

Statistical adaptation and optimal estimation in movement computation by the blowfly visual system

R.R. de Ruyter van Steveninck, W. Bialek, M. Potters*, R.H. Carlson
NEC Research Institute, 4 Independence Way, Princeton NJ 08540

INTRODUCTION

Here we combine ideas from two lines of research in visual information processing. The first is the study of the fundamental limits to the reliability of selected visual tasks as set by noise at the photoreceptor level. The second is the attempt to describe certain visual tasks, thought to be of critical importance, in a functional mathematical framework.

Probably the clearest example of the latter approach can be found in the work of Reichardt and collaborators. Starting with a formulation for the fundamental interactions involved in movement detection (Reichardt 1957), this line of thought was generalized to more complex visual computations, such as figure-ground discrimination (e.g. Reichardt and Poggio, 1976), and given a systematic mathematical formulation in terms of Volterra series (Poggio and Reichardt, 1976). These authors have always emphasized the fundamental role played by nonlinear interactions in solving nontrivial visual tasks.

Early examples of the first line of research are Hecht et al. (1942) and van der Velden (1944), who were interested in questions pertaining to the detectability of single photons, and in de Vries (1943) and Rose (1948) who analyzed the effects of photon shot noise on simple perceptual tasks. The work in this field was extended to measuring the reliability of nerve responses by Barlow and Levick (1969). Research into factors limiting the accuracy of vision has been very productive, and over the years the results have lent more support to the notion that visual signal processing may approach optimal performance under the right conditions (e.g. Savage and Banks 1992). Generally speaking, however, this line of research has been constrained to relatively simple stimuli where the total energy could be well-defined, such as light flashes, presentations of gratings for limited intervals and the like. The detection task in such cases ultimately consists of estimating a (positive or negative) excess photon count over a linearly weighted area of the retina. Although we certainly gain insight into the statistical efficiency of various processing levels in the

visual system, many interesting higher visual functions are not probed in this way.

To our knowledge, the case of nonlinear interactions between noisy photoreceptor signals has received little attention. In the following we therefore revisit the problem of a well-defined nonlinear computation, namely movement detection. Instead of assuming noiseless input as in Reichardt's original approach, we will explicitly incorporate noise in the photoreceptors that provide the input to the computation. For simple movement stimuli such as sudden displacement steps of wide field patterns, this problem is tractable. Comparing measurements of the statistical efficiency of the blowfly movement-sensitive neuron known as H1 with the performance of an ideal observer who uses realistic photoreceptor signals (de Ruyter van Steveninck 1986, de Ruyter van Steveninck and Bialek 1992), we find that over a 10 ms time interval the neuron approaches the ideal observer closely. It thus seems that under these simple conditions, the fly tries to compute movement in an optimal way on a behaviorally relevant time scale. One of the questions prompted by this result is whether optimal processing can be generalized to more complex visual tasks, and if so, what role adaptational processes may play in this.

AN INTRODUCTORY EXAMPLE

It has been known for some time that the response dynamics of H1 is highly adaptive (Zaagman et al. 1983, Maddess and Laughlin 1985, de Ruyter van Steveninck et al. 1986) over a wide range, with characteristic response decay times ranging from over 300 ms to about 10 ms. The time constant can be set by various dynamic parameters of the stimulus. In the case of stepwise movement, it adapts to the time interval between steps, independent of step size or contrast (Zaagman et al. 1983). Also, sinusoidal flicker of the full field (Borst and Egelhaaf 1987), and random flicker of independent pixels as well as random pattern movement (de Ruyter van Steveninck, unpublished), influence the time constant. Further, in a situation where moving gratings are used as a stimulus, the time constant is set over a significant range by the velocity

* Present address: Istituto di Fisica, Università di Roma, Roma, Italy

of the pattern, independent of its spatial wavelength or contrast (de Ruyter van Steveninck et al. 1986). Finally, the time constant is independent of the direction of movement, and is set locally in the visual field.

