Real-Time Performance of a Movement-Sensitive Neuron in the Blowfly Visual System: Coding and Information Transfer in Short Spike Sequences

R. De Ruyter Van Steveninck; W. Bialek

Real-time performance of a movement-sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences

BY R. DE RUYTER VAN STEVENINCK[1]† AND W. BIALEK[2]

[1] *Department of Biophysics, Laboratorium voor Algemene Natuurkunde, Rijksuniversiteit Groningen, Westersingel 34, 9718 CM Groningen, The Netherlands*
[2] *Departments of Physics and Biophysics, University of California at Berkeley, Berkeley, California 94720, U.S.A.*

We develop model-independent methods for characterizing the information carried by particular features of a neural spike train as it encodes continuously varying stimuli. These methods consist, in essence, of an inverse statistical approach; instead of asking for the statistics of neural responses to a given stimulus we describe the probability distribution of stimuli that give rise to a certain short pattern of spikes. These 'response-conditional ensembles' contain *all* the information about the stimulus that a hypothetical observer of the spike train may obtain. The structure of these distributions thus provides a quantitative picture of the neural code, and certain integrals of these distributions determine the absolute information in bits carried by a given spike sequence.

These methods are applied to a movement-sensitive neuron (H1) in the visual system of the blowfly *Calliphora erythrocephala*. The stimulus is chosen as the time-varying angular velocity of a (spatially) random pattern, and we consider segments of the spike train of up to three spikes with specified spike-intervals. We demonstrate that, with extensive analysis, a single experiment of roughly one hour's duration is sufficient to provide reliable estimates of the relevant probability distributions.

From the experimentally determined probability distributions we are able to draw several conclusions. (1) Under the conditions of our experiment, observation of a single spike carries roughly 0.36 bits of information, but spike pairs carry an interval-dependent signal that can be much larger than 0.72 bits; estimates of the total information capacity are in rough agreement with the maximum possible capacity given the signal-to-noise characteristics of the photoreceptors. (2) On average a single spike signals the occurrence of a velocity waveform that is positive (movement in the excitatory direction) at all times before the spike, whereas spike pairs can signal both positive and negative velocities, depending on the inter-spike interval. (3) Although inter-spike intervals are crucial in extracting all the coded information, the code is robust to

† Current address: Institute of Audiology, University Hospital Groningen, P.O. Box 30.001, NL 9700 RB Groningen, The Netherlands.

several millisecond errors in the estimate of spike arrival times. (4) Short spike sequences give reliable information about specific features of the stimulus waveform, and this specificity can be quantified. (5) Our results suggest approximate strategies for reading the neural code – reconstructing the stimulus from observations of the spike train – and some preliminary reconstructions are presented. Some tentative attempts are made to relate our results to the more general questions of coding and computation in the nervous system.

1. INTRODUCTION

The sensory systems of even relatively primitive creatures receive information from the environment at an enormous rate, and much of this sensory information must be processed in real time. Seldom does an organism have the luxury of expecting many precise copies of a stimulus to be transmitted in sequence, and the time scales for behavioural decision-making are often comparable to the timescales that characterize variations in the stimulus itself. This real-time mode of operation of the nervous system should be contrasted with the conventional outlook in neurophysiology. In most studies, especially on spiking neurons, stimulus–response relations are characterized by responses averaged over a large number of presentations of identical stimuli. This averaging procedure provides an experimentally convenient description of neural responses allowing, for example, the definition of 'receptive fields' and other characteristics of filtering and signal transduction in the nervous system. It should be clear, however, that these average responses are of limited relevance for an organism that does not itself have the opportunity to average.

This study is an attempt to characterize the spike train of a single neuron in a way that conforms more directly to the problem faced by the organism. The aim is to understand what inferences can be made about a time-dependent input signal by a hypothetical observer of the ongoing spike activity. Because we are dealing with a stochastic system this must necessarily be done in probabilistic terms. Specifically, the goal of our experiment is to determine the probability distributions that describe how short stretches of a continuously varying stimulus are encoded in short spike sequences. Given these distributions, a variety of theoretical results can be used to analyse the code quantitatively.

(i) We can determine the absolute information content, in bits, of particular spike sequences without any *a priori* assumptions about the informational significance of different features of the spike train.

(ii) We can characterize the way in which the space of all possible stimuli is divided into regions represented by distinct spike sequences. Do variations, for example, in inter-spike intervals correspond to variations in specific temporal features of the stimulus waveform?

(iii) We can define the degree of precision in spike timing that is required if the maximum possible information is to be extracted from the neuron, or alternatively the extent to which the neural code is robust to timing errors.

(iv) We can, at least approximately, combine information from successive spike sequences literally to read the code, providing an algorithm for real-time stimulus reconstruction using only the spike train as input.

We wish to emphasize a key feature of this approach, namely that it is not tied to any model of the coding process. On the contrary, if one has reason to suppose that particular features of the spike train carry special significance then one can actually use our methods to test this hypothesis.

In one sense our method of analysis is the opposite of the conventional one. Rather than considering the set of responses induced by a certain type of stimulus, we are interested in the set of stimuli leading to a certain type of response; it is this inverse relation that allows us to interpret the response as representing certain particular stimuli. In practice the stimulus is a computer-generated random wave-form, and we will average continuous short stretches of this waveform that immediately precede identical short spike sequences. The computation of such a conditional mean stimulus waveform is in the spirit of de Boer's reverse correlation method (de Boer 1967; de Boer & de Jongh 1978); reviews on this and related methods are given by Marmarelis & Marmarelis (1978) and by Eggermont *et al.* (1983). A crucial extension is that here the events forming the conditions are stretches of the neuronal response spread out in time, instead of single spikes. Furthermore, we will characterize not only the conditional mean waveform but rather the conditional probability *distribution* of waveforms. By estimating these distributions we learn not only what waveform is most likely to have generated a given spike sequence, but also the extent to which this most likely waveform can be trusted as a reliable estimate of the actual stimulus. This quantitative notion of reliability is essential for a meaningful characterization of the neural code.

It is clear that experiments of the sort described here must be done on rather carefully chosen preparations. Obvious requirements, as we are trying to answer intrinsically statistical questions, are long stable recordings that allow for large and well-behaved data sets. More fundamentally, we would like to study a system in which enough is known about both physiology and behaviour to be certain that we are looking at biologically meaningful stimuli and at neural responses on behaviourally relevant timescales. We have chosen to study H1, a wide-field movement-sensitive neuron in the visual system of the blowfly *Calliphora erythrocephala*. This neuron is directionally selective, its preferred stimulus being horizontal inward (back to front) movement, and has a receptive field almost covering one hemisphere (Hausen 1982). Vertical movement leaves the neuron virtually unaffected, whereas movement in the horizontal outward (null) direction suppresses its spontaneous activity. In the experiment described below, the fly watches a wide-field pattern that moves randomly. The H1 neuron is described as 'movement-sensitive', but of course other parameters of the stimulus such as contrast and spatial structure also influence its response. These parameters are kept constant throughout the experiment, so that for most of the analysis we characterize the stimulus by one function of time, the angular velocity of the pattern. In §6c we return to the consideration of other stimulus variables in connection with adaptation and coding ambiguity.

The starting point in our approach is the observation that the H1 neuron must convey its messages rapidly, because it acts as a link in the optomotor course-control loop (Hausen & Wehrhahn 1983; Hausen 1984). Typical response times in visually induced behaviour of free-flying flies are of the order of 30 ms (Land & Collett 1974), during which time the H1 neuron typically fires at most a few spikes.

The neuroanatomy of the fly visual system (Hausen 1984) suggests that there is little, if any, opportunity for averaging over several neurons, and the behavioural timescales tell us that there is no opportunity to average over many spikes. This system thus typifies the problem of real-time processing to which we alluded at the outset.

In the following sections we present the theoretical foundations for our probabilistic discussion, the methods required for the experiment itself, our approach to processing the raw experimental data and finally the results and implications derived from a specific experiment on H1.

2. THEORETICAL FOUNDATIONS

(a) A short review of earlier work

After Shannon (1948, 1949) published his classic work on information theory, the concepts of coding and information transfer have been widely used in both theoretical and experimental studies of neuronal signal transport. Following von Neumann's (1956) emphasis on the severity of reliability problems in computational systems the notions of reliability and redundancy in neural coding have also been widely discussed.

A comprehensive review of the concepts involved in neural coding is provided by Perkel & Bullock (1968). The authors give an exhaustive survey of possible coding schemes, and illustrate various possibilities with experimental data on a wide variety of preparations. The issue of reliability is treated by Bullock (1970), who presents a critique of the view (see, for example, Burns 1968) that the brain operates in a probabilistic manner with unreliable components. One of his important points is that neurons perform with a *degree* of reliability that must be quantified experimentally. In fact a quantitative experimental description of neural reliability has been rather elusive.

McKay & McCulloch (1952) consider the maximum amount of information that a spiking neuron can carry. Their analysis is based on interval coding, in which the information transfer is limited both by the refractory time and by the smallest discriminable time-interval between successive spike firing times. For typical values of these parameters maximum information rates of the order of 1000 bits s^{-1} are computed. In a comparable model, supplemented by experimental data, Rapaport & Horvath (1960) estimate a maximum information capacity of 4000 bits s^{-1}.

Stein (1967) challenges these results by stating that in most cases spiking neurons convey their information by a rate code, i.e. the information is contained in the number of spikes fired within a time window substantially longer than the average spike-interval. This is an assumption, however, which lacks direct experimental support. There are two more reasons why this assumption is unattractive. First, the information transfer is made to depend on an arbitrary choice of the duration of the time window, and it seems very difficult to make an independent estimate of this duration. Secondly, only the total number of spikes in the window is considered, so that all information carried by the structure of the spike train within the window is discarded. A related problem is that the analysis is

restricted to static stimuli. This difficulty is removed in two subsequent papers (Stein & French 1970; Stein *et al.* 1972), where the information transfer of two neuronal models is studied, by using the coherence function to describe the linear part of the coupling between stimulus and response. The basic problem remains, however, that these theoretical studies begin with definite postulates about the nature of the neural code rather than providing a guide for experimental exploration of the code.

In most experimental studies, assumptions are also made regarding the encoding of stimulus signals by the nerve cell. Grüsser (1962) and Grüsser *et al.* (1962) treat information transfer in cat retinal-ganglion cells by using a spike-counting model as in Stein's work, considering the discrimination of different steady-state inputs based on the number of spikes in a certain time-interval. An approach along the same lines is taken by Werner & Mountcastle (1965) in their study on mechano-receptive cutaneous afferents. A somewhat more sophisticated approach, but still based on a frequency-coding model, is offered by Walløe (1970), who investigates information transfer by stretch receptors. Finally, Färber (1968) considers the possibility of distribution coding, in which the information is carried by more statistical parameters than the mean rate alone, but again the analysis is limited to static stimuli. The main drawback of these approaches is that they offer no insight into the encoding of time-varying stimuli, and of course the results depend on the particular choice of the encoding model.

A procedure in which no assumptions are made regarding the encoding is described by Eckhorn & Pöpel (1974, 1975) in a study on the information capacity of cat lateral geniculate nucleus cells. In their method, an input–output joint distribution is computed from experimental data. Both the input, a pseudorandom train of light flashes, and the input, a spike train, are considered binary functions of time. As long as stationarity is a good approximation the method provides a complete description of the joint distribution over limited time intervals. From this the information transfer can be computed directly. The distribution's dimensionality grows linearly with the length of this time interval, hence the analysis is restricted to very short stretches of stimulus and response signal. Furthermore, the method is designed to avoid the question of encoding, so that the end result is essentially only a number representing the average information transfer, and no insight is obtained into the translation between stimulus signals and spike sequences. Finally, binary strings describe at best a limited class of stimuli that may not be typical of the cat's sensory ecology.

A very different class of experiments was pioneered by Fitzhugh (1957) and by Barlow & Levick (1969). These authors examine the reliability of neurons in forced-choice discrimination tasks between pairs of known stimuli. Such experiments are direct analogues of psychophysical experiments on humans, and they provide a rather direct definition of neuronal reliability; we have in fact generalized these ideas and applied them to H1 (de Ruyter van Steveninck *et al.* 1985; de Ruyter van Steveninck 1986; Bialek *et al.* 1986). It should be clear that the problem of discriminating among a discrete set of known possibilities is quite distinct from the more natural problem of real-time signal processing and estimation that we consider here.

The study of neural coding has by now accumulated a rather long history, and many specific experimental and theoretical results have been obtained. We hope that this short review has made clear, however, that there is still a need for a more direct experimental attack on the problem.

(b) Response-conditional ensembles and their information content

The neuronal response is a stimulus-dependent stochastic signal. This dependence is partial because the system is not noise-free, one obvious source of randomness in the visual system being photon shot noise (see §6d). The aim of the following analysis is to give a probabilistic description of the relation between both signals. To this end, neuronal responses are ordered in categories defined by the inter-spike intervals. More precisely we imagine taking a snapshot of the spike train at some observation time t_{obs}. At this instant a time t_0 has elapsed since the last spike, while the second-to-last spike occurred a time t_{-1} further in the past, and so on; see figure 1 below. Let us call this particular snapshot of the spike train R, because it is some particular response of the neuron. We would like to know what the observation of this snapshot R tells us about the stimulus, which in the H1 experiments is the angular velocity $v(t)$ of a pattern as it moves across the visual field.

