International Journal of Ethics in Engineering & Management Education
Website: www.ijeee.in (ISSN: 2348-4748, Volume 9, Issue 7, July 2022)

An Analog VLSI Motion Sensor Based on the Fly Visual Systems

Dr. Deepak
Professor, Dept. of ECE
MIST, Tamilnadu
India

Abstract: **The human brain is a daunting goal for biologists and engineers alike. Our brain takes several years to fully develop, and contains between 10^{10} and 10^{11} neurons (nerve cells), each communicating with 10^{3} other cells, on average. Brains of other animals (particularly invertebrates) are much smaller but still perform remarkably complex computations. Insect brains, for example, typically contain between 10^{5} and 10® neurons. As we shall see in the following chapter, insects perform sophisticated information-processing tasks rapidly and efficiently. In this body of work, we have attempted to extract computational principles from the visual system of the fly and apply these principles to an engineered system—an integrated, low-power visual motion sensor. As our engineering tool we use very-large scale integration (VLSI) of silicon circuits—the most advanced information-processing substrate available today. In particular, we explore continuous-time (un clocked), continuous-value (analog) circuit architectures.**

Keywords: **Sensors, VLSI, Fly, Light, Linear, Power, Control**

1. SENSORY SYSTEMS OF THE FLY

The fly is an attractive target for biologically-inspired approaches to engineering. Its brain and sensory systems have been studied for decades, so much is known about their operation. Of course, we are still decades (or centuries) away from understanding the entire system, but a wealth of behavioral and electrophysiological data has led to the development of several models of information processing.

Flies possess a diverse array of organs for sensing their environment. In addition to the familiar sense of vision, flies employ Coriolis-force "gyroscopes," polarized light sensors, and body proprioception to aid in navigation. In this chapter, we will discuss these sensory systems.

The Visual System of the Fly

Vision is a vitally important sense for flying insects. This information processing begins at the sensor—the retina. Despite the multi-lens construction of the compound eye, the pattern projected onto the underlying retina is a single image of the visual scene. Photoreceptors in the retina adapt to the ambient light level, and signal temporal deviations from this level. These signals are passed on to the next layer of cells, the lamina. Lamina cells generally show transient or highpass responses, emphasizing temporal change (Weckstrom et ah, 1992). The next stage of processing is the medulla, a layer of cells that are extremely difficult to study directly due to their small size. Indirect evidence suggests that local measures of motion (i.e., between adjacent photoreceptors) are computed here. These are the best-studied cells in the fly visual system, and much is known about their properties.

The Vestibular Sense

Dipterans (true flies and mosquitos) possess a remarkable evolutionary specialization for measuring angular velocity. The hind wings of these animals evolved from flight surfaces into dedicated angular rate "gyroscopes." These halteres, as they are called, resemble small balls at the end of sticks (see Figure). The halteres beat up and down antiphase to the wings at the wingbeat frequency (about 150 Hz in *CaMiphora;* over 200 Hz in the smaller *Drosophila*). They move at nearly constant velocity during each upstroke and downstroke, covering nearly 180° (Nalbach, 1993).

While body rotations produce centrifugal forces on the halteres, these forces are oriented radially and tangentially, and for typical maneuvers are several orders of magnitude smaller than the radial centrifugal forces due to halteres oscillation. Centrifugal forces are proportional to the square of angular velocity and thus provide no information on the direction of rotation. A more useful signal is the Coriolis force, which is proportional to the cross product of the instantaneous haltere velocity and the axis of body rotation. Coriolis forces acting normal to the plane of haltere oscillation are detected by about 335 campaniform sensilla organized in five groups at the haltere base. These sensory cells are embedded in the flexible exoskeleton, and act as strain gauges.

By integrating Coriolis force information over the haltere's 180° sweep, and by combining signals from the two non-coplanar halteres, the fly can measure angular rotation about all three axes. Vestibular information from the halteres system is critical for maintaining stable flight; when a fly's halteres are removed it quickly falls to the ground. (With only one haltere removed, flight is still possible.)

Free flight seems to be controlled by a combination of visual and vestibular sensors. Little is known about the interactions between these two sensory modalities. Recent experiments by Dickinson and colleagues have shown that visual interneurons stimulate small haltere control muscles that exert force in the same direction as Coriolis forces (Chan et ah, 1998). They propose a model where the halteres control flight equilibrium in a fast feedback loop and the slower visual interneurons steer the animal by tugging on the halteres to

International Journal of Ethics in Engineering & Management Education
Website: www.ijeee.in (ISSN: 2348-4748, Volume 9, Issue 7, July 2022)

create a vestibular illusion.

Ocelli

In addition to the two compound eyes, flies have three other photosensitive organs called ocelli. These sensors are located between the eyes on the dorsal region of the head. Each ocellus consists of a single circular lens approximately 75 pm in diameter which focuses light onto a low-resolution retina containing approximately 220 photoreceptors. The image produced on the retina is a wide- angle, underfocused view of the surroundings above and lateral to the fly. There is rapid convergence of the photoreceptors onto 4-6 interneurons that seem to measure mean brightness.

The ocelli seem to contribute to the *dorsal light response* observed behaviorally, where flies align the top of their head with the center of brightness (Schuppe and Hengstenberg, 1993). This often corresponds with the zenith in outdoor environments. Dragonflies appear to use their ocelli as horizon detectors to provide information on their head attitude relative to the ground (Stange, 1981). The ocelli seem to work in concert with the compound eyes to position the head.

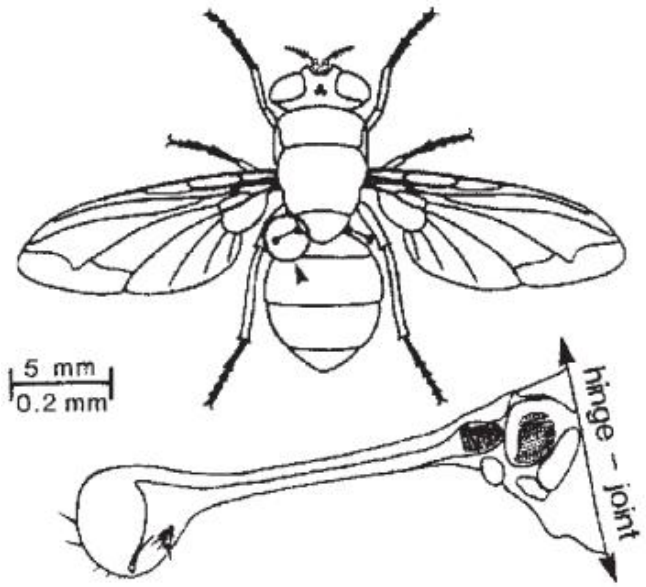

Figure 1: The halteres of the blowfly *Calliphora*.