Given the robustness and the magnitude of the effects, it would be highly surprising if this form of adaptation did not serve an important purpose. One possibility is that the effects seen in H1 are the result of adaptive low-pass filtering in retinotopic columns. The idea is that the statistical structure of the signals entering a photoreceptor is determined by the statistics of the visual scene, the optical properties of the photoreceptor, and the angular speed of the photoreceptor relative to the scene. Consequently, a measurement of pattern speed and salient contrast parameters provides the animal not only with data for direct use, but also gives it statistical information about the signals it could expect next. Specifically, if a pattern moves across the retina at angular velocity v , the power density spectrum of the optical signal transmitted by the photoreceptor optics is:

$$S(f) = \frac{C(f/v) \cdot M(f/v)}{v} \quad (1)$$

with $C(k)$ the contrast power density of the pattern, $M(k)$ the modulation power transfer function, k the spatial frequency (in cycles/degree), and f the induced temporal frequency in Hz. Thus a spatial frequency k will be mapped to a temporal frequency $f=vk$. Recently, power spectra of natural scenes, woods in this case, were measured (Ruderman and Bialek 1994), and shown to have power-law behavior over at least three decades of spatial frequency. Converted to the one-dimensional case relevant here, $C(k)$ is given as:

$$C(k) \propto \left(\frac{k_0}{k}\right)^{1-\eta},$$

with $\eta=0.19$ and formally $k_0=1$ cycle/ $^\circ$ to make the term in brackets dimensionless. Further, $M(k)$ is usually modeled as a Gaussian:

$$M(k) = \exp[-(2\pi k)^2 \sigma_s^2]$$

with a spatial width $\sigma_s=0.51^\circ$. Next we must take into account the signal transfer properties and the noise power spectrum of the individual photoreceptor. These were measured by intracellular recording, under the same illumination conditions that were used in the experiments on H1 (de Ruyter van Steveninck 1986). For our purposes, signal transfer is described

adequately by a linear contrast power transfer function $T(f)$, while the noise is Gaussian and given by the power density spectrum $N(f)$. A good descriptive fit to the ratio of these functions is:

$$N_{eff}(f) = \frac{N(f)}{T(f)} \approx 1.25 \cdot 10^{-4} \left[1 + \left(\frac{f}{f_c} \right)^{2.7} \right] \quad (2)$$

which describes the equivalent noise power spectrum referred back to the stimulus. f_c is about 18 Hz. We now assume the following:

- The visual system *measures* the speed and contrast variance of the pattern.




Fig. 1 Dots: values of the characteristic decay time of H1's response to pattern steps, after preadapting the neuron with movement at a velocity given by the abscissa. Square-wave patterns were used with contrasts as shown in the figure. Solid lines are filter time constants computed as described in text.

- The cutoff frequency of a filter within the visual system is adapted to filter out frequencies where the signal to noise ratio (SNR) of the photoreceptor input is too low.

The cutoff frequency f_0 can be found by assuming the form of $C(k)$ given above, scaling it with a prefactor to make the contrast variance equal to the measured value, measuring the speed and using knowledge of the effective noise power spectrum. Figure 1 presents the result of the computation of f_0 , and compares it to data obtained from H1 published earlier (de Ruyter van Steveninck et al. 1986). The results are expressed as the measured quantity $\tau=1/2\pi f_0$, which can be thought of as the integration time of a filter, adjusted to eliminate frequencies above f_0 . Obviously, the match between theory and experiment is far from perfect. But we must remember that the fly tries to solve a complicated problem, of which our model description can at best be a crude approximation. Also, there are no free parameters in our calculation. In several important respects there is agreement, however: First, for most of the range the orders of magnitude predicted for τ are within a factor of three from the measured values. Second, as also found for H1, τ should depend only weakly on contrast for low adaptation velocities. This is mainly due to the fact that at low velocities the optical signal is low-pass filtered with a very steep slope induced by the modulation transfer function. Third, the point of breakaway from the trend of decreasing τ with increasing velocity, does depend on contrast, which again is in line with the observations on H1. What we mean to demonstrate here is that the adaptive behavior of H1 is consistent with a mechanism that could serve an important purpose, namely suppressing high frequency components for which the SNR drops below a criterion value. This is a very useful thing to do prior to combining signals in nonlinear interactions, because such interactions mix frequencies. For example, the low-frequency output of a second-order interaction results from the convolution of the two input spectra. Suppose that the inputs contain useful signals only at low frequencies and that broad band noise is added. Then, even if the low frequency SNR of the inputs to the computation is high, the output of the computation may be completely dominated by the noise, because in the interaction the noise-noise cross product is integrated over the full bandwidth. This argument is quite general in the sense that the statistics of photon arrival induces broad band noise in the photoreceptors. Therefore, prefiltering signals prior to nonlinearities is important. Further, prefiltering should be adaptive if

the computations are to span a reasonable dynamic range.