Everything that we know about $v(t)$ by virtue of observing R is by definition contained in the conditional probability distribution $P[v(t)|R]$. This distribution measures the likelihood that any particular stimulus $v(t)$ gave rise to the observed segment of the spike train. More precisely we might thus refer to $P[v(t)|R]$ as a distribution functional as it describes the probability that the function $v(t)$ has a certain form. From $P[v(t)|R]$ we can find the most likely waveform and we can estimate the expected fluctuations around this most likely signal. The problem is thus to determine $P[v(t)|R]$ experimentally.

Imagine that we have performed a very long experiment, presenting the system with some randomly chosen, continuously varying $v(t)$. In such a long experiment, the response R will have occurred many times; there will be many observation times t_{obs} such that the last spike occurred at $t_{obs}-t_0$, and so on. Looking backwards in time from each such t_{obs} we will see some particular waveform $v(t_{obs}+\tau)$, where τ is negative. Let us imagine keeping a list of all these waveforms that preceded the response R. What we have done in making this list is to choose out of the statistical ensemble of all possible waveforms in our experiment a particular sub-ensemble, the *response-conditional ensemble*. This sub-ensemble in fact consists of waveforms that are chosen randomly out of the distribution $P[v(t_{obs}+\tau)|R]$, so that our experiment has given us a sort of Monte Carlo sampling of this distribution. This sampling constitutes all our experimental knowledge of the response-conditional ensemble. To parametrize this knowledge we can compute various moments of the distribution, with more samples of course required if we are to estimate the higher moments reliably.

Although we can continue to discuss abstract distributions, some practical comments are in order. In principle, $P[v(t_{obs}+\tau)|R]$ defines the probability distribution for a continuous function of time; in practice, of course, we must take discrete time samples if our experimental data are to be processed digitally. This

means that $v(\tau)$ is a finite dimensional vector measuring the velocity at a discrete set of times preceding t_{obs}. These distributions $P[v(t_{obs} + \tau) \,|\, R]$ are then probability *densities* in that $P[v] \mathrm{d}v_1 \, \mathrm{d}v_2 \ldots v_N$ is the probability that the N samples v_i of the waveform takes values in small regions of width $\mathrm{d}v_i$. Note that unless we build correlations into the stimulus itself, a spike sequence consisting only of events earlier than some observation time can, by causality, only inform us about stimuli before this time. In the experiment described below we will in fact use uncorrelated (white-noise) stimuli. Secondly, we clearly cannot attempt to evaluate all the moments of the response-conditional distributions. In practice we will represent this distribution as a multi-dimensional Gaussian, characterized by the mean velocity vector $w_R(\tau)$ and the covariance matrix $C_R(\tau_1, \tau_2)$, with τ the discrete time index and R denoting the particular response that forms the condition. We will check that this is a good approximation by computing a subset of the third and fourth moments.

Consider again our hypothetical observer monitoring the spike train in real time. Given that a response of a certain category has just occurred, the best estimate he can make about the immediate past of the stimulus is the mean waveform conditional on that response category. A measure of how good his estimate is, i.e. how closely he can expect the actual waveform to match the estimate, is given by the covariance of the conditional sub-ensemble. This can be seen by observing that the covariance is the multi-dimensional analogue of the variance, describing the 'width' of a distribution in the space of all possible stimulus waveforms. The concept of coding thus emerges in a natural way: loosely speaking the message conveyed by a particular spike sequence is the mean waveform conditional on that sequence, and the uncertainty in the message is given by the conditional covariance.

In the absence of any observations on the spike train all we know is that the stimulus waveform was chosen from some *a priori* probability distribution $P_0[v(\tau)]$ determined by the experimental or environmental conditions. The observation of a certain response changes our statistical knowledge from the *a priori* ensemble to the conditional ensemble: some classes of stimuli are judged more probable and others less probable than before the observation of R. A useful measure of this change is the Shannon (1948) information gained by observing such a response. For each response category this quantity can be computed directly from the two distribution functions $P[v \,|\, R]$ and $P_0[v]$. Let the information provided by the occurrence of a response R about the occurrence of a particular waveform $v(\tau)$ be denoted by $I[v(\tau); R]$. Then, with I in bits we have (Fano 1961):

$$I[v(\tau); R] = \ln\left(\frac{P[v(\tau) \,|\, R]}{P_0[v(\tau)]}\right). \tag{1}$$

The average information carried by R about the complete set of stimulus velocity waveforms is the expectation value of $I[v(\tau); R]$, found by integrating over all stimulus waveforms:

$$I(R) = \int \mathrm{D}v \, P[v(\tau) \,|\, R] \, I[v(\tau); R], \tag{2}$$

where the integration is over the entire space of waveforms, and Dv represents the infinitesimal volume element $dv(\tau_1)\,dv(\tau_2)\ldots dv(\tau_n)$.

In the case where the distributions are Gaussian $I(R)$ can be evaluated directly. Explicitly, we have the conditional distribution

$$P[v(\tau)\,|\,R] = [(2\pi)^n\,|C_R|]^{-\frac{1}{2}}\exp\left[-\tfrac{1}{2}(v(\tau_1)-w_R(\tau_1))^T C_R^{-1}(\tau_1,\tau_2)\,(v(\tau_2)-w_R(\tau_2))\right], \tag{3}$$

where $|C_R|$ is the determinant of the covariance matrix, C_R^{-1} is its inverse, $(\ )^T$ denotes the transpose of a vector $(\ )$, $w_R(\tau)$ is the mean velocity waveform conditional on the response R, and n is the dimension of the vectors. For conciseness of notation the time dependency will not be explicitly indicated in the following. Similarly, the *a priori* probability density is given by

$$P_0[v] = [(2\pi)^n\,|C_0|]^{-\frac{1}{2}}\exp\left[-\tfrac{1}{2}v^T C_0^{-1}v\right], \tag{4}$$

in which by construction C_0 is a diagonal matrix (there are no correlations in the *a priori* ensemble), and the diagonal elements of C_0 are the variance of the stimulus velocity. After some calculation we find that the integral defining $I(R)$ becomes

$$I(R) = \tfrac{1}{2}[\log_2(|C_0|/|C_R|) + \mathrm{Tr}\,(C_0^{-1}C_R - U) + w_R^T C_0^{-1}\,w_R], \tag{5}$$

with U the unit matrix and $\mathrm{Tr}(\ )$ denoting the trace of the matrix $(\ )$.

Through the first two terms in (5) the information content depends on the change in the covariance, whereas the third depends on the shape of the conditional mean waveform. The first term has a particularly simple interpretation. Because the determinants are measures of the volume in stimulus space associated with the two matrices, this term measures the information gained by the reduction in volume of the *a posteriori* stimulus space relative to that of the *a priori* distribution.

To summarize, we can characterize response-conditional ensembles by their moments, which are directly estimated from experiment. Having done this we can examine the structure of the ensemble, as described in detail below, and we can also compute the information conveyed by each spike sequence R.

(c) Reconstruction of the stimulus waveform

Our methods, as described above, provide precise data on the information conveyed by particular spike sequences R. Truly to read the neural code – provide an algorithm for real-time reconstruction of the stimulus based on the spike train alone – we must find a way to combine information from successive sequences. At present we have no experimental guidance on how to do this combination, so we will make what is admittedly a crude approximation, namely that successive spike sequences R are generated independently. Note that if by 'spike sequence' we simply mean a single spike, then we are assuming that neural firing is a Poisson process, which is indeed a strong assumption. But if 'spike sequence' includes groups of two or three spikes our assumption is progressively less stringent. Clearly as our definition of the responses categories R becomes more refined – as we look back further and further in time from t_{obs} – the assumption of independence is more and more accurate. In any event, it is our only possibility given our present data set, so it seems to be worth trying.

Mathematically the independence assumption is

$$P[R_1, R_2, ..., R_k \,|\, v] = \prod_{i=1}^{k} P[R_i \,|\, v]. \tag{6}$$

We are interested in finding the function $v(t) = v_{\text{est}}(t)$ that maximizes the conditional probability $P[v \,|\, R_1, ..., R_k]$, that is the most likely velocity given that the responses $R_1, ..., R_k$ have been observed. Now, by Bayes's theorem,

$$P[v \,|\, R_1, ..., R_k] = \frac{P[R_1, ..., R_k \,|\, v]\, P_0[v]}{P[R_1, ..., R_k]}, \tag{7}$$

from which, using (6):

$$P[v \,|\, R_1, ..., R_k] = \left(\prod_{i=1}^{k} P[R_i \,|\, v] \right) \frac{P_0[v]}{P[R_1, ..., R_k]}. \tag{8}$$

Again applying Bayes' theorem we have

$$P[R_i \,|\, v] = \frac{P[v \,|\, R_i]\, P[R_i]}{P_0[v]}, \tag{9}$$

which substituted in (8) yields

$$P[v \,|\, R_1, ..., R_k] = \left(\frac{1}{P[R_1, ..., R_k]} \prod_{i=1}^{k} P[R_i] \right) P_0[v] \left(\prod_{i=1}^{k} \frac{P[v \,|\, R_i]}{P_0[v]} \right), \tag{10}$$

in which the first factor is a constant, given that the particular combination of the R_i was observed, and the second factor is the *a priori* probability density. The third factor is a product of ratios of the probabilities of the relevant conditional velocity waveforms to those of the *a priori* waveforms. The probability densities involved in the second and third factor are the multi-dimensional Gaussians given by (3) and (4).

We wish to maximize expression (10) to find the velocity with maximum likelihood v_{est}. In the Gaussian approximation the probability distributions are especially simple, and we can find the maximum likelihood v as usual by demanding that the derivatives of $P[v \,|\, R_1, ..., R_k]$ with respect to each element of the vector v be zero. The result is that

$$v_{\text{est}} = \left[\sum_{i=1}^{k} (C_{R_i}^{-1} - C_0^{-1}) + C_0^{-1} \right]^{-1} \sum_{i=1}^{k} C_{R_i}^{-1}\, w_{R_i}, \tag{11}$$

where it should be remembered that v_{est} is a function of time. In the computations presented below, v_{est} at a certain moment in time includes contributions from the spike train up to 100 ms ahead. Through (11) all responses $R_i\,(i = 1, 2, ..., k)$ occurring in this time-window contribute to the calculation of v_{est}. Of course, the R_i occur at different instants during this window, and the corresponding velocity vectors w_{R_i} and covariances C_{R_i} must be shifted in accordance with these occurrence times. Note that if this procedure is truly done in real time then the observer may at a certain moment estimate the stimulus in the near past and then modify this estimate at some later time in response to new incoming information.

3. Methods

(a) Preparation and recording

In the experiment described here we used female *Calliphora erythrocephala* from the laboratory stock. Spike activity from the H1 neuron was recorded extra-cellularly with a tungsten microelectrode, which penetrated the back of the head through a small hole cut in the integument. Spike-interval times were digitized in units of 500 μs. These digitized intervals were stored on disk for off-line analysis. Further details are described in de Ruyter van Steveninck *et al.* (1986).

(b) Stimulus generation and presentation

The stimulus was a pattern displayed on a cathode ray tube (CRT). It consisted of 2048 vertical lines with a 0.029° spacing, set dark or bright at random with equal probability. When viewed through the photoreceptor aperture this is an excellent approximation to Gaussian white noise. The mean radiance was $32 \text{ mW m}^{-2} \text{ sr}^{-1}$ at a frame repetition frequency of 800 Hz. The fly viewed the screen via a square $(30.5°)^2$ diaphragm.

To the X-deflection voltage on the CRT a pseudo-random signal was added, generated by a Data-General MP/200 microcomputer. Each 500 μs, i.e. synchronized with the spike-interval timer clock, a random-valued movement step was selected from a sequence previously downloaded into core-memory. In this way the entire pattern jittered by means of small movement steps in the horizontal direction, with the time between steps much less than the integration time of the fly's photoreceptors (typically 13 ms under the conditions of this experiment, see de Ruyter van Steveninck (1986)). The sizes of these steps were integer multiples of 0.066°. The movement signal was cyclic, with a full period of 2^{16} timer ticks, or about 33 s. During the second half of this period, the pseudo-random signal repeated the sequence of the first half, except for a reversal of sign of the move-ment. In this way, the stimulus pattern executed a random walk with mean velocity zero. During the experiment the stimulus sequence was repeated until a sufficient number of spikes was recorded. In the experiment described here, 100 complete sequences were presented during approximately 53 min.

The sizes of the movement steps were identically distributed and were set independently at each clock pulse. The sizes were chosen from seven values, generated with the probabilities listed in table 1. The mean value as well as the skewness and the excess of this distribution are zero, and the variance is one unit. The distribution was obtained by transforming a uniformly distributed sequence generated by combining the outputs of a congruential generator and a shift register (Marsaglia *et al.* 1973), as implemented in the PORT library routine UNI (Fox *et al.* 1976).

