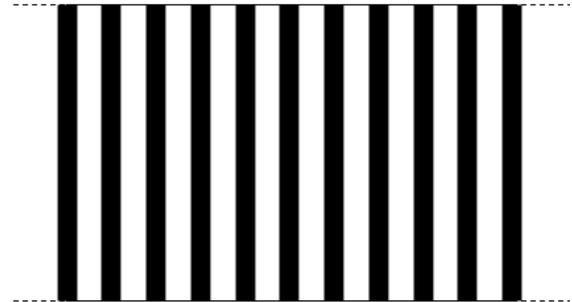
## VI. References

- [1] Eckert, H (1980). Functional properties of the H1-Neurone in the third optic ganglion of the blowfly, *Phaenicia*. *Journal of Comparative Physiology A* 135, 29-39.
- [2] Rieke F, Warland D, de Ruyter van Steveninck R, Bialek W. (1997). *Spikes: Exploring the Neural Code*. MIT Press.

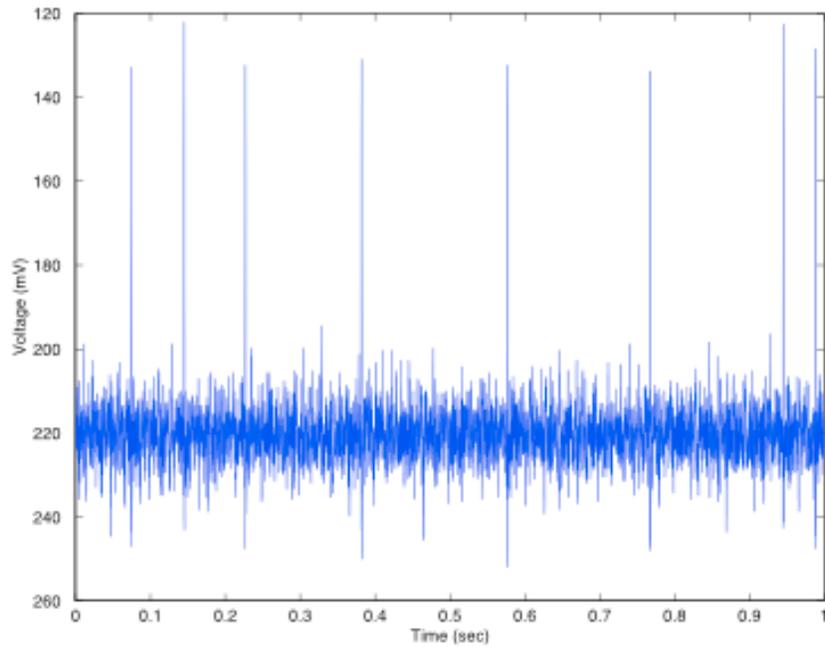
- [3] Tumer E (2002). *Wide field visual information encoding in the fly*. Physics 271: Biophysics Laboratory. UC San Diego.
- [4] Brotz T, Egelhaaf M, Borst A. (1995) A preparation of the blowfly (*Calliphora erythrocephala*) brain for in vitro electrophysiological and pharmacological studies. *Journal of Neuroscience Methods* 57: 37-46.
- [5] Dayan P, Abbott L. (2001) *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. The MIT Press.
- [6] Borst A, Theunissen F. Information theory and neural coding. *Nature Neuroscience* 2:11, 947-957.
- [7] Bialek W, Rieke F, de Ruyter van Steveninck RR. (1991) Reading a neural code. *Science* 252, 1854-1857.
- [8] Reinagel P, Reid RC (2000) Temporal coding of visual information in the thalamus. *Journal of Neuroscience* 20:14 (5392-5400).