These observations have prompted us to investigate the question in a more general and perhaps more natural situation, namely continuous estimation of the velocity of a wide-field pattern subject to Brownian motion. The full mathematical treatment of this problem is beyond the scope of this paper, but is given elsewhere (Potters 1994, Potters and Bialek 1994). Here we present a more intuitive approach which can be understood from first principles. Some of the predictions are presented and compared to results from experiments on the H1 neuron.

OPTIMAL VELOCITY ESTIMATION

The desideratum is to make the best least-squares estimate of the velocity of a wide-field pattern, given the photoreceptor signal-to-noise characteristics. Mathematically, the formulation is that of a statistical estimation problem in which the trajectory of a continuous random motion signal is estimated. One implication is that qualitative aspects of the computation should change when the signal-to-noise ratio at the photoreceptor level changes from a high to a low value. A rather dramatic demonstration of this is the prediction that the pure velocity estimator in the limit of high SNR should approach the gradient scheme (Limb and Murphy 1975), whereas at low SNR, the velocity estimator should shift its computation to a form of the original Reichardt correlator, with a multiplicative interaction between two linearly filtered versions of the photoreceptor signals. The latter situation is forced by a tradeoff between statistical and systematic errors: At high SNR, the optimal estimator can afford to compute high-order terms between photoreceptor signals, thus improving its velocity estimate by reducing systematic error. At low SNR, however, the statistical error in higher order terms in the computation grows faster than in lower order terms. Therefore, at very low SNR only the lowest-order term that contains any movement information, the second-order one, should be used.

The transition between the two regimes is smooth, and the particular version of the gradient model predicted can be seen as a suitably normalized version of a Reichardt correlator, where the normalization adapts to the SNR. A simplified version of this adaptive movement sensor is given by the following relation:

$$v_{est}(t) \equiv \frac{-\int dx S_x \cdot S_t}{D^{-1} + \int dx S_x \cdot S_x} \quad (3)$$

on contrast increases with decreasing light levels. This is expected from the simple picture: at lower light levels, the relative SNR decreases, thus widening the contrast range over which correlation operations rather than gradient schemes are optimal. In the experiment presented here, the range should scale to first approximation inversely with intensity, which is roughly correct. Although the response at the lowest light level increases slower than at the higher levels, its value at maximum contrast is not that different from the other cases. This is an interesting phenomenon, and it could point to a mechanism in which the dependence H1's response on contrast is gradually suppressed as the SNR increases. The following experiment was designed to test this more explicitly.

Adaptation to contrast

A more detailed way of looking at contrast dependence is to probe adaptation to contrast. As an adaptation stimulus, in this experiment the fly saw a Gaussian random pattern move at constant velocity (8 omm/s) in the neuron's null direction. (Adaptation movement in the preferred direction gave qualitatively the same results, although there were differences in detail resulting from the tonic high firing rate induced by such an adaptation stimulus.) After 3 seconds of adaptation movement, the velocity changed sign, and the contrast changed to a test contrast value. This test phase lasted for 1 second, whereupon the adaptation stimulus resumed. The estimated values of adaptation contrast at the photoreceptor level in this experiment were: 0.0, 0.02, and 0.2.

We probe H1's response to the test contrast, given a certain value of the adaptation contrast. To minimize the effect of transients induced by the contrast switch from adaptation to test, the movement response was computed over the last 500 ms of the presentation of the test stimulus. The results are shown in Fig. 3 at two test contrast resolutions. The top figure shows that the test response reaches a plateau as a function of test contrast, and the plateau values do not depend very much on adaptation contrast. On the other hand, the contrast gain of the test response depends strongly on the adaptation contrast in the region of low test contrasts, as can be seen from the bottom figure. This graph also shows that in the low contrast region, H1 shows typical correlator-like behavior, in that the movement response is a quadratic function of contrast. However, the gain of the correlator is very sensitive to the adaptation contrast. In this case, the gain varies from $2.5 \cdot 10^4$ to $0.3 \cdot 10^4$ spikes/[s·(rms contrast)²] when rms adaptation contrast changes from 0 to 0.2.