TABLE 1. Probability of occurrence of movement steps in the
pseudo-random stimulus

(The unit step size in the experiment is 0.066°, and steps were presented at 500 μs intervals.)

value (units)	−3	−2	−1	0	1	2	3
probability	1/216	12/216	51/216	88/216	51/216	12/216	1/216

4. DATA ANALYSIS AND DEFINITIONS

(a) *Accumulation of the conditional moments*

The result of the experiment is a time-sequence of spike occurrences, together with a stimulus velocity signal, both stored on disk. The principle employed in the subsequent off-line analysis is shown in figure 1. The top of the figure shows a part of the velocity signal $u(t)$, together with the sequence of spikes it induces. For the purpose of illustration the stimulus shown is a smoothed version of the actual stimulus. A time pointer (arrow) shows the instant up to which the analysis has proceeded. At this moment, t_{obs}, the last occurrence of a spike was 20 ms earlier, and the spike immediately preceding the latter occurred 30 ms ago. The immediate
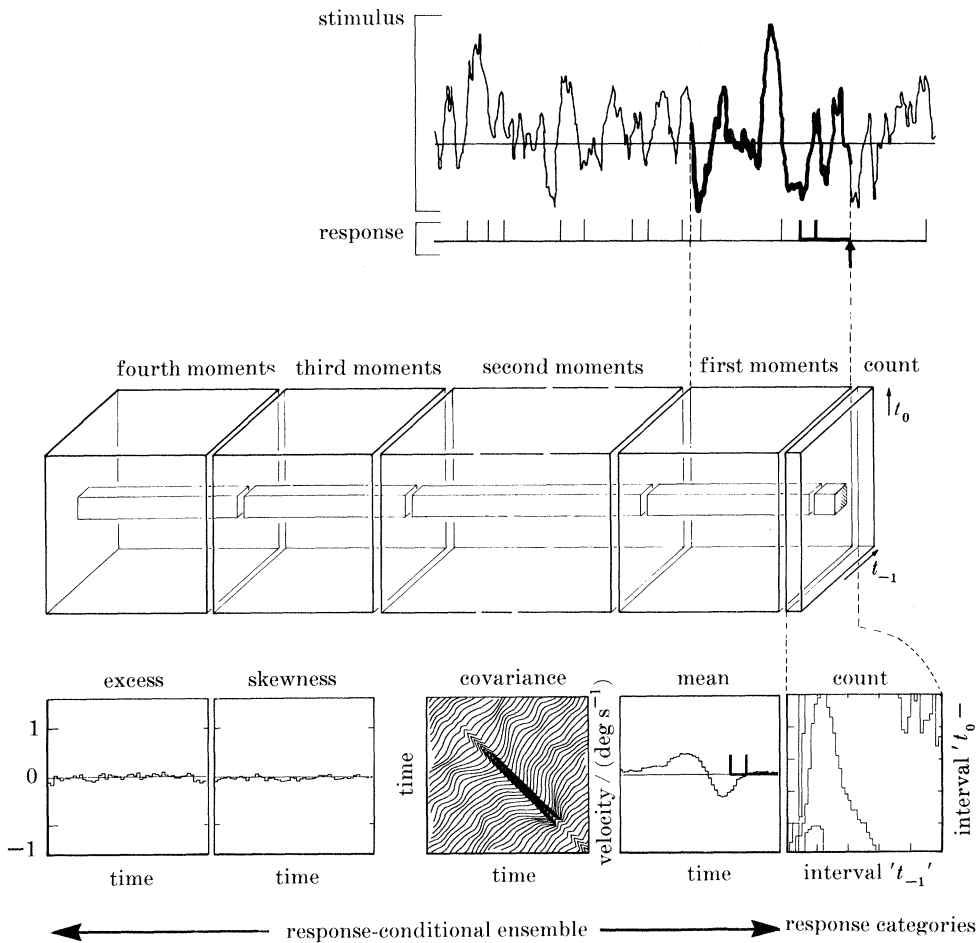
FIGURE 1. Explanation of the procedure followed in off-line analysis. The two traces at the top show a random stimulus (in this case the angular velocity in rigid movement of a wide-field pattern) together with the neuron's response, both as functions of time. In the large block at the centre, the first through fourth moments of the stimulus waveform are summed, conditional on different neuronal firing patterns. The bottom trace presents an example of the end result for one response condition after the full experiment has been analysed.

history of the neuron's signal can be approximated by two numbers: the time since the firing of the last spike, designated t_0, and the spike-interval between this spike and the preceding one, t_{-1}. In the example $t_0 = 20$ ms, and $t_{-1} = 10$ ms. This part of the immediate past of the response is drawn in a heavy line.

Once the response category is determined, the counter for this category is incremented by 1, while the velocity values $u(t_{obs} + \tau)$ of the stimulus waveform in a time-window immediately preceding t_{obs} (drawn in a heavy line) are added to a vector labelled by the response category. (In all results shown here τ runs from -100 to 0 ms.) The counter and the velocity vector are represented by the small cube and the bar respectively, drawn at the right inside the large block in figure 1. In addition, the cross-products $u(t_{obs} + \tau_1) u(t_{obs} + \tau_2)$ of all pairs of velocity values in the window are added into a matrix, and the third and fourth diagonal moments $u^3(t_{obs} + \tau)$ and $u^4(t_{obs} + \tau)$ are added to two vectors. This matrix and both these vectors are also labelled by the response category; they are represented in figure 1 by the continuation of the bar to the left. All the labelled vectors are composed of 50 time bins, each 2 ms wide. Because of its symmetry, the labelled cross-correlation matrix can be stored in an array of 1275 (2 ms)^2 elements instead of the full 2500. When all parameters are updated, the time pointer moves one bin ahead. In this next time bin, the pointer may or may not encounter a spike. If it does not, the procedure described above is repeated for a response category with an interval t_0 incremented by one time bin and t_{-1} remaining the same. If instead a spike occurs at the new t_{obs}, the new value of t_{-1} takes the value of the previous t_0, the new t_0 is set to 0 and the cycle repeats, starting with this new response category.

In the analysis each of the dimensions (each interval) of the response category is subdivided into 50 bins, each 1 ms wide, so that the total number of response categories is 2500. Along each dimension the 50 bins are supplemented by an extra bin in which the moments of all responses are summed for which the corresponding interval lengths exceed the time-span of the 50 bins.

(b) Computation of the response-conditional ensemble

If the data of the complete experiment are processed in the way described above we obtain, for each response category R, a count of the number of its occurrences in the spike train, together with the corresponding time-dependent first through fourth moments about the origin, summed over all occurrences of R. Each such collection will be called a response-conditional sum (RCS). From every RCS the following functions, acting as the parameters of a multidimensional distribution, can be computed: (i) the mean velocity $(w_R(\tau) = \langle u(\tau) \rangle_R)$; (ii) the covariance matrix $(C_R(\tau_1, \tau_2) = \langle u(\tau_1) u(\tau_2) \rangle_R - w_R(\tau_1) w_R(\tau_2))$; (iii) the diagonal skewness $(s_R(\tau) = \langle u_M^3(\tau) \rangle_R / [C_R(\tau, \tau)]^{\frac{3}{2}})$; and (iv) the diagonal excess $(e_R(\tau) = \langle u_M^4(\tau) \rangle_R / [C_R(\tau, \tau)]^2 - 3)$, where $\langle \ \rangle_R$ denotes averaging the quantity over all occurrences of the response category R during the experiment, and u_M^3 and u_M^4 are the third and fourth moments about the mean. The latter can be found from the first through third, and the first through fourth moments about the origin, respectively (Abramowitz & Stegun 1965). The results of these computations are shown at the bottom of figure 1. Together they constitute a parametric approximation

representing a probability density of waveforms in stimulus space. Such a set of waveforms with its associated probability distribution will be called a response-conditional ensemble (RCE), which is a sub-ensemble of the *a priori* stimulus ensemble E_0. For each response category, the mean and the covariance of the RCE describe a conditional multidimensional Gaussian approximation to the probability density in the vector space V of waveforms $v(\tau)$, as defined in (3) above.

The diagonal skewness and excess are computed to check that the measured distribution is a reasonable approximation to a Gaussian. If this approximation is correct, these two should be zero which, as can be seen from figure 1, is approximately true. Examples (except the skewness) for other response categories are shown in §5(*b*). In no case did we find large deviations of either the skewness or the excess, so we can be confident that the Gaussian approximation is reasonable.

In principle more complex response categories could be defined by taking account of more intervals. This is precluded, however, both by limited storage capacity and by limited experimental time. The latter limitation arises because a sufficient number of occurrences (of the order of 300) in each response category is required if we are to make a reliable estimate of the information content, as noted below.

(c) *Response categories and their notation*

The response category R shown in figure 1 can be symbolized by $R = ['t_{-1}'t_0{}^-]$, where the ordering of the superscripts (') and ($^-$) provides a pictorial notation for the type of firing pattern. The subscripted interval times denote the durations of both parts of the response category. For specific cases these are given in ms, so that here $R = ['10'20^-]$. The large block in the middle of figure 1 contains the RSCs with the response categories $['t_{-1}'t_0{}^-]$ for all values of t_{-1} and t_0. From this entire set of RCSS we can compute subsets. For example, the RCS for the category $R = ['t_0{}^-]$ – a single spike occurring t_0 ms in the past – is found by summing all RCSS with $R = ['t_{-1}'t_0{}^-]$ over all values of t_{-1}. In principle, response categories of the form $R = ['t_{-1}'t_0']$ – responses ending with a spike – can also be found from the set of RCSS with $R = ['t_{-1}'t_0{}^-]$ by taking the differences of RCSS with two adjacent values of t_0, but in practice it is easier to extract these cases directly from the spike train. More complicated subsets could be constructed, e.g. one in which two spikes occur at defined positions and a third spike at an unspecified time in between, but these cases will not be treated.

Table 2 lists the symbolic notations for all response categories that are constructed, together with the way in which they are computed. Where appropriate, pictorial notations for the response categories are used in the figures throughout, as in the mean waveform depicted in figure 1. These notations are also given in table 2. In specific cases the response acting as the condition on the mean velocity waveform and the covariance is given in square brackets. For example, $w['10']$ and $C['10']$ are the mean velocity and the covariance conditional on a 10 ms spike-interval.

TABLE 2. SYMBOLIC AND PICTORIAL NOTATIONS FOR STRETCHES OF NEURONAL
RESPONSE ACTING AS CONDITIONS IN COMPUTING THE RESPONSE-CONDITIONAL
SUMS (RCSS), TOGETHER WITH THE STRATEGY FOR THEIR COMPUTATION

notation		
symbolic	pictorial	computation
RCS $('t_{-1}'t_0^-)$		directly from experiment
RCS $(^-t_{-1}'t_0^-)$		$\sum_{\alpha > t_{-1}} \text{RCS}('\alpha't_0^-)$
RCS $('t_0^-)$		$\sum_{\alpha > 0} \text{RCS}('\alpha't_0^-)$
RCS $(^-t_0^-)$		$\sum_{\alpha > t_0} \text{RCS}('\alpha^-)$
RCS (0^-)		$\sum_{\alpha > 0} \text{RCS}('\alpha^-)$
RCS $('t_{-1}'t_0')$		directly from experiment
RCS $(^-t_{-1}'t_0')$		$\sum_{\alpha > t_{-1}} \text{RCS}('\alpha't_0')$
RCS $('t_0')$		$\sum_{\alpha > 0} \text{RCS}('\alpha't_0')$
RCS $(^-t_0')$		$\sum_{\alpha > t_0} \text{RCS}('\alpha')$
RCS $(0')$		$\sum_{\alpha > 0} \text{RCS}('\alpha')$

(d) The information content of the RCES

Equation (5) defines the information content of an RCE in the Gaussian approximation. Naïve attempts to evaluate this function by using experimental estimates of the covariance matrix led to significant difficulties. The problem is that the determinant of the covariance matrix is a strongly nonlinear function of the matrix elements, so that random statistical errors in the matrix elements lead to *systematic* errors in the estimate of the determinant. We have examined this statistical problem both by analytical methods and by numerical simulations, and found that to insure negligible systematic errors in the 50-dimensional problem treated here requires at least 300–500 samples of any particular response category. For several categories we have more than this number, and in these cases we noticed a very important feature of the covariance matrix, namely that only a small number, typically one or two, or its eigenvalues are different from the eigenvalues of the *a priori* covariance matrix. These one or two distinguished eigenvalues have associated eigenvectors that are smooth functions of time, whereas the eigenvectors belonging to higher eigenvalues are not. (Some examples of eigenvalue spectra and eigenvectors are shown in figures 8–10.) These results suggest that we should approximate the determinant in all cases by the product of the two lowest eigenvalues with all other eigenvalues replaced by their *a priori*

values (one, in suitable units). Again we emphasize that for those R with large samples, we know experimentally that this approximation is excellent.

