

J. B. M. Goense · R. Ratnam

Continuous detection of weak sensory signals in afferent spike trains: the role of anti-correlated interspike intervals in detection performance

Received: 28 February 2003 / Revised: 20 June 2003 / Accepted: 2 July 2003 / Published online: 14 August 2003
© Springer-Verlag 2003

Abstract An important problem in sensory processing is deciding whether fluctuating neural activity encodes a stimulus or is due to variability in baseline activity. Neurons that subservise detection must examine incoming spike trains continuously, and quickly and reliably differentiate signals from baseline activity. Here we demonstrate that a neural integrator can perform continuous signal detection, with performance exceeding that of trial-based procedures, where spike counts in signal- and baseline windows are compared. The procedure was applied to data from electrosensory afferents of weakly electric fish (*Apteronotus leptorhynchus*), where weak perturbations generated by small prey add ~ 1 spike to a baseline of ~ 300 spikes s^{-1} . The hypothetical postsynaptic neuron, modeling an electrosensory lateral line lobe cell, could detect an added spike within 10–15 ms, achieving near ideal detection performance (80–95%) at false alarm rates of 1–2 Hz, while trial-based testing resulted in only 30–35% correct detections at that false alarm rate. The performance improvement was due to anti-correlations in the afferent spike train, which reduced both the amplitude and duration of fluctuations in postsynaptic membrane activity, and so decreased the number of false alarms. Anti-correlations can be exploited to improve detection performance only if there is memory of prior decisions.

Keywords Continuous detection · Electroreception · Interspike interval correlations · Neural coding · Sequential hypothesis testing

Abbreviations B binomial · CV coefficient of variation · EOD electric organ discharge · ELL electrosensory lateral line lobe · EPSP excitatory postsynaptic potential · ISI interspike interval · M_0 Markov order zero · M_1 Markov order one · N noise · OC operating characteristic · PDF probability density function · ROC receiver operating characteristic · S signal · SNR signal-to-noise ratio · S + N signal in noise

Introduction

Animals encounter a continuous stream of sensory information, which consists of both relevant and irrelevant information. From this stream of information they have to detect, extract and attend to the relevant information quickly and reliably. Here we focus on the detection task. This task is challenging because the stimulus must be detected in the presence of ongoing background activity, while neither the onset nor the properties of the stimulus are known. The problem is multiplied when considering the neural representation of stimuli. Weak but important signals tend to be obscured in a mix of irrelevant information from both the environment and intrinsic neural noise. The neural processing of sensory data may involve circuits that incorporate some form of online hypothesis testing to detect signals of relevance on a continuous basis. Physiological mechanisms for the continuous testing of hypotheses at the single neuron level are currently unknown, although there is some evidence linking single neuron activity to decision making (Kim and Shadlen 1999; Gold and Shadlen 2001). Here, we address the issue of signal detection in the weakly electric fish *Apteronotus leptorhynchus* (brown ghost knife fish). Based on the afferent input to the electrosensory lateral line lobe (ELL), we examine how a continuous detector located in the ELL may detect small changes in the afferent spike train.

J. B. M. Goense (✉) · R. Ratnam
Beckman Institute, University of Illinois at Urbana-Champaign,
IL 61801, USA
E-mail: goense@uiuc.edu
Tel.: +1-217-3337071
Fax: +1-217-2448371

J. B. M. Goense
Center for Biophysics and Computational Biology,
University of Illinois at Urbana-Champaign, IL 61801, USA

The electric fish's self-generated electric organ discharge (EOD) is modulated by objects that differ in electrical impedance from the surrounding water, which provides sensory cues that allow the fish to navigate and hunt in the dark (review: Bullock and Heiligenberg 1986). Probability coding, or P-type electrosensory afferents, that maintain a baseline activity of ~ 300 spikes s^{-1} in the presence of undisturbed EOD, respond to the modulations by increasing or decreasing their probability of firing. Based on studies of prey capture behavior, Nelson and MacIver (1999) estimated that at the time of detecting a prey, the transient modulation in P-type afferent activity is about 1%. This adds about one extra spike to the total of 60 spikes expected in the time taken for the prey to sweep across the afferent's receptive field (~ 200 ms). Such a weak signal can be difficult to detect given the intrinsic fluctuations in the spike count (Ratnam and Nelson 2000). Using a trial-based ideal observer scheme, Ratnam and Nelson (2000) showed that if only spike counts were used, then 2–3 extra spikes could be detected with 90% probability in a 100-ms window. While the results do not fully explain the predicted performance, they noted that the high degree of regularity of the P-type afferent improved detection performance, since for renewal (Poisson type) spike trains the detection probability was only 10–15%. The regularity is due to anti-correlations in the interspike interval sequence that serve to reduce the variability of the spike count distribution over long counting windows (Ratnam and Nelson 2000; Chacron et al. 2001). However, a trial-based approach does not explicitly utilize the information contained in the temporal sequence of spikes, because only the first-order statistics [mean and standard deviation (SD) of the spike count] are used, and not the correlation structure of the sequence. The anti-correlation structure of the spike train (second-order statistics) may provide additional information that can be used to improve detection performance. In particular, neurons in the ELL, that receive input from P-type afferents, may exploit such information, and so better approach the performance suggested by behavior.

Here we present a continuous detection procedure that is able to exploit the anti-correlations in the input spike train, and can serve as a model for a detection strategy that may be employed by the ELL. This procedure is based on sequential testing (Wald 1947) and marks a radical departure from fixed-sample or trial-based testing using, for instance, a post-stimulus time histogram (PSTH). It models the well-known properties of neural integration of excitatory postsynaptic potentials (EPSPs), and spike generation when the membrane voltage reaches a threshold. We assume that a detector neuron continuously performs a binary hypothesis test on the summed EPSPs. At every time instant the neuron decides whether the summed EPSP reflects unchanging background spiking activity in the presynaptic neuron or whether there was a signal embedded in it. Unlike trial-based testing, where the

beginning and duration of the testing interval are predetermined, in sequential testing the neuron continually tests the EPSPs, and terminates only when it determines that a signal was present. At this point, it generates a spike, signaling a hit. We show that because the input is continuously evaluated, as opposed to the evaluation of a single time interval in trial-based testing, the neuron can utilize memory of prior activity to improve the reliability of decision-making. Further, continuous testing requires a smaller sample size than trial-based procedures (see Wald 1947; Siegmund 1985), so the speed of decision-making is optimized.

We consider the baseline (spontaneous) activity of a sensory afferent fiber in the weakly electric fish, and perturb it at a random point in time by inserting a spike or by shortening an interspike interval (ISI). The coding of a stimulus with a single or few spikes and its detection, has been studied in several systems using spike counting or trial-based testing (reviews: Rieke et al. 1997; Parker and Newsome 1998; see also Fitzhugh 1957; Barlow et al. 1971; Relkin and Pelli 1987; Lee et al. 1993; Vallbo 1995; Tougaard 1999). In most cases however, the background rates are relatively low, which facilitates the detection task. It becomes considerably more difficult when the signal-to-noise ratio (SNR) is low, due to small changes in spiking activity and high background rates. Here, we address the task of detecting a small transient perturbation in the presence of significant background spiking activity, without the benefit of repeated trials, or knowledge of the time of occurrence. The main findings are: (1) a simple physiologically plausible neuron (the leaky integrator) can perform reliable signal detection on a continuous basis. The leaky integrator is used to illustrate the concepts, because of its simplicity, and because it is a common model of neural membrane activity, but more accurate models can also be used. (2) The performance of this detector is determined to a large extent by the statistical properties of the baseline spiking activity. Neural spiking activity is often modeled as a renewal process (e.g., a Poisson or Gamma process), and for such spike trains detection performance is relatively poor. However, the presence of negative correlations (anti-correlations) between adjacent ISIs as observed in the P-type afferent spike trains results in a dramatic improvement in the detectability of weak signals. In a negatively correlated spike train, the duration of a given ISI depends on the preceding ISI, such that long ISIs are followed by short ISIs and vice versa (see Figs. 5 and 7 in Ratnam and Nelson 2000). We also demonstrate that only sequential procedures can fully exploit such temporal correlations in spike trains, and increase detection performance beyond the performance achievable by traditional trial-based procedures. This is because in contrast to trial-based procedures (Ratnam and Nelson 2000; Chacron et al. 2001), sequential testing incorporates memory in the decision making process. Anti-correlations in spike trains have been observed in several sensory systems (Kuffler et al. 1957; Amassian et al. 1964; Bullock and Chichibu 1965;

Lowen and Teich 1992; Ratnam and Nelson 2000; Steuer et al. 2001) and may be more widespread than hitherto reported. Thus, while it is often assumed that spike trains have no memory (i.e., ISIs are independent), this study indicates that correlations in afferent spike trains can play an important role in signal encoding and detection.

Materials and methods

We consider a spontaneously discharging neuron. At some unknown point in time its spike train is perturbed due to a stimulus, which either adds one spike or shortens an ISI. Subsequent to the perturbation the neuron returns to its baseline activity. Since baseline spiking activity fluctuates, if the perturbations are weak and limited in duration, they can be obscured by the fluctuations. Our goal is to determine whether there are physiological mechanisms that can correctly detect perturbations while ignoring the intrinsic fluctuations.

Spike trains and signal generation

Extracellular recordings of spontaneous activity were made from P-type primary electrosensory afferents of the weakly electric fish *A. leptorhynchus*. The experimental procedures are described in Xu et al. (1996). *A. leptorhynchus* has a quasi-sinusoidal EOD waveform with a fundamental frequency that depends on the individual and ranges from 750–1,000 Hz. P-type units fire maximally once per EOD cycle and randomly skip cycles between successive spikes. On average, they fire on about one-fourth of the EOD cycles. This ratio is the per-cycle-probability of firing p . Information is encoded in the neural spike train as changes in p , and hence these units are called probability coders. Since p changes with stimulus intensity, stimulus amplitude is coded as a change in spike rate (rate-coding mechanism). Experimentally obtained spike trains from P-units (sampled at 13.89 kHz, 72- μ s interval) were resampled at the unit's EOD frequency, which corresponds to a sampling period of approximately 1 ms (Ratnam and Nelson 2000). The spike train was represented as a discrete binary valued sequence $x[n]$, where n is the number of elapsed EOD cycles since the start of the recording. Since units fire at most once per EOD cycle, $x[n]=1$ if there is a spike at cycle n , otherwise $x[n]=0$. To determine the effect of discretization, the detection performance for the resampled spike train was compared with the original spike train, sampled at 13.89 kHz.

To mimic the task a postsynaptic neuron faces when detecting a small change in spiking activity, two types of synthetic signals were added to the baseline spike train: (1) a spike was randomly added to $x[n]$ at a location that did not already contain a spike; (2) a randomly selected ISI was shortened by 1–3 EOD cycles. In the case of the afferent, shortening an ISI by k EOD cycles effectively adds k/\bar{I} spikes, where \bar{I} is the mean ISI. These signals represent small perturbations of the spike train and are motivated by experimental observations of small numbers of spikes being added due to weak stimuli (Fitzhugh 1957; Vallbo 1995; Tougaard 1999), or of ISIs being shortened (Goldberg and Fernandez 1971; Blanks et al. 1974; Tricas and New 1998; Ratnam et al. 2001).

Detection theory

The detection of a change in spiking activity in response to a signal is performed by a statistical test of hypotheses (reviews: Green and Swets 1966; Rieke et al. 1997; Gabbiani and Koch 1998). Testing requires a decision statistic, which is usually the spike count in a given time window, and a decision-making strategy. In trial-based detection as usually employed in neurophysiology (e.g., Relkin and

Pelli 1987; Shofner and Dye 1989; Lee et al. 1993; Ratnam and Nelson 2000; Chacron et al. 2001), the neuron's spike count in response to a stimulus (the alternate hypothesis, H_1) is measured under repeated trials, and compared to the baseline (spontaneous) activity (the null hypothesis, H_0). We test for the mean increase A in the spike count y due to the stimulus by counting in a fixed time window. This can be represented as:

$$\begin{aligned} H_0 : y &= y_n \text{ (noise only: Null hypothesis)} \\ H_1 : y &= y_n + A \text{ (signal + noise: Alternate hypothesis)} \end{aligned} \quad (1)$$

When the spike count crosses a specified threshold γ , it is assumed a signal is present, and H_1 is accepted, otherwise H_0 is accepted. The performance of the detector can be assessed from the probability distribution functions (PDFs) of the count under the two hypotheses, and the threshold (Fig. 1). Performance is determined by the probability of a correct detection (P_d), acceptance of H_1 when it is true; and by the probability of false alarm (P_{fa}), acceptance of H_1 when in fact H_0 is true. By systematically varying the threshold the relationship between P_d and P_{fa} can be depicted using a receiver operating characteristic (ROC) curve. This decision-making strategy is the Neyman-Pearson scheme and trials are assumed to be independent. A common parameter to characterize signal discriminability is d' , which refers to the degree of overlap between the PDFs. When the PDFs are Gaussian distributions with equal variance, $d' = \frac{A}{\sigma}$. Thus, both the SD (σ) and the mean increase in the spike count determine signal detection performance. The SNR is the square of d' .