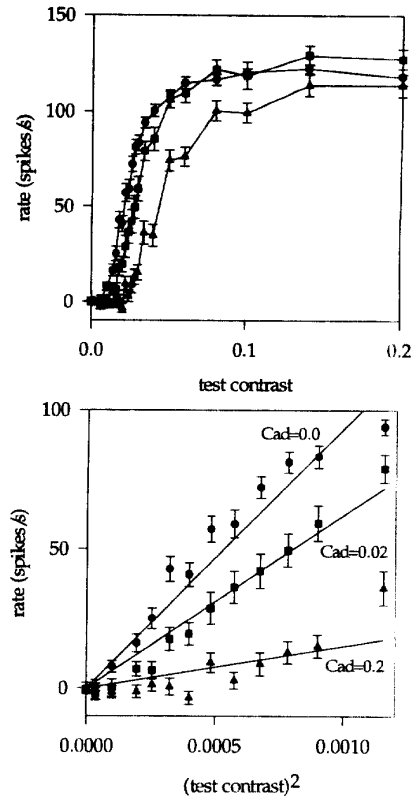


Fig 3. Dependence of H1's movement response on test contrast, when adapted to different adaptation contrasts. The data in the bottom figure are for the low contrast regime only, and the abscissa is quadratic. The fact that straight lines on this plot fit the data well means that for this test contrast range, H1 responds like a correlator.

Dependence on contrast and velocity

Here we probe the simultaneous dependence of the movement response on contrast and velocity. In order to keep the neuron in a reasonably steady state of adaptation we presented null-direction movement at a constant velocity for periods of 3 seconds, interleaved with preferred direction movement during 1 second. Meanwhile the contrast of the pattern changed in triangular fashion, with a period of 100 seconds so that we probed the response to preferred direction movement 25 times each period. The pattern consisted of 1 omm. wide bars, set dark or bright randomly. The result, expressed as the averaged firing rate during each preferred-direction stimulation as a function of both contrast and velocity, is shown in Fig. 4. Clearly the response plateaus as a function of rms contrast for contrast values above 0.1. However, even if the response as a function of contrast is well in the plateau region, the response still depends strongly on velocity

with D the diffusion constant parametrizing the whole-field Brownian motion of the pattern, and where S_x and S_t are spatial and temporal derivatives of the spatiotemporal function $S(x,t)$. This function is in effect the signal to noise ratio at the level of the photoreceptors. At the intensities used in laboratory conditions $S(x,t)$ is almost linearly related to the light intensities. In the full description the differentiation operations are implemented by biphasic filters, the shapes of which depend on stimulus parameters. Here we ignore these dependencies, as we are more interested in the qualitative behavior. For a rigidly moving pattern:

$$S(x,t) = S(x - v \cdot t), \quad (4)$$

and if noise were not present one could simply estimate velocity by taking the ratio of the time and space derivatives of S . This is what happens in the limit of (3) when the integral dominates the denominator. At low SNR, where D^{-1} dominates, the estimated velocity is the product of two linearly filtered versions of the spatiotemporal light intensity, which is in essence the classical Reichardt correlator. If this type of flexibility is built in to the fly visual system, we should be able to adapt the fly's brain in different ways to regimes of different SNR and probe its properties. The experiments described below are an attempt to do so.

METHODS

The experiments were performed on female wild type *Calliphora vicina*, caught outdoors. The fly was put in a plastic tube, its wings, thorax and head fixed with wax. Care was taken to leave the proboscis free so that the animal could be fed occasionally. The H1 neuron was recorded by an extracellular tungsten electrode reaching the lobula plate through a small hole cut in the back of the head.

