Visual Coding in the Blowfly H1 Neuron: Tuning Properties and Detection of Velocity Steps in a new Arena

Jeff Moore and Adam Calhoun
TA: Erik Flister
UCSD Imaging and Electrophysiology Course, Prof. David Kleinfeld

1. Introduction

The blowfly visual system is frequently used in vision experiments, and is attractive because of its simple preparation and small scale. H1 is well-suited for these experiments because of its large size and its ease to record from. Because it is four synapses away from the retina, it may be seen as similar to cells in visual cortex. It responds primarily to horizontal full-field motion, and is directionally tuned to regressive motion, from the back to the front of the eye [1]. This important system allows the fly to remain stable and avoid objects while in flight.

We sought to characterize various properties of the cell, both to test the capabilities of our arena and the capabilities of the fly. Stimuli were presented by changing three parameters: movement speed, spatial frequency, and refresh rate. These experiments should expose some of the limits of the fly's ability to perceive these events.

2. Methods

2.1 Fly preparation

The blowfly *Calliphora* was ordered from Ward's Natural Science (Rochester, NY). The pupae were housed in clear plastic water containers. They were placed in a small dish to let hatch. A dish of water and a dish of sugar were also placed in the containers. One end of the container was cut off and replaced with pantyhose. Cutting off part of the pantyhose allowed easy access to the flies without letting them escape, while still letting them breathe. The last several batches of flies failed to hatch for unknown reasons, despite changing many of the housing elements.

When a fly was needed, it was trapped in a test tube with a small air hole. It was then anesthetized by placing it in the freezer for eight minutes. Upon removing the fly, its legs, wings, and antennae were quickly removed. These wounds were covered in wax to prevent movement from vestigial remnants of the limbs. The fly was then waxed to the platform with its head positioned off of it. The head was pressed forward so that the back of the head was pointing vertically, and then waxed down with the proboscis.

The back of the head behind the eye was then exposed. After initially taking out the mucous-like substance, we found we were better able to record from H1 when we left it in. A small incision was made on the back for the reference wire. While in the rig, a saline solution was intermittently placed in the opening. The solution is described in [2].

2.2 Stimulus delivery

The stimulus was delivered through a system of LEDs. Ninety-six columns of LEDs were arranged in an arc encompassing 270 degrees around the center of the fly platform. This was controlled by three NI-DAQ cards (two NI-DAQ 6224 cards and one NI-DAQ 6229 card, National Instruments, Austin, TX). The output of the three cards was fed through a custom-designed box to switch the signal from six cables to two.

These cards were controlled from a MATLAB program (Natick, MA) through the NI-DAQmx C library. All digital writes to the cards were synchronized. We broke every frame of our stimulus down into one hundred subframes to digitally control perceived brightness. If we wanted an LED to display a brightness at, say, half maximum, the LED would be turned on for approximately fifty randomly chosen subframes. Since the NI-DAQ cards were capable of 1MHz output, this allowed us to show frames of stimuli at 10kHz.

Subframes were calculated using an algorithm that converted from arbitrary spatial and temporal waveforms to desired LED intensities. The brightness was gamma-corrected with a gamma of 2.2. Our software allows for individual trial lengths ranging from ~1ms to ~15s (memory buffer limit).

2.3 Recording

We used the MATLAB Data Acquisition Toolbox to acquire data, using two analog inputs. One input was wired into the rightmost LED. This LED was non-functional, but a signal was sent to it at the start of every trial. In this way, we could ensure synchronization of the input from the electrode with the start of each trial. The other input came from an amplifier/filter, which amplified the voltage between the active and reference electrodes (filters were set all the way open). Trials with movement artifacts were removed manually from the dataset.

To find H1, a sine grating stimulus was used. The grating moved with a constant speed of 100 deg/s, switching direction every second. The fly was placed on the rig tilted upward and in the direction of the side of the head being recorded from, allowing us access to the cell. The electrode (3 MOhm Tungsten, MicroProbe, Inc) was then inserted until a characteristic spiking sound was heard from motion in the regressive direction to the contralateral eye.

2.4 Stimuli

Several stimuli were used. All had a characteristic sine wave (in space) which was moved at constant speed. To test flicker we presented a stimulus moving at 50 deg/sec with 10 cycles, at artificial frame rates of 5Hz, 10Hz, 30Hz, 70Hz, 100Hz, 300Hz, 700Hz, 1000Hz, and 10000Hz.