We modified the standard trial-based procedure in two ways. First we replaced the counting procedure with leaky integration. This makes the decision statistic more realistic as it better reflects the membrane voltage of the postsynaptic neuron. Then we modified the decision-making strategy, replacing trial-based testing with sequential (continuous) testing.

Decision statistic: the neural integrator output

Counting spikes in a fixed time window (boxcar counting, or trial-based testing) can be performed such that successive window locations are non-overlapping (Fig. 2A), or by sliding the window forward one EOD cycle (or sampling period) at a time (overlapped counting, Fig. 2B). A non-overlapping window counts each spike only once, while an overlapping window counts each spike T times, for the duration of the overlap. Each spike causes an elevation of the filter output for a time T . Thus, the effect of a single spike persists for a time T in the output of the counter, and the counts from overlapping windows are correlated over a time T .

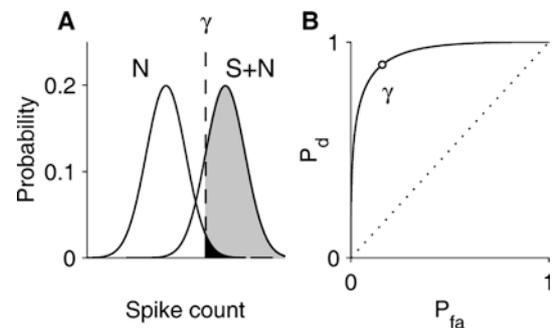


Fig. 1A,B In trial-based testing, the presence of a signal in noise ($S+N$) versus the presence of noise alone (N) can be tested using a binary hypothesis test. **A** The probabilities of detection (P_d) and false alarm (P_{fa}) are assigned based on spike count distributions of the N and $S+N$ situations. The area under the probability density functions (PDFs) to the right of the threshold γ determines P_d (light gray) and P_{fa} (dark gray). **B** The receiver operating characteristic (ROC) curve shows the trade-off between P_d and P_{fa} as a function of threshold. As the threshold decreases, both P_d and P_{fa} increase

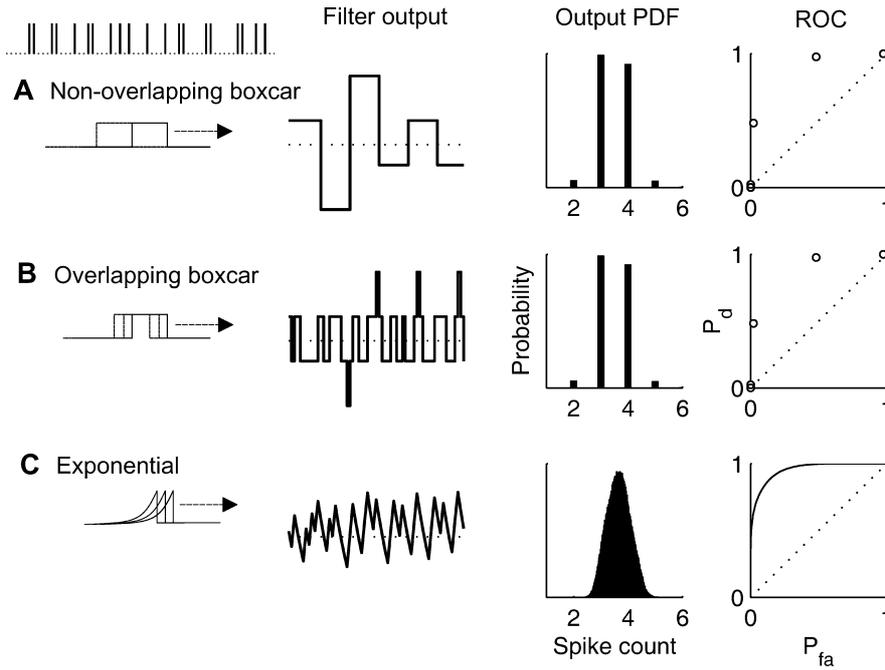


Fig. 2A–C Spike counting implemented using filters. Spike sampling rate is the electric organ discharge (EOD) frequency. **A**, **B** Spikes can be counted in a fixed time window of duration T by multiplying the spike train (top left) with a unit-amplitude boxcar window (first column) and summing. The second column shows the output of the filter; the dotted line indicates the mean filter output. In **A** successive counts obtained by sliding the window forward without overlap results in output every T EOD cycles. In **B** successive counts obtained by sliding the window forward in steps of one EOD cycle result in output every EOD cycle. **C** A leaky integrator can be implemented as a filter with exponential window shape, and time constant T . The output of a leaky counter is continuously valued in contrast to the discrete values obtained from the boxcar window. The third column shows the PDF obtained by binning the filter output. In **A** and **B** for large sample size (here $\sim 3 \times 10^6$ EOD periods) the PDFs are identical. The fourth column shows the ROC when the task is to detect an added spike. The dotted line is chance performance. The ROCs for boxcar filtering are discrete, with only a few possible P_d s and P_{fa} s (circles), while the ROC for the exponential filter is continuous valued

Overlapped counting is a linear filtering operation. The boxcar window is the impulse response $h[n]$ of the filter, which is the output of the filter in response to a single spike. If $x[n]$ is the input spike train, the output $y[n]$ of the filter is the convolution of the window with the spike train according to $y[n] = x[n] * h[n]$. For the overlapping boxcar window where counting is performed in a window of length T , $h[n] = 1$ for $0 \leq n < T$, and zero otherwise. It can be seen that the non-overlapping boxcar filter output is a subset of $y[n]$ sampled at intervals T apart (Fig. 2A, B). While spike counting is often employed in rate coding models, a more realistic description should consider real membrane characteristics. While there are many models of neural integration, a commonly used mechanism is leaky integration, which is a low-pass filter with an exponentially decaying impulse response (Fitzhugh 1957; van Rossum 2001). The decay rate is governed by the time constant T of the membrane. The filter has exponentially weighted memory and gives greatest weight to the most recently occurring spike (Fig. 2C). This is unlike boxcar counting where memory is perfect but finite and all spikes in the window are weighted equally. The impulse response function for the leaky integrator is $h[n] = e^{-n/T}$, where $n \geq 0$ is an integer. For large T (≥ 10), the area under the filter is approximately T , and so the mean output for a leaky integrator becomes equal to the area under a boxcar window of length T . It should be noted that for the

exponential filter, $h[n]$ takes non-integer values, and the output $y[n]$ can assume values between 0 and T . When a spike is convolved with the filter function, the output is the spike broadened in time according to the impulse response function h . Thus, the filter output samples will be correlated, and the extent of broadening can be defined by the correlation time:

$$\tau_c = 1 + \sum_{k=1}^{\infty} \rho_k \quad (2)$$

$k = 1, 2, 3 \dots$, where ρ_k are the serial correlation coefficients of the output. For a process in which the input samples are independent, the correlation time of the output after filtering is equal to the time constant of the filter.

P-type baseline activity $x[n]$ was filtered using a leaky integrator with time constant T , and the PDF was obtained by binning the filter output $y[n]$. The PDF describes the possible values that $y[n]$ can assume and their probability of occurrence, but it does not provide information about dependencies (i.e., correlations) in the time series. Figure 2 shows the PDFs for the three filters. In addition to being a more realistic description of neural integration, the leaky integrator has the advantage that the PDF is continuous valued. The properties of the PDFs determine the ROC, since P_{fa} and P_d are calculated from the area of the PDFs under H_0 and H_1 .

To determine the detectability of small changes to the spike train, a single spike (the signal) was added to the spike train. The spike train was filtered and the detectability of the signal in the output $y[n]$ examined. If a spike is added at $n = m$, and the spike train is exponentially filtered, the effect of the spike for $n \geq m$ is given by the impulse response $h[n - m] = e^{-(n-m)/T}$. For large T (≥ 10), it can be shown that the mean increase A in filter output in a window of length T_s beginning at $n = m$ is $A \cong \frac{T_s}{T} (1 - e^{-T_s/T})$. In other words, the effect of the added spike persists in the output by increasing its mean for some time. Thus, to test for the presence of a signal, we define a “signal window”, the time following the added spike during which the effect of the spike is appreciable. Since the correlation time of the filter output is T , the time constant is a natural candidate for the signal window duration, thus we set $T_s = T$ for the remainder of this work.

Signals can also be weaker than a single added spike, such as those that induce a change in spike timing. This is modeled by shortening an ISI. If an interval is shortened by k EOD cycles, the mean change in y is kA/\bar{I} . If $k < \bar{I}$, the perturbation in $y[n]$ is less than the addition of a spike. Thus shortening intervals allows for

more graded responses to a stimulus. Since the effect of the shortened interval is not noticed until the filter encounters the spike terminating the interval that is shortened, the signal window corresponds to the duration T from the spike terminating the interval. Other possible signals, for instance combinations of spike additions and deletions, or lengthening and shortening of intervals can be treated in a similar manner.

Decision-making strategies

The Neyman-Pearson test can be extended so that the filter output is examined continuously. This means that decisions are made every sampling instant, or EOD cycle. The filtered spike train $y[n]$ was used as the decision statistic to perform the binary hypothesis test. Earlier we defined a signal window, since filtering (integrating) the input spike train smears out a spike. For a given threshold γ , whenever the filter output $y[n]$ crosses the threshold within the signal window, it was considered to be a correct detection, otherwise it was considered to be a false alarm. There are several ways in which the hypothesis test can be performed (Fig. 3).

Scheme 1: trial-based testing

Consider an exponential filter with time constant T that produces filter output $y[n]$. If only every T th sample is retained (i.e., subsample; Fig. 3B), this is equivalent to the non-overlapping counting scheme (since the correlation time of the filter is T). Hypothesis testing using these samples is similar to the trial-based procedures employed in neurophysiology, where spikes are counted in windows of fixed duration over repeated trials (Relkin and Pelli 1987; Shofner and Dye 1989; Lee et al. 1993). It should be noted that for a signal window of length T only one decision is made in the signal window. P_{fa} and P_d were calculated by dividing the number of hits by the number of trials, in the absence and presence of a signal, respectively.

Scheme 2: testing every sampling instant

The trial-based detection scheme can be extended by considering all samples $y[n]$ (Fig. 3A). The number of possible samples during the signal window is equal to T , and so T decisions are made in the signal window. This scheme calculates the per sample probability of detection or false alarm, and differs from scheme 1 in which only one sample per window of duration T is considered. Although this scheme can be implemented sequentially, all decisions are independent. That is, while decisions are made sequentially, only the total number of threshold crossings is considered, without reference to prior decisions. Thus, only the information contained in the PDF of $y[n]$ is used.

We are interested in improving schemes 1 and 2 because, in addition to their lack of biological realism, they suffer from other drawbacks. In scheme 2, if some of the T samples in the signal window do not exceed threshold, then the number of misses is increased and so the detection probability is lowered. Further, both schemes fail to incorporate important information afforded by temporal information in the spike train, in particular, negative correlations between adjacent ISIs. We can improve on these decision strategies by using two different kinds of sequential decision-making strategies.