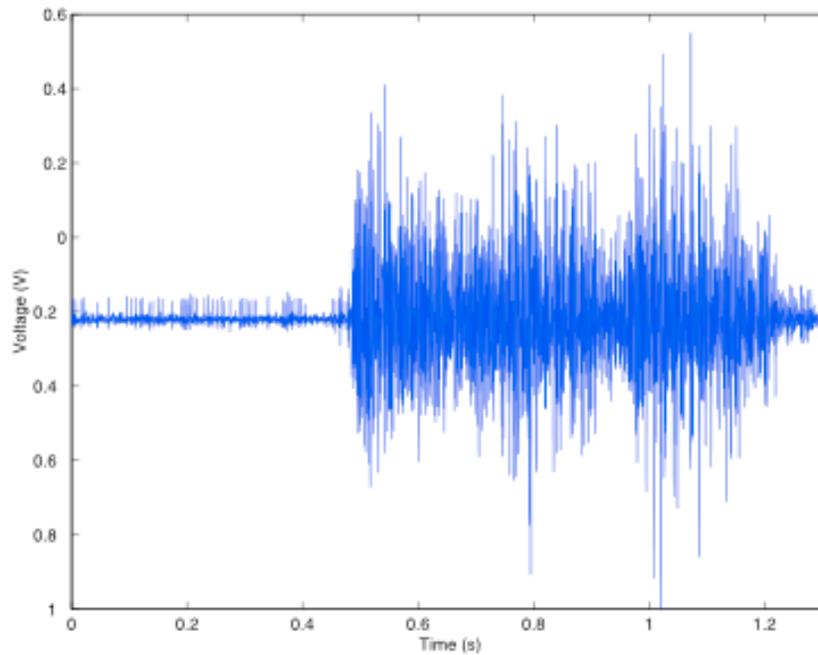
**Figure 1:** A diagram of the fly arena adapted from [3]. There are 96 vertical LED columns covering 288 horizontal degrees of visual scene.



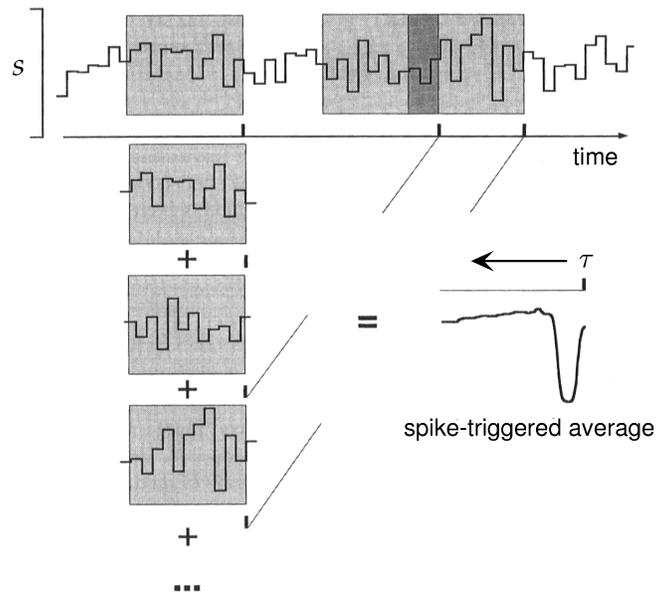
**Figure 2:** A diagram of the LEDs illuminated during all of the experiments. The brightness of every other LED column is modulated according to the input pattern. Every other LED light always remains off.



