

## Short-Term Synaptic Plasticity Can Enhance Weak Signal Detectability in Nonrenewal Spike Trains

Niklas Lüdtke

*Niklas.Ludtke@manchester.ac.uk*

Mark E. Nelson

*m-nelson@uiuc.edu*

*Beckman Institute, University of Illinois at Urbana-Champaign, Urbana, IL 61801, U.S.A*

We study the encoding of weak signals in spike trains with interspike interval (ISI) correlations and the signals' subsequent detection in sensory neurons. Motivated by the observation of negative ISI correlations in auditory and electrosensory afferents, we assess the theoretical performance limits of an individual detector neuron receiving a weak signal distributed across multiple afferent inputs. We assess the functional role of ISI correlations in the detection process using statistical detection theory and derive two sequential likelihood ratio detector models: one for afferents with renewal statistics; the other for afferents with negatively correlated ISIs. We suggest a mechanism that might enable sensory neurons to implicitly compute conditional probabilities of presynaptic spikes by means of short-term synaptic plasticity. We demonstrate how this mechanism can enhance a postsynaptic neuron's sensitivity to weak signals by exploiting the correlation structure of the input spike trains. Our model not only captures fundamental aspects of early electrosensory signal processing in weakly electric fish, but may also bear relevance to the mammalian auditory system and other sensory modalities.

### 1 Introduction ---

In response to a growing body of experimental studies on short-term synaptic plasticity (Tsodyks & Markram, 1997; Zucker & Regehr, 2002; Xu-Friedman & Regehr, 2004; Blitz, Foster, & Regehr, 2004), models of nonlinear synaptic transmission have emerged that emphasize the functional importance of the relative timing of presynaptic action potentials (Maass & Zador, 1999; Markram, 2003; Abbott & Regehr, 2004). Sensitivity of synaptic transmission to the history of presynaptic spikes becomes especially relevant when interspike intervals (ISIs) are correlated. Such input correlations have been reported in various sensory modalities (Lowen & Teich, 1992; Teich, Turcott, & Siegel, 1996; Bahar et al., 2001) and are particularly prominent in

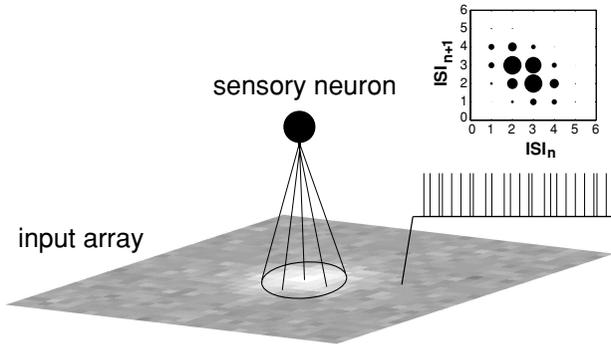


Figure 1: Schematic representation of a sensory neuron receiving afferent input from an array of spiking units. The stimulus intensity is encoded in the afferent firing rates. The noisy pattern of gray-scale values in the input array depicts spike counts obtained within a fixed time window. The insets show a spike train of an individual afferent and the corresponding joint ISI histogram, revealing negative ISI correlations. Short ISIs tend to be followed by longer ISIs and vice versa. Since the firing activity is stochastic, the spike counts exhibit variability. For temporally correlated ISIs, the spike count variability depends on the count window length.

the active electrosensory system of weakly electric fish (Ratnam & Nelson, 2000).

The negative ISI correlations observed in the electrosensory system have been successfully modeled using an integrate-and-fire type mechanism with threshold fatigue (Chacron, Maler, & Longtin, 2001; Brandman & Nelson, 2002). In these models, the firing threshold is elevated following an action potential and subsequently decays toward a baseline level. With appropriate parameters, the ISI sequence exhibits negative correlations, that is, short ISIs tend to be followed by longer ISIs and vice versa.

**1.1 Detection of Weak Signals.** In this article we explore the interplay between ISI correlations and fast synaptic plasticity in sensory systems that respond to extremely weak stimuli by detecting small changes in afferent spike activity.

Figure 1 shows a generalized scheme of a detector neuron receiving input from a receptor array via afferent fibers. The gray-scale intensity represents stimulus intensity encoded as a noisy, localized pattern of afferent firing rates above a baseline activity. Throughout this article, we assume there are no interconnections among the afferent fibers and hence no spatial correlations in the activation pattern other than those due to the stimulus. We are particularly interested in situations where the signal-to-noise ratio is low, such that the stimulus-induced change in activity is small compared to

intrinsic fluctuations of sensory afferent activity. Under these circumstances, the detector neuron has to perform spatiotemporal averaging to enhance the signal relative to the stochastic background activity.