Patterns were generated using a Digital Signal Processor board (Ariel) based on a Motorola 56001 processor, and consisted of frames of nominally 200 vertical lines, written at a frame rate of 500 Hz. Thus the patterns were essentially 1-dimensional, but extended in the vertical direction. They were displayed on a Tektronix 608 monitor (phosphor P31), at a radiance of $165 \text{ mW}/(\text{sr}\cdot\text{cm}^2)$. Taking spectral and optical characteristics of the photoreceptor lens-wave guide into account, we estimate a flux of effectively transduced photons of about $4 \cdot 10^4 \text{ s}^{-1}$ (de Ruyter van Steveninck, 1986).

Prior to each set of experiments the distance of the fly to the screen was adjusted so that 4 lines on the screen

corresponded to 1 horizontal interommatidial angle, which for the fly is about 1.35° . All velocity values in the subsequent experiments are therefore given in units of ommatidia of the fly's horizontal sampling raster per second, i.e. omm/s.

RESULTS

Contrast and light intensity

In this experiment the fly watched a pattern moving in the preferred direction of H1, at a speed of 8 omm/s. The pattern consisted of vertical lines with intensities set independently from the others according to a Gaussian distribution with rms value 0.5 and mean 1.0. However, because the line width was set at one quarter of the interommatidial angle, the lines were substantially blurred by the photoreceptor optics, resulting in an estimated rms contrast value of 0.2. The experiment was done for three values of the light intensity, by having the fly look at the screen either directly, or through a $D=1$ or $D=2$ neutral density filter. Figure 2 shows the result of this experiment for the three light levels. For $D=0$ the response of H1 increases steeply as contrast increases and reaches a plateau value at relatively low contrast values. A noteworthy feature of the curve is also that it is not symmetric: the response to the downward contrast flank is different from that to the upward one. As the stimulus period was 300 seconds, this means that long-term adaptation effects play a role in shaping the response. For lower light intensities, it can be seen that the range over which the response depends more or less quadratically




Fig 2. Response of H1 to movement in the preferred direction, at moderate speed (about 8 ommatidia/s). During movement the contrast of the random bar pattern was ramped from 0 to 0.2 and back, as shown. Three levels of illumination were used: low (black curve), middle (dark gray curve), and high (light gray curve). Illumination levels were spaced by factors of ten.

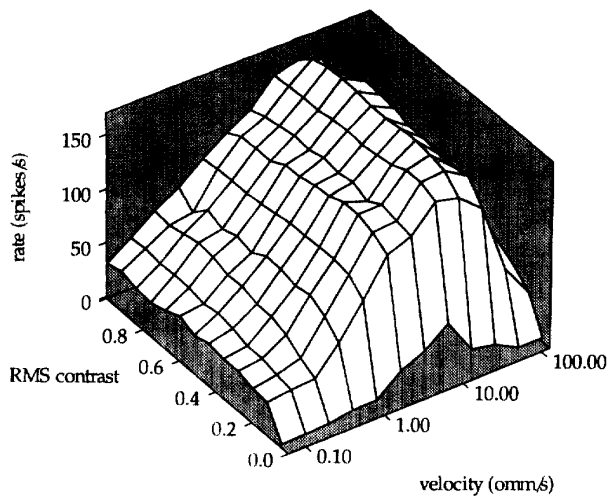


Fig. 4: Firing rate of H1 as a function of both contrast and velocity. Velocity was presented alternately in null direction (3 s, $v=8$ omn/s, and preferred direction (1 s, v as given by axis). Rate in this figure is the average response to the 1 second test stimulus.

over a wide dynamic range. So, at low contrast the output of H1 is ambiguous in the sense that it depends both on contrast and velocity. At higher contrasts, however, H1 seems to have normalized contrast away, and is sensitive to velocity only. This is in line with the behavior predicted by Eq. 3, but we should of course be careful interpreting these results, as there are other parameters that influence H1's response, such as the spatial structure of the stimulus.