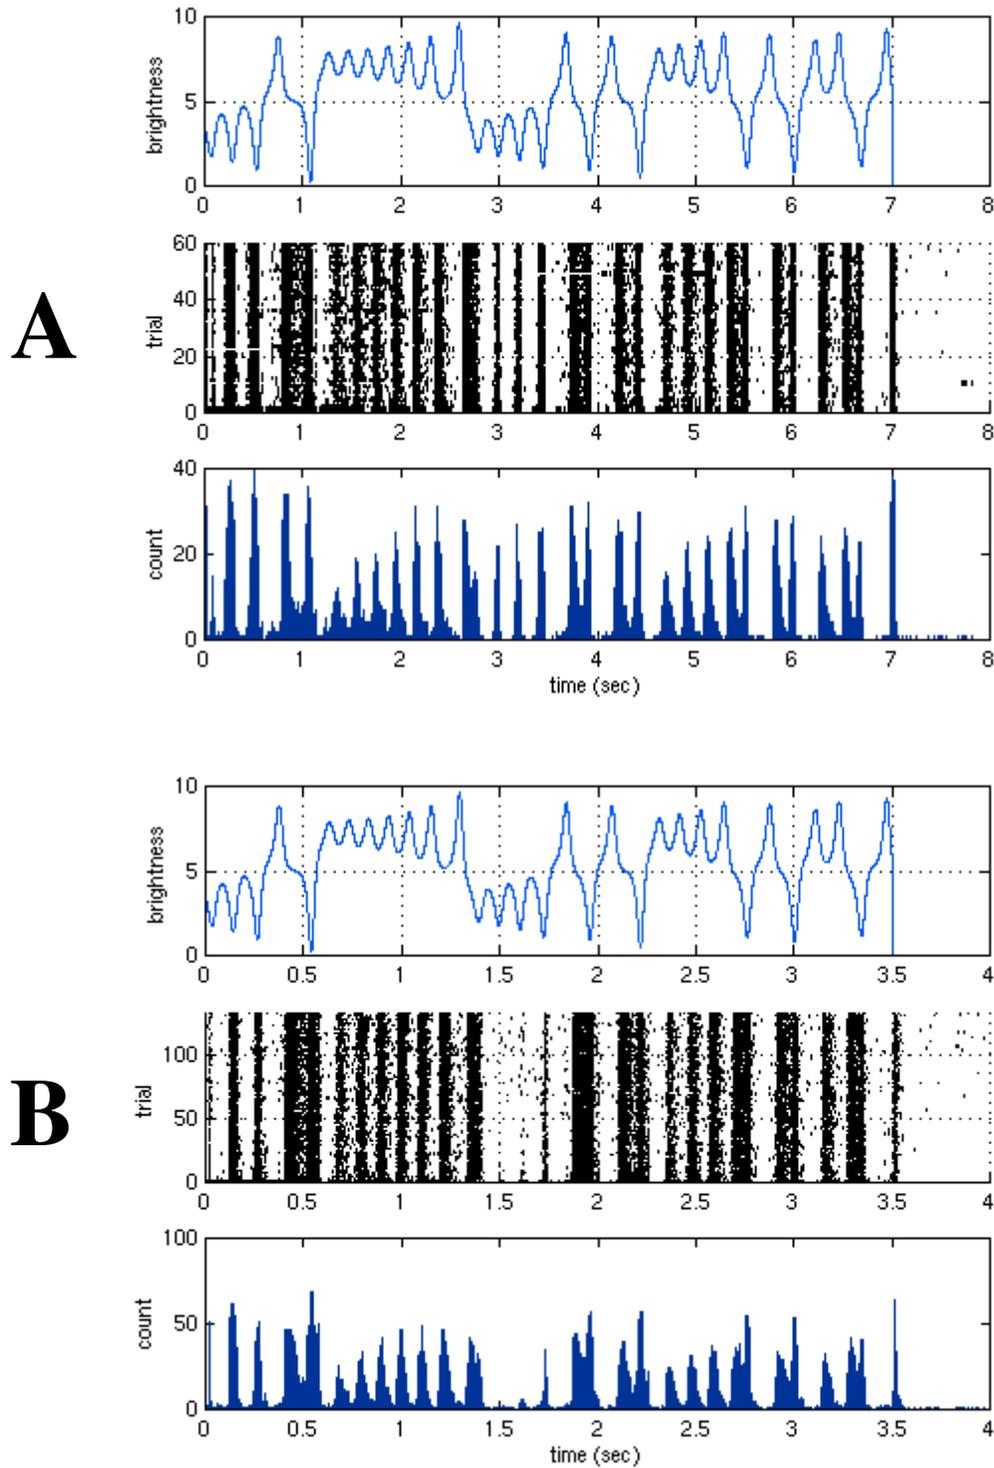
**Figure 3:** A portion of a time series from an extra-cellular recording of an H1 neuron during visual stimulation [3]. Note the large spikes amongst lots of noise and even other spiking activity.