The halteres evolved from hindwings but no longer serve any aerodynamic function. Located in the "waist" between the thorax and the abdomen, these devices beat up and down (i.e.. in and out of the plane of the page) antiphase to the wings. Groups of mechanoreceptors at the base measure Coriolis forces produces by the angular rotation of the animal. Adapted from Nalbach, 1993.

Polarized Light Detection

The dorsal regions of insect eyes contain polarization-sensitive photoreceptors. Bees and desert ants have been shown to use skylight polarization patterns as a compass and can infer their heading even when the sun and much of the sky is obscured by clouds. These specialized photoreceptors are most sensitive to ultraviolet light, which is more scattered and polarized by the atmosphere than longer-wavelength "visible" light (Welmer. 1987). Polarization-sensitive cells have been found in flies (Wolf et ah. 1980). though polarization-sensitive behaviors have not been investigated in detail.

Linear Acceleration

Presumably, flying insects can also detect linear acceleration. While there are no known organs dedicated to this sense, the legs and neck of the fly are equipped with proprioceptive sensors that measure both position and strain. Presumably, flies can sense the inertia of their head and limbs and infer acceleration.

In the remainder of this dissertation, we will focus primarily on visual motion sensing. Visual motion perception underlies many interesting behaviors in the fly and could be applied to useful engineering applications. The following chapter introduces motion detection algorithms, including the model commonly used to explain early vision in the fly.

In Chapter 9, we will revisit vestibular sensing and discuss how Coriolis-force sensors might be integrated with visual motion sensors

2. MOTION DETECTION - ALGORITHMS AND VLSI IMPLEMENTATION

During the past 15 years, many analog, digital, and hybrid VLSI motion sensors have been developed and tested. Most of these designs incorporate photodetection and motion computation on the same chip. These focal-plane processors typically cannot achieve the high pixel density of dedicated CMOS imagers or CCDs, but rather trade off density for functionality. By extracting motion information at the level of light detection instead of using an external microprocessor, large savings in size, power, and system complexity is achieved.

Nearly every motion detection algorithm devised has been implemented in VLSI in some form. Motion detection algorithms can be divided into two broad classes: *feature-tracking* or *token-based* algorithms, and *intensity-based* algorithms. Models of motion detection in the fly represent a special case of intensity-based algorithms. In this chapter, we will review the principles of motion detection commonly used in both hardware implementations and biological models.

Feature-Based Motion Detection

Algorithms of this type use feature detectors to identify salient points in the raw image. Binary- tokens indicating the absence or presence of a feature are then passed on to a velocity-estimation stage. Two types of feature detectors have been used in silicon motion sensors: spatial feature detectors and temporal feature detectors

3. VLSI REIEHARDT DETECTOR DESIGN

We developed two distinct circuit architectures for VLSI Reiehardt motion detectors. Both circuits operate in

International Journal of Ethics in Engineering & Management Education
Website: www.ijeee.in (ISSN: 2348-4748, Volume 9, Issue 7, July 2022)

continuous time with analog signals, and incorporate light sensing and information processing on the same chip. The first circuit described is largely a *current-mode* design. That is, signals are represented as currents throughout the majority of the circuit. The second circuits is called the *voltage-mode* design since signals are represented a voltages, although the final output is produced as a current. Both circuits make use of the weak-inversion, or subthreshold region of operation of the AI OS transistor for micropower operation.

Current-Mode Design

Each elementary motion detector (EMD) uses photodiodes as light sensors. We use a four-transistor adaptive photoreceptor circuit developed by Delbriick (Delbriick and Mead, 1996) that produces a continuous-time output voltage proportional to the logarithm of light intensity. This circuit has a temporal lowpass characteristic with a cutoff frequency that can be set with a bias voltage. The photoreceptor is connected to a temporal derivative circuit (Mead, 1989), which has a highpass behavior. Transient firing, characteristic of a temporal highpass response, has been observed in fly laminar cells that receive input from retinal photoreceptors (Weckstrom et ah, 1992). Together, the lowpass filtering of the photoreceptor and the highpass filtering of the temporal derivative circuit form a bandpass filter which improves performance by eliminating dc illumination (which contains no motion information), and attenuating high-frequency noise such as the 120 Hz flicker of ac incandescent lighting. These bandpass filters were set to attenuate frequencies below 2.8 Hz and above 10 Hz.

The temporal derivative circuit relies on a high-gain differential amplifier in a negative feedback configuration to keep the voltage on the capacitor equal to the input voltage. As the capacitor charges and discharges to maintain this equality, the currents through the two source follower transistors (labeled "sf" in Figure) may be measured. The outputs of the temporal derivative circuit are these two unidirectional currents, which are proportional to the positive and negative components. This two-channel representation is useful for current-mode circuits, since the following translinear circuits work only with unidirectional currents. It should be noted that the use of ON and OFF channels introduces nonlinearities into the circuit that are not accounted for in the simple model described in Chapter 3.

We use the phase lag inherent in a first-order lowpass filter as a time delay. The currents from the temporal derivative circuit are passed to current-mode first-order lowpass filter circuits (Himmelbauer, 1996). These are log-domain filters that take advantage of the exponential behavior of field-effect transistors (FETs) in the subthreshold (weak inversion) region of operation. Note that two filters are needed for each EMD—one for the ON channel, and one for the OFF channel, which are processed in parallel. The time constant of the filters is controlled with a bias current that can be set externally. This time constant can be changed to tune the EMD to a specific optimal temporal frequency. We fixed this time constant to 40 ms, which gave our chip a maximum temporal frequency sensitivity of 4 Hz, similar to motion-sensitive neurons in flies (O'Carroll et ah, 1996).

To correlate the delayed and non-delayed signals for motion computation, we use a current-mode multiplier circuit. This circuit also takes advantage of the exponential behavior of subthreshold FETs to perform a computation. The weighted sum of these voltages is computed with the capacitive divider on the floating gate of the output transistor, and this transistor exponentiates the summed voltages into the output current, completing the multiplication. Any trapped charge remaining on the floating gates from fabrication is eliminated by exposing the chip to ultraviolet light, which imparts sufficient energy to the trapped electrons to allow passage through the surrounding insulator. This circuit represents one of a family of floating-gate MOS translinear circuits developed by Minch that are capable of computing arbitrary power laws with current-mode signals (Minch et ah, 1996b).