CONCLUSION

Probably the most important conclusion we can draw from our results, is that the visual system of the fly seems to be very adaptive in its computations. Furthermore, we tentatively understand some of the observed behavior as the result of incorporating statistical knowledge about stimulus structure into neural computations, in order to optimize the processing of signals that are ultimately derived from noisy photoreceptor voltages. If our interpretations are correct, adaptation of the kind we observe here should be much more widespread. Another important lesson is that experimental conditions can be manipulated such that H1 shows distinct features of either a correlator, or of a pure velocity sensor. Thus, proponents of the correlator scheme, and those in favor of gradient models (see e.g. Buchner 1984) may really be looking at two sides of the same coin, namely the adaptive movement sensor described by Eq. 3.

ACKNOWLEDGEMENTS

We thank Al Schweitzer for his wonderful technical support.

REFERENCES

- Barlow, H.B. and Levick, W.R. 1969 Three factors limiting the reliable detection of light by retinal ganglion cells of the cat. *J. Physiol.* **200**:1-24.
- Bialek, W., Rieke, F., de Ruyter van Steveninck, R.R. Warland, D. 1991 Reading a neural code. *Science* **252**, 854-1857.
- Borst, A., Egelhaaf, M. 1987 Temporal modulation of luminance adapts time constant of fly movement detectors. *Biol. Cybern.* **56**, 209-215
- Buchner, E. 1984 Behavioural analysis of spatial vision in insects. In: *Photoreception and vision in invertebrates* M.A. Ali ed., pp. 561-622. New York, NY: Plenum Press.
- Hecht, S., Schlaer, S., Pirenne, M.H. 1942 Energy, quanta, and vision. *J. Gen. Physiol.* **25**,819-840.
- Limb, J.O., Murphy, J.A. 1975 Estimating the velocity of moving objects in television signals. *Comp.Graph.Image.Proc.* **4**,311-327
- Maddess T., Laughlin S.B. 1985 Adaptation of the motion sensitive neuron H1 is generated locally and governed by contrast frequency. *Proc. R. Soc. Lond. B* **225**,251-275
- Poggio, T. Reichardt, W. 1976 Visual control of orientation behaviour in the fly, Part II. Towards the underlying neural interactions. *Q. Rev.Biophys.* **9**,377-438.
- Potters, M. 1994 Toward a predictive theory of neural computation: motion estimation in the fly's visual system. Thesis, University of Princeton.
- Potters, M., Bialek, W. 1994 Statistical Mechanics and Visual Signal Processing. Submitted to *J. de Physique*.
- Reichardt, W. 1957 Autokorrelations-Auswertung als Funktionsprinzip des Zentral-nervensystems. *Z. Naturf.* **12b**,448-457.
- Reichardt, W., Poggio, T. 1976 Visual control of orientation behaviour in the fly, Part I. A quantitative analysis. *Q. Rev.Biophys.* **9**,311-375.
- Rose, A. 1948 The sensitivity performance of the human eye on an absolute scale. *J. Opt. Soc. Am.* **38**,196-208.
- Ruderman, D., Bialek W. 1994 Statistics of natural images: Scaling in the woods. *Phys. Rev. Letters* **73**,814-818
- de Ruyter van Steveninck, R.R. 1986 Real-time performance of a movement-sensitive neuron in the blowfly visual system. Thesis, Rijksuniversiteit Groningen, The Netherlands.
- de Ruyter van Steveninck, R.R., Zaagman, W.H., Mastebroek, H.A.K. 1986 Adaptation of transient responses of a movement-sensitive neuron in the visual system of the blowfly *Calliphora erythrocephala*. *Biol. Cybern.* **54**,223-236.
- de Ruyter van Steveninck, R.R., Bialek, W. 1992 Statistical reliability of a blowfly movement-sensitive neuron. In: *Advances in Neural Information Processing Systems* Moody, J.E., Hanson, S.J., Lippmann, R.P. eds., pp 27-34. Morgan Kaufmann, San Mateo, California.
- Savage, G.L., Banks, M.S. 1992 Scotopic visual efficiency: Constraints by optics, receptor properties and rod pooling. *Vision Res* **32-4**, 645-656.
- de Vries, Hl. 1943 The quantum character of light and its bearing upon threshold of vision, the differential sensitivity and visual acuity of the eye. *Physica* **10**,553-564.
- van der Velden, H.A. 1944 Over het aantal lichtquanta dat nodig is voor een lichtprikkel bij het menselijk oog. *Physica* **11**,179-189.