**1.2 Sensory Neurons as Likelihood Ratio Detectors.** The accumulation and evaluation of evidence conveyed by the activities of multiple afferents can be assessed using a key concept from statistical detection theory: the likelihood ratio. This quantity compares the probability of the afferent activity in the presence of a stimulus (baseline plus signal) and without a stimulus (pure baseline) by computing the ratio

$$\frac{p(\text{afferents; signal + baseline})}{p(\text{afferents; baseline})}.$$

Using a more formal notation, we introduce the detection-theoretic definition of an optimal detector. Let  $\vec{a}$  be the input data vector (afferent activities). The ratio of data likelihoods under assumptions of signal presence (hypothesis  $H_1$ ) and absence (null hypothesis  $H_0$ ), respectively, is compared to a threshold  $\gamma$ :

$$\Lambda = \ln \frac{p(\vec{a}; H_1)}{p(\vec{a}; H_0)} > \gamma. \quad (1.1)$$

The detector decides in favor of  $H_1$  if the threshold is exceeded. The choice of threshold value determines the probability of a false alarm. According to the Neyman-Pearson theorem (Kay, 1998), the likelihood ratio detector is optimal in the sense that it achieves maximum probability of detection for a given false alarm probability.

Although the Neyman-Pearson theorem holds for any monotonic function of the likelihood ratio, one usually takes the logarithm. Under the independence assumption, the ensemble activity can be expressed as a product of individual afferent probabilities, and the logarithm of the product of probabilities then transforms into a sum of log likelihoods. Moreover, the logarithm of the likelihood ratio can be written as the difference of the log likelihoods.

Temporal evidence accumulation can be accomplished by summing the likelihood ratios obtained at different time instances. This procedure is commonly referred to as the cumulative sum (CUSUM) algorithm (Page, 1954) and is a repeated sequential likelihood ratio test (Wald, 1948). The resultant quantity is the cumulative log likelihood ratio, denoted by  $\Lambda_{\text{cum}}$  and defined in a recursive fashion:

$$\Lambda_{\text{cum}}[k + 1] = \max\{\Lambda_{\text{cum}}[k] + \Lambda[k], 0\}; \quad \Lambda_{\text{cum}}[0] = 0. \quad (1.2)$$

Note that the cumulative log-likelihood ratio  $\Lambda_{\text{cum}}$  undergoes rectification. The detector decides  $H_1$  at time step  $k$  if

$$\Lambda_{\text{cum}}(k) > \tilde{\gamma},$$

where  $\tilde{\gamma}$  is a threshold that determines the false alarm rate.

Intuitively, the rectification in equation 1.2 seems advantageous, since it limits the accumulation of negative evidence, allowing a faster recovery of the cumulative likelihood ratio once  $\Lambda[k]$  turns positive. In addition, it has been formally proved (Moustakides, 1986) that this scheme is indeed optimal in the sense that for a given false alarm rate, the CUSUM algorithm exhibits the shortest average detection delay.

It has been demonstrated that neurons have the ability to carry out likelihood ratio computations (Gold & Shadlen, 2001). In the context of decision making and motion perception, evaluating the logarithm of a likelihood ratio is equivalent to calculating the difference in firing rate of two neurons with opposite preferred directions of motion, provided the neural responses are described by normal, Poisson, or exponential densities (Gold & Shadlen, 2001). Furthermore, if the responses are independent and identically distributed over time, temporal accumulation of evidence, as in the CUSUM procedure, can be accomplished by an integrate-and-fire neuron (Gold & Shadlen, 2002).

**1.3 Spike Train Statistics and Detection Performance.** To assess theoretical detection performance limits of an individual sensory neuron for different input spike train statistics, we first derive a CUSUM detector model based on equation 1.2 for afferents with renewal spike statistics. Though designed for a renewal process, this type of detector can also operate with temporally correlated input. Detection performance has been shown to improve in the presence of negative ISI correlations, since a detector can passively benefit from the reduced spike count variability caused by the correlations (Ratnam & Nelson, 2000; Chacron et al., 2001; Goense & Ratnam, 2003).

The main contribution of this article is a more sophisticated detector model that actively utilizes the redundancy in temporally correlated input. Rather than matching the average firing probability, this detector operates with an estimate of the current firing probability at each time instance. The interdependence of adjacent ISIs requires the use of conditional likelihoods dependent on the spike train history, which poses a computational challenge and raises the question of how conditional firing probabilities could be represented in neural systems. We suggest that a record of afferent spike train history can be kept implicitly in the form of short-term synaptic plasticity and demonstrate that such synaptic plasticity would enable a sensory neuron to robustly track the fluctuations of the presynaptic firing probabilities.