We tested spatial frequencies moving at 100 deg/sec with a refresh rate of 10 kHz. The number of cycles appearing on the screen was 3, 5, 8, 20, 40, and 60. To test preferred speed, we displayed a sine wave with a spatial frequency of 10 cycles and a refresh rate of 10 kHz. Speeds used were 10 deg/sec, 20 deg/sec, 40 deg/sec, 60 deg/sec, 75 deg/sec, 100 deg/sec, 200 deg/sec, 300 deg/sec, 500 deg/sec, and 700 deg/sec.

3. Results

3.1 Raw Data and Spike Sorting

Stimuli presented to the fly consisted of 20 consecutive trials of progressive motion followed by regressive motion (top right). A sample recording is shown in Fig 3.1(top left). The same recording is separated into trials (middle left), and sample spikes are shown (middle right). To get spike times, we applied a threshold voltage and counted upward threshold crossings. In several recordings for each electrode placement in each fly, we checked plots of peak vs. trough spike amplitude with different thresholds (Fig 3.2, top left and right), and chose thresholds so that only one cluster was seen in each plot (right). We also checked the spike autocorrelation to make sure there was a dip near 0 lag lasting several milliseconds, corresponding to the refractory period (bottom).

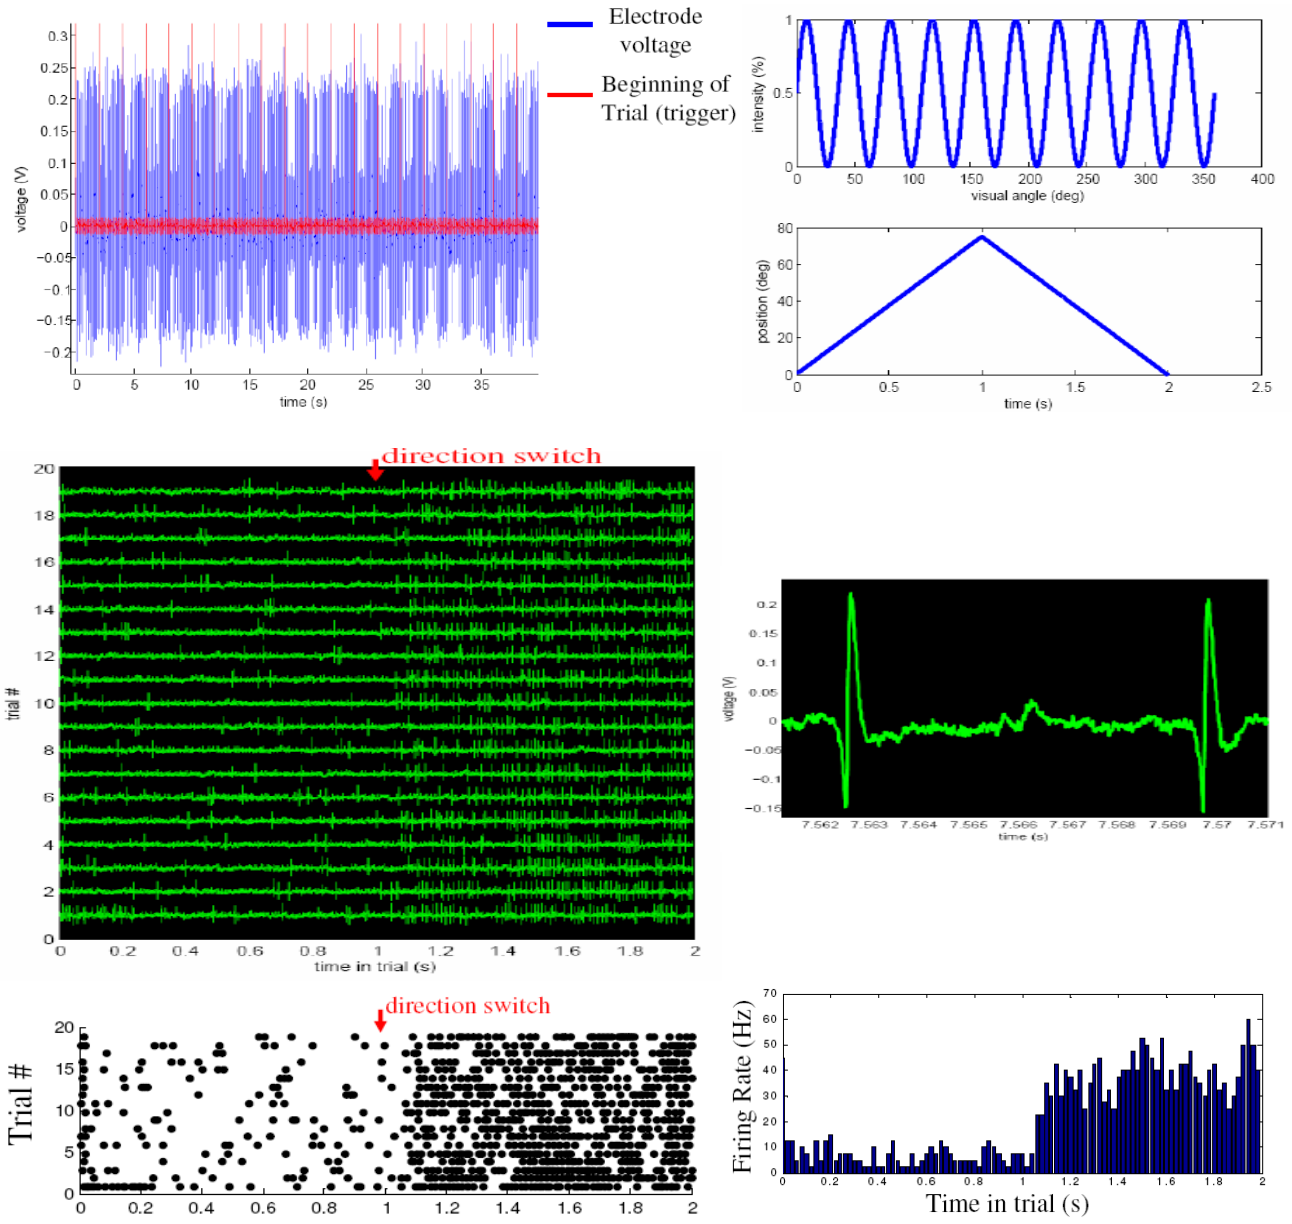

Fig 3.1. Sample recording: 75 deg/s set. (top) raw recording and triggers. (middle left) separated by trials. (middle right) single spikes. (bottom left) raster plot of spike times and (bottom right) spike histogram.

3.2 Flicker vs. Motion

Because the flies were presented with apparent motion stimuli that were varied by changing the position of a pattern of lit LED panels, we verified that the fly could actually see full-field motion and was not simply responding to the flickering at frame transition boundaries. This was done by artificially decreasing the frame rate and evaluating the resulting H1 response. To decrease the frame rate, we held the stimulus constant for multiple frames, effectively down-sampling to the desired frame rate. We tried frame rates varying from 5Hz to our maximum 10 kHz (see Methods). Examples of our down-sampled frames are shown in Fig 3.3 (left).

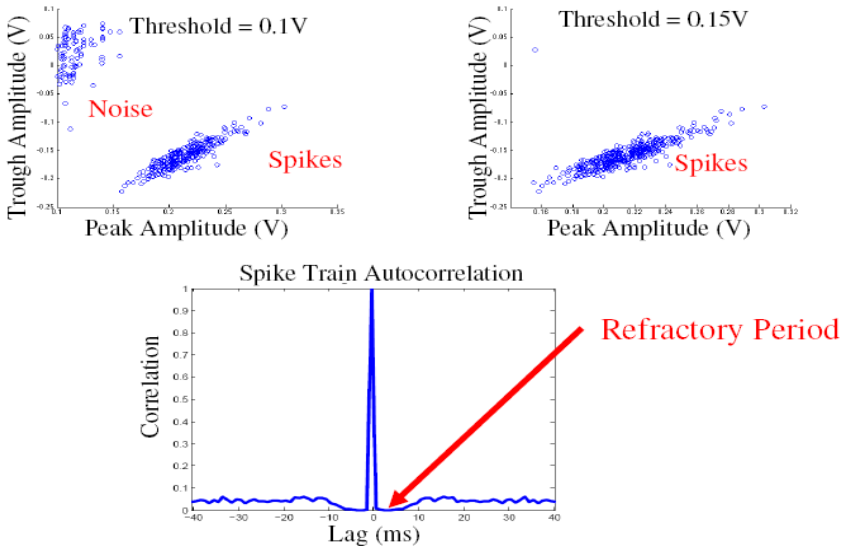

Fig 3.2. Spike Sorting. Appropriate threshold yields a single cluster of spikes (top). With appropriate threshold, no spikes occur in the refractory period in the autocorrelation function (bottom).