Scheme 3: sequential testing

Here again, every sample point is tested. At the first instant that $y[n] \geq \gamma$ (H_1 is favored) a hit is scored. The hit is a correct detection if it occurs within the signal window, otherwise it is a false alarm. After a threshold crossing the test is terminated. If $y[n] < \gamma$, H_0 is favored, and testing continues. This is the classical sequential

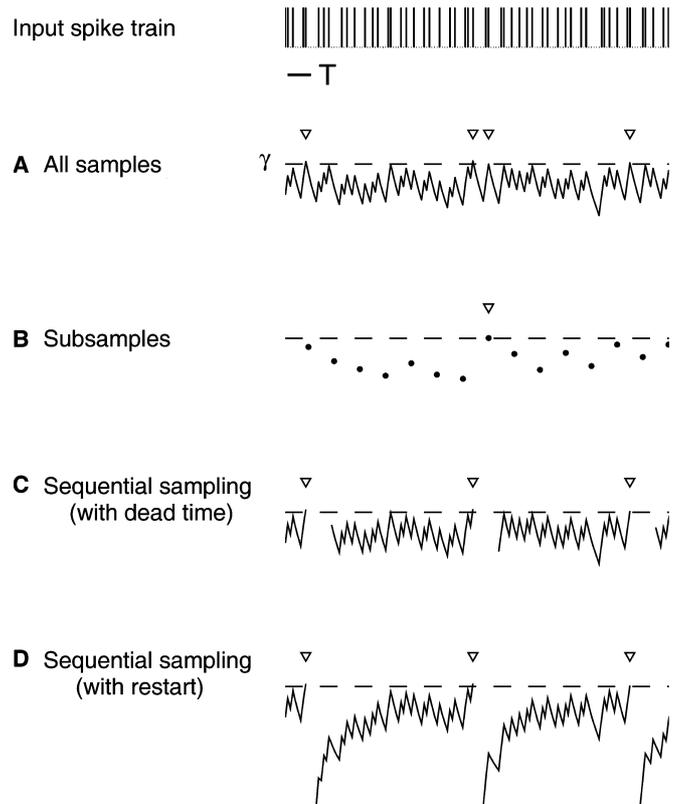


Fig. 3A–D Illustration of the detection schemes. The input spike train ($267 \text{ spikes s}^{-1}$, $p = 0.35 \text{ spikes/EOD cycle}$) is filtered with an exponential filter with time constant T , and decisions are made using four different schemes (*horizontal bar* is $T = 10$ EOD cycles, 13 ms). **A** Filter output (mean 3.66 ± 0.42 spikes) which is evaluated every EOD cycle. When the filter output exceeds the threshold $\gamma = 4.36$ a hit is generated (*triangle*). **B** Trial-based testing using non-overlapping samples is performed by using every T^{th} sample of the filter output (since the correlation time of the filter equals T). The sampled output points (*dots*) are a subset of samples shown in **A** separated by time T . Note that only one of the hits generated in **A** is sampled in **B**. Neither scheme takes into consideration the temporal correlation between counts. **C** Sequential testing, where every output sample is tested as in **A**, except that testing terminates when the threshold is exceeded, and restarts T EOD cycles later. **D** The counter is reset following a hit and restarted on the next EOD cycle, which makes continuous sequential testing possible. This is the leaky integrate-and-fire neuron. The advantage of testing strategies **C** and **D** compared to **A** is that groups of correlated hits are excluded (the second and third hits in **A** that are separated by a time interval $< T$ register as only one hit in **C** and **D**). The counting strategy in **C** and **D** is similar to a random walk towards a barrier and is sensitive to the history or temporal sequence of counts

testing procedure of Wald (1947). The sequential scheme can be extended by restarting after a dead-time (Fig. 3C), or by resetting the integrator (Fig. 3D) immediately following a threshold crossing. These modifications to schemes 1 and 2 result in a radical departure from trial-based testing, with important consequences. The sequential scheme overcomes the deficiency of scheme 2, since one hit in the signal window is sufficient to trigger detection. Another advantage is that it incorporates memory of prior decisions. However, a drawback is that the PDF is truncated, because testing stops when the threshold is exceeded, and filter output values above the threshold are non-existent, which makes the PDF dependent on the threshold. Because of this P_{fa} cannot be calculated and so the ROC cannot be constructed. Instead the false

alarm rate (number of false alarms per second) is used, and performance is characterized using an operating characteristic (OC) curve of P_d versus the false alarm rate.

When using sequential testing with dead-time (Fig. 3C), testing is terminated after each hit and restarted after the dead-time T_d . T_d is set equal to T since the correlation time of the fluctuations in the filter output is $\sim T$, and the output will have drifted back to its normal values. The drawback of this scheme is that it is not biologically realistic, but it has the advantage of being more analytically tractable. For sequential testing with reset (Fig. 3D), the integrator is reset to zero. Upon exceeding the threshold at some time n , we restart the integration by setting $y[n]=0$ (the resting potential of the leaky integrator), and let decision making continue on subsequent samples. This is the familiar integrate-and-fire model. Resetting of the integrator after a threshold crossing erases the memory of the prior events. Since the integrator has to return to its baseline values, there is a time period when no threshold crossings are possible (the absolute refractory period), and a further period when threshold crossings are less likely (relative refractory period).

Comparisons with low order models

Ratnam and Nelson (2000) and Chacron et al. (2001) showed that regularity of the spike train, and particularly negative ISI correlations, decreased the SD of the counts in trial-based detection, and thus increased the detectability of a signal. However, trial-based detection does not explicitly use the temporal correlations in the spike train, since only the information in the PDF is used. Sequential testing differs in that it is sensitive to the serial dependencies in the input spike train. Thus, it is expected that the temporal structure of the spike train affects signal detectability when using sequential tests. To test this, the afferent spike train was compared with models that reproduce the lower order statistical features of the data, and incorporate temporal correlations to varying degrees. The afferent was compared with the binomial (B) and the zeroth-order Markov (M_0) renewal process models, and with a non-renewal model, the first-order Markov (M_1) process. These models were used since they are point-process models with well-defined statistical properties. These properties, including ISIs, joint ISIs and serial correlations are described in Ratnam and Nelson (2000). Note that they do not model the membrane characteristics or the driven response of the neuron, and more detailed models of P-type afferents exist (Kashimori et al. 1996; Nelson et al. 1997; Kreiman et al. 2000; Chacron et al. 2000). The models by Nelson et al. (1997) and Kreiman et al. (2000) are renewal models, while the model by Chacron et al. (2000) exhibits serial dependencies. The binomial spike train was generated by shuffling $x[n]$, in order to make the samples independent. Only the per-cycle probability of firing p is matched to the afferent. The M_0 spike train was generated by shuffling the ISIs, and preserves both p and the ISI distribution. The M_1 spike train was generated by constrained shuffling of the ISI sequence, so that the joint distribution of adjacent ISIs, the ISI distribution and p are preserved. M_1 exhibits non-zero ISI correlations. The afferent and model spike trains were subjected to the same analysis throughout this work.

Results

Spontaneous P-type afferent activity was recorded from 49 fibers from six *A. leptorhynchus*. The electric organ of *Apteronotus* is neurogenic, and continually discharges throughout the fish's life, as well as under neuromuscular blockade. Therefore we also refer to the baseline spike rate (the firing rate in the absence of a stimulus) as spontaneous activity. The EOD frequencies were

between 750 and 1000 Hz. The afferent firing rates ranged from 65 to 563 spikes s^{-1} (250 ± 117 spikes s^{-1}), with a mean ISI of 4.3 ± 2.3 EOD cycles. The statistical properties of baseline spike trains and the lower order statistical models have been characterized in detail in Ratnam and Nelson (2000).

Filtering of the spike train

The spike train was integrated, or counted using boxcar and exponentially weighted filters. The increase in the filter output $y[n]$ due to a spike is the convolution of the spike with the filter impulse response $h[n]$. Figure 2 shows the comparison of the filter output, the PDFs and the ROCs for three filters: a non-overlapping boxcar of length $T=10$ EOD cycles, which is the same as counting in fixed time windows (Fig. 2A), an overlapping or sliding boxcar (Fig. 2B), and an exponential filter with time constant $T=10$, which models the postsynaptic membrane characteristics (Fig. 2C). The filter output and the count distributions for the boxcar filters are discrete, resulting in an ROC that consists of only a few points. Because the output of the non-overlapping boxcar filter is a subset (every T^{th} sample) of the overlapping boxcar, the PDFs and ROCs are identical for a large number of samples. They differ in that for the non-overlapping boxcar only one output sample is taken every T EOD cycles, and thus all output samples are uncorrelated, while this is not the case for the overlapping boxcar. The discrete ROC prohibits a finely tuned detection algorithm, i.e., because there are so few points on the curve, it is likely that the only choice would be between a sensitive detector with an unacceptable false alarm rate (large P_d and P_{fa}), or an insensitive detector with low false alarm rate (small P_d and P_{fa}). Thus counting spikes online provides little flexibility in the implementation of a neural detector.

When using leaky integration, for large T (≥ 10) the mean filter output (3.66 ± 0.42) is almost the same as for boxcar counting (3.48 ± 0.59), and thus spike counting and leaky integration are equivalent. To determine whether the discretization of the spike train introduces sampling errors, we filtered the spike trains resampled at different rates with a leaky integrator and compared the ROCs. The spike train was discretized at the maximum possible sampling rate, the sampling rate at which spikes were acquired (13.89 kHz, 72 μs resolution), and at the EOD rate (763 Hz, 1300 μs). The worst-case error in the detection probability was 2.3%. Thus, given the limitations on spike sampling rates that are available in experimental recordings, this discrepancy is an upper bound on the discretization error. For P-type afferents, which encode modulations only in their mean firing rate and not in the phase of their firing with respect to the EOD period, increasing the sampling rate beyond the EOD rate does not provide improved detection performance. For the remainder of the study, we use spike

trains discretized at the EOD rate, and use only the leaky integrator.

There are several advantages to using a neural integrator, such as a leaky integrator, instead of boxcar counting: (1) membranes modeled as a leaky integrator are biologically more realistic; (2) the output of the integrator can take any value, which leads to continuous valued PDFs, and so the operating point on the ROC curve can be set for the entire range of P_d (or P_{fa}) values (Fig. 2C); (3) the integrator can respond to changes in spike trains more quickly and accurately because most recent samples are given more weight, whereas boxcar counting weighs all EOD cycles in the window equally; and (4) the noise fluctuations are smaller for the leaky integrator, and thus the SD of the output is lower (for example, for a binomial spike train the SD of the leaky integrator output is a factor of $\sqrt{2}$ less than the SD of a boxcar counter). Thus, despite the ubiquity of count analysis in the literature, more realistic neural integration has many advantages, particularly for coding and signal detection. The leaky integrator used here can be replaced with more accurate neural models, or Hodgkin-Huxley type neurons. This increases the amount of detail in model detector, but the following methods apply in the same manner.

Effect of filter time constant on SD of the filter output

Signal detectability is determined by the signal strength (the added spike), and the SD of the PDF (the noise fluctuations in the filter output). While the signal strength only depends on the number of spikes added, the SD of the integrator output depends on the time over which spikes are integrated. For instance, for a renewal spike train the SD grows in proportion to \sqrt{T} (Cox 1962). Even though for non-renewal processes such as electrosensory afferent spike trains, this is not necessarily the case (see below, and Ratnam and Nelson 2000), the choice of a time constant remains crucial, and we wish to determine an integrator time constant that gives best detection performance.

Figure 4A shows the SD of the filter output as a function of time constant for the afferent and the models B, M_0 and M_1 for a representative fiber. The representative fiber in Fig. 4A will be used throughout this work. Fluctuations in the filter output are smallest for the afferent over a large range of time constants. For most afferents the curve of the SD as a function of T stays flat for a large range of T values, and often it shows a minimum (inset). On the other hand, for the models there is a large increase in the SD with increasing filter time (see Ratnam and Nelson 2000 for a similar analysis of the coefficient of variation, CV). Intuitively, it can be understood that the minimum is a result of regularity (i.e., negative ISI correlations), by considering the case in which a long ISI is always followed by a short interval and vice versa (Ratnam and Nelson 2000). If the

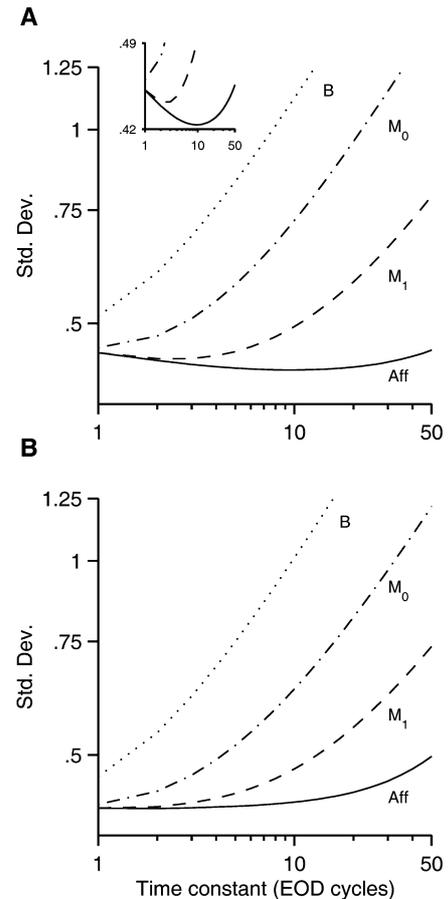


Fig. 4A,B Standard deviation (SD) of the count PDF as a function of the leaky integrator time constant T (in EOD cycles) for afferents (*Aff*) compared to their matched models (binomial: *B*, zeroth-order Markov: M_0 , first-order Markov: M_1). **A** SD as a function of T for a representative afferent and its matched models. The SD increases with increasing T for the models, but for the afferent the curve exhibits a minimum at $T=10$ (*inset: expanded view*). Nearly two-thirds of the afferents exhibited a minimum in the curve. **B** Mean SD curves for the population. The SD for the afferents stays relatively flat up to about 20 EOD cycles (~ 25 ms) and thereafter increases slowly

counting window now spans two intervals, the variability of the counts in this window will be decreased. This is the case for the M_1 model (Fig. 4A inset) that preserves only the adjacent afferent ISI order. Afferent spike trains, on the other hand, are regularized over much longer time scales (Ratnam and Nelson 2000), and thus the minimum in the SD curve occurs later. Half of the afferents showed a clear minimum (such as seen in Fig. 4A) that could be fit with a polynomial, which occurred between 3 and 13 cycles (95% confidence limit, mean of 8 EOD cycles), corresponding to a range of 3–17 ms. Figure 4B shows the mean SD over the population of fibers. There is no clear minimum because not all fibers showed a minimum (not shown). The minimum is shallow or not present when the PDF cannot be approximated well by a Gaussian distribution, or changes its shape for different time constants. Nevertheless, the mean SD of the afferent population is

relatively flat for time constants up to ~ 20 EOD cycles (~ 25 ms).