After the multiplication stage, the currents from the ON and OFF channels are summed, and the final subtraction of the left and right channels was done off-chip. Due to transistor mismatch, there was a gain error of approximately 2.5 between the left and right channels that was compensated for manually. It is interesting to note that there is no significant offset error in the output currents from each channel. This is a consequence of using translinear circuits which typically have gain errors due to transistor mismatch, but no fixed offset errors.

One entire EMD (left and right channels) consists of 31 transistors and 25 capacitors with 8.0 pF of total capacitance. Most of the capacitors were small devices (8 pm x 8 pm or less) associated with the floating-gate multiplier circuits. Each EMD takes 0.044 mm^2 of silicon area in a 2.0-pm CMOS process, including the integrated photoreceptors. By operating most of the transistors in the subthreshold regime, we achieve extremely low power dissipation (approximately 7.5 pW per elementary motion detector).

We built a simple model of the HS cell by constructing a one-dimensional array of 13 complete EMDs and linearly summing their outputs. This is easily achieved due to the current-mode nature of the EMD output signals; we simply tie all the wires together.

Voltage-Mode Design

Our voltage-mode version of the Reichardt motion detector offers several advantages over the currentmode design, including superior matching characteristics and reduced contrast dependence. To the best of our knowledge, this is the closest approximation to this biological motion sensor that has been built.

Circuit Architecture

As in the current-mode design, we measure light intensity with an adaptive photoreceptor circuit developed by Delbriick and Mead (Delbriick and Mead, 1996). This four-transistor circuit uses a substrate photodiode and source follower (Mi) to convert incident light into a logarithmically encoded voltage. A high gain amplifier (M_2 and M3) and feedback network *(Ci* and C_2) amplify the voltage signal by a factor of 18. The adaptive element (M4) acts as a nonlinear feedback element

International Journal of Ethics in Engineering & Management Education
Website: www.ijeee.in (ISSN: 2348-4748, Volume 9, Issue 7, July 2022)

that conducts only if the voltage across it exceeds several hundred millivolts. This allows the photoreceptor to adapt to large changes in illumination. Thus we maintain a large dynamic range over a wide operating range. At low bias current levels, the bandwidth of the photoreceptor is limited by the parasitic output capacitance C_p. For a detailed discussion of this circuit, see (Delbriick and Mead, 1996).

The adaptive photoreceptor signal is sent to a $g_m C$ highpass filter. We use a source follower to provide a low-impedance driver, but in future designs we will leave this out and compensate for the increased output capacitance by increasing the photoreceptor bias current I_{pr}. We use a highpass filter for two reasons. First, the ac coupling eliminates any systematic offsets caused by device variation in the adaptive photoreceptor. Second, by fixing the dc component of the signal to V_a, we can eliminate any common-mode effects later in the circuit.

The delay is accomplished with a first-order $g_m C$ lowpass filter. The bias transistor in the circuit was made several times minimum size to improve time constant matching across the chip. By operating this circuit at low current levels, we can achieve time constants useful for motion detection (10-100 ms) with reasonably sized capacitors (on the order of 1 pF).

4. VLSI MOTION DETECTOR CHARACTERIZATION

The voltage-mode elementary motion detectors described in the previous chapter demonstrated superior matching characteristics for similar pixel sizes. In this chapter, we characterize in detail the behavior of the voltage-mode EMD to both simple and complex visual stimuli.

Methodology

All of the experiments in this chapter were carried out on a 1 x 22 array of motion sensors fabricated on a 2.2 mm x 2.2 mm die in a standard 1.2 pm CMOS process. A 2.6 mm focal length lens was mounted directly over the chip, giving a 35° field of view across the entire array. The angle $4>$ between adjacent photoreceptors was 1.5°, comparable to the eyes of many flying insects (Land, 1997). The chip was biased to an appropriate operating range, and the bias settings were unchanged during all experiments, except where explicitly stated.

For experiments involving spatial integration over many sensors, the individual output currents were summed on two wires, one for the rightward-facing half-receptors (i.e., the toi signal in Figure), and one for the leftward-facing half-receptors (i.e., the $m2$ signal in Figure). The currents were measured with off-chip sense amplifiers. The two opponent signals were subtracted to yield a direction selective response.

We presented computer-generated visual stimuli on a standard monitor (Sony Multiscan 17se II) with a refresh rate of 72 Hz. Our software was able to update the screen at approximately the same rate. The bandwidth of the adaptive photoreceptors was set sufficiently low to attenuate screen refresh artifacts by 20 dB. This also prevented the photoreceptors from responding to the 120 Hz signal in ac incandescent lighting.

We generated visual stimuli with spatial resolution far exceeding the motion sensor array resolution. We used a 64-value gray scale to generate sinusoidal gratings and other complex stimuli of varying contrasts.

Direction Selectivity

Figure shows the output of a single Reichardt motion sensor and the summed output of the 22-element sensor array in response to a sinusoidal grating drifting along the sensor axis. The sensorChip testing methodology. We mounted a lens directly over the chip to focus an image on the photoreceptor array. Moving patterns were generated on a standard computer monitor. The temporal bandpass filters in each EMD blocked the 72 Hz refresh rate signal from the monitor.

array is highly direction selective, giving responses of opposite sign to motion in opposite directions. The individual sensor shows a high degree of pattern dependence superimposed on a dc direction selective response. Much of this pattern dependence is caused by device mismatch in the Gilbert multiplier. If the differential pairs are not perfectly matched, the output contains components of the raw input signals.

Pattern dependence is greatly reduced by spatially integrating over a small group of motion sensors that see different phases of the stimulus. Pattern dependence has also been observed in motion-sensitive cells in flying insects, where it is also reduced by spatial integration (Reinhardt and Egelhaaf. 1988; Single and Borst. 1998). In principle, pattern dependence could also be removed through temporal integration (averaging over time), but this would limit the response time of the sensor. We chose spatial integration, which sacrifices resolution, but maintains temporal bandwidth. The transients observed at the onset of motion are also observed in biological motion-sensitive cells and have been shown to be a consequence of summing many EMDs which see different phases of a periodic stimulus (Egelhaaf and Borst. 1989).