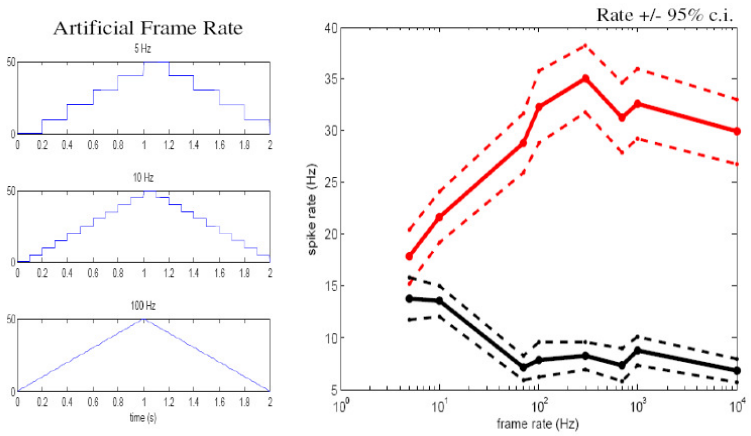

Fig 3.3. Effect of frame rate on firing rate. (left) Stimulus position vs. time plots. (right) firing rate vs. frame rate. Stimuli consisted of one second of progressive motion followed by one second of regressive motion repeated 20 times at each frame rate. Red line indicates preferred direction, black line indicates null direction.

We examined the average firing rate of the H1 neuron (1 fly) as a function of frame rate using one second of progressive motion followed by one second of regressive motion at 50 deg/s. The firing rates in the null and preferred directions are shown in Fig 3.3 (right). Firing reached a steady value at ~100 Hz frame rate.

We also examined whether the timing of the spikes aligned with frame transition times. To do this, we computed the average of all of the pairwise cross-correlation histograms for the 20 trials at each frame rate. We also computed 95% confidence intervals for what the cross-correlation histogram assuming spike times are completely random within the time periods over which motion occurred in a single direction. 100 sets of 20 trials with random spike times were generated, and confidence intervals were determined empirically. The results are shown in Figure 3.4. At 10 Hz frame rate we see significant peaks in the cross-correlation at a period of 1/10Hz. At 100 Hz frame rate and above we do not see such peaks. This confirms that above 100 Hz, spikes are not related to frame transitions. In the results to follow, we present stimuli at a frame rate of 10 kHz.

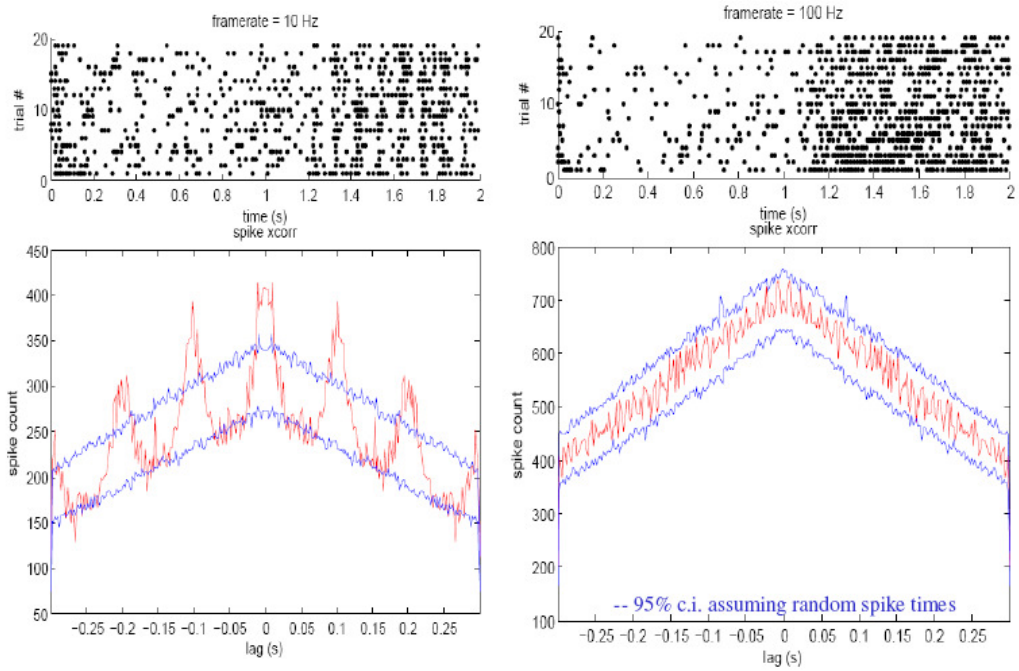

Fig 3.4 Cross correlations of flicker response. Red line indicates the actual cross-correlation, while blue indicates 95% confidence intervals.