Since the mean shift in the PDF under H_1 is independent of the time constant for long T values, detection performance is determined by the SD, and is best at the minimum of the SD curve. The relatively constant SD over a range of time constants for the afferent indicates that detection performance will be close to optimal, and will not vary much over this range of time constants. The above results suggest an optimal range of integrator time constants of 3–17 ms. This range is in good agreement with the known time-constant for the post-synaptic neuron. The E-type (excitatory) pyramidal neurons in the ELL of *A. leptorhynchus*, which receive input from the electrosensory afferents, have mean time constants of about 16 ms (Berman and Maler 1998). Since each ELL neuron receives input from multiple afferents we choose a single time constant to represent a hypothetical ELL detector cell. We used a time constant of ten EOD cycles (~ 13 ms, the optimal time constant for the representative afferent) for the remainder of this work. For a time constant $T=10$, the mean filter output $y[n]$ for the afferent in Fig. 4A is 3.66 ± 0.42 , while the SD is 1.12 for the binomially shuffled spike train, 0.72 for M_0 and 0.50 for M_1 . Table 1 shows the SD of the filtered spike trains for $T=10$ for the population.

Single spike detection performance—trial based detection

To illustrate the important issues in online (sequential) detection, we first present results for the two simple and equivalent detection schemes (schemes 1 and 2, see Materials and methods). These are the trial-based detection (scheme 1), and testing every sampling instant (scheme 2). Although the filter output samples in scheme 2 are not independent, the performance of schemes 1 and 2 are the same, because the decisions are made independently without reference to prior decisions. In scheme 2, T decisions are made in the signal window, and P_d is calculated by dividing the number of hits by T . In scheme 1, only one sample is considered during T , which is the same as choosing any of the T samples from scheme 2 at random. It can be shown that the

Table 1 Discriminability d' and standard deviation (SD) of the leaky integrator output for trial-based detection

	d'	SD $y[n]$
Afferent	1.58 ± 0.26	0.42 ± 0.07
M_1	1.43 ± 0.32	0.48 ± 0.11
M_0	1.13 ± 0.41	0.63 ± 0.20
B	0.65 ± 0.12	1.01 ± 0.13

Shown are the mean \pm SD over the afferent population ($n=49$) and the matched models. *B* binomial; M_1 Markov order one; M_0 Markov order zero. Improved discriminability for afferents over the models is a result of the smaller SD of the filter output. The mean filter output $y[n]$ for the population of afferents was 2.99 ± 1.32

detection probabilities are the same for a sufficiently large number of samples. The same applies to the calculation of P_{fa} . Hence, the two schemes are equivalent. Intuitively, this also follows because non-overlapping counting is a subset of the overlapped counting employed in scheme 2.

Both schemes result in identical detection performance and the ROC is shown in Fig. 5 for the representative afferent and matched models, filtered using the leaky integrator. The detectability of the change in filter output $y[n]$ due to the extra spike is much greater for the afferent than for the matched models. This is also evident in the discriminability d' which is 1.53 for the representative afferent, whereas for the models d' is 0.53 (B), 0.88 (M_0) and 1.31 (M_1). Table 1 shows d' for the population. The improved signal detectability for the afferents compared to the models is due to the lower SD of the PDF of the filter output (Fig. 4), a result of the regularity and the strong negative correlations in the ISI sequence (Ratnam and Nelson 2000; Chacron et al. 2001).

Single spike detection performance—sequential detection

The filter-based detection schemes above are equivalent to trial-based hypothesis testing that is commonly encountered in the literature. The most significant drawback of these is that the decisions are made independently, without reference to prior decisions. If on the other hand, the detector has memory of prior decisions, it is likely to improve performance. For instance, we can

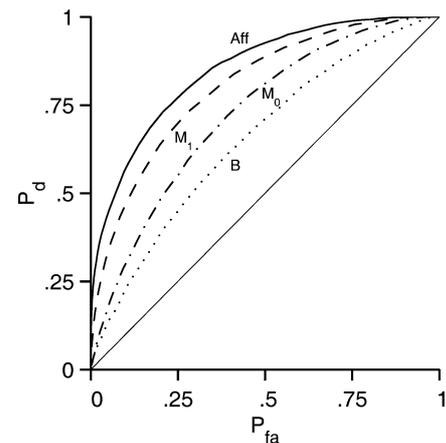


Fig. 5 The ROC for trial-based detection of a randomly added spike for the representative afferent in Fig. 4A, and its matched models (afferent: solid line, M_1 : dashed line, M_0 : dash-dot line, B: dotted line, chance: diagonal line). The spike train was filtered with a leaky integrator (time constant $T=10$) and the task was to detect an added spike. P_d is the probability of a threshold crossing in each EOD cycle within a signal window of duration $T_s=10$ after the added spike (correct detection). P_{fa} is the probability of a threshold crossing for each EOD cycle outside the signal window (false alarm). Afferent performance is markedly superior to the renewal models B and M_0

make use of the knowledge that a burst of threshold crossings within a time T of the first threshold crossing is most likely due to the same event (either due to the inserted spike or a random threshold crossing) since each spike affects the filter output for a duration T . This means that threshold crossings are no longer independent, and performance improvements can result if the detector refers to prior decisions. Thus we need a decision-making system with memory, like continuous (sequential) detection.

For example, if we consider scheme 2, to obtain $P_d = 1$ the filter output needs to be above threshold in *each* EOD cycle within the signal window. Even a single miss can lower the probability of detection. Since adding a spike elevates the filter output for T samples (the convolution of the spike with the filter) the likely event will be a burst of hits within a time T of the added spike. Therefore, if this burst is detected, or even one threshold crossing occurs shortly after the spike insertion point, it is sufficient to determine that the signal has been detected (and the remaining hits can be ignored). Consequently, we need to look for only one threshold crossing instead of T in the window $T_s = T$ following the added spike.

This procedure can also be applied to the false alarms, although in this case there is no prior knowledge of their time of occurrence. But the same argument holds, and we assume that threshold crossings within T EOD cycles after the first hit are correlated, and constitute one and the same event. After a threshold

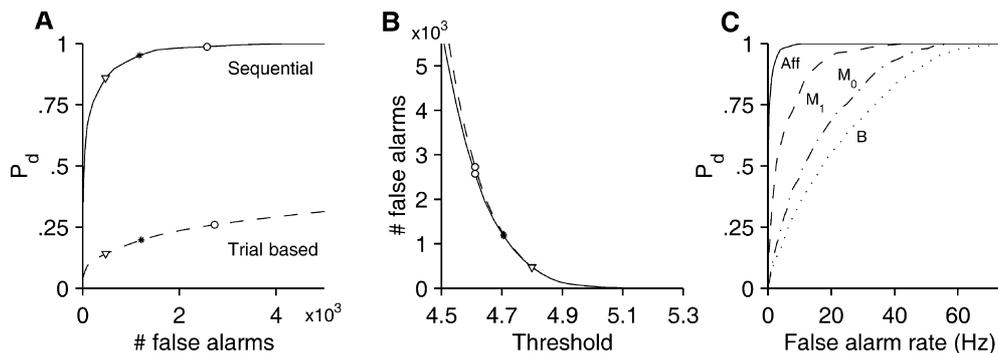
crossing the test is stopped and resumed T EOD cycles later (i.e., impose a dead-time). This truncates the upper end of the PDF (values $>$ threshold), and so P_{fa} cannot be calculated. P_d however can be calculated by the usual method. Instead of P_{fa} it is more appropriate to work with the false alarm rate, or the number of false alarms. In the appendix we present an approximation that allows us to calculate P_{fa} and d' , and compare the performance based on these metrics.

Figure 6A shows the dramatic improvement in detection performance for the afferent when using the sequential method compared to the trial-based method with testing every sampling instant (scheme 2). For a given threshold P_d , and thus the detection performance, is much higher for the sequential method (Fig. 6A) than for the trial-based method, even though the number of false alarms at high thresholds is nearly the same in both methods (Fig. 6B). Figure 6C shows the detection probability as a function of the false alarm rate for the representative fiber compared to the matched models. Performance gains due to the afferent regularity are most dramatic at low false alarm rates. For example, at a rate of one false alarm per second, P_d for detecting a spike is 0.83 for the afferent, 0.06 for the B model, 0.07 for M_0 and 0.27 for M_1 . This corresponds to a performance improvement for the afferent over the models by a factor 14, 12, and 3, respectively. At lower false alarm rates the improvement of the afferent over the models is even larger.

Detection performance for spike train perturbations smaller than a single spike

Given this remarkable detection performance for a single added spike (note that the baseline rate is ~ 300 spikes s^{-1}) it is relevant to ask whether a smaller perturbation can be detected. Preliminary experiments indicate that when very weak amplitude-modulated (AM) stimuli were applied to the animal (Ratnam et al. 2001), the number of spikes added was less than one. This corresponds to a shortening of intervals, and may represent the weakest signal that can be encoded. Interval shortening has also been observed in the elasmobranch electrosensory system (Tricas and New 1998) and in the vestibular system (Goldberg and Fernandez 1971; Blanks et al. 1974). Figure 7 shows the detection

Fig. 6A–C Signal detection performance when a single spike is added to the baseline spike train. **A** Comparison of the operating characteristic (OC) curve for the representative afferent of Fig. 5, using the trial-based (*dashed line*) method (testing every EOD cycle, see Fig. 3A), and the sequential (*solid line*) method (leaky integrator with dead-time, see Fig. 3C). The OC curve shows P_d as a function of the absolute number of false alarms in the spike train recording (~ 400 s). The sequential method provides vastly improved performance over trial-based detection. **B** Number of false alarms as a function of threshold, for the trial-based (*dashed line*) method (scheme 2) and the sequential method (*solid line*). The symbols at three example thresholds correspond to those in **A**. **C** OC showing P_d versus the false alarm rate using the sequential method, compared to the performance for the matched models (afferent: *solid line*, M_1 : *dashed line*, M_0 : *dash-dot line*, B: *dotted line*). The leaky integrator time constant was $T = 10$ EOD cycles, and the sequential method employed a dead-time of 10 EOD cycles



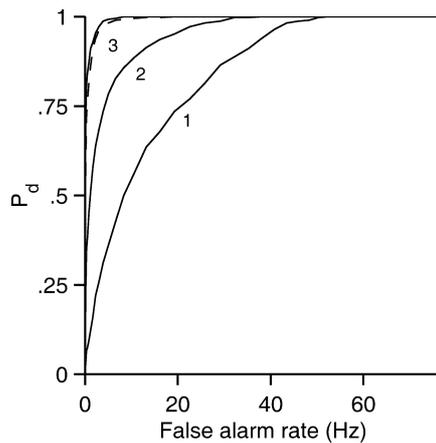


Fig. 7 Signal detection performance when interspike intervals (ISIs) are shortened. OC curve for the representative afferent of Fig. 5 for ISIs shortened by 1, 2, and 3 EOD cycles, analyzed using the same detector as in Fig. 6C. The mean ISI for the afferent was 2.9 EOD cycles, and shortening an interval by 1 EOD cycle corresponds to adding 0.34 spikes. Shortening an interval by 3 EOD cycles is approximately the same as adding a single spike (*dashed line*, as in Fig. 6C). Shortening an interval by fewer cycles produces a graded response. Leaky integrator time constant was $T = 10$

performance when a single ISI was shortened in steps of 1, 2, and 3 EOD cycles. The mean ISI for this afferent was 2.9 EOD cycles, so shortening an interval by one EOD cycle corresponds to the addition of 0.34 spikes. The OC shows that shortening an interval by three EOD cycles is approximately equivalent to adding a spike (*dashed line*), and that graded responses are obtained when intervals are shortened by smaller values.