Thus, the incoming evidence can be evaluated in terms of current conditional log-likelihood ratios, resulting in more efficient detection of weak signals.

**1.4 Relation to Electrosensory Signal Processing.** Our simplified detector model is motivated by the study of electrosensory prey detection in weakly electric fish (Nelson & MacIver, 1999). The model captures the fundamental aspects of the feedforward pathway at the first stage of electrosensory processing, in the electrosensory lateral line lobe (ELL), which receives electrosensory afferent input.

Typical electrosensory stimuli induced by small prey are localized and extremely weak perturbations of the fish's self-generated electric field. The resultant minute changes in transdermal potential are sensed by cutaneous electroreceptors and encoded in the activities of primary electrosensory afferent nerve fibers. Due to the small signal amplitude and the variability of afferent firing, the fish must solve a challenging detection task. It is estimated that small prey (such as *Daphnia magna*), at a typical detection distance of 2 cm from the skin, will cause only about one extra spike above a background of 60 spikes within a 200 ms interval (Ratnam & Nelson, 2000).

Our simulations show that detection performance is enhanced through the interplay between ISI correlations and synaptic plasticity in a model neuron. We speculate that a neural correlate of this mechanism could be implemented by short-term plasticity at ELL excitatory afferent synapses onto ELL pyramidal neurons.

**1.5 Outline.** First, we formulate the prey detection task within the framework of statistical detection theory and establish a link between the key concept of the likelihood ratio and the integrate-and-fire model of a sensory neuron. Next, we develop two alternative sequential likelihood ratio detector models—one for uncorrelated and the other for correlated spike trains. We then specify the electrosensory signal of a prey-like stimulus in a simplified cylindrical geometry, derive the corresponding models for a putative electrosensory detector neuron, and compare their performance in the presence or absence of negative correlations in the input spike trains. Finally, we discuss the implications of our results with regard to a possible role of short-term synaptic plasticity in the enhanced detection of weak signals encoded in correlated spike trains.

## 2 A Log-Likelihood Ratio Detector for Binomial Spike Trains

Let  $a_i[k]$  be the spike state of the  $i$ th input fiber at the current time step  $k$ , and  $\vec{a}[k]$ , the spike state of an ensemble of  $n$  fibers:

$$a_i[k] \in \{0, 1\}; \quad \vec{a}[k] \in \{0, 1\}^n.$$

Assuming independence, firing probabilities for individual afferents can be multiplied, and the logarithm of the likelihood ratio is then given by

$$\begin{aligned} \Lambda &= \log \frac{P(\vec{a}[k]; H_1)}{P(\vec{a}[k]; H_0)} = \log \prod_{i=1}^n \frac{P(a_i[k]; H_1)}{P(a_i[k]; H_0)} \\ &= \sum_{i=1}^n \log P(a_i[k]; H_1) - \log P(a_i[k]; H_0). \end{aligned} \quad (2.1)$$

In electrosensory afferents, the independence assumption is well justified due to the absence of interconnections between afferent ganglion cells (Maler & Berman, 1999).

A simple spiking model is the binomial probability encoder, which represents the signal amplitude by a proportional change in firing probability at each time step. The resultant spike train is a renewal process: the ISIs are independent random variables. The likelihood of an individual afferent spike state  $a_i$  is

$$\log P(a_i[k]; H_1) = \begin{cases} \log[r_{\text{base}}\Delta t + s_i g \Delta t] & \text{if } a_i[k] = 1 \\ \log[1 - r_{\text{base}}\Delta t - s_i g \Delta t] & \text{if } a_i[k] = 0 \end{cases}, \quad (2.2)$$

where  $r_{\text{base}}$  is the baseline firing rate,  $s_i$  the signal strength at the  $i$ th receptor site,  $g$  the gain, and  $\Delta t$  the duration of a time step.