3.3 Speed and Spatial Frequency Tuning

Speed tuning was measured by presenting consecutive 20 trials of a sinusoidal grating (10 cycles/360 degrees) which moves progressively for one second and then regressively for one second. Sets of trials were collected at different rates varying from 5 deg/sec to 700 deg/sec (see Methods). The mean firing rate in the preferred and null directions is shown in Fig 3.5 (left). The neuron appears to be tuned to a peak speed of ~ 75 deg/s.

Spatial frequency tuning was measured in the same manner by presenting sinusoids with different spatial frequencies while holding the speed at 100 deg/s. The mean firing rates are shown in Fig 3.5 (right). Due to the fact that each LED subtends 3 degrees of visual angle, direction of motion is ambiguous at 60 cycles/360 degrees so the null and preferred directions should not be distinguishable.

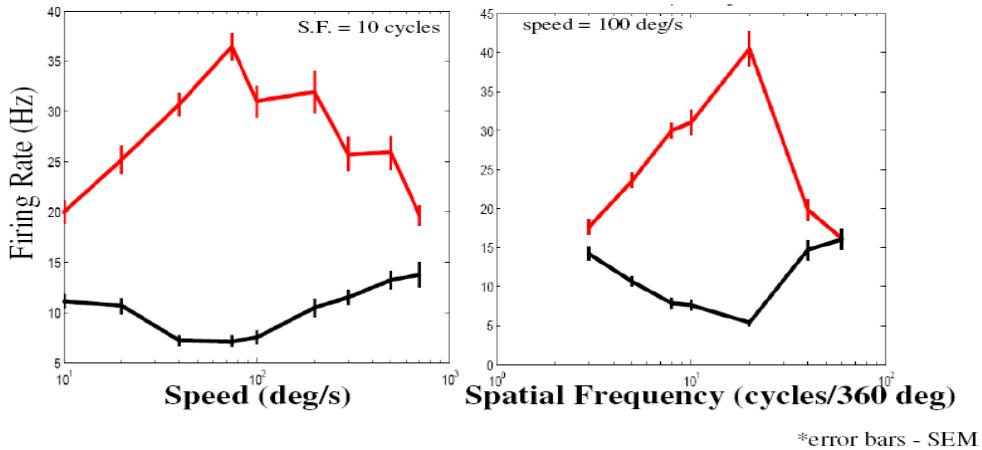

Fig 3.5 Speed and spatial frequency tuning. Red indicates preferred direction, black indicates null direction. (left) Firing rate vs. speed. (right) firing rate vs. spatial frequency.

3.4 Spike Train Statistics

To determine whether H1 firing follows Poisson statistics, we computed the coefficient of variation and Fano-Factor at the various speeds tested above, and found confidence intervals by resampling trials (Fig 3.6). In the preferred direction, the CV was lower than would be expected from a Poisson process (CV =1). This could be due to spike intervals timed to visual edges reaching certain locations in the visual field, or due to the refractory period of the neuron. However, the Fano-Factor was not significantly different from 1. The neuron's ability to transmit information is shown in Fig 3.7, calculated using the Strong method [3]. Information rate appears to follow the speed and spatial frequency tuning profiles. In subsequent analysis, we assume a rate code for stimulus speed.