Spatiotemporal Frequency Tuning

Next, we varied the temporal and spatial frequencies of the sinusoidal gratings. Figures show the mean response of the sensor array, as well as the standard deviation of the signal over ten temporal cycles of the stimulus. The error bars give an indication of the magnitude of residual pattern dependence. Noise levels were far below the deterministic pattern-dependent fluctuations observed. Theoretical fits are plotted as dashed lines. These fits use Equation, including parameters from

International Journal of Ethics in Engineering & Management Education
Website: www.ijeee.in (ISSN: 2348-4748, Volume 9, Issue 7, July 2022)

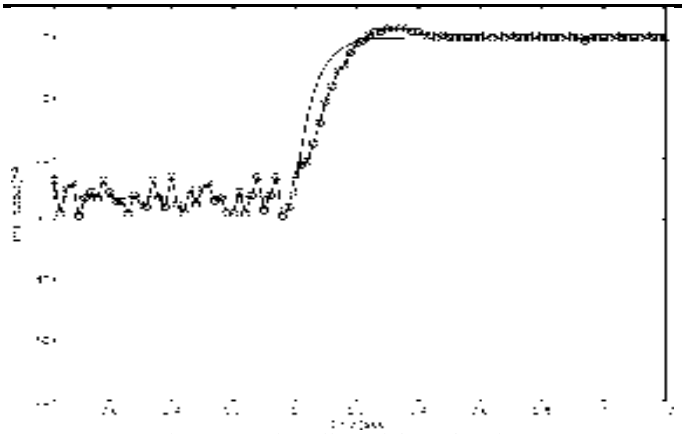

Figure 2: Motion offset: Rightward motion.

Chip output is plotted against Equation 5.33 for th = 200 ms and r = 80 ms. Only the scaling constant was varied to fit the data. The first-order temporal lowpass behavior of the adaptive photoreceptor ($T_p u_0 to$ = 30 ms) and the first-order temporal highpass filter that follows (tr = 200 ms).

The same parameters were used for all fits in Figures: r = 80 ms, $<f>$ = 1.5°, and a fixed constant of proportionality. The circuit behaves as a Reichardt motion detector over a wide range of spatial and temporal frequencies. Similar temporal frequency tuning is observed in the motion-sensitive lobular plate neurons of flies.

Spatial aliasing should produce response reversals at $N/2<j>$, where N = 1,2,3, Indeed, the first reversal can be seen near f_s = 1/2$<j>$ « 0.33 cpd. This reversal is also observed in flies, and has been used to measure their interommatidial angle $<j>$ (Gotz, 1965). The effect of aliasing at higher spatial frequencies is reduced by the finite photoreceptor size as well as by slightly defocused optics. Both effects attenuate at high spatial frequencies.

Contrast Dependence

One of the biggest disadvantages of the Reichardt motion detector is its strong (quadratic) dependence on contrast. This not only confounds contrast with spatiotemporal frequencies in the response, it also greatly amplifies the effect of high-contrast features, while attenuating low-contrast features. Studies of natural scenes have shown that low contrasts are much more common (Ruderman and Bialek, 1994), so we do not want them to be underrepresented by a motion sensor.

Interpixel Variation

As discussed in Chapter 4, good interpixel matching is essential before large arrays of local motion sensors will be feasible. We measured the matching characteristics of our motion sensors across the 22-sensor 1-D array that spanned 1.3 mm. We used a sinusoidal grating of fixed spatial frequency, and measured the mean response of each sensor in the array as we varied its velocity from —200 deg/s to +200 deg/s. Figure shows the mean and standard deviations of the 22 responses measured across the chip. The ratio of standard deviation to maximum mean response varied between 0.10 and 0.25, depending on the temporal frequency. The individual sensors perform similarly, indicating good matching of gains, dc levels, and time constants.

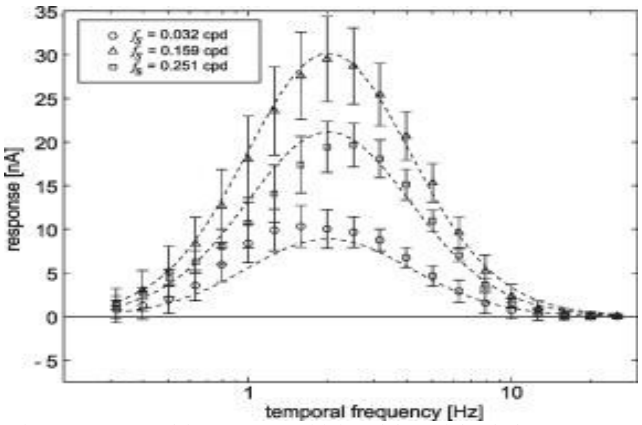

Figure 3: Temporal frequency sensitivity of the Reichardt detector array

In this figure and in those following, error bars show one standard deviation of the time response computed over 10 stimulus periods, and represent residual deterministic pattern dependence such as that seen in Figure. Noise levels were small by comparison. Dashed lines show fits to Equation

Response to Naturalistic Stimuli

While useful for initial evaluations, sinusoidal gratings are simple artificial stimuli that a creature is rather unlikely to encounter while navigating through the real world. (We have also repeated the above experiments with square-wave gratings, and the results are very similar.) We would like to characterize the performance of our sensor with real-world stimuli to test its robustness in the face of more complex visual scenes.

One of the difficulties in measuring "real-world robustness" is that complex stimuli may be hard to define and standardize. If we use a "cluttered office environment" for a visual stimulus, how does another group in a different lab reproduce this stimulus to evaluate the relative robustness of their sensor? Of course the real world is always the ultimate acid test for robustness, but we propose a useful middle ground: generating random stimuli that conform to the statistics observed in the natural environment, and using these to test sensors.