Note also that, because most of the power of the stimulus signal is above the cutoff frequency of the visual system, one expects at the outset that the covariance matrix contains meaningful components only in the low-frequency region. Because of this, the off-diagonal elements can be smoothed to decrease the influence of statistical errors in the computation of the eigenvalues themselves. In the results presented in §5 the procedure followed is to subtract C_0 from C_R, then smooth $(C_R - C_0)$ with a symmetrical two-dimensional Gaussian filter of standard deviation 4 ms, and finally add C_0 again. Note that 4 ms is much smaller than the photoreceptor integration time (13 ms) under the conditions of our experiments (de Ruyter van Steveninck 1986).

5. RESULTS

(a) General

Here we present the results of one experiment only. Although the results from experiments done under different conditions differ in detail, the essential features remain the same. For reference an overview is presented of the amount of data processed, and some other global parameters of the experiment. The response conditions of the form $['t_{-1}'t_0']$, i.e. with a spike occurring at t_{obs} (closed intervals), and of the form $['t_{-1}'t_0']$, i.e. with no spike occurring at t_{obs}, (open-ended intervals) are worked out separately. Because the computations in the case of open-ended intervals are much more time-consuming, these are done on a smaller body of experimental data than for the closed intervals.

The experiment lasted 53 min, during which time 100 stimulus sequences were presented. A total of 123 077 spikes were recorded, with an average firing rate of 39 spikes s^{-1}. The full amount of spikes was processed for the closed intervals, and about half this number for the open-ended intervals. Figure 2 shows the count
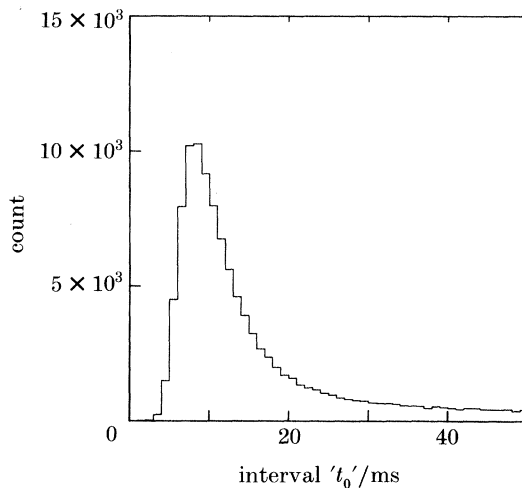
FIGURE 2. Count histogram of the occurrences of closed spike intervals $['t_0']$ over the entire experiment. The bin width is 1 ms.

histogram for single intervals $['t_0']$, which peaks at $t_0 = 9$ ms. The two-dimensional count histogram for $['t_{-1}'t_0']$ is presented as a contour plot in figure 3. The figure is symmetrical under the interchange of the two intervals, but successive intervals are not generated independently. This is demonstrated by figure 4, which shows interval histograms for the second interval, conditional on three different choices for the length of the first. Clearly, there is a preference for short intervals to be followed (or, because of the time symmetry, preceded) by relatively short ones and conversely for the long intervals. A possible mechanistic interpretation of this behaviour is that the input signal of the spike generator has a correlation time of the order of the mode of the interval distribution.
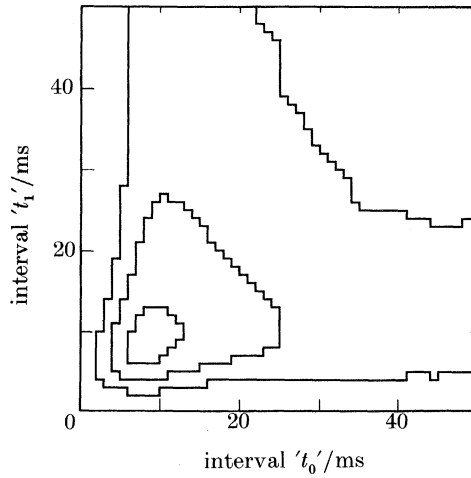

FIGURE 3. Two-dimensional count histogram of interval pairs $'t_{-1}'t_0'$ in $(1 \text{ ms})^2$ bins. The data are represented as a logarithmic contour graph with a distance between successive contour lines of one decade.
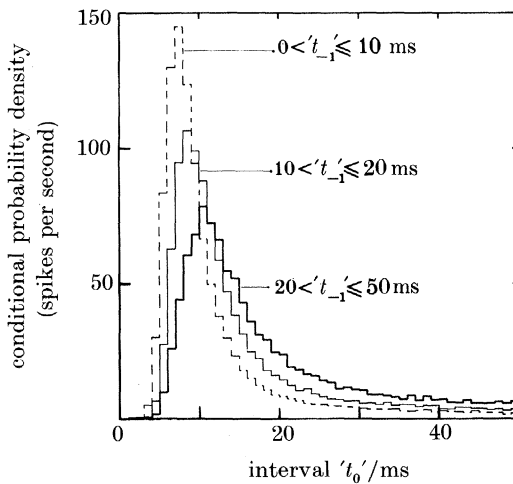

FIGURE 4. Conditional interval histograms $P('t_0' \,|\, 't_{-1}')$ for three different regions of $'t_{-1}'$. We see that, on average, short intervals are followed by short intervals, and long by long.

(b) *Examples of RCEs for different response conditions*

Representations of response-conditional ensembles (RCEs) for a selection of different responses are provided by figure 5 *a–h*. Two simple conditions: a single spike [0'] and a 50 ms period of non-firing [$^-$50$^-$], are depicted in figure 5 *a* and *b*. The mean stimulus waveform conditional on a single spike is a smooth function of time peaking 25 ms before the spike occurs. Its peak height is 16° s^{-1}. The multiple-line plot represents the covariance matrix, scaled by the *a priori* stimulus variance. The elements on the top-left to bottom-right diagonal are set to zero to provide a reference. In reality these elements are almost equal to the value of the nearest off-diagonal elements plus one. Positive (negative) values are represented by deviations in top-left (bottom-right) direction. The figure shows that the off-diagonal elements have a region of negative covariance, centred at about 35 ms before the spike occurs. The fact that the off-diagonal values are negative means that the waveforms that constitute the RCE are constrained to deviations from the mean that have less low-frequency power than waveforms from the *a priori* ensemble E_0. To put it in probabilistic terms, if a waveform deviates from the mean in a certain sub-region close to $\tau = -35$ ms, the probability associated with the waveform in the conditional ensemble decreases. Because the off-diagonal elements are negative this decrease can be partly compensated by a neighbouring deviation of reversed sign. Loosely speaking, this means that the H1 neuron is not sensitive to fluctuations in the stimulus with a period smaller than the width of the region of negative covariance.

The waveform conditional on the long silent interval (figure 5*b*) is negative, meaning that responses of this type are preferentially generated by movements in the neuron's null direction. As in the single-spike case, the off-diagonal covariances are negative in a certain region of the covariance matrix. Here the region of negative covariance is stretched out in time, much as the conditional mean waveform is stretched.

Figure 5*c* shows the RCE for $R = ['30^-]$. The mean velocity exhibits a distinct cross-over from positive to negative velocity. It is interesting to compare its shape to $w[0']$, with the latter waveform shifted to align the spikes for both cases. For times earlier than 40 ms preceding the spike in each case the mean velocity waveforms $w[0']$ and $w['30^-]$ are equal within our experimental error. At about 35 ms before the spike, $w['30^-]$ begins to deviate from $w[0']$, i.e. the extra condition of non-firing during a 30 ms interval starts to have its effect. Correspondingly the covariance matrix $C['30^-]$ has a somewhat larger region of negative covariance than $C[0']$.

Starting from a given open-ended interval we can ask what happens if a spike is subsequently fired. This is shown in figure 5*d* for the case $R = ['30']$, in which the mean waveform shows an additional cross-over to positive velocities. Again, a comparison can be made with the previously shown cases. If the first spike in ['30'] is aligned with the spike in ['30$^-$] the mean waveforms associated with both conditions are approximately equal up to 55 ms before the second spike in ['30']. If $w[0']$ is aligned with the second spike of $w['30']$ we see that, looking backward in time, the waveforms are equal up to the peak in velocity that precedes the spike.
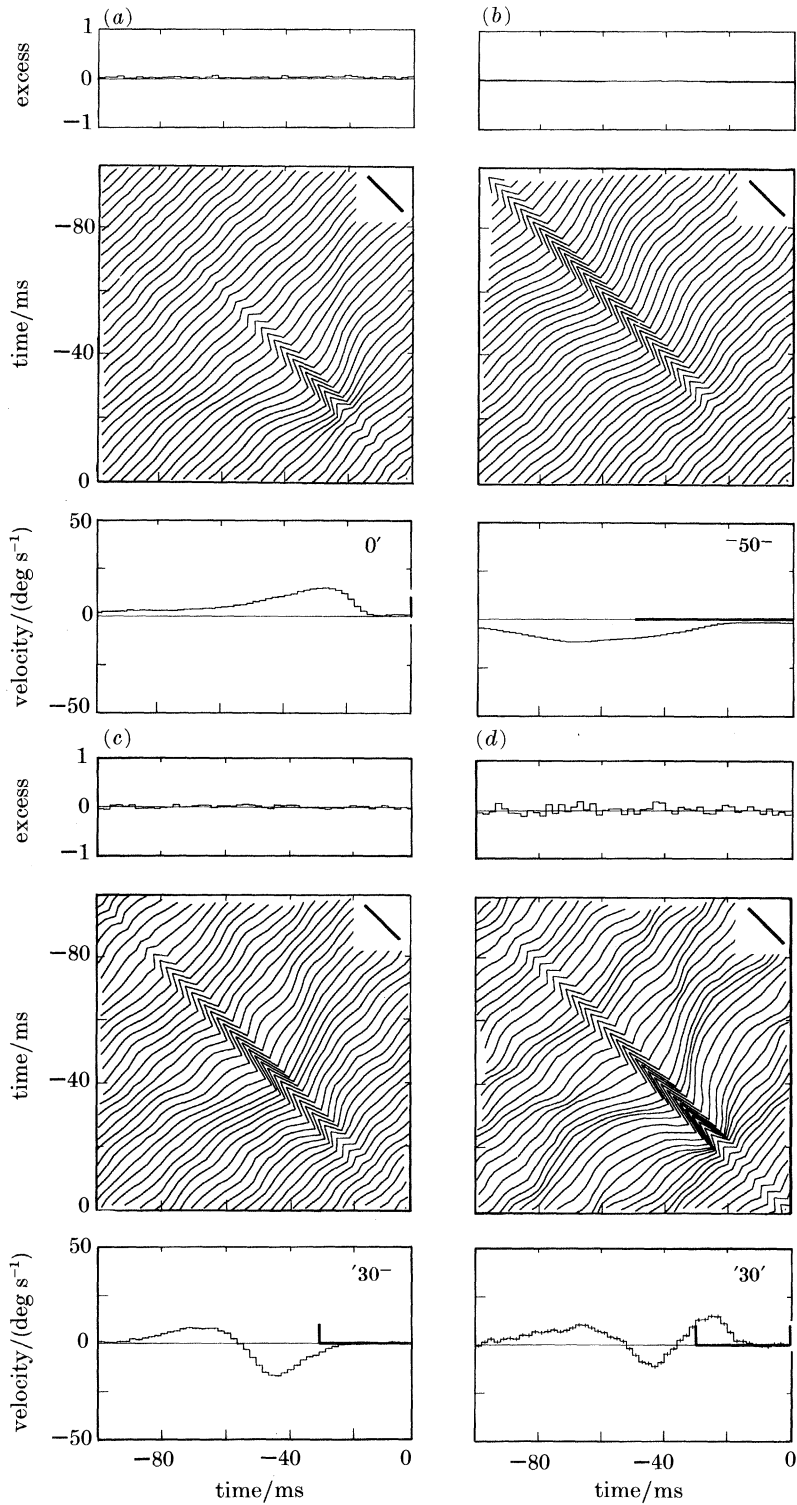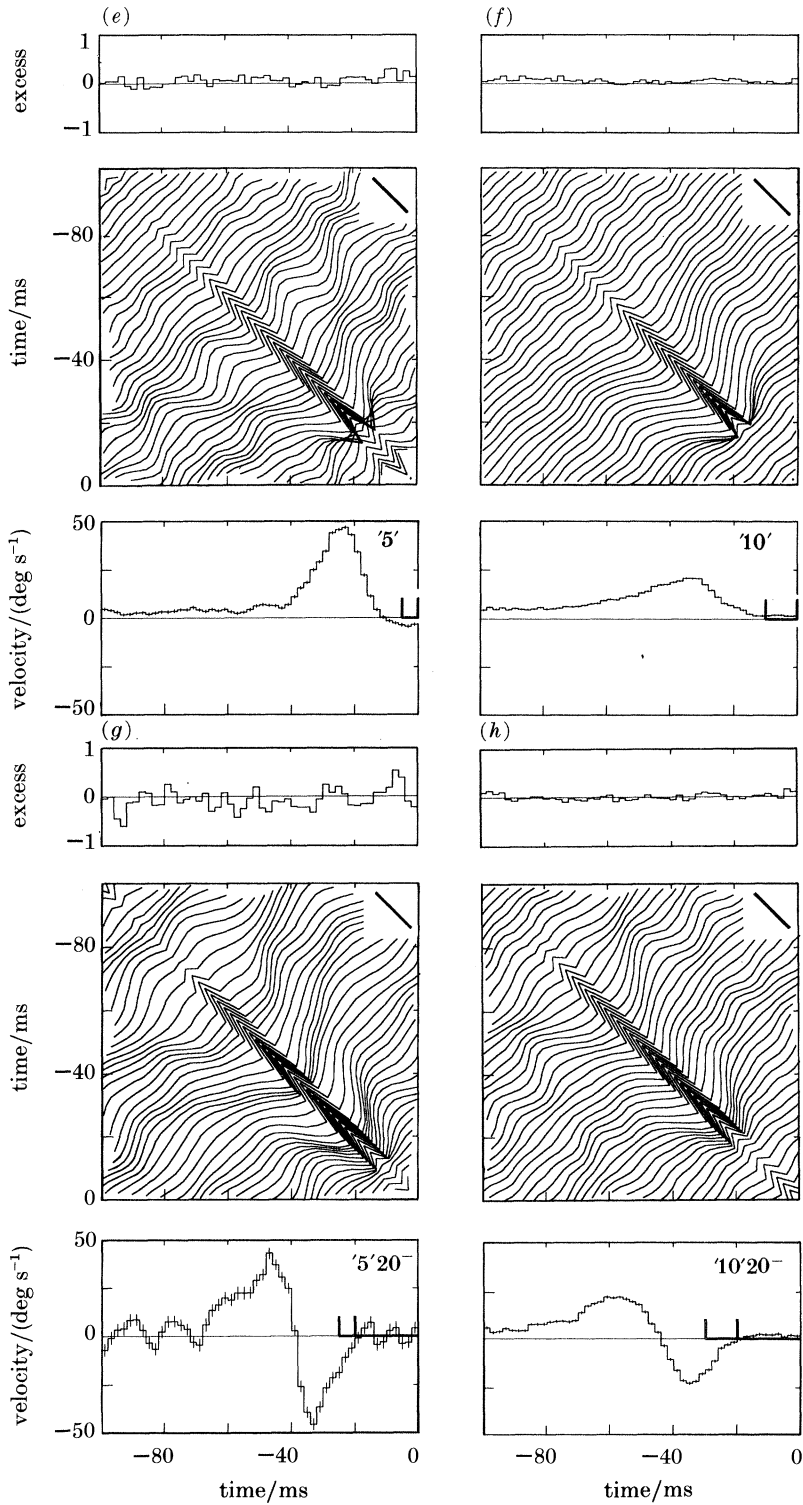
FIGURE 5 *a–d*. For description see page 398.