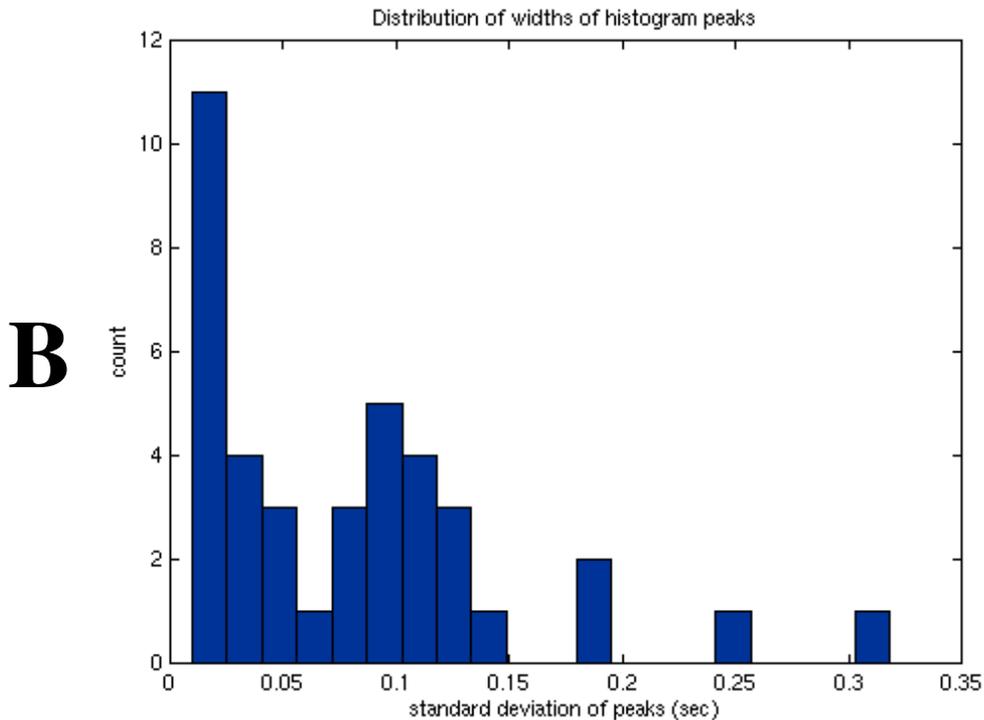
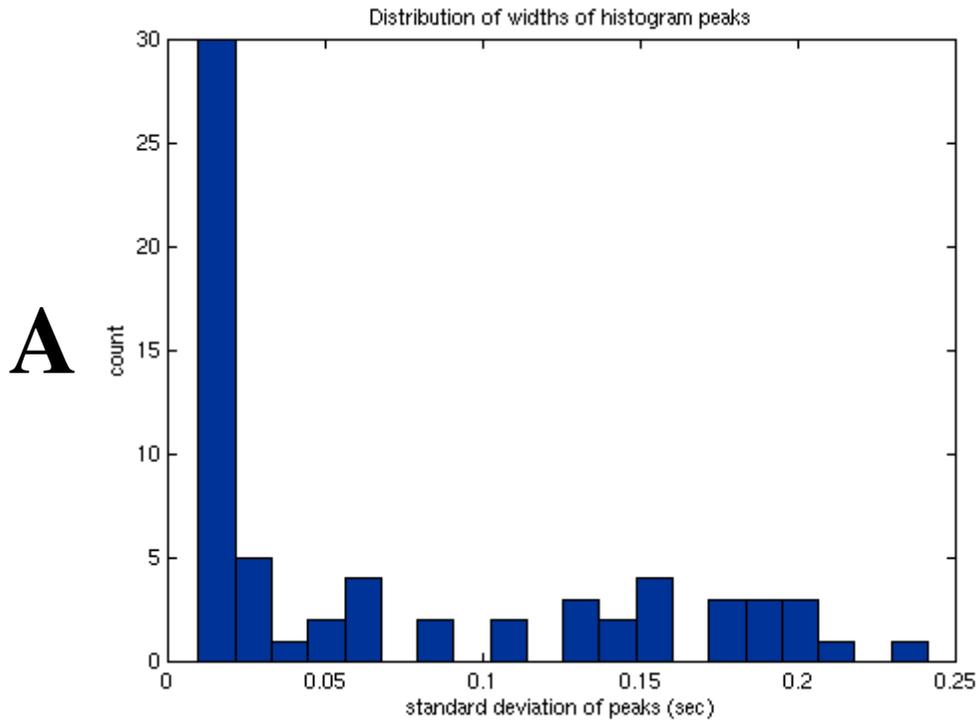
**Figure 4:** Movement artifacts. When the fly moves, the spikes (on the left) are dwarfed by the disruption caused by the physical movement of the electrode (on the right) [3].