If the increment in firing rate caused by the signal,  $s_i g$ , is small compared to the baseline firing rate, the logarithm in equation 2.2 can be approximated using a first-order Taylor expansion:  $\log(x + x_0) \approx \log x_0 + x/x_0$ . Furthermore, one can combine the two rows by weighting the entries with  $a_i$  and  $1 - a_i$ , respectively. Hence,

$$\begin{aligned} \log P(a_i[k]; H_1) &\approx a_i[k] \left[ \log(r_{\text{base}}\Delta t) + \frac{g\Delta t s_i}{r_{\text{base}}\Delta t} \right] \\ &+ (1 - a_i[k]) \left[ \log(1 - r_{\text{base}}\Delta t) - \frac{g\Delta t s_i}{(1 - r_{\text{base}}\Delta t)} \right]. \end{aligned} \quad (2.3)$$

Accordingly, one can obtain the corresponding probability for the null hypothesis by setting the signal amplitude to zero ( $s_i = 0$ ):

$$\log P(a_i[k]; H_0) \approx a_i[k] \log(r_{\text{base}}\Delta t) + (1 - a_i[k]) \log(1 - r_{\text{base}}\Delta t). \quad (2.4)$$

Subtracting equation 2.4 from 2.3 yields:

$$\begin{aligned} &\log P(a_i[k]; H_1) - \log P(a_i[k]; H_0) \\ &= a_i[k] \frac{g\Delta t s_i}{r_{\text{base}}\Delta t} - (1 - a_i[k]) \frac{g\Delta t s_i}{(1 - r_{\text{base}}\Delta t)} \end{aligned}$$

$$\begin{aligned}
&= A a_i[k] \underbrace{g \Delta t \tilde{s}_i \left[ \frac{1}{r_{\text{base}} \Delta t} + \frac{1}{1 - r_{\text{base}} \Delta t} \right]}_{w_i} - A \underbrace{\frac{g \Delta t \tilde{s}_i}{(1 - r_{\text{base}} \Delta t)}}_{b_i} \\
&= A [a_i[k] w_i - b_i], \tag{2.5}
\end{aligned}$$

where  $A$  is the stimulus amplitude and  $\tilde{s}_i$  denotes the normalized signal strength at receptor site  $i$ , in the sense that  $A = \max\{s_i\}$  and  $s_i = A\tilde{s}_i$ . Therefore, the likelihood ratio of the entire ensemble state at time step  $k$  is given by

$$\log \frac{P(\bar{a}[k]; H_1)}{P(\bar{a}[k]; H_0)} = A \left[ \sum_{i=1}^n a_i[k] w_i - b \right], \text{ where } b = \sum_{i=1}^n b_i. \tag{2.6}$$

The resultant detector constitutes a linear filter with a fixed spatial receptive field, since the afferent inputs  $a_i$  are weighted according to the expected relative signal strength  $\tilde{s}_i$  at the associated receptor sites.

Note that linearity in  $a_i[k]$  is not caused by the linearization in signal strength. The Taylor approximation in  $s_i$  merely makes it possible to factor out the signal amplitude. Hence, the detector responds preferentially to a stimulus with particular spatial characteristics, and the absolute stimulus intensity determines the detection delay. Therefore, in all of our simulations, we focus on the spatial aspects of detection and restrict our analysis to stimuli with constant spatial characteristics and instantaneous onset. Adapting the detector to more complex stimuli with varying intensity and spatial extent would require a varying threshold and dynamic weights matched to the expected time course of the signal intensity at the receptor locations.

### 3 Spike Trains with Negative ISI Correlations

---

Correlations among neighboring ISIs can influence the spike count statistics on timescales well beyond that of the mean ISI. For instance, spike trains with negative ISI correlations can exhibit a lower spike count variability than a surrogate renewal process with identical ISI distribution or a Poisson process. Such long-term regularization has been observed in the auditory nerve (Lowen & Teich, 1992) and is particularly prominent in electrosensory afferents, where the Fano factor (variance to mean ratio) of the spike count can be reduced by more than an order of magnitude at behaviorally relevant time scales of 100 to 200 ms (Ratnam & Nelson, 2000). It has been suggested that such regularization enhances the encoding and detectability of weak signals (Ratnam & Nelson, 2000; Chacron et al., 2001; Brandman & Nelson, 2002; Goense & Ratnam, 2003). Therefore, it is expected that a naive detector based on equation 2.6 would passively benefit from the higher signal-to-noise ratio of regularized input spike trains. However, there is a potential additional benefit: ISI correlations imply a certain degree of

predictability (Raciot & Longtin, 1997), since individual spike events are statistically dependent on previous ISIs.