In the set of all possible images a computer monitor can display, the subset of these images that do not look like random noise is vanishingly small. It has been found that natural images exhibit a predictable statistical structure (Field. 1987; Ruderman and Bialek. 1994; Dong and Atick. 1995). These statistics hold for images of natural as well as man-made objects. Static natural scenes exhibit

International Journal of Ethics in Engineering & Management Education
Website: www.ijeee.in (ISSN: 2348-4748, Volume 9, Issue 7, July 2022)

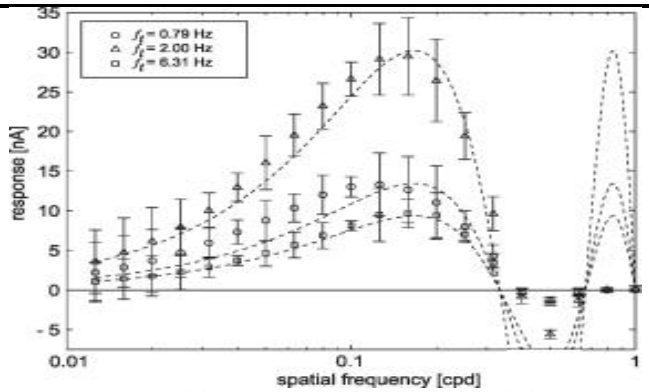

Figure 4: Spatial frequency sensitivity of the Reichardt detector array

In this figure and in those following, error bars show one standard deviation of the time response computed over 10 stimulus periods, and represent residual deterministic pattern dependence such as that seen in Figure. Noise levels were small by comparison. Dashed lines show fits to Equation 3.19. Spatial aliasing is reduced by the finite photoreceptor size as well as by slightly defocused optics.

Speed Tuning

A common criticism of the Reichardt motion detector is that it is not a true speed sensor; even with contrast saturation, the output is dependent on the spatial frequency of the stimulus. Recently, a variation of the Reichardt motion detector has been proposed that greatly reduces the spatial frequency dependence and makes it more speed sensitive (Zanker et ah, 1999).

By using half motion detectors with no opponency, we achieve a more speed-tuned response. However, the lack of opponency reduces directional selectivity. Flying insects seem to use both systems: a highly directional, temporal frequency dependent optomotor system for stabilizing flight and a speed-tuned, non-directional system for navigation (Srinivasan and Zhang, 1993; Srinivasan et ah, 1999). It remains unclear whether this second system uses a correlation-type of motion computation.

Power Dissipation

One entire Reichardt motion sensor consumes 50 nW of power at Vdd = 2.5 V under normal indoor illumination of 10 cd/m^2. Photocurrent contributes significantly to overall power consumption. Under increased illumination of 2500 cd/m^2, the power consumption increased to 110 nW. The circuit consumes approximately the same amount of power with or without motion present. This is the lowest power requirement of any motion sensor we are aware of. A 100 x 100 array of 2-D Reichardt motion sensors would consume less than 1 mW. Results described in this chapter also were published in Harrison and Koch, 2000a

5. STIMULUS RECONSTRUCTION

In this chapter, we evaluate our sensor's ability to encode information about the velocity of a simple stimulus. We also compare our sensor's encoding ability with the encoding ability of the HS cell in the lobular plate of the blowfly *Calliphora erythrocephala,* which has been previously measured by other researchers. We replicated the experiments of Haag and Borst, 1997, which use *stimulus reconstruction techniques* to measure encoding fidelity. We shall first describe the techniques and then discuss the experimental results.

Stimulus Reconstruction Techniques

One way of measuring how well a sensor encodes a stimulus is to determine how well we can reconstruct an unknown stimulus from the sensor's response. The stimulus reconstruction technique used in the following experiments finds the linear filter that transforms the sensor response into an estimate of the stimulus that is optimal in the least-squares sense. That is, the reconstruction filter minimizes the square error between the estimate and the actual stimulus. This linear reconstruction represents a lower bound on encoding ability. Nonlinear filters could of course generate an estimate with a lower error, but we compare the sensor's response variability under repeated experiments with identical stimulus conditions to generate an upper bound on encoding ability. A more detailed description of these techniques may be found in Borst and Theunissen, 1999.

Suppose we stimulate a time-invariant system with a set of i Gaussian stimuli s, (f) and record responses $r\hat{}t)$. Given the frequency-domain representation of the stimuli $Si(f)$ and the responses $Ri(f),$ the optimal reverse reconstruction filter is given by the average cross-correlation normalized by the average autocorrelation of the responses: that respond to horizontal motion in the visual held. These cells exhibit relatively symmetric hyperpolarization and depolarization on the order of ±10 mV to bidirectional motion. During stimulus presentation, which lasted 40 seconds, the intracellular potential was recorded using a sampling rate of 2 kHz. Each 40-second stimulus and response $[s(t)$ and $r(t)]$ was broken down into 4-second segments $[si(t)$ and $ri(t)]$. The 40-second stimulus was presented multiple times as well.

Using the cell's response to the random velocity stimuli, a reverse reconstruction filter was calculated using Equation 6.1. This filter was used to generate the estimated stimulus $sest(t)$. Figure shows a 600 ms excerpt of the stimulus with the estimated stimulus superimposed. Based on the linear reconstruction derived from the intracellular response, it is clear that the HS cell encodes slow changes in pattern velocity, but fails to encode rapid changes.

This frequency-limited encoding can be quantified by computing the coherence function $7^2(/)$, as shown in Figure. A coherence of 1 indicates that the stimulus is reconstructed perfectly by the reverse reconstruction filter at a certain frequency. The coherence plot indicates that the cell encodes velocity fluctuation information up to approximately 10 Hz.

There are two possible reasons for a coherence value less than one. First, there may be noise added to the signal in the cell. Second, the encoding may be nonlinear, which the linear reconstruction filter could not account for. We know from our analysis of Reichardt motion detectors in that models of the HS cell do not predict a linear encoding of velocity. In order to determine which effect (noise or nonlinear encoding),

International Journal of Ethics in Engineering & Management Education
Website: www.ijeee.in (ISSN: 2348-4748, Volume 9, Issue 7, July 2022)

dominates the loss of coherence. Haag and Borst estimated the noise independently.

Pattern Velocity Estimation by Silicon EMD Array

We repeated the above experiment using a silicon 1 x 13 EMD array. The experiment was performed using the same experimental apparatus and control software as in the fly experiments. The current mode output of the chip was converted to a voltage by a sense amplifier, and this voltage was sampled at 2 kHz during stimulus presentation.

Figure shows a segment of the stimulus and the estimation derived from the chip response. Comparing this with Figure we see that the chip performs similarly to the HS cell; it follows the stimulus at low frequencies, but misses high-frequency information. The coherence function for the silicon system is displayed in Figure. along with the expected coherence function. Figure shows the signal and noise spectra for the chip. From this data, one can see that the relative noise level in our chip is several times lower than that in the fly's HS cell. Most of the missing accuracy- in chip is due to coding nonlinearities, not noise.