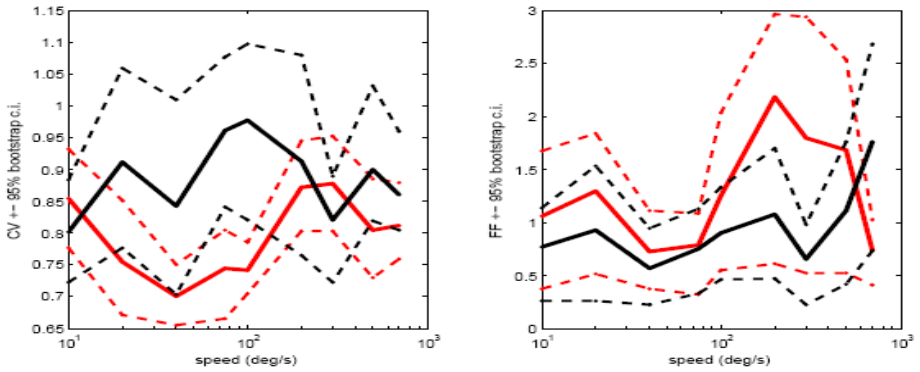

Fig 3.6. Spike Train Statistics (left) CV. (right) Fano Factor. Black indicates null direction, red indicates preferred direction.

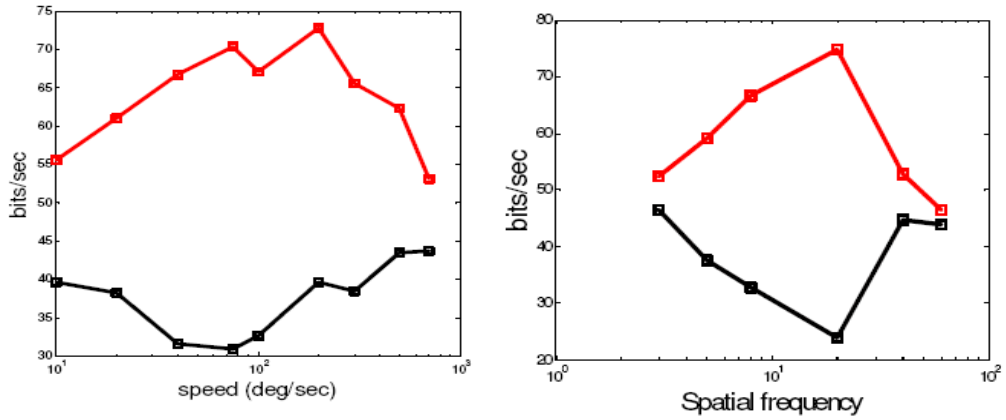

Fig 3.7 Information transmission rates. (left) Strong Information rate vs. speed. (right) Strong information vs. spatial frequency. Red is preferred direction, black is null direction.

3.5 Time to detect direction switches

The stimuli presented above consist of instantaneous (within 1us) steps in velocity from the null to the preferred direction, and vice versa. However, there is a delay in the neural response to the step (Fig 3.1, bottom left). This delay is potentially relevant for determining how quickly the fly can respond to perturbations in the flight path. In this section, we examine how quickly the fly can determine that a switch in direction occurred, assuming the only information about the stimulus is in the rate. In other words, we assume the fly is simply counting its H1 spikes as they come in.

With this assumption, we determine how well an ideal observer can tell the difference between the distributions of preferred rates and null rates as a function of how much time following a switch we allow the observer gather information over. Specifically, we determine the optimal probability of correctly assigning a single new observation of firing rate as coming from one of the two directions.

Let us denote the time window following the switch as T. Then, for a given set of trials at a certain speed, we compute the distribution of firing rates (across trials) in the one second period during which the null direction was presented, and the distribution of firing rates in the time window T during which the preferred stimulus occurred, immediately after a switch. As we slide the criterion threshold (for determining which of the two distributions an observation comes from) from zero to the maximum firing rate, we compute the probability of a correct detection as a function of the probability of a false positive. The integral of this function is the optimal probability of detection. Fig 3.8 shows this probability as a function of T for switches from the null to preferred direction (left) and preferred to null direction (right). We then assign an arbitrary threshold probability (85% correct) and find the time it takes to reach this level as a function of speed (Fig 3.9). For switches to the preferred direction, we also show the average number of spikes to reach criterion.