**3.1 Toward a Correlation-Sensitive Detector.** The nonrenewal property of the input spike trains should be accounted for in the calculation of the cumulative likelihood ratio. Since the firing probability is determined by previous spike events, one would have to consider conditional probabilities of spike states dependent on the spike train history. In the simplest case, the current firing probability depends on only the previous ISI, and the likelihood ratio would be of the form

$$\log \frac{P(\bar{a}[k] | I[k-1]; H_1)}{P(\bar{a}[k] | I[k-1]; H_0)},$$

where the conditional probabilities could be obtained from joint ISI histograms. However, since experimentally observed Markov orders of electrosensory afferent ISIs are at least five or greater (Ratnam & Nelson, 2000), one is faced with a dilemma: while it may be technically feasible to obtain estimates of conditional firing probabilities using higher-order joint histograms from baseline spike trains of sufficient length, such an approach would clearly be biologically implausible and thus provide no further insight into a possible physiological mechanism. Therefore, we propose a solution based on the nonlinear dynamics of the spike generating process.

**3.2 A Nonlinear Adaptive Threshold Model.** We introduce a generalization of the time-discrete afferent model by Brandman and Nelson (2002). As in most other models of electrosensory afferents, the firing threshold is raised following an action potential and decreases as long as no spike is generated. This type of threshold adaptation leads to negative ISI correlations, since the threshold level at each time step depends on its previous value, thus creating a memory of the spike train history. The equations of the Brandman-Nelson model are:

$$v[k] = c s[k] + W[k] \quad (3.1)$$

$$a[k] = \begin{cases} 1 & \text{if } v[k] > \theta[k] \\ 0 & \text{otherwise} \end{cases} \quad (3.2)$$

$$\theta[k+1] = \theta[k] - \beta/\alpha + a[k] \beta, \quad (3.3)$$

where  $v$  is the membrane potential,  $s$  the signal amplitude,  $W \sim \mathcal{N}(0, \sigma)$  the intrinsic gaussian white noise component,  $\theta$  the firing threshold, and  $c$  the gain. Beginning with an arbitrary initial value  $\theta[0]$ , the threshold decays by a fixed amount of  $\beta/\alpha$  during a time step and, if an action potential has been generated, is raised by amount  $\beta$ . For more biological realism, the

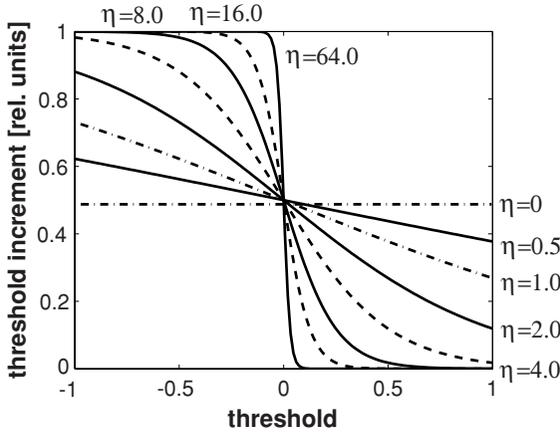


Figure 2: The sigmoid function governing the threshold increment for several different values of  $\eta$ . The increment is normalized by the factor  $2\beta$ . For extremely large  $\eta$ , the sigmoid approaches a step function; for  $\eta = 0$ , it reduces to a constant.

linear decay could be replaced with an exponential one (Chacron, Longtin, St. Hilaire, & Maler, 2000). Since  $\beta$  relates an action potential at step  $k$  to the firing threshold at step  $k + 1$ ,  $\beta$  controls the degree of correlation between subsequent ISIs.

In our modified version of the above model, we introduce a new refractory term in equation 3.3. Instead of a constant boost following an action potential, the threshold is now raised by a variable amount dependent on its current value  $\theta[k]$ :

$$\theta[k + 1] = \theta[k] - \beta/\alpha + a[k] \beta g\{\theta[k]\},$$

$$\text{where } g(\theta) = \frac{2e^{-\eta\theta}}{1 + e^{-\eta\theta}}. \tag{3.4}$$

If the threshold is low, it can be boosted by a maximum amount determined by the parameter  $\beta$ . If the threshold has reached a high level, the boost is smaller than  $\beta$ ; at a very high threshold level, virtually no further increase is possible and the decay exceeds the boost. The increase is governed by the sigmoid function  $g(\theta; \eta)$ . The saturation parameter  $\eta$  controls the degree to which the increase depends on threshold level (see Figure 2), thus also affecting ISI correlations. In the special case that  $\eta = 0$ , the sigmoid in equation 3.4 reduces to a constant,  $g(\theta) \equiv 1$ , resulting in the same constant boost as in the linear adaptive threshold model.