Figure overlays the measure coherence functions of the HS cell and the chip. The fly encodes velocity fluctuations in the 5 Hz to 20 Hz range with higher accuracy than our chip. We believe that parasitic capacitances in the current-mode lowpass filter begin introducing nonlinearities at these frequencies, resulting in a less accurate encoding.

We used information-theoretic techniques to quantify upper and lower bounds on the mutual information encoded by the chip (Borst and Theunisseri, 1999). The lower bound and upper bound.

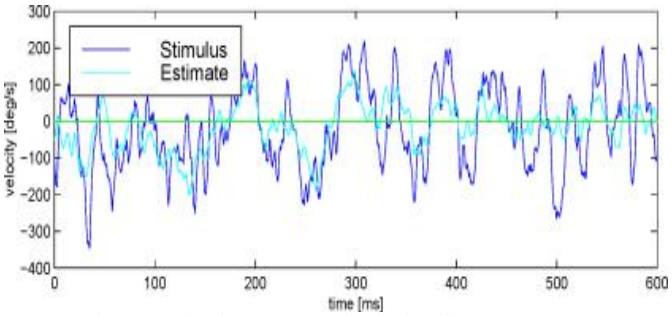

Figure 5: Stimulus reconstruction in the silicon EMD array

Using an upper frequency limit $\mathbf{f_{max}}$ of 50 Hz. the lower-bound information rate of the chip was calculated to be 27 bits/sec. compared with 37 bits/sec in the HS neuron using identical stimuli and upper frequency limit (Haag and Borst. 1998). The channel capacity (upper bound) of our sensor was 137 bits/sec compared with 110 bits/sec in the HS neuron. This is not surprising considering our chip dissipated many orders of magnitude more power than the fly's visual system. The true measure of transinforniation probably lies closer to the lower bound. Our chip used far fewer EMDs than an HS neuron, and was thus subject to greater pattern dependence which obscured velocity information. Parasitic effects and second-order nonlinearities also prevented our sensor from encoding velocity information as efficiently as the fly's HS neuron.

6. OPTOMOTOR CONTROL

In Chapter 2, we introduced the optomotor response, where visual motion information is used as a feedback control signal to estimate and cancel self-rotation. This sensorimotor loop is perhaps the best-studied visually-guided behavior of the fly. We will first describe optomotor experiments performed with flies and then describe real-time experiments where our sensor was compared directly against flies.

Experiments Previously Performed on Flies

Warzecha and Egelhaaf recently characterized the optomotor behavior of the fly under closed-loop conditions (Warzecha and Egelhaaf, 1996). A female sheepfly *(Lucilia cuprina,* Calliphoridae) was rigidly attached to a meter that measured yaw torque produced while the fly attempted to turn in response to visual stimuli, reducing the fly's behavior to a single degree of freedom. Vertical bars were presented to a large region of the fly's visual field, and could be drifted clockwise or counterclockwise. In closed-loop experiments, the fly's yaw torque was measured in real time and scaled by a constant gain term to yield angular velocity. This simulates the observed dominance of air friction in determining the instantaneous angular velocity in flies (Reichardt and Poggio, 1976). The fly's simulated angular velocity was subtracted from the angular velocity imposed by the experimenter. The resulting signal was used to control the drift rate of the visual stimulus. This simulated free-flight conditions, and allowed evaluation of the optomotor system performance.

The imposed motion schedule consisted of 3.75 s of zero imposed motion, then 7.5 s of clockwise rotation at 44 deg/s. Figure shows the torque data and resulting stimulus position for an individual trial. Figure shows the averaged data over 139 trials in a total of five animals. See Warzecha and Egelhaaf, 1996, for details on the experimental protocol.

The fly is able to stabilize its flight and cancel out most of the imposed motion. Simulation results suggest that the nonmonotonic temporal frequency response of Reichardt motion detectors results in greater stability for the optomotor control system (Warzecha and Egelhaaf, 1996). The individual trials show an oscillatory component to the torque response around 2 Hz. This oscillation is not phase-locked to the stimulus since it is not present in the average torque trace. Oscillations are not observed under open-loop conditions, suggesting they arise from optomotor feedback (Geiger and Poggio, 1981; Warzecha and Egelhaaf, 1996). Notice that despite the large amplitude of the torque oscillations, the position trace is not dominated by this effect. This fluctuation amplitude, in terms of number of photoreceptors, is close to the amplitude observed in human microsaccades (Warzecha and Egelhaaf, 1996). Poggio and colleagues observed similar oscillations in closed-loop experiments and proposed that they arose from the 60-75 ms synaptic delay inherent in the fly-visual system (Geiger and Poggio, 1981; Poggio and Reichardt, 1981).

International Journal of Ethics in Engineering & Management Education
Website: www.ijeee.in (ISSN: 2348-4748, Volume 9, Issue 7, July 2022)

Application to Autonomous Vehicle Control

Optic flow patterns produced by self-motion are one of the richest sources of navigation information available to a mobile creature (Gibson, 1950). As an animal moves through its environment, images of the outside world move across its retina in predictable ways. Objects being approached grow larger; objects left behind grow smaller. When moving forward, images of nearby objects move across the retina faster than images of distant objects. If a creature rotates in place, the entire visual scene moves across its retina at a rate that is independent of object distance. Much information can be gained from patterns of visual motion, even if no explicit object recognition is performed (Duchon et ah, 1998). Indeed, motion parallax information is immune to camouflage that can defeat even the most sophisticated static pattern recognition scheme when object and background have similar textures. Humans have no difficulty detecting the structure of randomly patterned objects against identically patterned backgrounds from motion cues alone.

Using egomotion-induced optic flow for robot navigation is a computationally demanding sensory task. By its very nature it must be done in real time. Most object recognition tasks are performed on static images, and often one can tolerate latencies of several seconds. But optic flow is available only while the robot is moving, and relevant information must be extracted in real time and fed back to the motor control system to steer the robot in the right direction. The rate of computation needed depends on the rate of robot motion, but typical real-world situations require times on the order of tens or hundreds of milliseconds.

Robot Optomotor System

As mentioned above, the fly uses visual motion information to stabilize its flight. Mismatch of body components or environmental disturbances may impart rotation on the animal, but sensory feedback is used to produce compensatory torque responses. This sensorimotor feedback is known as the optomotor system, and is one of the best-studied behaviors of the fly (Gotz, 1975; Warzecha and Egelhaaf, 1996).