FIGURE 5*e–h*. For description see page 398.

A comparison of figure 5c and d demonstrates a fundamental characteristic of the neuronal code, i.e. that the occurrence of an action potential at a single point in time contributes to a description of the stimulus over a certain more or less extended *region* of time. Consequently, on the observation of an action potential the best estimate of the stimulus changes discontinuously from one function of time to another.

RCES conditional on the same firing patterns, but with interval lengths of 5 ms and 10 ms respectively are shown in figure 5e and f. Contrary to the 30 ms case, the mean waveforms shown here represent unitary events in the sense that there are no zero-crossings (see also figure 6). The mean waveform $w['10']$ can be approximately described as the sum of two waveforms conditional on a single spike, with one of the two shifted backward in time by 10 ms. The waveform $w['5']$ can certainly not be constructed in this manner: its peak value is close to $50°$ s^{-1}, i.e. three times as high as the peak of $w[0']$. This implies that the stimulus is
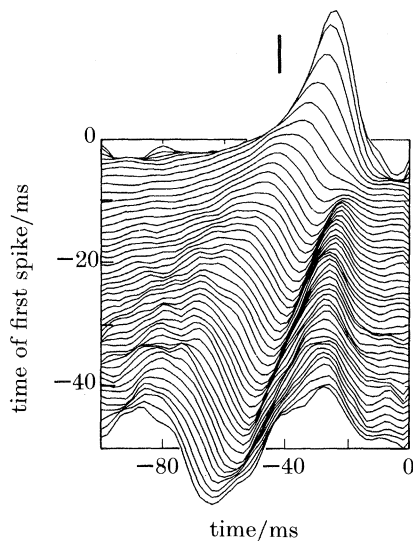

FIGURE 6. Compilation of conditional mean waveforms for different lengths of closed intervals $'t_0'$. The time of occurrence of the first spike relative to the second, i.e. $-t_0$, is given by the ordinate. The smallest value of t_0 shown is 4 ms. The calibration bar on top represents 10 deg. s^{-1} on the velocity scale.

FIGURE 5 (a–h). Response-conditional ensembles (RCES) for a selection of eight different response categories. For each category the RCE is represented by the conditional average waveform (bottom) and the covariance (middle); the diagonal excess (top) is shown to illustrate that the Gaussian approximation is reasonable. The abscissae represent time with respect to t_{obs} (the last point of the response). Covariance matrices are represented by consecutive sections normal to the top-left bottom-right diagonal. The elements on this diagonal are set to zero to provide a reference. The calibration bars in the top-right corners have a length equal to 0.05. Positive (negative) values are in the top-left (bottom-right) direction.

The error bars on the average waveforms (in some of the cases too small to be visible) represent standard errors of the mean in each individual time bin to provide an indication of how accurately the waveforms are determined by the experiment.

encoded in a distinctly nonlinear way. A further aspect of this nonlinearity is that in the negative region of the covariance matrices the values are more negative than for the previous cases, which means that the contribution of the covariance to the information content is larger. Also, the border of this region tends to shift somewhat toward t_{obs}. The covariance matrix for the 5 ms case appears to have considerable structure, but most of this is due to statistical errors associated with the finiteness of our experiment.

It is interesting to note that 10 ms is very close to the most probable interval (compare figure 2), and it is only here that we observe simple linear superposition of the single spike waveforms. Evidently the degree of nonlinearity in stimulus encoding is larger with shorter intervals, which corresponds to the neuron deviating further from its most probable firing pattern. This result is echoed in the interval-dependence of the information content, as discussed below.

Finally, figure 5 g, h shows RCEs in which the previous conditions are extended by an extra empty 20 ms interval. As before, the tails of $w['t_{-1}'t_0{}^-]$ agree reasonably well with the shape of $w['t_{-1}']$ if the latter is shifted by t_0 ms. The more recent history of both conditional waveforms is characterized by a crossover to negative velocities. A striking feature of these latter parts of the waveforms is that the crossover has an amplitude roughly equal to that of the preceding positive peak. Presumably this is related to the correlation among successive intervals demonstrated in figure 4. The interpretation is that if a short interval occurs, the next interval is preferentially shorter than average. If this preference is due to a relatively long correlation time in the signal feeding the spike-generator, the stimulus must be more strongly negative to compensate for this and hence generate a long second interval.

A compilation of conditional waveforms for $R = ['t_0']$ is shown in figure 6. The functions plotted represent waveforms as a function of time along the abscissa, parametrized by the interval length given by the ordinate. It is clear that for long intervals we can distinguish three different phases in the average velocity waveform: two positive peaks occurring 25 ms before each of the two spikes, and a trough in between. As the interval becomes shorter than 10–15 ms, these separate phases merge into a single peak that becomes very high for the shortest intervals. Interestingly, this transition occurs for intervals with a length of the order of the photoreceptor integration time (13 ms) (see de Ruyter van Steveninck (1986)). A functional interpretation is that, loosely speaking, structure in stimulus events on a timescale below the photoreceptor integration time cannot be detected. However, the H1 neuron can generate intervals of shorter duration. These intervals may therefore be used to encode higher stimulus amplitudes.

(c) *Information content of single intervals and rate of information transfer*

The information content of the RCEs conditional on $['t_0{}^-]$ and $['t_0']$ as a function of t_0 is given by figure 7. The values were computed from (5), approximating the determinant as described in §4(*d*) above.

For t_0 close to zero with $R = ['t_0{}^-]$ (heavy line) we obtain the information associated with a single spike, which amounts to 0.36 bits. If the empty interval becomes longer, the information associated with R first decreases, reaching a

minimum of 0.23 bits for $t_0 = 11$ ms, and then rises again when the interval becomes longer. The rising segment has a slope of 17 bits s^{-1} which represents the asymptotic rate of information transfer associated with the *absence* of spikes. The dip in the information content occurring at $t_0 = 11$ ms may seem paradoxical at first sight. After all, the occurrence of a longer open interval is necessarily preceded by the occurrence of a shorter one, and additional information can be provided during the extra stretch of time. Note, however, that this need not be true in general. A simple counter-example is one in which the message provided by the additional length is in conflict with what was received earlier. This increases the uncertainty, and thus decreases the amount of information conveyed.
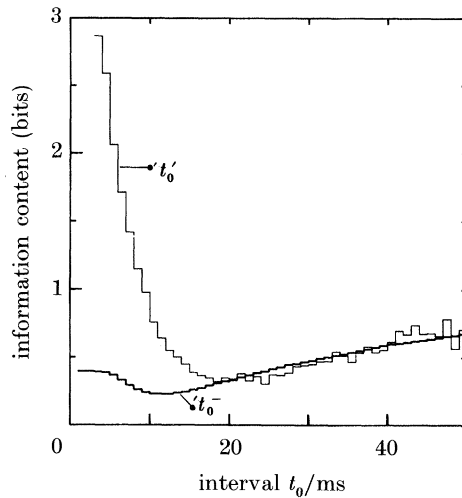

FIGURE 7. Information content of open intervals (heavy line) and closed intervals (thin line) as a function of interval length.

The information carried by short closed intervals considerably exceeds that of the corresponding open intervals. For intervals shorter than 10 ms the full information content is greater than the sum of its parts: two spikes considered independently carry 0.72 bits, and the stretch of empty interval between the spikes contributes 0.17 bits at most (which is the value for a 10 ms interval computed from the 17 bits s^{-1} asymptotic rate mentioned above), whereas short closed intervals carry more than 1 bit. This demonstrates that an important part of the information is carried by *interactions* between spikes. The importance of these short intervals in the total transfer of information can be appreciated by observing that intervals shorter than 8.5 ms, just 16% of the total population of intervals, are responsible for half the average rate of information transfer as computed below.

The information contents of open and closed intervals are equal for interval lengths exceeding 17 ms. Given that a long empty interval has occurred, the occurrence of an extra spike after 17 ms does not add to the information content although it codes for a different event in stimulus space (cf. figure 5c, d). In other

words, the occurrence of the second spike codes for a different average event but the precision of the estimate is the same as that of the empty interval. One reasonably expects that if the intervals become longer the generation of the second spike becomes independent of the time since the previous spike was fired, so that no information is contained in the cross-interaction of the two spikes. This seems to be the case for intervals exceeding 17 ms (see also figure 11).

The intervals in the tail of the interval distribution (see figure 2) are distributed exponentially with a characteristic time of $\tau = 35$ ms. The self information (Fano 1961) of each set of interval lengths lying within a small time window dt is equal to the logarithm of the probability. For the exponential distribution this is just $I = [(t/\tau - \ln (dt/\tau)]/\ln 2$. The time derivative of this, $dI/dt = (\tau \ln 2)^{-1}$, measures the increase of information content with increasing interval length, amounting to 41 bits s^{-1} here. Because of the properties of the exponential distribution this result also holds for open intervals. This rate is more than twice as large as the observed value of 17 bits s^{-1}, which means that there is a significant stimulus-independent component in the generation of these long intervals. Intervals longer than 17 ms contribute only 10 % of the total information transfer rate.

We estimate the information rate carried by spikes in isolation by multiplying the information contained in a single spike by the spike rate. The result is 14 bits s^{-1}. Neglecting correlations in the coding among successive intervals we can compute the average rate of information transfer based on single intervals under the stimulus conditions used in the experiment. The rate of information transfer associated with a certain interval is simply the information content of that interval divided by its length. Averaging this quantity over all intervals, and weighting by their probability of occurrence, we obtain a value of 87 bits s^{-1}, substantially more than the value found for single spikes. Note that the information rate associated with the different intervals varies drastically, and that it can be quite high; a closed interval of 5 ms conveys its information at a rate of 500 bits s^{-1}. In principle the rate of information transfer could be computed along the same lines on the basis of double intervals, but there are too few occurrences, especially of short interval pairs, to allow a reasonably accurate estimate.

(d) Eigenvalues and eigenvectors of the covariance matrix

As noted earlier, the ratio of the determinant of the conditional covariance matrix to that of the *a priori* covariance matrix describes a contraction in the space of stimulus waveforms associated with the response. In the experiment, the waveforms that have given rise to a particular response belong to an ensemble governed by a multi-dimensional Gaussian centred on the mean waveform which is narrower in the space of stimulus waveforms than the *a priori* distribution. A convenient way of representing such a distribution is by a hypersurface of constant probability density. For a Gaussian this is a multi-dimensional ellipsoid, the principal axes of which are given by the eigenvectors of the covariance matrix. We will, for convenience, scale the covariance matrix to the *a priori* stimulus variance, so that the linear contraction in stimulus space along the principal axes is given simply by the square roots of the corresponding eigenvalues. Because the distribution in stimulus space associated with a certain response is narrower along

certain eigenvectors of C_R, the remaining part of the *a priori* stimulus ensemble must be taken up by distributions associated with all other response conditions. Thus the eigenvectors of C_R should point from the mean waveform w_R in the direction of mean waveforms $w_{R'}$ conditional on other responses. This is demonstrated below for some simple cases.