Leaky integrate-and-fire neuron

In the sequential scheme described above, testing was terminated following a threshold crossing and restarted after a time T . This corresponds to a ‘blinking’ period or dead-time following a threshold crossing. Such a long refractory time would not only miss potentially relevant information about the stimulus, but also the absolute refractory period is generally much shorter than that considered here. How then is the previous scheme to be incorporated in a biologically realistic setting? If instead of a fixed dead-time, the leaky integrator resets after a threshold crossing, and begins integrating again from a resting value (typically $y=0$), then further threshold crossings are unlikely until the membrane is recharged, at roughly a time T later (Fig. 3D). A major benefit of resetting the integrator is that it resets the memory of prior events, and thus, correlated hits are less likely to occur. Figure 8A shows the OC for the decision scheme that stops and restarts (dead-time of length T , same as Fig. 6C) overlaid with several points for an integrate-and-fire neuron that resets to zero. At low false alarm rates, the number of false alarms and detections determined using the integrate-and-fire model is practically equal to the number calculated using the detector with

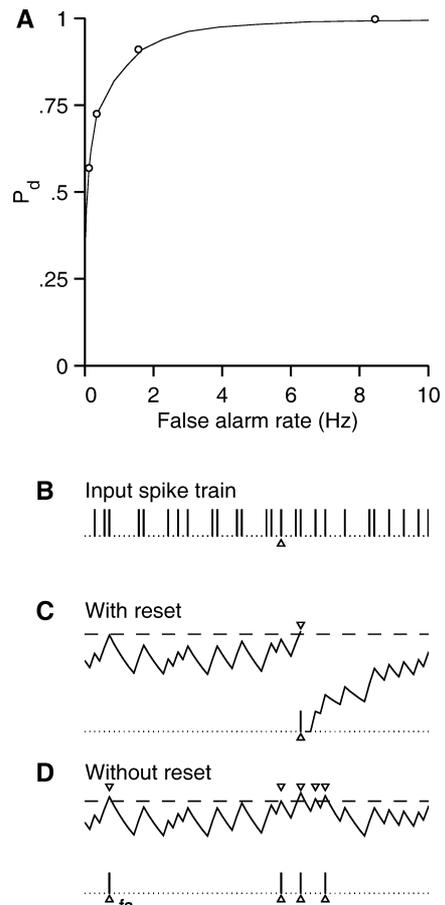


Fig. 8A–D The leaky integrate-and-fire neuron as a sequential detector (Fig. 3D). **A** Comparison of the OC for the integrate-and-fire neuron (*circles*) with the sequential (*solid line*, same as Fig. 6C) detector with dead-time (Fig. 3C) for a representative afferent. The integrate-and-fire neuron resets to zero after each threshold crossing and begins integrating anew. Its performance is almost the same as for sequential detection with dead-time. **B–D** show the effect of the reset on bursting activity. **B** Input spike train with spike added at $t=0$ (*triangle*). **C** Output of a leaky integrator neuron that resets to zero after a threshold crossing. The high threshold (4.82, *dashed line*) results in a false alarm rate of ~ 1 Hz and a P_d of 0.83. A single spike (*bar*) shows the response of the detector neuron, which fires a spike typically a few EOD cycles after the location where the added spike in the input occurred. **D** Output of a leaky integrator neuron that does not reset after a threshold crossing but has an absolute refractory period of 3 EOD cycles. The *bars* show the response of the detector neuron, which fires an isolated spike in response to a false alarm (*bar labeled ‘fa’*), while correct detections typically result in bursts of a few spikes (*unlabeled bars*). The lower threshold (4.56, *dashed line*) results in a P_d of 0.995 and a false alarm rate of ~ 10 Hz. The first detection occurs usually at $t=0$ or 1. If a neuron postsynaptic to the detector neuron can distinguish between bursts of hits and isolated hits, the false alarm rate can be lowered, while P_d will not change significantly (Fig. 9). Reset values and refractory periods between these two cases will result in intermediate levels of bursting

dead-time. Thus, for the afferent the performance of the integrate-and-fire model is very similar to the sequential scheme with dead-time, which indicates that very high detection performance is possible in a physiologically plausible setting.

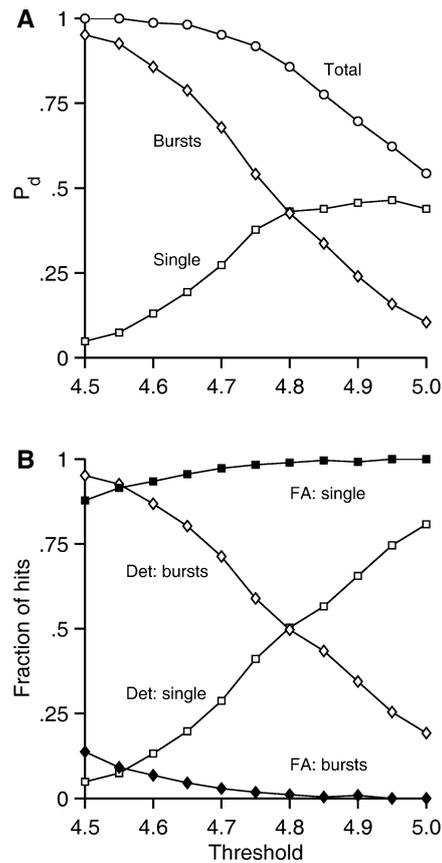


Fig. 9A,B Burst firing can enhance signal detectability. **A** The detection probability as a function of threshold shows that the total probability of detection decreases with increasing threshold (circles). The total probability was divided into the proportion that resulted from a single hit (squares) or a burst of hits (lozenges). At low thresholds, most detections are a result of burst hits, whereas at high thresholds the proportion of single hits increases. **B** Comparison of relative number of burst hits (lozenges) and single hits (squares) for detections (open symbols) and false alarms (filled symbols). Most single hits are due to false alarms and most burst hits are due to correct detections of a signal

The properties of such a detector neuron depend on the threshold and the reset value. The threshold determines the operating point on the OC curve. Figure 8C, D illustrates how different thresholds and resets affect the properties of the detector. A detector neuron that has a high threshold and resets to zero, results in a low false alarm rate. The neuron fires a single spike in response to a detection or false alarm (Fig. 8C). The spike is detected within a few EOD cycles after it is inserted. Partial reset values (Lánský and Musila 1991; Bugman et al. 1997) or no reset, result in different behaviors of the integrate-and-fire neuron. If the integrator is partially reset the refractory period becomes shorter, and bursts of threshold crossings can occur. Further, if the reset value is high, the memory is not completely erased. Figure 8D illustrates the case where there is no reset. A lower threshold is chosen here, and thus the false alarm rate is higher. There are now multiple, closely spaced threshold crossings (the neuron

fires a burst of spikes) in response to the added spike, while a false alarm results only in an isolated hit. This is due to the anti-correlations in the baseline spike train, which limit the duration and amplitude of fluctuations from the mean baseline activity (see below). Detection usually occurs immediately or on the first or second EOD cycle after the inserted spike (Fig. 8D). For values of reset intermediate to those between a full reset (Fig. 8C) and no reset (Fig. 8D) the properties are intermediate.

The presence of burst hits and isolated hits can improve detection performance if bursts and isolated hits can be differentiated. In Fig. 9A, we consider the probabilities of bursts and isolated hits within the signal window. It was observed that for low thresholds detections almost always resulted in a burst of hits, but as the threshold increases, isolated hits become more likely and bursts become less common. If all hits are classified as isolated hits or bursts of hits, it becomes clear that for a given threshold, the proportion of hits that occurs as a burst is different for false alarms and detections (Fig. 9B). At the relevant thresholds, only a small fraction of false alarms occurs as a burst of hits, while particularly at the lower thresholds, a significant fraction of the detections occurs as a burst. That false alarms tend to result in single hits is due to the negative correlations in the spike train, because although burst firing due to a false alarm is uncommon for the afferent, it is common for renewal models where ISIs are uncorrelated.

If a neuron postsynaptic to the detector can distinguish between isolated spikes and small bursts of spikes, even higher detection performance than is reported here may be possible. While it is beyond the scope of this work to suggest a neural detection mechanism that can take into account bursting activity, bursting seems to provide a convenient mechanism for distinguishing correct detections from false alarms, and hence raise detection probability beyond those reported here. These results are also true for a related species, *Apteronotus albifrons* (Goense et al. 2003). Burst firing has received much interest recently since it has been observed that bursts carry more information than isolated spikes (Kepecs et al. 2002; Gabbiani et al. 1996; Metzner et al. 1998).

Since a burst of hits in response to a false alarm is uncommon for the afferent, even for lower thresholds, the number of false alarms is relatively insensitive to properties of the integrate-and-fire model, such as the reset value or the length of the absolute refractory period. For instance, the false alarm rate did not change when the model was reset to zero (Fig. 8C) or to the mean filter output (not shown). The insensitivity of the false alarm rate to resets and the refractory period would allow for a robust implementation of the detector. Presumably real neurons with more complex filtering properties than the simple leaky integrator can exploit the statistical properties of the incoming spike trains to improve detection performance.

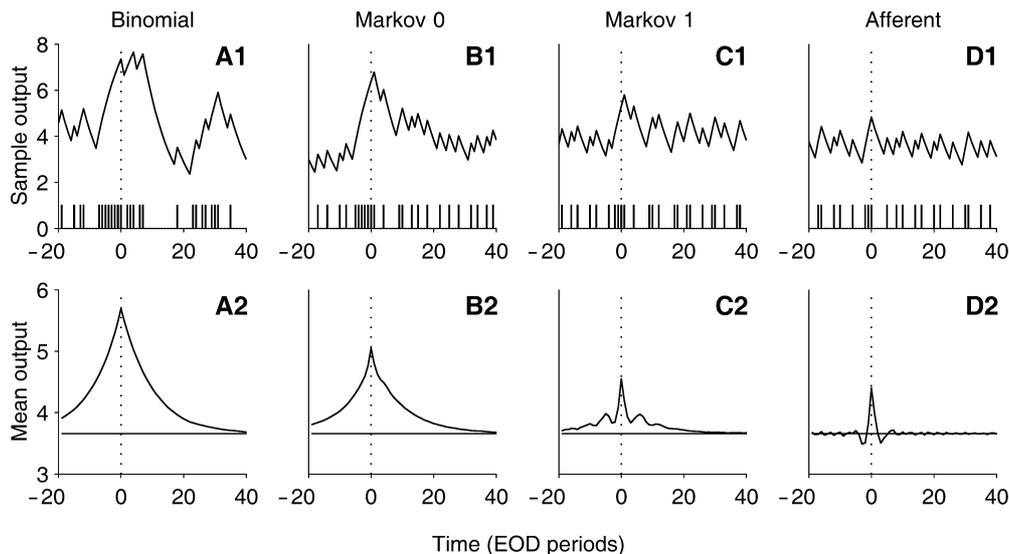


Fig. 10 Integrator output in the vicinity of a false alarm for the representative fiber (*D*) shown in Fig. 5, and its matched models (*A* binomial, *B* M_0 , *C* M_1). The spike trains were filtered using a leaky integrator with time constant $T=10$. *Panels A1–D1*: example segments of the filter output along with the corresponding stretch of spike train around a false alarm (at $t=0$). The threshold was set so that $P_{fa}=0.1$ for all cases. Fluctuations in the filtered output are larger for the models. *Panels A2–D2*: the mean filter output in the neighborhood of a false alarm averaged over all false alarm segments in the spike train. The panels demonstrate that both the amplitude and the duration of the false alarms for the afferent (*D2*) are greatly reduced in comparison with the models (*A2–C2*)

Correlations in the input and output

Since the output of the neural integrator is expected to remain elevated for a time T due to the properties of the integrator (the correlation time), the response to a false alarm is also expected to persist for a duration T . The detector without reset is thus expected to respond to a false alarm with a burst of hits. Surprisingly, while this is true for the models, it is not so for the afferent as we showed above. This has important implications for signal detection since it suggests that there are mechanisms built into neural activity that can accurately reject false alarms. This can be seen by comparing the filter outputs for a false alarm for the afferent and the models. Figure 10 shows a sample trace of the filter output in the neighborhood of a false alarm (top row), and the mean of such traces aligned at the time of the threshold crossing (bottom row). For the binomial model (panel A2) the mean filter output did indeed decay with a time constant T . But for the afferent, the mean trace (panel D2) decays much faster than the filter time constant. Note also that for the afferent, the peak amplitude of a fluctuation is much smaller than that for the models (panels B2–C2).

This faster decay for the afferent is also apparent if we examine the auto-covariance function. This is shown in Fig. 11, which shows the auto-covariance function of the baseline spike train (top row) and of the filter output (bottom row), for the representative afferent (Fig. 11,

panel D) and the models (panels A–C). The covariance function for the binomial spike train (Fig. 11, panel A1) shows that the samples are independent, as expected. The filter output covariance (Fig. 11, A2) has width T (10 EOD cycles), and a correlation time τ_c of 10.3 EOD cycles as predicted from theory. The afferent however, shows an almost identical covariance function for both the spike train (Fig. 11, panel D1) and filter output (Fig. 11, panel D2). For the afferent the correlation time τ_c of the filter output is 1.1 EOD cycles. This indicates that correlations persist for a very short time, much shorter than the time constant of the filter. The M_0 and M_1 models exhibit varying degrees of correlation. For M_0 , τ_c for the filter output is 9.7 EOD cycles, and for M_1 , which is more regular and shows more narrowing of the output covariance function, τ_c is 5.2 EOD cycles. Median correlation times for the population are 1.41 EOD cycles for the afferent (inter-quartile range 1.54), 10.5 for B (iqr 0.29), 8.24 for M_0 (iqr 2.19), and 4.09 for M_1 (iqr 2.70).