7. NONLINEAR SPATIAL INTEGRATION

In the previous chapters, we have described arrays of Reichardt elementary motion detectors with linear spatial integration—currents from each EMD are tied to a single wire and sum linearly. In this chapter we will explain why linear spatial integration is undesirable and introduce a circuit architecture of nonlinear integration based on the properties of wide-held lobular plate neurons in the fly's optic lobe.

Real-world optic flow fields are sparse. Natural images have areas of little or no contrast such as blank patches of sky. The lack of spatial detail in these areas leads to "holes" in the otherwise full-field patterns of optic flow produced by rotation. When estimating self-rotation, for example, one would like to extract information based on wide-held motion that is robust against these gaps.

Gain Control in Fly Tangential Neurons

Flies have developed a remarkably elegant method for dealing with optical flow sparseness. The optic lobe in the brain of the fly contains several wide-held motion-sensitive neurons that integrate motion information from many elementary motion detectors (EMDs) in large receptive fields to produce estimations of self-rotation (Krapp and Hengstenberg, 1996). These neurons have been studied for decades, and much is known about their response properties. One property exhibited by some of these cells is called gain control, and seems to make the sensory response robust against gaps in the optical flow field.

Gain control describes the saturating response of these motion-sensitive cells with increasing stimulus size. As the extent of the stimulating pattern across the visual receptive field increases linearly, the response of the cell saturates, but it saturates at different levels for different stimulus velocities. This cannot be explained by a simple saturating output channel. The wide-held motion-sensitive neuron is integrating motion information spatially, but this integration is nonlinear. This size-dependent saturation assures that at reasonably high levels of stimulation, the cell is not sensitive to gaps in the optic flow field. (Featureless parts of the visual scene decrease the effective stimulus size.) The cell now encodes the stimulus velocity, which in this case may represent a measure of self-rotation, largely independent of the visual sparseness of the environment. Featureless areas Features.

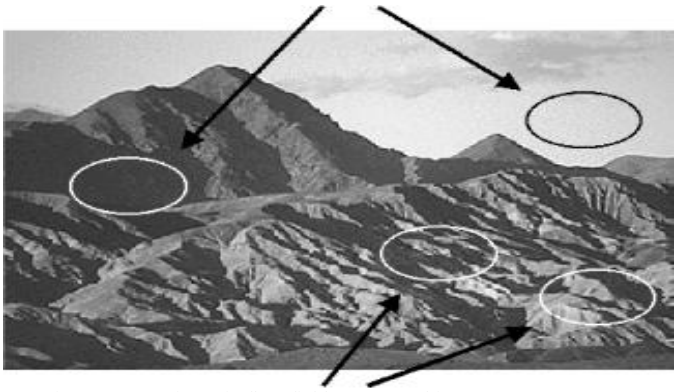

Figure 6: Gaps in optic flow fields. In natural images, some areas are featureless, and thus convey no local motion information.

They effectively decrease the pattern size. A wide-field motion sensor ideally should integrate optic flow from a large field of view, but the presence of some featureless areas should not affect its output. ("Ignore the blank parts.")

Algorithm and Biological Architecture

The neural architecture and biophysical mechanisms underlying gain control in the fly are now understood (Borst et ah. 1995; Single et ah. 1997). Simple linear models of spatial integration result in an output that is linearly dependent on stimulus size. Size-dependent saturation comes about if we use a more accurate model of the wide-field motion-sensitive neuron. In this model, the EMD outputs are not directly conveyed to the wide-field neuron. Instead, the EMDs modulate synapses, which are modeled as conductances

International Journal of Ethics in Engineering & Management Education
Website: www.ijeee.in (ISSN: 2348-4748, Volume 9, Issue 7, July 2022)

between the intracellular potential and a fixed ion reversal potential. Depending on the type of ion involved, the reversal potential can be above or below the resting potential of the cell, creating excitatory or inhibitory synapses. We connect preferred-direction EMDs to excitatory synapses and null-direction EMDs to inhibitory synapses. Each EMD modulates its corresponding ion channel conductance, which acts to pull the cell away from its resting potential, where it is held by the fixed leakage conductance.

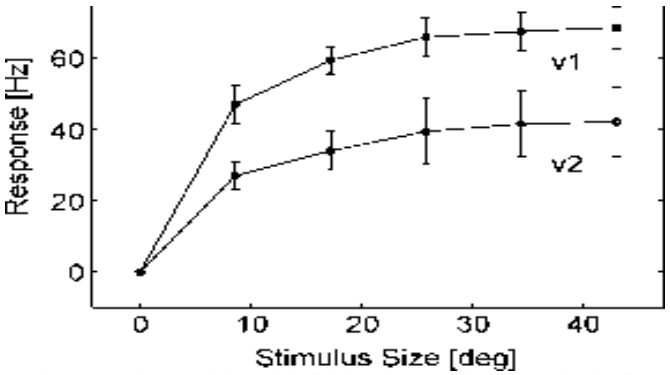

Figure 7: Gain control in a wide-field motion-sensitive neuron in the fly

The cell response saturates with increasing pattern size, and saturates at different levels depending on the stimulus velocity *(vl* = 72°/sec; *v2* = 360°/sec). Stimulus was a sinusoidal grating with spatial wavelength of 24° and 29% contrast. Data represent mean ± SKA I of extracellular recordings from the lobula plate spiking neuron HI of four different female blowflies *(Calliphora erythrocephala).* Data reprinted from (Single et ah, 1997).

8. SYSTEM INTEGRATION

As discussed in Chapter 2, flies use a wide variety of sensors to execute tasks. In most autonomous systems, sensors of many types will be present. How can the motion sensors described in the previous chapters be integrated with sensors of other types, particularly vestibular sensors? This chapter addresses possible strategies for sensory fusion that use the strengths of one sensor to compensate for deficiencies in another.

Visual Motion Sensors and Their Limitations

Visual motion sensors have a significant limitation which must be taken into account when using their signals at the system level: In the absence of a patterned visual stimulus, it is impossible to compute optic flow. If the sensor is pointed towards a blank wall or if the environment is completely dark, the motion sensor will have a "zero optic flow" output regardless of the relative motion between the sensor and the external environment. In this case, we should ignore the motion sensor. On the other hand, if the sensor is looking at a high-contrast, patterned stimulus and reports "zero optic flow," that conveys a great deal of information. In order to combine visual and vestibular information reliably, we need some type of "confidence measure" from the visual motion sensor that conveys some information about the contrast of patterns in the field of view.