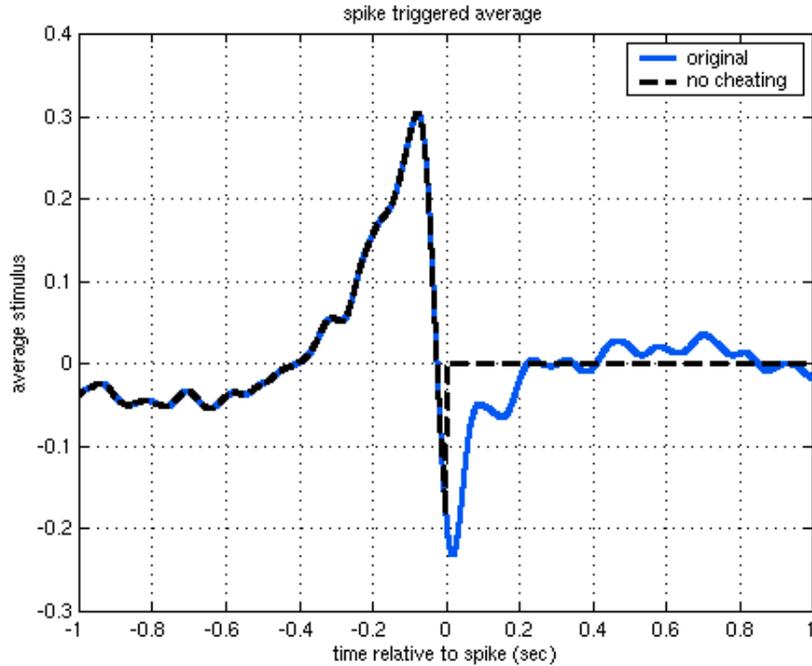
**Figure 5:** Spike Triggered Average. A schematic diagram for the calculation of the spike triggered average adapted from [5].



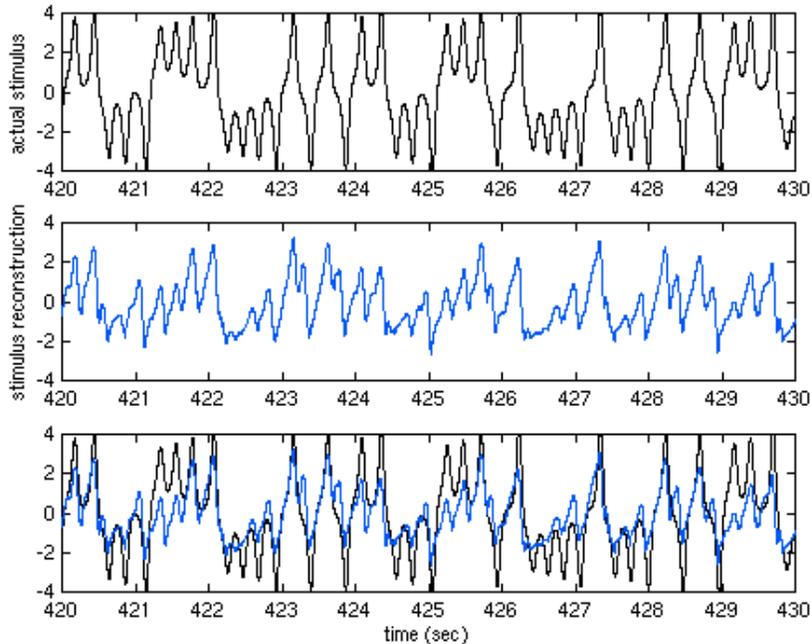
**Figure 6:** Spike timing precision. In two sets of experiments labeled *A* and *B*, the H1 neuron was presented multiple trials of a brightness-modulated visual stimulus. The visual stimulus lasted 5 seconds (plus 1 second of darkness). The brightness over time  $s(t)$  sampled from the Lorenz system is plotted in the top panels. The middle plots display the corresponding raster diagrams signifying the filtered spiking events that occurred during each trial (60 trials in *A* and 120 trials in *B*). The bottom plots show the corresponding histogram of the raster diagram. The bin resolution of the histogram is 2 msec.