Figures 10 and 11 indicate that the filter output is seemingly decorrelated for the afferent. (The Appendix treats some of the theoretical aspects of the effect of the decorrelation on sequential detection performance.) This decorrelation has the effect that when there is a false alarm, the integrator output returns to its baseline values very rapidly, and noise fluctuations are more limited in duration. Both the amplitude and duration of the noise fluctuations are smaller for the afferent than for the models. The smaller amplitude fluctuations are reflected in the SD (Table 1 and Fig. 4), and have also been observed by Ratnam and Nelson (2000) and Charon et al. (2001). However, in addition to the reduced SD, making use of the property that the noise fluctuations are also of shorter duration results in an improvement in the detection performance (Fig. 6) that goes beyond those achieved by non-sequential procedures. These properties, of a limited duration and amplitude of the false alarms, suggest that biology may have evolved a strong noise suppression ability (presumably using anti-correlations).

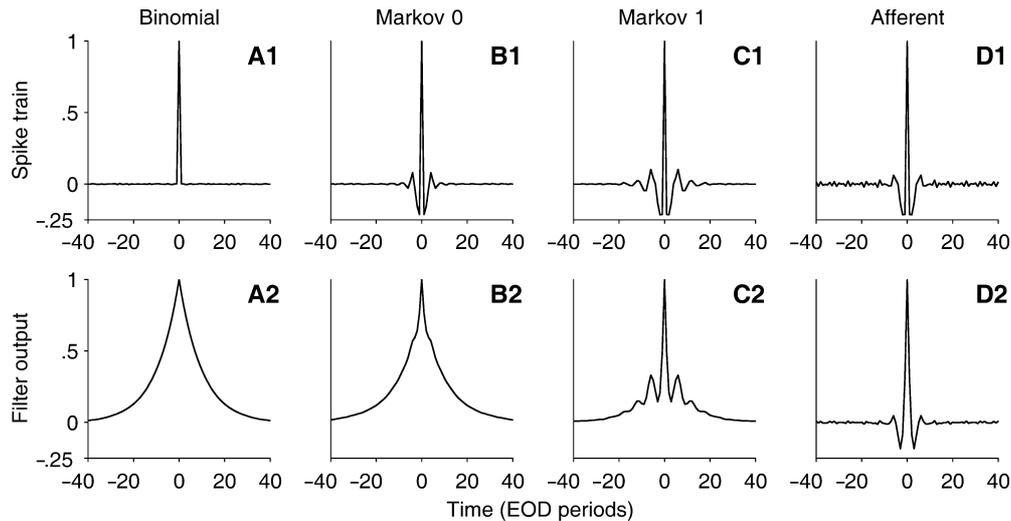


Fig. 11 Auto-covariance functions for the input spike trains (*top row*) and filtered output (*bottom row*), shown for afferent (*D*) and matched models (*A* binomial, *B* M_0 , *C* M_1). The afferent is the same as in Fig. 5. Spike trains were filtered using a leaky integrator with time constant $T = 10$. The samples of the binomial input spike train are independent and identically distributed, and so the covariance function is zero except at the origin (*panel A1*), whereas the auto-covariance of the filter output has width T (*panel A2*). The M_0 and M_1 models exhibit varying degrees of correlation (*panels B1*, *C1*). Note that for M_0 the ISIs are independent, but the spike train demonstrates negative correlations. The output of the integrator shows a more or less narrowing of the covariance function (*panels B2*, *C2*). The afferent, however, shows an almost identical covariance function for both input (*panel D1*) and output (*panel D2*), i.e., integrating a strongly anti-correlated spike train does not substantially increase the correlation time

Discussion

The methods of detection theory were applied in a continuous (sequential) framework to P-type electro-sensory afferent spike trains, to study the detectability of small perturbations in the afferent firing rate. We show that (1) a postsynaptic neuron can function as a sequential detector, which integrates the incoming spike train and at each instant performs a continuous (sequential) test on the output, to test for the presence of a signal. An integrate-and-fire neuron can function in this manner, although other integration schemes, such as the Hodgkin-Huxley model can also be used; (2) a continuous detector can detect small increases in firing rate of the input neuron reliably (single spikes), even when the small stimulus-induced changes are superimposed on a high baseline firing rate; and (3) that the temporal patterns in the spike train (particularly negative ISI correlations) are central to this ability. They improve detection performance both by decreasing the amplitude of noise fluctuations in the filtered spike train, as well as by limiting the duration of noise fluctuations. We note that negative correlations in spike trains are ubiquitous, although it is only recently that they have generated greater interest with respect to signal detection (Ratnam and Nelson 2000; Chacron et al. 2001).

Sequential detection

All information about the environment, available to an animal, is encoded in the form of neural spike trains. The mechanism of encoding information in the spike train, and the amount of information contained in spike trains, both in the firing rate and in the precise timing of spikes, is well studied (Reich et al. 1997; reviews: Rieke et al. 1997; Borst and Theunissen 1999; Buračas and Albright 1999). However, how the next higher neuron decodes the signal has received little attention. Since the first step is detection of a signal, we applied the methods of detection theory and sequential analysis to the problem of detecting a small perturbation in a spike train. That this problem is behaviorally relevant has been reported in several studies on detection of small changes in spiking activity (Fitzhugh 1957; Bastian 1981; de Ruyter van Steveninck and Bialek 1995; Vallbo 1995; Tougaard 1999; Ratnam et al. 2001; VanRullen and Thorpe 2001b).

In weakly electric fish, nearby objects modulate the per-cycle firing probability of P-type afferents. When the object is small or far away the modulations are extremely weak due to the low electrical contrast with the surrounding water, and so small changes in the afferent firing rate can be obscured by the intrinsic fluctuations in the baseline firing rate. Previous studies showed that the behavioral threshold for *Apteronotus* is $< 1 \mu\text{V cm}^{-1}$ (Knudsen 1974), and in this range the expected change in firing rate of P-type afferents is about 1 spike s^{-1} , superimposed on a baseline firing rate of about $300 \text{ spikes s}^{-1}$ (Bastian 1981; Ratnam et al. 2001). Nelson and MacIver (1999) showed that small water fleas (*Daphnia magna*, 2–3 mm diameter) could be detected at a distance of about 2 cm from the fish. Based on computer reconstructions, they estimated that the peak change in the per-cycle probability of firing of an afferent at the time the prey is detected, is about one added spike in a 200-ms window. Ratnam and Nelson (2000) simulated the detectability of added spikes using a trial-based approach and

showed that 2–3 extra spikes in a 100-ms window could be detected with 90% reliability. While this performance greatly exceeds those attained by renewal spike trains (such as Poisson spike trains), it is not adequate to explain the predicted behavioral performance. This study shows that the trial-based procedure limits detection performance, and if a biologically realistic sequential procedure is used, then a detector neuron located in the ELL may achieve the predicted performance (i.e., single spike detection). The single spike superimposed on the baseline activity of the electro-sensory afferent can be detected reliably, even though the baseline is highly variable, as judged by the CV of the ISI (Kreiman et al. 2000; Ratnam and Nelson 2000). The detectability of such small changes agrees with the observed sensitivity of the electric fish. We show that despite the high baseline firing rate, it is possible to detect a single spike in the afferents in a time that allows for fast behavioral responses (<10 ms). The range of optimal membrane time constants found also agrees with the measured time constant of the ELL E-cell (Berman and Maler 1998). A factor that is not taken into account here, but that may further enhance signal detectability is afferent convergence. The typical ELL cell receives input from 10–20 afferents (Shumway 1989b), and the averaging this provides for can enhance detectability if multiple afferents carry a signal. Detecting weak signals is only one of the tasks of the ELL. The ELL cells are variable in their morphology and properties (Bastian and Nguyenkim 2001) and it has been shown there are three spatial maps in the ELL (Heiligenberg and Dye 1982; Shumway 1989a, 1989b). Furthermore, the spatiotemporal tuning of the cells has been shown to depend on the type of stimulus (Bastian et al. 2002; Chacron et al. 2003). Thus, how the ELL processes the information present in the afferents depends on many factors, and although this study shows that it is theoretically possible to extract one additional spike in a single afferent using a simple neural mechanism, whether and where in the ELL this happens is still an open question.

The implications of this study extend to neural signal detection at large. Detection theory has been applied to neurons since the 1950s (Fitzhugh 1957; Relkin and Pelli 1987; Shofner and Dye 1989; Lee et al. 1993; Celebrini and Newsome 1994). However, the SNR is often relatively high, i.e., either the baseline activity is low, the changes in response due to the stimulus are high, or both. Another drawback of the trial-based method that is typically used is that it does not suggest how such a scheme may be implemented in a neuron. Furthermore, the long counting periods that are often used are not compatible with the speed of neural processing (Gautrais and Thorpe 1998). Finally, they typically rely on spike counting, and so most of the information contained in the temporal pattern of the spikes is not used. Some methods have expanded on this by also using temporal information in the spike

train (Geisler et al. 1991; de Ruyter van Steveninck and Bialek 1995), and this does improve detection performance. But we are not aware of a method in which this has been shown in a biologically realistic setting. In this work the detection task was phrased as a sequential decision making task. The decision statistic is the output of a leaky integrator with the spike train as input. Sequential analysis is sensitive to the temporal pattern (correlation) of the samples. Note also, that in contrast to spike counting, neural integration preserves the temporal information in the spike train, since the most recent spikes are weighted more heavily. This sequential detection scheme can exploit the temporal information in the spike train, and thereby improve detection performance. Furthermore, we have shown that this is readily extended to a biologically realistic neuron, with no loss in performance.

Sequential detection is relevant in many neural systems, since the detection of a small perturbation in a spike train within a short time of its occurrence is a general problem in the nervous system. This is important, since the representation of relevant information, like a sensory signal, by a single spike or very few spikes is common (reviews: Rieke et al. 1997; Parker and Newsome 1998). It has been shown to occur in the visual system (Fitzhugh 1957; Barlow et al. 1971; Lee et al. 1993; de Ruyter-van Steveninck and Bialek 1995), the mechanosensory system (Vallbo 1995), the auditory system (Relkin and Pelli 1987; Tougaard 1999), and in cortical neurons. The benefit of sequential detection, that decisions are made continuously, becomes important when quick detection is required, since the time scales of biologically relevant signals and response times are often short. The response time, or the time to perception, has been shown to be on the order of 20–200 ms (de Ruyter van Steveninck and Bialek 1995; Thorpe et al. 1996; Rolls et al. 1999; VanRullen and Thorpe 2001a). Sequential detection (continuous signal detection) may be advantageous when the temporal dynamics of decisions or tracking of the decision statistic is required, since it may allow for earlier detection than is possible using trial-based testing with fixed window size (Kim and Shadlen 1999; Gold and Shadlen 2000; Hernandez et al. 2002).

Role of regularity, correlations and temporal structure

The sensitivity reported here relies on temporal correlations in the spike train, particularly negative correlations between adjacent ISIs. Although it has been suggested that a refractory period introduces negative correlations (de Ruyter van Steveninck and Bialek 1995; Berry and Meister 1998; Goldberg 2000; Panzeri and Schultz 2001) and negative correlations, or a high degree of regularity have been observed in a number of systems (Kuffler et al. 1957; Hagiwara and Morita 1963; Amassian et al. 1964; Goldberg and Fernandez 1971; Blanks

et al. 1974; Lowen and Teich 1992; Tricas and New 1998; Goldberg 2000; Steuer et al. 2001), there is relatively little data on their relevance to neural coding. It has been shown that negative correlations can improve detectability of weak signals (Ratnam and Nelson 2000; Chacron et al. 2001), and can increase the information content of the spike train (Stein 1967; Chacron et al. 2001; Panzeri and Schultz 2001), but here we establish a direct connection between negative correlations in the ISIs and signal detectability.

Correlations can improve detectability in two ways. First, they decrease the SD of the spike count (Ratnam and Nelson 2000; Chacron et al. 2001), or in our case the SD of the neural integrator output (Fig. 4). The smaller SD of the spike count however does not necessarily require negative correlations in the spike train or non-renewality, and thus does not fully explain their possible function; a renewal process can have an equally low variability. Here we show that correlations in the spike train can improve detectability by a second mechanism, beyond the gain that results from a lower SD. The negative correlations result in an effective decorrelation of the filter output, as a result of which random fluctuations from the mean also have a shorter duration (Fig. 11), and so a false alarm is less likely. The shorter correlation time of noise fluctuations makes signals easier to detect (DeWeese 1996). Furthermore, under certain conditions, a true detection tends to result in burst firing of the detector neuron, whereas the response to a false alarm tends to be an isolated spike. Thus, higher-order neurons could use burst firing to further improve detection performance. In weakly electric fish, burst-firing in ELL neurons can increase the information carrying capacity (Gabbiani et al. 1996; Metzner et al. 1998; Kepecs et al. 2002), and as suggested here, this may be exploited for improving signal detection performance. The short correlation-time of intrinsic fluctuations resulting from the anti-correlations in the spike train cannot be exploited in a trial-based scheme, but only in the more biologically plausible scheme of sequential detection. Thus, the work establishes a connection between correlations and sequential detection performance, and suggests a biological basis supporting sequential detection.