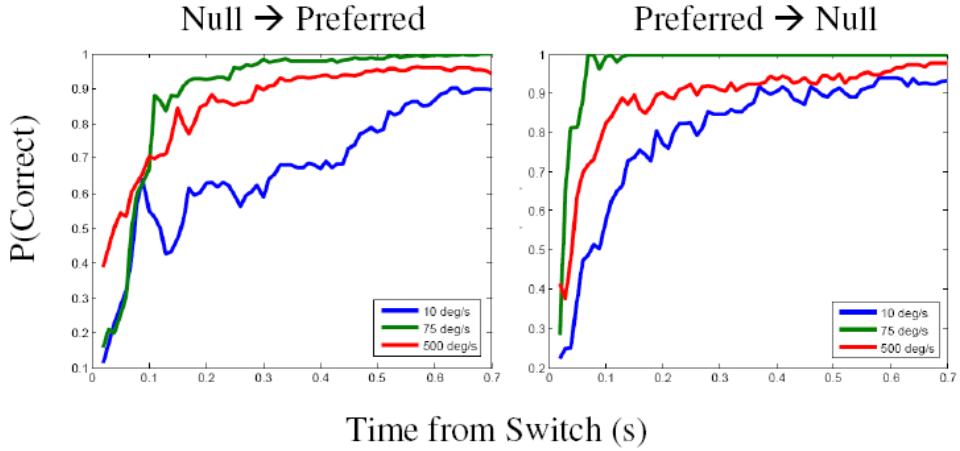

Fig 3.8 ROC Analysis. Probability of correctly distinguishing between the preferred and null direction, given one second of observation of the first direction, and T seconds of observing the post-switch direction (T is on the horizontal axis). Probability values below 0.5 indicate that the average rate in the time window T was lower than the null rate.

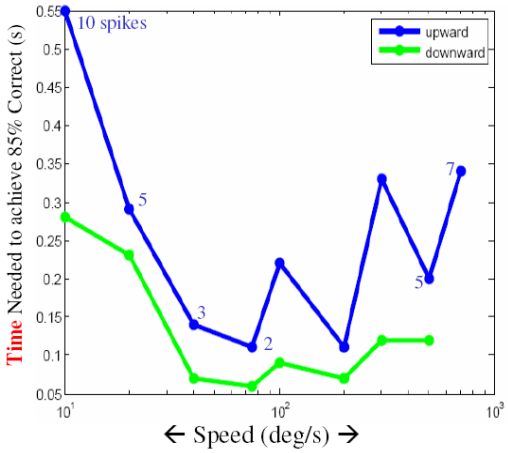

Fig 3.9 Time (and number of spikes) required to detect a switch as a function of stimulus speed.

The curves roughly follow the speed tuning curve of the neuron, and the time to detect a switch to the null direction is faster than a switch to the preferred direction. This difference may be due to the time scale of inhibition from the contralateral H1 (since stimuli were presented to both eyes simultaneously). For a switch to the preferred direction there is a delay of ~100 ms at the peak speed, and for a switch to the null direction there is a delay of ~50 ms.

4. Discussion

We have implemented a new hardware setup (designed by Erik Flister) for a fly arena capable of presenting graded intensity visual signals at a fast frame rate, and we have written the corresponding software to compute and present stimuli with arbitrary intensity profiles with arbitrary velocities. We demonstrate that this arena setup is capable of presenting stimuli that are similar to smooth motion patterns as seen by the H1 neuron (at least at the speeds and spatial frequencies described above). We have also confirmed that the H1 neuron is tuned to different speeds as well as spatial frequencies, as described in [1]. We then use a ROC analysis to predict the delay time with which the fly would be able to respond to instantaneous direction switches (at the contrast levels of the rig and limited to the spatial frequency tested). The tumbling speed of the blowfly can reach 1200 deg/s [4], and velocity changes can exceed 30 Hz in natural flight [5] (stimulus courtesy of RR de Ruyter van Steveninck), so the delays observed in our recordings would appear to be detrimental. However, instantaneous shifts in direction do not usually occur in natural flight, and the delay may be due to inhibition from the contralateral H1 that we allowed to build up over an entire second. The fly also has a vestibular system (the halteres) that may mediate the fast perturbations in flight trajectory.

We had also planned to compare the H1 response to natural flight stimuli to noise stimuli with different bandwidths. We designed and displayed these stimuli, but were unable to obtain recordings due to lack of flies.

References

[1] Eckert H. J. J. Comparative Physiology, 1979.

[2] T. M. Brotz, M. Egelhaaf, and A. Borst. J. Neurosci. Meth.,1995.

[3] Strong et al. Phys. Rev. Letters, 1998.

[4] Fairhall A, Lewen G, Bialek W, de Ruyter van Steveninck. Nature 2001.

[5] Nemenman I, Lewen G, Bialek W,de Ruyter van Steveninck RR. PLoS Computational Biology 2008.