Figure 8 shows the three eigenvectors with the lowest eigenvalues for a covariance matrix conditional on a single spike. The (scaled) eigenvalues are 0.69, 0.89 and 0.90 respectively. Of the three eigenvectors, only the one with the lowest eigenvalue seems to have a clear structure, and this clearly corresponds to the largest reduction in variance from the *a priori* value of 1. Thus this particular neural event provides us with reliable information only about stimulus variations along this one particular direction in the space of all possible stimuli. If the stimulus deviates from the mean waveform $w['0']$ in the direction of this eigenvector, the RCE tells us that that this deviant waveform has a lower probability of generating a spike at t_{obs}. On the other hand, if the deviations are orthogonal to this eigenvector they make little difference to the probability of a spike at t_{obs}. In this sense observation of a single spike provides us with substantial information about certain specific temporal features of the stimulus waveform, but essentially no information about other features.
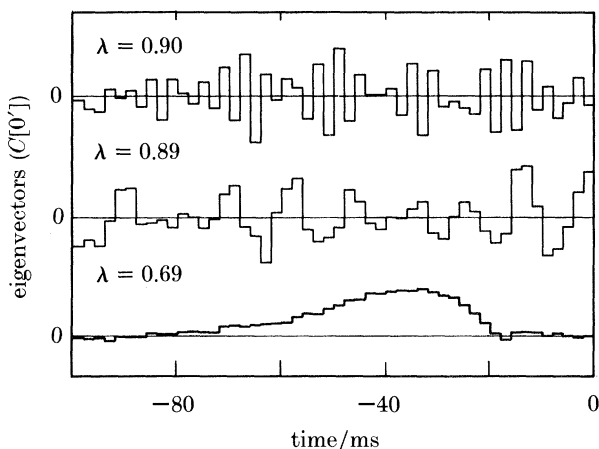

FIGURE 8. Eigenvectors corresponding to the three lowest eigenvalues of the covariance matrix conditional on a single spike.

We can try to understand the encoding of particular 'stimulus directions' in more conventional physiological terms. For negative-going deviations from the mean waveform this is quite easy; if a stimulus waveform occurs with a lower amplitude than the mean conditional waveform, the stimulus is simply less effective. For positive-going deviations the interpretation is slightly more complicated: if the stimulus waveform has a higher amplitude than $w[0']$, the probability of generating spikes goes up. Therefore, the probability of firing a spike shortly *before* t_{obs} also increases. Because of the neuron's refractoriness, however, the

probability of generating an additional spike at t_{obs} drops again, and on the average this may overcompensate the effect of the higher stimulus amplitude. In this connection it is probably significant that the peak of the eigenvector occurs 35 ms before t_{obs}, i.e. 10 ms before the peak in w_R (see figure 5a).

When the response is extended to a spike followed by an open interval, an interesting transition occurs. For open intervals shorter than 5 ms the picture is basically the same as for the single spike case, with the mean waveform and the eigenvector shifted in time. But if the empty interval becomes longer than 6 ms, a second eigenvector with an associated lower eigenvalue develops. This transition can be seen from figure 9, where the lowest two eigenvalues for $R = ['t_0^-]$ are shown as a function of t_0.
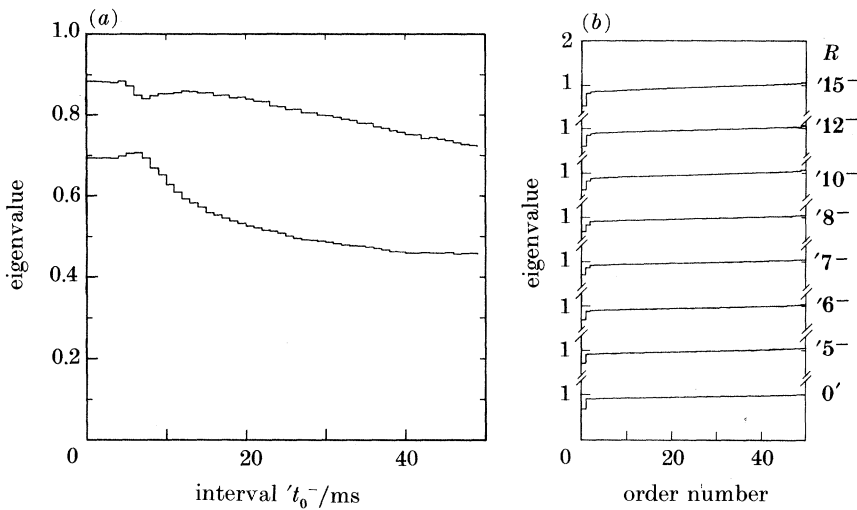

FIGURE 9. (a) The two lowest eigenvalues of the covariance matrix conditional on open intervals of varying length. (b) Complete eigenvalue spectra conditional on open intervals for a selection of interval lengths.

As an example, figure 10a shows the shape of the three eigenvectors of $C['10^-]$ with the lowest eigenvalues. Here we find two eigenvectors with a meaningful structure. The one with the lowest eigenvalue (0.66) has roughly the same shape as the one found for the single spike. The next eigenvector (eigenvalue 0.86) has a different structure, exhibiting a zero-crossing 38 ms before t_{obs}. Now, various combinations of these two eigenvectors map on to different responses. This is shown in figure 10b–d for three classes of response. In these figures, the contributions of the eigenvectors to the differences in the mean conditional waveforms are found by a least-squares fit, or equivalently as the inner products of the normalized eigenvectors with the velocity difference measured in units of the *a priori* stimulus standard deviation. Figure 10b shows the difference waveform $(w['5'5^-]-w['10^-])$ in a heavy line, together with the fit. The contributions of the two eigenvectors are 1.04 (lowest eigenvalue) and 1.45 respectively. For the case $(w['10']-w['10^-])$ the weights are 0.64 and 0.36, and for the case $(w['16^-]-$
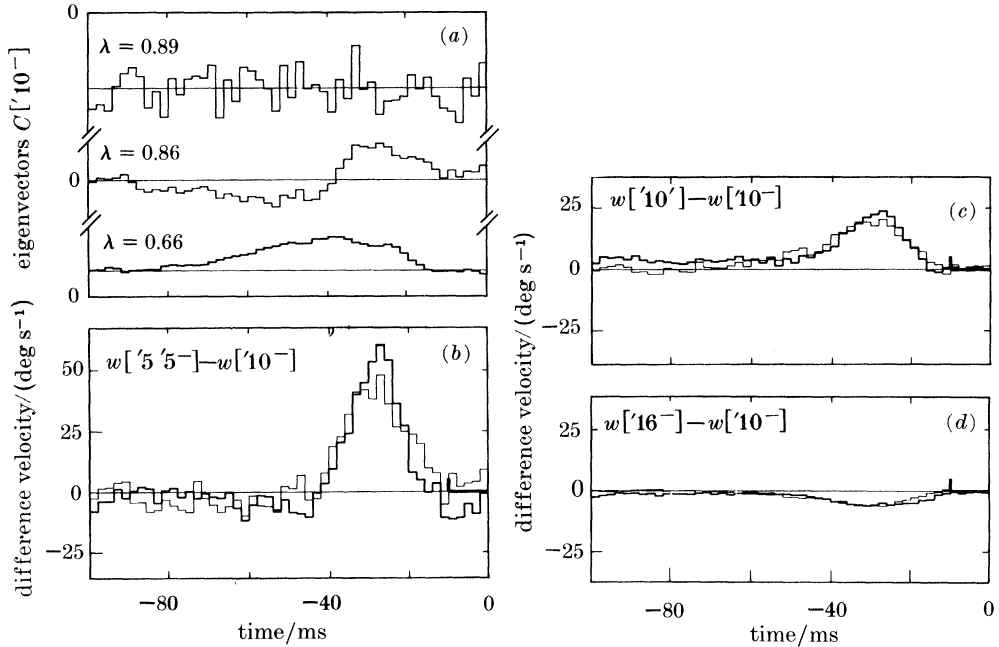
FIGURE 10. (a) Eigenvectors corresponding to the three lowest eigenvalues of the covariance matrix conditional on a 10 ms open interval $['10^-]$. (b–d) Heavy lines are the differences among the mean waveforms $w['5'5^-]-w['10^-]$, $w['10']-w['10^-]$, and $w['16^-]-w['10^-]$, respectively. Thin lines are the fits to these difference waveforms using the first two eigenvectors in (a).

$w['10^-])$ these are -0.26 and -0.12 respectively. In each case the fit is very good, so that these two eigenvectors indeed provide a good set of 'coordinates' for describing the changes in stimulus waveform encoded by different spike sequences.

(e) Discriminability of waveforms conditional on different closed intervals

Within the framework of our interpretation, an interesting question concerns the precision with which a hypothetical observer must make his observations – how accurately must he time the neuronal signal – to obtain the maximum possible information. Clearly if two RCEs corresponding to different inter-spike intervals, for example, are essentially indistinguishable, the observer loses very little by lumping these two intervals in one bin. We can measure the distinguishability of two RCEs R_1 and R_2 by using the discriminability parameter d' familiar from psychophysics (Green & Swets 1966; see also Van Trees 1967):

$$d'(R_1, R_2) = [(w_{R_1} - w_{R_2})^T C_{R_1}^{-1} \cdot (w_{R_1} - w_{R_2})]^{\frac{1}{2}}. \tag{12}$$

Quantitatively, if we were to be presented with the event R_1 or the event R_2, d' is related to the probability that we could distinguish these events by looking at the stimuli that gave rise to them. If it is not possible to make this distinction then d' is near zero and from an informational point of view we might as well consider R_1 and R_2 to be the same event. If the distinction is very easy then d' is large; the crossover defining 'reliable distinction' is conventionally taken as $d' = 1$, and in

figure 11 we show the interpolated contour line of the values of $(t_0[R_2] - t_0[R_1])$ with $d' = 1$, as a function of $t_0[R_1]$. The results are for closed single intervals, i.e. both responses are of the form $['t_0']$. The figure shows that the observation of a 5 ms interval corresponds to a stimulus event that can be discriminated from an event associated with a 7.5 ms interval, $t_0[R_2] - t_0[R_1] = 2.5$ ms. As $t_0[R_1]$ becomes longer, the required timing becomes less precise: the mean waveform for $t_0[R_1] = 10$ ms can be discriminated only from mean waveforms corresponding to intervals that are longer by 8 ms. For large values of $t_0[R_1]$, d' reaches unity for intervals $t_0[R_2]$ that differ from $t_0[R_1]$ by about 17 ms.
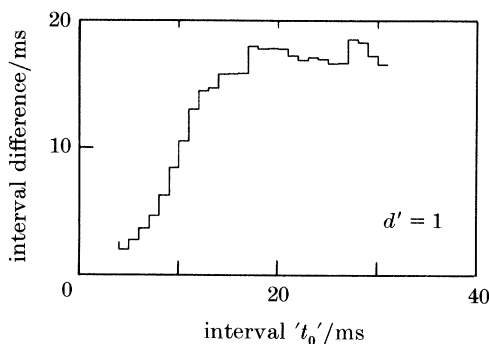

FIGURE 11. Difference interval as a function of the interval length for which the corresponding mean waveforms can be discriminated with $d' = 1$. See text for details.

The conclusion from these results is that, once a spike is fired, the precision with which the observer must remember its position as time proceeds should be high shortly after the spike has occurred. As time goes on and no subsequent spike is fired, his time resolution can be less and less precise. After about 17 ms, the mutual timing of the two spikes is no longer important, and they can be considered as representing independent events. In this event the observer may forget that there was a spike, the only salient feature of the history being that no spike was fired during the past 17 ms. Because the precision required for optimal information extraction is not terribly great, we may say that the neural code is substantially robust to timing errors, even in a single neuron.

(f) Reconstruction of the stimulus waveform based on the neuronal response

We now turn to the real-time estimation problem referred to in the Introduction. Our general task is to reconstruct the stimulus waveform based on information carried by the spike train alone. What we know from the experiments described above concerns the information carried by short spike sequences, that is by generalized neural events in isolation, so we shall have to make some approximation in order to combine information from successive events. As discussed above, the only natural assumption not requiring us to fix some arbitrary parameters is the assumption that, given a fixed input to the cell, these events are generated independently. Again we emphasize that, although approximate,

this reconstruction strategy is defined completely by the structure of the RCEs determined above.