Behavioral aspects of the detection task

We proposed a biologically plausible method that is able to exploit correlations in a spike train to achieve high detection performance. This is a mainly conceptual study that highlights important features of the problem, and as such did not address many of the properties that determine whether such a stimulus leads to a perception or behavioral response. The detectability of a signal at the level of the afferent does not necessarily mean it is also behaviorally relevant. This study shows that a single spike in the input can in principle be extracted, but whether this actually happens also depends on many

factors. For instance, the detection performance depends on the reliability of spike propagation and transmission, i.e., failures in spike propagation or neurotransmitter release. The spatiotemporal properties of the neural filters at higher levels in the pathway also influence detection performance. This is illustrated in the case of the mechanosensory afferent neuron in the human hand, where a single spike was shown to lead to a perception (Vallbo and Johansson 1976; Vallbo 1995). However, whether an afferent spike led to a perception depended on the region of the hand, in some regions one spike sufficed while in others multiple spikes were needed, illustrating the importance of convergence and higher order spatial filters. We therefore do not presume that a single spike will always lead to a behavioral response and further experiments are needed to address this question in the weakly electric fish. Some other factors that determine whether such a change in neural spiking leads to a percept are the acceptable false alarm rate, and which features of the stimulus are important. There is a trade-off between detection probability and false alarm rate (Fig. 6), between accuracy and speed (Reddi and Carpenter 2000), sensitivity and resolution (Schiller and Logothetis 1990), and which properties are important depends on the neural system, the level, and the purpose of the system. The relevant false alarm rate is also likely to be different for different neurons and systems. For example, if low false alarm rates are required, it means signals will be missed, but if false alarms can be eliminated at a later processing stage, such as by discriminating between burst firing and isolated spikes, higher false alarm rates may be tolerable. Sensitivity can often be increased and false alarms eliminated by convergence (Shadlen and Newsome 1998), but this requires some degree of redundancy, although in sensory systems, convergence can decrease (spatial) resolution. Besides the detection problem, neurons have to solve an estimation problem. The requirements are not necessarily the same for both tasks, and again there may be a compromise. Further studies will be needed to address these issues.

Summarizing, physiologists have focused on trial-based testing using aggregate spike counts. This procedure is not only biologically unrealistic, but also fails to exploit the mechanisms of dynamic noise-suppression, like the presence of correlations in the spike train. In behaving animals the demands of decision making are far more stringent, and a framework that can exploit the statistical properties of spike trains, is the continuous detection method suggested here. The method is realistic, and although real neurons have more diverse integrative mechanisms, its broad principles and advantages are likely to remain the same.

Acknowledgements The work was conducted in the laboratory of Dr. Mark E. Nelson (University of Illinois). We would like to thank him for his support, and the comments and suggestions he provided at various stages of the project. We also thank Noura Sharabash and Dr. Zhian Xu for help in collecting data, and Dr. Pim van Dijk and Dr. Rob de Ruyter van Steveninck for

comments on the manuscript. This work was supported by grants from NIMH (R01-MH49242) and NSF (IBN-0078206) to M.N. The experiments comply with the *Principles of animal care*, publication No. 86-23, revised 1985, of the National Institute of Health and also with the current laws of the United States of America.

Appendix

Performance comparison between sequential and trial-based testing

Here we compare the performance of trial-based detection and sequential detection with dead-time, and discuss the effect of correlations in the spike train on the performance of the sequential detector. As indicated earlier, a direct comparison of the performance of trial-based testing and sequential testing is not possible. The decision-making processes are different, since in sequential testing resetting the detector when a threshold is exceeded and restarting the process, truncates the PDF under H_0 at the threshold value. Consequently, calculating d' , σ^2 and P_{fa} is not possible in sequential testing. This is in contrast to trial-based testing where shape of the PDFs under both hypotheses is independent of the threshold value. The class of problems where testing is stopped and restarted after the threshold is reached, or where the integrator is reset, are random walk level-crossing problems. The theory underlying these problems has wide applicability and has received much attention in physics, engineering and statistics (Blake and Lindsey 1973; Leadbetter et al. 1983), but closed-form solutions are not available, even for relatively simple cases such as renewal spike trains. Given these difficulties, our interest is to ask whether there is an equivalent but simpler process that captures the performance of the sequential scheme, and allows us to calculate the SD of the PDF and P_{fa} . This may allow us to explain the performance improvement obtained in sequential testing and the role of temporal dependencies in the spike train on these improvements.

To simplify the sequential testing scheme shown in Fig. 3C, we divide $y[n]$ into non-overlapping blocks of length T and examine whether there is at least one hit within this block. If the block contains a signal then the hit is a correct detection otherwise it is a false alarm. Let P_d^* denote the detection probability and P_{fa}^* the false alarm probability for the simplified scheme. To calculate the probability of detection or false alarm within a block, we only have to examine the maximum value of the T samples within the window. Thus, we can construct a sequence $y_T[k]$, where k ($k \geq 0$) is the number of blocks of length T as follows:

$$\begin{aligned} H_0 : y_T[k] &= \max_{0 \leq i < T} (y[kT + i]) \quad (\text{noise only}) \\ H_1 : y_T[k] &= \max_{0 \leq i < T} (y[kT + i] + A) \quad (\text{signal + noise}) \end{aligned} \quad (3)$$

where i are the individual samples in the block. Signal detectability will be improved if the SD of the PDFs for $y_T[k]$ is smaller than for $y[n]$, since A is the same for both cases.

Now we can construct the PDFs of $y_T[k]$. It can be shown that if x_i is a sequence of independent random (i.i.d.) variables where $0 \leq i < T$, the distribution of the random variable $x = \max_{1 \leq i < T} \{x_i\}$ is $F(x) = \prod_{i=1}^T F_i(x)$, where $F_i(x_0)$ is the probability that $x < x_0$ for the i^{th} sample and $F(x_0)$ is the probability that $x < x_0$ for at least one of the samples within the block. The right tail $P(x) = 1 - F(x)$. Therefore, if $P_{fa,i}$ and $P_{d,i}$ are the probabilities of a false alarm or detection in a single EOD cycle i within the block, then the probability of detection P_d^* is:

$$P_d^* = 1 - \prod_{i=1}^T (1 - P_{d,i}) \quad (4)$$

The probability of a false alarm is:

$$P_{fa}^* = 1 - \prod_{i=1}^T (1 - P_{fa,i}) \quad (5)$$

For a given threshold, it can be shown that the probabilities (P_d^* and P_{fa}^*) of detecting at least one threshold crossing in a block, are larger than the probabilities of detecting a threshold crossing in one cycle ($P_{fa,i}$ and $P_{d,i}$).

Figure 12A shows the OC for the representative afferent and its matched binomial model, for the sequential decision method (solid line), and the simplified block-based method (dashed line). For the afferent the OC is the same in both cases, whereas for the binomial model the simplification underestimates the performance slightly. For renewal spike trains, the deviation arises due to the presence of bursts of false alarms. If a burst of threshold crossings occurs, then in the sequential scheme, the dead-time is always triggered on the first threshold crossing. This is not the case for the block-based scheme and thus bursts are more likely to spread out over two blocks, and hence the number of false alarms is somewhat higher. For the afferent, bursts of false alarms only occur when P_d^* is very close to 1 (Fig. 9), but for the models, particularly the binomial model, bursts occur more frequently even for small P_d^* values. Hence, for renewal spike trains, the block-based methods underestimates P_d^* for a given P_{fa}^* , and is not as good an approximation of the sequential scheme.

With this simplification we can compare the performance of the block-based method with the trial-based method. Figure 12B, C shows the ROCs for the representative afferent and binomial model, where P_d^* and P_{fa}^* follow Eqs. 4 and 5 (solid lines), compared to the trial-based method (dash-dotted line). The figure shows that for the afferent there is a large improvement in performance compared to the trial-based method. For example, at a P_{fa} of 0.013, which corresponds to a false alarm

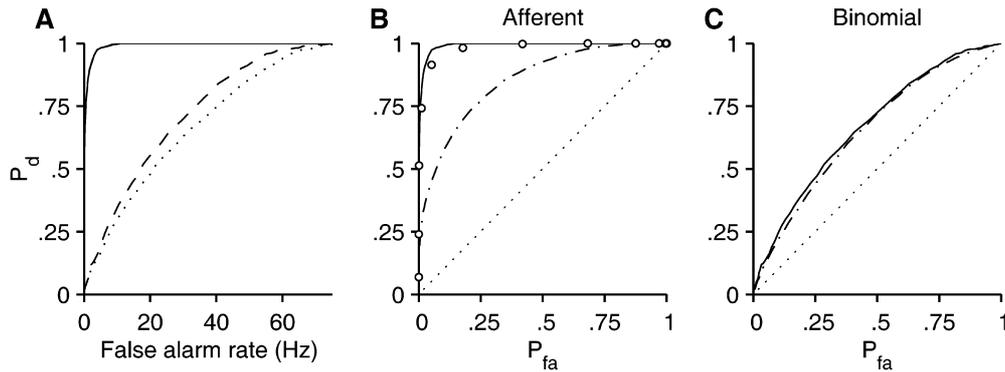


Fig. 12A–C Comparison of trial based and sequential detection schemes. **A** The sequential scheme (Fig. 3C) is compared to a simplified block based scheme to test the validity of the simplification. The figure shows the representative afferent in Fig. 5 and its matched binomial model (afferent sequential: *solid line*, afferent block-based *overlaps afferent sequential*, binomial sequential: *dashed line*, binomial block-based: *dotted line*). The simplification yields similar performance for the afferent (the curves are superimposed), but the binomial model shows some differences. The performance of the M_0 and M_1 models is intermediate (not shown). **B** ROCs for the simplified scheme (*solid line*) compared with the ROCs obtained from trial-based detection (*dash-dot line*) for the afferent (*dotted line*: chance). Sequential detection provides a large performance improvement over trial-based detection. The *open circles* show the performance for the case in which all output samples are independent (shuffled filter output), and the performance is nearly the same as for the afferent, i.e., the filter output samples are seemingly decorrelated (Fig. 11). **C** Results for the binomial model (as in **B**). The ROCs are nearly identical for the binomial model. Thus, sequential detection results in improved performance when input spike trains are negatively correlated, but may not provide much benefit when spike trains are binomial (Poisson)

rate of 1 Hz, P_d for detecting a spike is 0.83 for the afferent using the sequential scheme while the P_d for the trial-based scheme is 0.29. For the binomial model there is hardly any improvement. For this fiber the SD of $y_T[k]$ is 0.21 (afferent), 1.04 (B), 0.67 (M_0) and 0.34 (M_1). The discriminability d' is 3.25 (afferent), 0.63 (B), 1.01 (M_0) and 2.03 (M_1). The population means of the SD and d' are shown in Table 2.

To examine the role of the correlations in the filter output, the filter output was shuffled to remove serial dependencies. Since the filter introduces positive correlations in its output, the $y[n]$ are in general not independent within the time T , and therefore $P_{d,i}$ and $P_{fa,i}$ are not identical and not independent. If, on the other hand, the $y[n]$ are shuffled so as to remove dependencies, then the samples are independent and identically distributed, and the probabilities $P_{d,i}$ of detection for each cycle are equal (P_d), and P_d^* simplifies to:

$$P_d^* = 1 - (1 - P_d)^T \quad (6)$$

Similarly for P_{fa}^* :

$$P_{fa}^* = 1 - (1 - P_{fa})^T \quad (7)$$

Table 2 Discriminability d' and the SD of the PDF that is obtained by taking the maximum value attained by the leaky integrator output in each block of size T

	Filter output		Shuffled filter output
	d'	SD of $y_T[k]$	
Afferent	3.37 ± 1.06	0.22 ± 0.05	0.22 ± 0.05
M_1	2.56 ± 1.16	0.32 ± 0.11	0.28 ± 0.09
M_0	1.65 ± 1.07	0.56 ± 0.24	0.45 ± 0.19
B	0.72 ± 0.10	0.94 ± 0.08	0.66 ± 0.04

Shown are the mean \pm SD of the afferent population ($n=49$) and their matched models. The improved discriminability for the afferents over the models is due to the smaller SD. The right column shows the SD when the filter output is shuffled prior to dividing into blocks. For the models shuffling decreases the SD, whereas for the afferent shuffling did not result in a change in the SD, i.e., the filter output for the afferent is apparently decorrelated, as also seen in Fig. 11

The ROC for the shuffled afferent filter output (Eqs. 6 and 7, Fig. 12B, circles) is almost the same as the ROC for the block based afferent output (Eqs. 4 and 5, Fig. 12B, solid line). This suggests the afferent filter output samples behave as if they were uncorrelated. The SD of $y_T[k]$ for the shuffled output is 0.22 (representative afferent), the same as without shuffling (0.21). For the models, the SD decreases as a result of shuffling, as expected. The SD for the shuffled output is 0.69 (B), 0.51 (M_0), and 0.30 (M_1). Table 2 shows the population means. The improvement in the ROC (Fig. 12B) is reflected in the smaller SD of $y_T[k]$ compared to the SD of $y[n]$. This follows directly from Eqs. 6 and 7, from which it can be shown that the distribution of the maximum value of a set of i.i.d. random variables will always have a smaller variance than the original distribution.