Level-dependent threshold boosting is also a feature of the (time continuous) model of Chacron, Pakdaman, and Longtin (2003). The difference is that instead of a linearly increasing function, in our model the threshold boost is governed by a nonlinear monotonically decreasing function of threshold level. Our assumption is biologically plausible, since physiological firing thresholds cannot increase arbitrarily and must eventually saturate. Though not biophysically detailed, our model exhibits spike train statistics that closely resemble those of actual afferents. Figure 3 compares the ISI statistics of the model to those of an extracellular afferent recording in the absence of electrosensory stimulation due to external objects. We use the term *baseline* rather than *spontaneous* activity since the electroreceptors are always driven by the continuously oscillating field generated by the fish's electric organ, which remains intact under anesthesia and immobilization.

In appendix A, we provide a detailed analysis of the influence of parameter  $\eta$  on the dynamics of threshold sequences. The properties of the iterative map defined by equation 3.4 play a crucial role in deriving an advanced likelihood ratio detector.

**3.3 The Firing Probability.** In the above model, the firing probability at time step  $k$  is given by the tail probability of the potential  $v$ , which is obtained by integrating the probability density function  $p(v)$ :

$$P[k] = P(v[k] > \theta[k]) = 1 - \int_{-\infty}^{\theta[k]} p(v) dv.$$

We will assume the noise to be gaussian of zero mean and variance  $\sigma$ . Since the potential  $v$  is the sum of the signal and the noise component, the pdf  $p(v)$  is also gaussian with the same variance but with mean  $\mu = cs$ . Equivalently, one can integrate a gaussian of zero mean and subtract the signal contribution from the threshold,

$$P[k] = \frac{1}{2} - \frac{1}{\sqrt{2\pi}\sigma} \int_0^{\theta[k]-cs[k]} e^{-v^2/2\sigma^2} dv,$$

where the gaussian integral

$$\frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^0 e^{-v^2/2\sigma^2} dv = \frac{1}{2}$$

has already been subtracted.

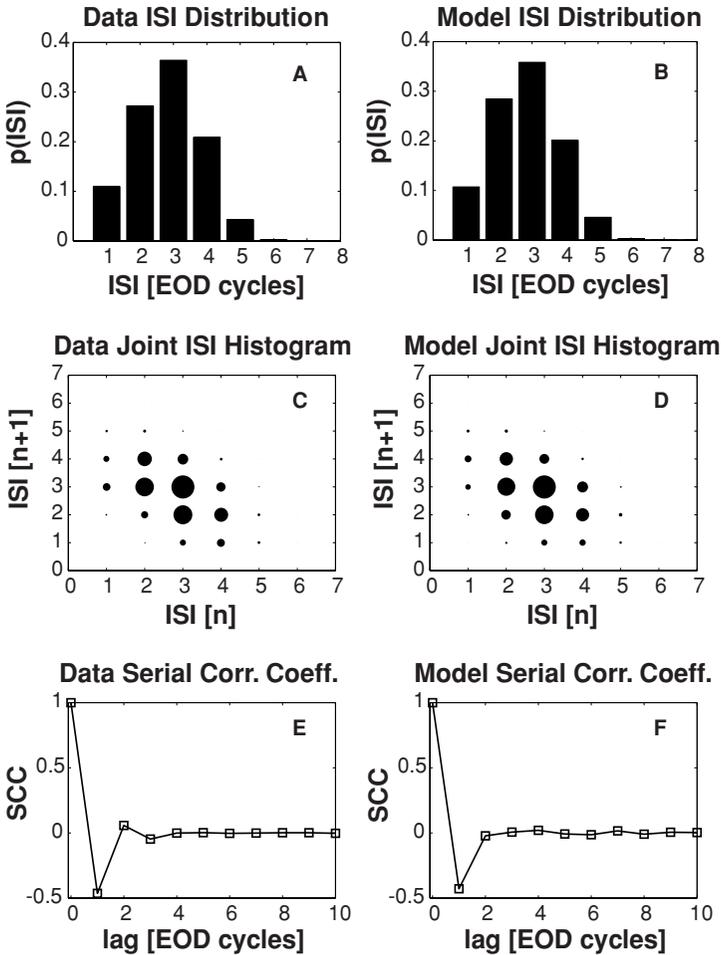


Figure 3: Comparison of the baseline interspike interval statistics of spike trains obtained from a tonic electroreceptive afferent (A, C, E) and the nonlinear adaptive threshold model (B, D, F). Model parameters were automatically optimized by maximum likelihood estimation (see section 3.7.1). The timescale is measured in units of the electric organ discharge (EOD) cycle, which is roughly 1 ms. The afferent spike train was recorded extracellularly from a weakly electric fish (*Apteronotus leptorhynchus*; data by Rama Ratnam).