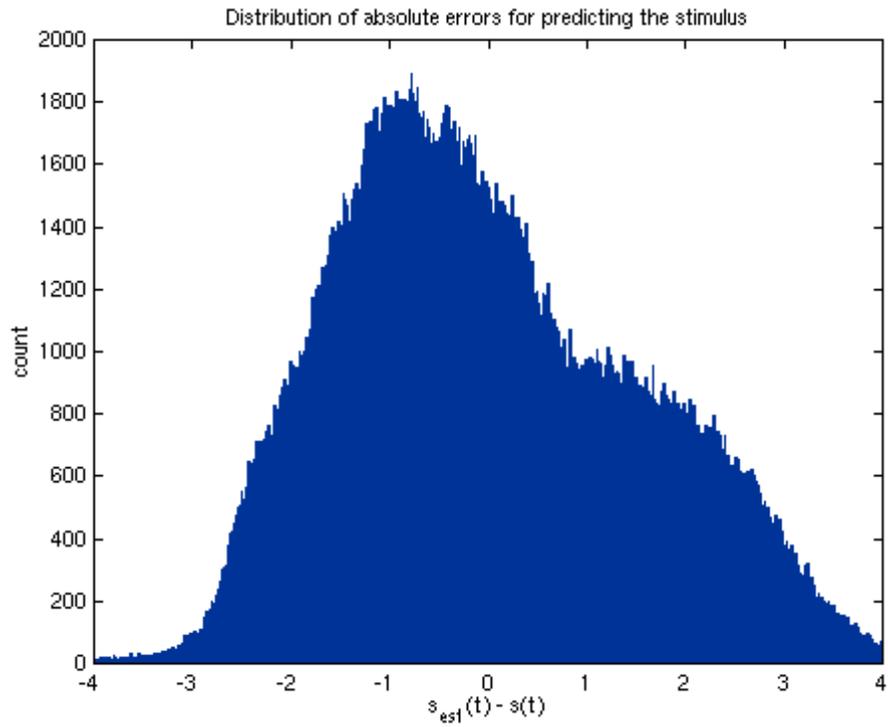
**Figure 7:** Spike timing precision. In the two sets of experiments labeled A and B in figure 6, the peaks of the histograms on the bottom roughly correspond to individual spiking “events.” The widths of these peaks, as measured by the standard deviation associated with a Gaussian fit, roughly quantifies the spiking precision of the neuron [6]. A distribution of these standard deviations of the peak widths for experiments A and B are shown above. The means of both of these distributions are  $\sim 20$  msec. A peak is identified as any event that goes above one standard deviation from the mean of the histograms in figure 6.



**Figure 8:** Spike triggered average. The spike triggered average  $K(t)$  is calculated as described in the Analysis section. This spike triggered average is plotted with respect to a spike occurring at  $t=0$ . Also,  $K(t)$  is reversed with respect to  $t$  for clarity. Hence,  $K(t)$  for  $t < 0$  represents the stimulus that occurred *before* the spike occurred and  $K(t)$  for  $t > 0$  represents the stimulus that occurred *after* the spike occurred. The blue line represents the original spike triggered average. The black line represents the causal, “no cheating” version in which all correlations after the occurrence of the spike are artificially zeroed. The “no cheating” version of the kernel is used to generate the plots below.



**Figure 9:** Reconstruction of the stimulus. The top and middle panel display a portion of the actual and reconstructed stimulus,  $s(t)$  and  $s_{est}(t)$ , respectively. The bottom panel is an overlay of the two.



**Figure 10:** A distribution of errors between actual and reconstructed stimuli. The mean of this distribution is  $-1.6e-3$  and the standard deviation is 1.5.