References

- Amassian VE, Macy J, Waller HJ, Leader HS, Swift M (1964) Transformations of afferent activity at the cuneate nucleus. In: Gerard RW, Duyff J (eds) Information processing in the nervous system. Excerpta Medica Foundation, Amsterdam, pp 235–254

- Barlow HB, Levick WR, Yoon M (1971) Responses to single quanta of light in retinal ganglion cells of the cat. *Vision Res [Suppl]* 3:87–101
- Bastian J (1981) Electrolocation. I. How the electroreceptors of *Apteronotus albifrons* code for moving objects and other electrical stimuli. *J Comp Physiol A* 144:465–479
- Bastian J, Nguyenkim J (2001) Dendritic modulation of burst-like firing in sensory neurons. *J Neurophysiol* 85:10–22
- Bastian J, Chacron MJ, Maler L (2002) Receptive field organization determines pyramidal cell stimulus-encoding capability and spatial stimulus selectivity. *J Neurosci* 22:4577–4590
- Berman NJ, Maler L (1998) Inhibition evoked from primary afferents in the electrosensory lateral line lobe of the weakly electric fish (*Apteronotus leptorhynchus*). *J Neurophysiol* 80:3173–3196
- Berry MJ II, Meister M (1998) Refractoriness and neural precision. *J Neurosci* 18:2200–2211
- Blake IF, Lindsey WC (1973) Level-crossing problems for random processes. *IEEE Trans Inform Theory* 19:295–315
- Blanks RHI, Estes MS, Markham CH (1974) Physiologic characteristics of vestibular first-order canal neurons in the cat. II. Response to constant angular acceleration. *J Neurophysiol* 38:1250–1268
- Borst A, Theunissen FE (1999) Information theory and neural coding. *Nat Neurosci* 2:947–957
- Bugmann G, Christodoulou C, Taylor JG (1997) Role of temporal integration and fluctuation detection in the highly irregular firing of a leaky integrator neuron model with partial reset. *Neural Comput* 9:985–1000
- Bullock TH, Chichibu S (1965) Further analysis of sensory coding in electroreceptors of electric fish. *Proc Natl Acad Sci USA* 54:422–429
- Bullock TH, Heiligenberg W (1986) *Electroreception*. Wiley, New York
- Buračas GT, Albright TD (1999) Gauging sensory representations in the brain. *Trends Neurosci* 22:303–309
- Celebrini S, Newsome WT (1994) Neuronal and psychophysical sensitivity to motion signals in extrastriate area MST of the macaque monkey. *J Neurosci* 14:4109–4124
- Chacron MJ, Longtin A, St-Hilaire M, Maler L (2000) Suprathreshold stochastic firing dynamics with memory in P-type electroreceptors. *Phys Rev Lett* 21:5328–5343
- Chacron MJ, Longtin A, St-Hilaire M, Maler L (2001) Negative interspike interval correlations increase the neuronal capacity for encoding time dependent stimuli. *J Neurosci* 21:5328–5343
- Chacron MJ, Doiron B, Maler L, Bastian J (2003) Non-classical receptive field mediates switch in a sensory neuron's frequency tuning. *Nature* 423:77–81
- Cox DR (1962) *Renewal theory*. Methuen, London
- DeWeese M (1996) Optimization principles for the neural code. *Network* 7:325–331
- Fitzhugh R (1957) The statistical detection of threshold signals in the retina. *J Gen Physiol* 40:925–948
- Gabbiani F, Koch C (1998) Principles of spike train analysis. In: Koch C, Segev I (eds) *Methods in neuronal modeling: from ions to networks*. MIT, Cambridge, pp 313–360
- Gabbiani F, Metzner W, Wessel R, Koch C (1996) From stimulus encoding to feature extraction in weakly electric fish. *Nature* 384:564–567
- Gautrais J, Thorpe SJ (1998) Rate coding versus temporal order coding: a theoretical approach. *Biosystems* 48:57–65
- Geisler WS, Albrecht DG, Salvi RJ, Saunders SS (1991) Discrimination performance of single neurons: rate and temporal-pattern information. *J Neurophysiol* 66:334–362
- Goense JBM, Ratnam R, Nelson ME (2003) Burst firing improves the detection of weak signals in spike trains. *Neurocomput* 52–54:103–108
- Gold JI, Shadlen MN (2000) Representation of a perceptual decision in developing oculomotor commands. *Nature* 404:390–394
- Gold JI, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. *Trends Cogn Sci* 5:10–16
- Goldberg JM (2000) Afferent diversity and the organization of central vestibular pathways. *Exp Brain Res* 130:277–297
- Goldberg JM, Fernandez C (1971) Physiology of peripheral neurons innervating semicircular canals of the squirrel monkey. III. Variations among units in their discharge properties. *J Neurophysiol* 34:676–684
- Green D, Swets J (1966) *Signal detection theory and psychophysics*. Wiley, New York.
- Hagiwara S, Morita H (1963) Coding mechanisms of electroreceptor fibers in some electric fish. *J Neurophysiol* 28:784–799
- Heiligenberg W, Dye J (1982) Labeling of electroreceptive afferents in a gymnotic fish by intracellular injection of HRP: the mystery of multiple maps. *J Comp Physiol* 148:287–296
- Hernández A, Zainos A, Romo R (2002) Temporal evolution of a decision-making process in medial premotor cortex. *Neuron* 33:959–972
- Kashimori Y, Goto M, Kambara T (1996) Model of P- and T-electroreceptors of weakly electric fish. *Biophys J* 70:2513–2526
- Kepecs A, Wang X-J, Lisman J (2002) Bursting neurons signal input slope. *J Neurosci* 22:9053–9062
- Kim J-N, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2:176–185
- Knudsen EI (1974) Behavioral thresholds to electric signals in high frequency electric fish. *J Comp Physiol* 91:333–353
- Kreiman G, Krahe R, Metzner W, Koch C, Gabbiani F (2000) Robustness and variability of neuronal coding by amplitude sensitive afferents in the weakly electric fish *Eigenmania*. *J Neurophysiol* 84:189–204
- Kuffler SW, Fitzhugh R, Barlow HB (1957) Maintained activity in the cat's retina in light and darkness. *J Gen Physiol* 40:683–702
- Lánský P, Musila M (1991) Variable initial depolarization in Stein's neuronal model with synaptic reversal potentials. *Biol Cybern* 64:285–291
- Leadbetter MR, Lindgren G, Rootzén (1983) *Extremes and related properties of random sequences and processes*. Springer, Berlin Heidelberg New York
- Lee BB, Wehrhahn C, Westheimer G, Kremers J (1993) Macaque ganglion cell responses to stimuli that elicit hyperacuity in man: detection of small displacements. *J Neurosci* 12:1001–1009
- Lowen SB, Teich MC (1992) Auditory-nerve action potentials form a nonrenewal point process over short as well as long time scales. *J Acoust Soc Am* 92:803–806
- Metzner W, Koch C, Wessel R, Gabbiani F (1998) Feature extraction by burst-like spike patterns in multiple sensory maps. *J Neurosci* 18:2283–2300
- Nelson ME, MacIver MA (1999) Prey capture in the weakly electric fish *Apteronotus albifrons*: sensory acquisition strategies and electrosensory consequences. *J Exp Biol* 202:1195–1203
- Nelson ME, Xu Z, Payne JR (1997) Characterization and modeling of P-type electrosensory afferent responses to amplitude modulations in a wave-type electric fish. *J Comp Physiol A* 181:532–544
- Panzeri S, Schultz SR (2001) A unified approach to the study of temporal, correlational and rate coding. *Neural Comput* 13:1311–1349
- Parker AJ, Newsome WT (1998) Sense and the single neuron: probing the physiology of perception. *Annu Rev Neurosci* 21:227–277
- Ratnam R, Nelson ME (2000) Non-renewal statistics of electrosensory afferent spike trains: implications for the detection of weak sensory signals. *J Neurosci* 20:6672–6683
- Ratnam R, Goense JBM, Nelson ME (2001) The response of P-type electrosensory afferents to weak prey-like stimuli. 6th Int Congr Neuroethol, Bonn
- Reddi BA, Carpenter RH (2000) The influence of urgency on decision time. *Nat Neurosci* 3:827–30
- Reich DS, Victor JD, Knight BW, Ozaki T, Kaplan E (1997) Response variability and timing precision of neuronal spike trains *in vivo*. *J Neurophysiol* 77:2836–2841

- Relkin EM, Pelli DG (1987) Probe tone thresholds in the auditory nerve measured by two-interval forced-choice procedures. *J Acoust Soc Am* 82:1679–1691
- Rieke F, Warland D, Ruyter van Steveninck RR de, Bialek W (1997) *Spikes: exploring the neural code*. MIT Press, Cambridge
- Rolls ET, Tové MJ, Panzeri S (1999) The neurophysiology of backward visual masking: information analysis. *J Cogn Neurosci* 11:300–311
- Rossum M van (2001) A novel spike distance. *Neural Comput* 13:751–763
- Ruyter van Steveninck R de, Bialek W (1995) Reliability and statistical efficiency of a blowfly movement-sensitive neuron. *Philos Trans R Soc Lond B* 348:321–340
- Ruyter van Steveninck R de, Lewen GD, Strong SP, Koberle R, Bialek W (1997) Reproducibility and variability in neural spike trains. *Science* 275:1805–1808
- Schiller PH, Logothetis NK (1990) The color-opponent and broad-band channels of the primate visual system. *Trends Neurosci* 13:392–398
- Shadlen M, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896
- Shofner WP, Dye RH (1989) Statistical and receiver operating characteristic analysis of empirical spike count distributions: quantifying the ability of cochlear nucleus units to signal intensity changes. *J Acoust Soc Am* 86:2172–2184
- Shumway CA (1989a) Multiple electrosensory maps in the medulla of weakly electric gymnotiform fish. I. Physiological differences. *J Neurosci* 9:4388–4399
- Shumway CA (1989b) Multiple electrosensory maps in the medulla of weakly electric gymnotiform fish. II. Anatomical differences. *J Neurosci* 9:4400–4414
- Siegmund D (1985) *Sequential analysis: test and confidence intervals*. Springer, Berlin Heidelberg New York
- Stein RB (1967) The information capacity of nerve cells using a frequency code. *Biophys J* 7:797–826
- Steuer R, Ebeling W, Russell DF, Bahar S, Neiman A, Moss F (2001) Entropy and local uncertainty of data from sensory neurons. *Phys Rev E* 64:061911, 1–6
- Thorpe SJ, Fize D, Marlot C (1996) Speed of processing in the human visual system. *Nature* 381:520–522
- Tougaard J (1999) Receiver operating characteristics and temporal integration in an insect auditory receptor cell. *J Acoust Soc Am* 106:3711–3718
- Tricas TC, New JG (1998) Sensitivity and response dynamics of elasmobranch electrosensory primary afferent neurons to near threshold fields. *J Comp Physiol A* 182:89–101
- Vallbo AB (1995) Single-afferent neurons and somatic sensation in humans. In: Gazzaniga MS (ed) *The cognitive neurosciences*. MIT Press, Cambridge, pp 237–252
- Vallbo AB, Johansson (1976) Skin mechanoreceptors in the human hand: neural and psychophysical thresholds. In: Zotterman Y (ed) *Sensory functions of the skin*. Pergamon Press, Oxford, pp 185–199
- VanRullen R, Thorpe SJ (2001a) The time course of visual processing: from early perception to decision-making. *J Cogn Neurosci* 13:454–461
- VanRullen R, Thorpe SJ (2001b) Rate coding versus temporal order coding: what the retinal ganglion cells tell the visual cortex. *Neural Comput* 13:1255–1283
- Wald A (1947) *Sequential analysis*. Wiley, New York
- Xu Z, Payne JR, Nelson ME (1996) Logarithmic time course of sensory adaptation in electrosensory afferent nerve fibers in a weakly electric fish. *J Neurophysiol* 76:2020–2032