After applying the transformation  $\tilde{v} = v/\sqrt{2}\sigma$ , the firing probability can be expressed in terms of the complementary error function (**erfc**):

$$P[k] = \frac{1}{2} - \frac{1}{\sqrt{\pi}} \int_0^{\frac{\theta[k] - c s[k]}{\sqrt{2}\sigma}} e^{-\tilde{v}^2} d\tilde{v} = \frac{1}{2} \text{erfc}\left(\frac{\theta[k] - c s[k]}{\sqrt{2}\sigma}\right). \quad (3.5)$$

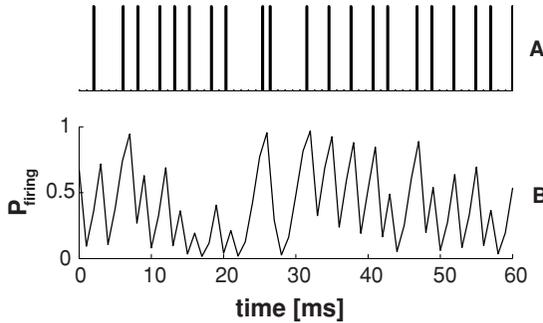


Figure 4: (A) Spike train generated by our nonlinear adaptive threshold model and (B) the corresponding firing probability. The time course of the firing probability resembles that of the postsynaptic current at a depressing synapse. Note that the very short timescale is not an inherent feature of the model; parameters were chosen to match the high baseline firing rates of electrosensory afferents.

Figure 4 shows a model spike train and the (discretized) time course of the corresponding firing probability, which bears a striking resemblance to the postsynaptic current at a depressing synapse, except at a much shorter timescale. Typical cortical time constants of short-term synaptic depression range in the hundreds of milliseconds (Zucker & Regehr, 2002). However, the timescale of the model is flexible and can be controlled via the parameter  $\alpha$  in equation 3.4, which determines the mean ISI and hence the baseline firing rate. In Figure 4, the baseline firing rate is set to approximately 300 Hz in order to match typical baseline firing rates of electrosensory afferents. Choosing a lower baseline firing rate would increase the timescale of the course of the firing probability accordingly. Therefore, we hypothesize that a form of short-term synaptic plasticity could enable a postsynaptic neuron to track the varying firing probabilities associated with its afferent input.

**3.4 The Log-Likelihood Ratio for Correlated Spike Trains.** In a similar manner as in section 2, one can derive an expression for the logarithm of the likelihood ratio for correlated model afferent spike trains. However, there is no longer a constant baseline firing rate, since the threshold is a dynamic variable influenced by previous spike events. Using equation 3.5, one can calculate the likelihood of individual afferent spike states  $a_i[k]$  under the two alternative hypotheses (denoted by  $H_x$ ). Instead of constant firing probabilities  $P(a_i[k]; H_x)$ , the likelihood ratio consists of conditional probabilities  $P(a_i[k] | \theta_i[k]; H_x)$ , and the threshold value  $\theta_i[k]$  contains a record of the entire spike train history up to time  $k$ . Thus, the logarithm of

the likelihood ratio has the following form:

$$\Lambda(\vec{a}[k]) = \sum_{i=1}^n \log P(a_i[k] | \theta_i[k]; H_1) - \log P(a_i[k] | \theta_i[k]; H_0). \quad (3.6)$$

The conditional probabilities  $P(a_i[k] | \theta_i[k]; H_x)$  replace more complicated higher-order Markov models with transition probabilities of the type

$$P(a_i[k] | a_i[k - 1], a_i[k - 2], \dots, a_i[k - m]; H_x),$$

where the spike train history is explicitly taken into account. In the subsequent sections, we describe a mechanism that enables the detector to track the changing thresholds  $\theta_i$  and thus the firing probabilities. By iteratively updating the estimate of the current threshold values,  $\theta_i[k]$ , arbitrary Markov orders can be taken into account implicitly, using the same formalism without increasing the model complexity. The entire spike train history can thus be absorbed into one variable.