We propose a simple method to return a confidence measure from the motion sensor. We measure the local spatial derivative in the direction of each motion detector, then sum the absolute value of all these spatial derivatives over the receptive field of the sensor. If this value is low, we know there is no significant image contrast, and we should ignore the output of the motion sensor.

We use an "antibump" circuit to estimate the absolute value of the spatial derivative at each pixel. This compact, five-transistor circuit, originally described by Delbriick, produces a current that approximates the absolute value of the differential voltage input over a range of about $8Ut$ (Delbriick, 1993a). Figure shows the antibump circuit. Its output current I_{out} is given by

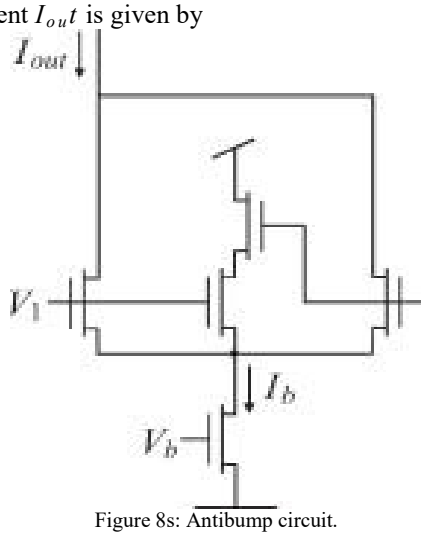

Figure 8s: Antibump circuit.

When lj and V2 are similar, most of the bias current If, flows through the middle leg. When lj and Vj are very different, most of the bias current will flow through either the left or right outer leg, increasing the output current I_{ou}

where W/L is the width-to-length ratio of each MOS transistor.

Figure shows the measured response of such a circuit fabricated in a 1.2 /mi process with S = 16. The theoretical fit indicates a functional value of S = 50. This is likely due to the narrower effective channel width also observed by Delbriick in his original experiments with this circuit (Delbriick, 1993a).

If the differential voltage to the antibump circuit comes from two adjacent photoreceptors, the output will approximate the absolute value of the spatial derivative along an axis. Since the output of these circuits is a current, it is trivial to sum the signals over the entire EMD array, and thus generate a monotonic measure of image contrast.

Coriolis-Force Gyroscopes and Their Limitations

As we discussed in Section 2.2, flies rely on Coriolis-force sensors to provide angular velocity information during flight. Many commorcially-availablo gyroscopes today operate on the same principle: A mass is driven into oscillation, and the Coriolis force normal to the axis of oscillation is measured. For a review of small Coriolis-force gyroscopes, see

Sodorkvist, 1994.

The SNR of a gyroscope is proportional to angular velocity. Typical small Coriolis-force gyroscopes have sensitivities of around 1 deg/s, and operate on 35 mW of power. However, as we decrease power usage in MEMS gyroscopes to the microwatt range, their sensitivities will decrease. Both signal-to-noise ratio and bandwidth are proportional to power dissipation in analog systems (Vittoz, 1994). Visual motion sensors can be designed to detect very slow velocities of optic flow simply byincreasing their time constants or decreasing the angular spacing between adjacent photoreceptors. Perhaps wide-field optic flow sensors could fill in information about low angular velocities.

9. CONCLUSION

In this dissertation, we have presented VLSI circuits modeled after neural mechanisms in the visual system of the fly. We have characterized these sensors and compared them against their counterparts in biology. Our biologically-inspired sensors operate at sub-milliwatt power levels—the lowest of any VLSI motion sensor we are aware of—and are capable of perceptual discriminations in situations where the signal is weaker than the noise. While testing our circuits, we developed new methodologies for evaluating the robustness of visual motion sensors in natural-image conditions. We investigated several circuit architectures to improve the robustness of these motion detectors by incorporating computational strategies used by the fly. By increasing our system's biophysical fidelity, we increased its performance—an encouraging example of how biologically-inspired approaches to engineering can yield valuable results.

REFERENCES

[1]. Adelson and Bergen, 1985] Adelson, E. H. and Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. Journal of the Optical Society of America A, 2:284-299.

[2]. Andreou et al., 1991] Andreou, A. G., Strohbehn, K., and Jenkins, R. E. (1991). Silicon retina for motion computation. In Proceedings of the 1991 IEEE International Symposium on Circuits and Systems, pages 1373-1376.

[3]. Autrum, 1958] Autrum, H. (1958). Electrophysiological analysis of the visual systems in insects. Experimental Cell Research Suppliment, 5:426-439.

[4]. Barlow and Levick, 1965] Barlow, H. B. and Levick, W. R. (1965). The mechanism of directionally selective units in the rabbit's retina. Journal of Physiology, 178:447-504.

[5]. Barrows, 1998] Barrows, G. L. (1998). Feature tracking linear optic flow sensor for 2-D optic flow measurement. In Proceedings of the 1998 International Conference on Automation, Robotics, and Computer Vision (ICARCV'98), Singapore, pages 732-736.

[6]. Benson and Delbriick, 1991] Benson, R. G. and Delbriick, T. (1991). Direction selective silicon retina that uses null inhibition. In Lippman, R., Moody, J., and Touretzky, D. S., editors, Advances in Neural Information Processing Systems 3, pages 756-763. Morgan Kaufman, San Maieo. CA.

[7]. Borst, 1990] Borst, A. (1990). How do flies land? BioScience, 40:292-299.

[8]. Borst and Bahde, 1988] Borst, A. and Bahde, S. (1988). Spatio-temporal integration of motion. Naturwissenschaften, 75:265-267.

[9]. Borst and Egelhaaf, 1989] Borst, A. and Egelhaaf, M. (1989). Principles of visual motion detection. Trends in Neurosciences, 12:297-306.

[10]. Borst and Egelhaaf, 1990] Borst, A. and Egelhaaf, M. (1990). Direction selectivity of blowfly motion-sensitive neurons is computed in a two-stage process. Proceedings of the National Academy of Sciences USA, 87:9363-9367.

[11]. Borst et ah, 1995] Borst, A., Egelhaaf, M., and Haag, J. (1995). Mechanisms of dendritic integration underlying gain control in fly motion-sensitive interneurons. Journal of Computational Neuroscience, 2:5-18.