The response events treated here are single spikes, closed intervals, and closed double intervals. These cases are distinguished by the reconstruction depth: the first has a reconstruction depth of one spike, the second of two, and the third of three spikes, with the latter taken in a non-overlapping way, i.e. the last spike of the previous event coincides with the first spike of the present event, and so on. Because during the second half of each stimulus period the first half was repeated with a change of sign, we symmetrize the reconstruction by combining the responses of the neuron that occur exactly half a stimulus period apart. This means, in effect, that we reconstruct the stimulus by observing the responses from two independent H1 neurons with opposite directional selectivities. Translated to the fly, this procedure amounts to combining the responses of the two H1 neurons on either side of the head, provided that the velocities on the left and right side are the same, and that on both sides the stimulus intensity distributions have the same statistics. These conditions would occur in natural situations when the fly makes a turn in a static, statistically homogeneous environment.

Figure 12 presents stimulus reconstructions of 2 s of the same trace of the experiment for the three reconstruction depths, computed according to (11) (see §2c). The traces also show the stimulus waveform, and the spike sequences are depicted at the bottom. The number of spikes fired during this time window is 131 for one direction of movement and 48 for the reverse direction. It is clear that the reconstruction based on single spikes has much less structure than the other two. Another noteworthy property is that in the region of high spike activity the reconstruction generally overestimates the stimulus. The overestimate becomes less pronounced, however, when the reconstruction depth increases to three spikes. Therefore, the effect is most likely due to serial correlation in the spike train, the explanation being that events that occur at larger separation in time are less correlated. In other words, the assumption of statistical independence becomes a noticeably better approximation as the reconstruction depth increases. To be fair, it is not clear that a reconstruction depth of three spikes is sufficient to validate the independence hypothesis convincingly, but the quality of the reconstruction at this depth is already quite good. To our knowledge this is the first case in which a quantitative attempt has been made literally to read the neural code; although our reading is certainly not optimal, as there are obvious systematic errors in the reconstruction, we find the results of figure 12 encouraging.

A rough indication of the quality of the fit is obtained by computing the coherence function (Bendat & Piersol 1971) between actual stimulus and reconstruction, and from that the average information transfer. If we think of the true signal and the reconstructed signal as a pair of random (noise-like) waveforms, the coherence function is essentially the frequency-dependent correlation coefficient between these two random processes. As such, it measures any *linear* coupling of signal and reconstructed waveforms, but does not take account of nonlinear couplings. This implies that if our reconstruction procedure systematically distorts the waveform – e.g. by overshooting – these systematic effects will degrade the coherence function in the same way as would random noise in the reconstruction.
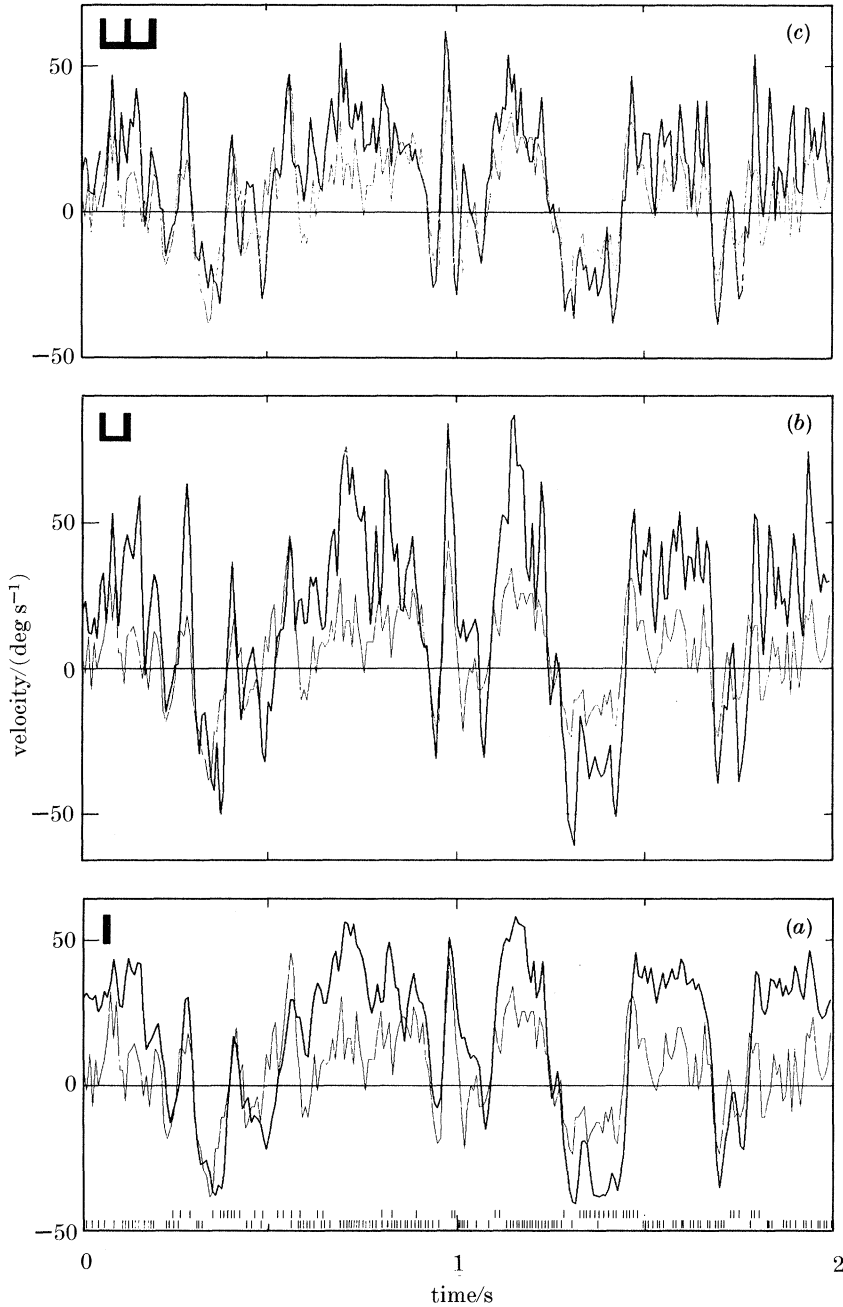
FIGURE 12. Reconstructions of the stimulus waveform (heavy lines) together with the actual waveform (thin lines) for a 2 s time window. The reconstructions are made on the basis of combined neural responses, symbolized by the small vertical bars at the bottom of (*a*), during two opposite phases of the stimulus. In the reconstruction we assume that (*a*) single spikes, (*b*) spike pairs or (*c*) spike triplets are generated independently, as described in the text.

The results of the coherence calculations, derived from the 32 seconds of signal of which figure 12 gives an example, are 16 bits s^{-1} for the single-spike reconstruction and 27 and 30 bits s^{-1} respectively for single- and double-interval reconstructions. From the results presented in §5c we would expect a maximum of 28 and 174 bits s^{-1} respectively for the case of single spikes and single intervals. Comparing this with the values obtained from the coherence function we see that our reconstruction procedure may still be far from optimal. More specifically, this comparison suggests that our procedure does indeed systematically distort the waveform, and there is thus some hope for systematic improvements in our reading of the code.

6. Discussion

(a) Review of the method

The method presented here offers the possibility of a systematic study of neural coding and provides insight into the question of how the signal from a given neuron is to be interpreted in real time. An important feature is that the method allows a definition of 'the code' that is rigorous and model-independent, but at the same time intuitively appealing and experimentally accessible. In addition, the information transferred by the neuron can be directly quantified in terms of Shannon's (1948) information measures.

It should be emphasized that, although we quantify the information transfer in short spike sequences, we do not have an estimate of the ultimate information capacity of the neuron; we have made no attempt to choose stimulus conditions that would maximize the transmitted information. On the contrary, we have used relatively low-contrast stimuli in the hope of making our Gaussian approximation (equation (3)) more reasonable.

It should be clear that our methods are generalizable to the study of simultaneous activity in several neurons. Conceptually this would involve extending our definitions of response categories to include combinations of firing patterns from more than one cell. In practice this approach may quickly become limited by the available recording times, because as we define more and more response categories we require more and more events to obtain good estimates of each response-conditional ensemble. It should certainly be possible to give a complete description based on combinations of single intervals from two neurons, but if a signal is shared among a very large number of cells it may be necessary to give a more approximate description of the encoding.

Based on the experimental description of coding by short sequences of spikes the problem of literally reading the neural code – reconstructing the stimulus waveform from the spike-train data alone – can be solved to a reasonable degree using the approximation that successive spike sequences are generated independently. It is clear that if by 'spike sequence' we mean a single spike, then our approximation is seriously limiting the quality of the reconstruction. By extending the depth of reconstruction to include spike triplets we substantially improve the comparison between reconstructed and actual waveforms. Nonetheless, visual inspection of figure 12 (and longer intervals as well) strongly suggests that even

our deepest reconstruction contains significant systematic errors, such as 'over-shoots' in regions of high firing rate, and our discussion of the coherence function also points towards the possibility of systematic errors in the reconstruction. If these impressions are confirmed by more detailed analysis we would have to conclude that our independence approximation limits the reconstruction quality even when the independent units are taken as spike triplets, or equivalently double intervals. The complete solution to the problem of reading the code requires finding the optimal way of combining the messages conveyed by successive spike sequences, and most importantly of showing how the parameters of this optimal combination can be determined experimentally. We hope to return to this issue in a subsequent publication.

(b) *Coding and refractoriness*

If we assume that neurons can mark the arrival of spikes with some fixed temporal precision, then the maximum rate of information transmission per spike would be achieved for neural firing patterns that most closely approach a Poisson process (Marko 1962); in this limit the information content of a given inter-spike interval would be proportional to its length (Färber 1968). For the cell studied here it is clear that the latter property obtains only at long intervals, and the interval distribution itself deviates strongly from the Poisson (exponential) form at short intervals. Most importantly, as shown in figure 7, the short intervals carry relatively more information, so that messages with high information content are transmitted very rapidly. These results indicate that, in this cell, coding has not been optimized for the highest possible average transmission rate. Rather, the emphasis seems to be on conveying relatively rare but important messages, in this case sudden large movements, as rapidly as possible.

One of the likely physiological mechanisms underlying this apparent coding strategy is the increase in firing threshold after a spike has been fired. In this way the phenomenon of relative refractoriness may play an important role in neural coding. If this interpretation is correct then similar coding strategies may be found in a wider variety of neurons, especially in those that fire irregularly. We also note that the correlations introduced into the spike train by refractoriness would aid in the detection and correction of errors (Brillouin 1962) associated with, for example, the failure of a synapse or interneuron to transmit a spike; we have no idea how significant such errors might be, but the concept of an error-correcting neural code remains intriguing.

(c) *Extrapolation to natural conditions*

An important caveat regarding our approach to the coding problem is that all our results are specific to the particular conditions of our experiment. We believe that this must be a general limitation for any experimental probe of neural coding. The hope is of course that some qualitative features of the code may be uncovered that may extrapolate to more general, and in particular 'natural', stimulus conditions.

The basic problem in extrapolating from a single set of experiments is that the response-conditional ensembles $P[v(t) | R]$ depend on the *a priori* stimulus distri-

bution $P_0[v(t)]$, as well as on other parameters of the sensory environment that we have not chosen to include in our definition of 'the stimulus'. If all these parameters are held fixed, then reading the neural code requires learning a limited set of probability distributions, essentially those we have measured or defined in our experiment. The fact that reliable estimation in general requires such detailed statistical knowledge of the possible stimuli and their representation in the nervous system is a special case of the often-discussed problem of prior probabilities (Jaynes 1968; Berger 1985). It is interesting that our difficulties in obtaining experimental estimates of the response-conditional ensembles provide at least an approximate view of the difficulties an organism would encounter if it had to learn to read its own neural code.

The fact that an organism does not live in a world with fixed stimulus conditions makes the coding problem much more difficult. In our specific case, the fly encounters in nature a wide variety of velocity flow fields in a still wider variety of images. By choosing ensembles of signals – random velocities against time and random images – we have made an important step towards the natural sensory environment, but in reality even these ensembles are subject to substantial changes; a clear example is the change in image statistics upon flying from an open field into a dense forest, or from clear air into a fog bank.

It is well known that sensory systems adapt to such changes in stimulus parameters; the photoreceptors themselves adapt to changes in light level, and there is movement-specific adaptation in the response of H1 (Maddess & Laughlin 1985; de Ruyter van Steveninck et al. 1986). Adaptation has generally been viewed as a means of preserving sensitivity to small changes in a steady background signal. With the approach taken here it is clear that adaptation is accompanied by a problem of coding ambiguity, where the same spike train will represent different stimuli if the organism finds itself in different states of adaptation. As a first step towards understanding these issues we would like to do the experiments described here under different conditions, using the notion of discriminability among RCEs (cf. §5e), to quantify the ambiguity of coding under different states of adaptation. In fact it may be possible to use such discriminability measures to quantify the 'states of adaptation' themselves, in effect defining adaptation by changes in coding strategy. One interesting possibility is that longer-timescale statistical properties of the firing pattern may contain enough information to resolve ambiguities in the interpretation of short spike sequences, so that our hypothetical observer need not have access to other cells in order to understand the signals conveyed by H1 itself.

(d) Comparison with the theoretical limit

Movement information is present at the level of the retina in the form of correlations among signals in the photoreceptor array, these correlations being induced by the rigid motion of the stimulus pattern across the visual field. Photoreceptor noise tends to destroy these correlations, so that the reliability of the movement signal measured at H1 is ultimately limited by the photoreceptor noise level. We have measured the signal transfer and noise properties of Calliphora photoreceptors under the same conditions of illuminance used in the H1 experiments. The main result (de Ruyter van Steveninck 1986) is that up to roughly

10 Hz the photoreceptor is limited by photon shot noise, whereas above this cutoff frequency the noise level begins to rise; this pattern of frequency dependence is consistent with a model in which the 'excess' noise is contributed largely by fluctuations in the latency of the single-photon response.

The photoreceptor noise data, together with the known geometry of the stimulus and the photoreceptor array, allow us to estimate the total amount of movement information available in the retinal signals (Bialek *et al.* 1986). For the conditions of our experiment (the spatial structure of the pattern, the number (2640) of photoreceptors illuminated, etc.) the result is 298 bits s^{-1}. In comparing this with our estimate of H1's information transfer rate (87 bits s^{-1}; cf §5*c*) we should keep several points in mind. (i) H1 is primarily responsive to movement in only one direction, so presumably it can achieve at best half the calculated maximum rate. (ii) The theoretical maximum is for a system that looks at a rigorously static pattern moving rigidly through small angular displacements. Although this describes our experiment, it is unrealistic to assume that the visual system has made this assumption; because of the need to continuously re-estimate the spatial pattern the true information rate will be reduced. (iii) We have made no effort to find special stimulus conditions under which the apparent efficiency of information transfer would be optimized. Granting these remarks, we would find it quite impressive if H1 approached optimal performance even within an order of magnitude, whereas in fact the approach is within a factor of four.

The hypothesis that sensory and neural systems might approach some fundamental theoretical limits to their performance dates at least from the turn of the century. Elsewhere one of us (Bialek 1987) has tried to bring the evidence on this issue up to date, gathering results from several different sensory modalities. For vision the approach to optimality has recently been re-emphasized by Barlow (1980, 1981). We would like to postpone a complete discussion of the approach to optimality in H1 until we can perform a more detailed and rigorous comparison between theory and experiment. In fact such a comparison is much easier for forced-choice discrimination experiments than for the real-time experiments discussed here. (See, for a preliminary account, de Ruyter van Steveninck (1986).) We cannot resist, however, pointing out certain tentative conclusions.

First, to approach optimality the fly must both use a near-optimal algorithm for the extraction of movement information and it must add very little noise as it carries out this computation. Although this may not be obvious from our derivation of the limiting information capacity (Bialek *et al.* 1986), it turns out that the optimal algorithm is essentially a computation of delayed near-neighbour correlations, as proposed many years ago (Reichardt 1957); in this case a variant of the simplest movement sensor is also the optimal movement sensor. Independently of Reichardt's arguments, the approach of H1 to optimal performance thus allows us to conclude that the neural circuitry that underlies movement computation in the blowfly must be functionally equivalent to a relatively simple spatio-temporal correlator. The evidence that a generalized Reichardt model is in fact applicable to H1 (see, for example, de Ruyter van Steveninck 1986) may therefore also be viewed as confirming the existence of computational elements required for the approach to optimality.

The detailed structure of the optimal computation depends on the *a priori*

stimulus ensemble, so that if the fly is to maintain optimality in a changing sensory environment the parameters of H1's dynamic response must change in a specific manner as the global stimulus parameters are varied. This adaptation of the dynamic response to maintain optimality is much more than a simple 'resetting of the operating point' to maintain sensitivity in the presence of a steady background signal. In fact, adaptation of dynamic response parameters has been reported for the H1 neuron (Maddess & Laughlin 1985; de Ruyter van Steveninck *et al.* 1986), and the observed behaviour is qualitatively in accord with the view of adaptation as a strategy for optimal computation in a dynamic environment. We are currently attempting to put this broader notion of optimal performance to a more quantitative test.

The preceding comments concern the algorithmic structure of the movement computation. We may also draw some tentative conclusions regarding the implementation of these algorithms in neural hardware. Because there is no evidence for significant anatomical redundancy in the pathway from retina to H1 (Strausfeld 1976), we must assume that each element in the movement computation is in fact carrying out its function reliably, adding very little noise to the photoreceptor signal. Specifically, in these and other experiments (de Ruyter van Steveninck 1986) we have evidence for near-optimal computation under conditions where the signal-to-noise ratio for the movement signal is near unity. But this signal is the coherent sum of signals from roughly 2500 photoreceptors, so that the elementary movement signals from each pair of correlated neighbouring photoreceptors are at signal-to-noise ratios of roughly $1/\sqrt{2500}$ (*ca.* 1/50). Each elementary correlator must thus preserve about 5–6 bits of accuracy in its output ($\ln{(50)} = 5.64$) in order that the subsequent averaging among parallel channels be effective in revealing the signal. We believe that this last argument is quite general. Although it is popular to appeal to 'square-root-of-N' as a means of increasing the signal-to-noise ratio, there are many cases, such as movement detection, where non-trivial computations must be carried out *before* the averaging is done. As shown here, in the absence of significant redundancy this may lead to rather stringent requirements on the computational hardware.

To summarize, our comparisons of H1 performance with the theoretical optimum lead to a tentative picture of the blowfly visual system as performing optimal, nearly noiseless computations, with much of the low noise level ascribed not to collective interactions among neurons but rather to the intrinsic precision of individual cells. It seems fair to say that this has not been a common view of the nervous system.

REFERENCES

Abramowitz, M. & Stegun, I. 1965 *Handbook of mathematical functions.* New York: Dover.

Barlow, H. B. 1980 The absolute efficiency of perceptual decisions. *Phil. Trans. R. Soc. Lond.* B **290**, 71–82.

Barlow, H. B. 1981 Critical limiting factors in the design of the eye and visual cortex. *Proc. R. Soc. Lond.* B **212**, 1–34.

Barlow, H. B. & Levick, W. R. 1969 Three factors limiting the reliable detection of light by retinal ganglion cells of the cat. *J. Physiol., Lond.* **200**, 1–24.

Bendat, J. S. & Piersol, A. G. 1971 *Random data: analysis and measurement procedures.* New York: Wiley.

Berger, J. O. 1985 *Statistical decision theory and Bayesian analysis*, 2nd ed. Berlin: Springer-Verlag.

Bialek, W. 1987 Physical limits to sensation and perception. *A. Rev. Biophys. biophys. Chem.* **16**, 455–478.

Bialek, W., de Ruyter van Steveninck, R. R. & Zaagman, W. H. 1986 *Reliability and information capacity of a neural code.* Institute for Theoretical Physics report 86-117, University of California, Santa Barbara.

de Boer, E. 1967 Correlation studies applied to the frequency resolution of the cochlea. *J. aud. Res.* **7**, 209–217.

de Boer, E. & de Jongh, H. R. 1978 On cochlear encoding: potentialities and limitations of the reverse correlation technique. *J. acoust. Soc. Am.* **63**, 115–135.

Brillouin, L. 1962 *Science and information theory*, 2nd ed. New York: Academic Press.

Bullock, T. H. 1970 The reliability of neurons. *J. gen. Physiol.* **55**, 565–584.

Burns, B. D. 1968 *The uncertain nervous system.* London: Edward Arnold.

Eckhorn, R. & Pöpel, B. 1974 Rigorous and extended application of information theory to the afferent visual system of the cat. I. Basic concepts. *Kybernetik* **16**, 191–200.

Eckhorn, R. & Pöpel, B. 1975 Rigorous and extended application of information theory to the afferent visual system of the cat. II. Experimental tests. *Biol. Cybernet.* **17**, 7–17.

Eggermont, J. J., Johannesma, P. I. M. & Aertsen, A. M. H. J. 1983 Reverse-correlation techniques in auditory research. *Q. Rev. Biophys.* **16**, 341–414.

Fano, R. M. 1961 *Transmission of information.* Cambridge, MA: MIT Press.

Färber, G. 1968 Berechnung und Messung des Informationsflusses der Nervenfaser. *Kybernetik* **5**, 17–29.

Fitzhugh, R. 1957 The statistical detection of threshold signals in the retina. *J. gen. Physiol.* **40**, 925–948.

Fox, P. A. (ed.) 1976 *The PORT mathematical subroutine library.* Murray Hall, NJ: Bell Telephone Laboratories Inc.

Green, D. M. & Swets, J. A. 1966 *Signal detection theory and psychophysics.* New York: Wiley.

Grüsser, O. J. 1962 Die Informationskapazität einzelner Nervenzellen für die Signalübermittlung im Zentralnervensystem. *Kybernetik* **1**, 209–211.

Grüsser, O. J., Hellner, K. A. & Grüsser-Cornehls, U. 1962 Die Informationsübertragung im afferenten visuellen System. *Kybernetik* **1**, 175–192.

Hausen, K. 1982 Motion sensitive interneurons in the optomotor system of the fly. II. The horizontal cells: receptive field organization and response characteristics. *Biol. Cybernet.* **46**. 67–79.

Hausen, K. 1984 The lobula complex of the fly: structure, function and significance in visual behavior. In *Photoreception and vision in invertebrates* (ed. M. A. Ali), pp. 523–559. New York and London: Plenum Press.

Hausen, K. & Werhahn, C. 1983 Microsurgical lesion of horizontal cells changes optomotor yaw responses in the blowfly *Calliphora erythrocephala*. *Proc. R. Soc. Lond.* B **219**, 211–216.

Jaynes, E. T. 1968 Prior probabilities. *IEEE Trans. Syst. Science Cybernet.* **4**, 227–241.

Land, M. F. & Collett, T. S. 1974 Chasing behaviour of houseflies (*Fannia canicularis*). A description and analysis. *J. comp. Physiol.* **89**, 331–357.

Maddess, T. & Laughlin, S. B. 1985 Adaptation of the movement-sensitive neuron H1 is generated locally and governed by contrast frequency. *Proc. R. Soc. Lond.* B **225**, 251–275.

Marko, H. 1962 Die Ausnutzbarkeit eines Telegraphiekanals zur Informationsübertragung. *Nachrtech. Z.* **15**, 451–466.

Marmarelis, P. Z. & Marmarelis, V. Z. 1978 *Analysis of physiological systems. The white noise approach*. New York: Plenum Press.

Marsaglia, G., Ananthanarayanan, K. & Paul, N. 1973 *Random number generator packet 'Super-Duper'*. School of Computer Science, McGill University.

McKay, D. M. & McCulloch, W. S. 1952 The limiting information capacity of a neuronal link. *Bull. Math. Biophys.* **14**, 127–135.

von Neumann, J. 1956 Probabilistic logics and the synthesis of reliable organisms from unreliable components. In *Automata studies* (ed. C. E. Shannon & J. McCarthy), pp. 43–98. Princeton, NJ: Princeton University Press.

Perkel, D. H. & Bullock, T. H. 1968 Neural coding. In *Neurosciences Research Program Summaries* (ed. F. O. Schmitt, T. Melnechuk, G. C. Quarton & G. Adelman), vol. 3, pp. 405–527. Cambridge, MA: MIT Press.

Rapaport, A. & Horvath, W. J. 1960 The theoretical channel capacity of a single neuron as determined by various coding systems. *Inf. Control.* **3**, 335–350.

Reichardt, W. 1957 Autokorrelations-Auswertung als Funktionsprinzip des Zentralnervensystems. *Z. Naturf.* **12**b, 448–457.

de Ruyter van Steveninck, R. R. 1986 *Real-time performance of a movement-sensitive neuron in the blowfly visual system*. Academisch Proefschrift, Rijksuniversiteit Groningen.

de Ruyter van Steveninck, R. R., Bialek, W. & Zaagman, W. H. 1985 Vernier movement discrimination with three spikes from one neuron. *Perception* **13**, A47–A48.

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. Cybernet.* **54**, 223–236.

Shannon, C. E. 1948 A mathematical theory of communication. *Bell. Sys. Tech. J.* **27**, 379–424 and 623–657.

Shannon, C. E. 1949 Communication in the presence of noise. *Proc. Inst. Radio Engrs* **37**, 10–21.

Stein, R. B. 1967 The information capacity of nerve cells using a frequency code. *Biophys. J.* **7**, 797–826.

Stein, R. B. & French, A. S. 1970 Models for the transmission of information by nerve cells. In *Excitatory synaptic mechanisms* (ed. P. Anderson & J. K. S. Jansen), pp. 247–257. Oslo University Press.

Stein, R. B., French, A. S. & Holden, A. V. 1972 The frequency response, coherence, and information capacity of two neuronal models. *Biophys. J.* **12**, 295–322.

Strausfeld, N. J. 1976 *Atlas of an insect brain*. Berlin: Springer-Verlag.

van Trees, H. L. 1967 *Detection, estimation, and modulation theory. Part I*. New York: Wiley.

Walløe, L. 1970 On the transmission of information through sensory neurons. *Biophys. J.* **10**, 745–763.

Werner, G. & Mountcastle, V. B. 1965 Neural activity in mechanoreceptive cutaneous afferents: response relations, Weber functions, and information transmission. *J. Neurophysiol.* **28**, 359–397.