Using equation 3.5, one can calculate the likelihood of the state of the  $i$ th afferent at time  $k$  under the signal hypothesis:

$$\begin{aligned} & \log P(a_i[k] | \theta_i[k]; H_1) \\ &= \begin{cases} \log \left[ \frac{1}{2} \operatorname{erfc} \left( \frac{\theta_i[k] - c s_i[k]}{\sqrt{2}\sigma_i} \right) \right] & \text{if } a_i[k] = 1 \\ \log \left[ 1 - \frac{1}{2} \operatorname{erfc} \left( \frac{\theta_i[k] - c s_i[k]}{\sqrt{2}\sigma_i} \right) \right] & \text{if } a_i[k] = 0 \end{cases}. \end{aligned} \quad (3.7)$$

Equation 3.7 can also be written as

$$\begin{aligned} \log P(a_i[k] | \theta_i[k]; H_1) &= a_i[k] \log \left[ \frac{1}{2} \operatorname{erfc} \left( \frac{\theta_i[k] - c s_i[k]}{\sqrt{2}\sigma_i} \right) \right] \\ &+ (1 - a_i[k]) \log \left[ 1 - \frac{1}{2} \operatorname{erfc} \left( \frac{\theta_i[k] - c s_i[k]}{\sqrt{2}\sigma_i} \right) \right]. \end{aligned} \quad (3.8)$$

Again, one can linearize the expressions, since the signal  $s_i[k]$  introduces only a small perturbation:  $c s_i[k] \ll \theta_i[k]$ . Hence,

$$\begin{aligned} & \log \left[ \frac{1}{2} \operatorname{erfc} \left( \frac{\theta_i[k] - c s_i[k]}{\sqrt{2}\sigma_i} \right) \right] \\ & \approx \log \frac{1}{2} + \log \operatorname{erfc} \left( \frac{\theta_i[k]}{\sqrt{2}\sigma_i} \right) - \frac{\operatorname{erfc}'(\theta_i[k]/\sqrt{2}\sigma_i)}{\operatorname{erfc}(\theta_i[k]/\sqrt{2}\sigma_i)} \frac{c s_i[k]}{\sqrt{2}\sigma_i} \\ & = \log \frac{1}{2} + \log \operatorname{erfc} \left( \frac{\theta_i[k]}{\sqrt{2}\sigma_i} \right) - \frac{\sqrt{2} c \exp(-\theta_i[k]^2/2\sigma_i^2)}{\sqrt{\pi} \sigma_i \operatorname{erfc}(\theta_i[k]/\sqrt{2}\sigma_i)} s_i[k], \end{aligned} \quad (3.9)$$

where the prime in the first row denotes the derivative of the complementary error function, which is defined as

$$\text{erfc}'(x) = \frac{d}{dx} \left\{ 1 - \frac{2}{\sqrt{\pi}} \int_0^x e^{-\tilde{x}^2} d\tilde{x} \right\} = -\frac{2}{\sqrt{\pi}} e^{-x^2}.$$

Analogously, one obtains a linearization of the second term in equation 3.8:

$$\begin{aligned} \log \left[ 1 - \frac{1}{2} \text{erfc} \left( \frac{\theta_i[k] - c s_i[k]}{\sqrt{2}\sigma_i} \right) \right] &\approx \\ \log \frac{1}{2} + \log \left[ 1 + \text{erf} \left( \frac{\theta_i[k]}{\sqrt{2}\sigma_i} \right) \right] & \\ - \frac{\sqrt{2} c \exp \left( -\theta_i[k]^2 / 2\sigma_i^2 \right)}{\sqrt{\pi}\sigma_i \left[ 1 + \text{erf} \left( \theta_i[k] / \sqrt{2}\sigma_i \right) \right]} s_i[k]. & \end{aligned} \tag{3.10}$$

For convenience, the regular error function **erf** is employed using the definition **erfc** = 1 - **erf**.

Inserting 3.9 and 3.10 into equation 3.8 yields the afferent log likelihood under the signal hypothesis. From this result, the log likelihood under the null hypothesis is obtained by setting the signal intensity to zero ( $s_i[k] = 0$ ). Given these log likelihoods, one obtains the logarithm of the likelihood ratio of the spike state vector of the afferent ensemble,  $\Lambda(\vec{a}[k])$ :

$$\begin{aligned} \Lambda(\vec{a}[k]) &= A \sqrt{\frac{2}{\pi}} \sum_{i=1}^n a_i[k] \\ &\quad \times \underbrace{\frac{c_i \tilde{s}_i[k]}{\sigma_i} \left( \frac{\exp \left( -\theta_i[k]^2 / 2\sigma_i^2 \right)}{1 - \text{erf} \left( \theta_i[k] / \sqrt{2}\sigma_i \right)} + \frac{\exp \left( -\theta_i[k]^2 / 2\sigma_i^2 \right)}{1 + \text{erf} \left( \theta_i[k] / \sqrt{2}\sigma_i \right)} \right)}_{w_i[k]} \\ &\quad - \underbrace{\frac{c_i \tilde{s}_i[k]}{\sigma_i} \frac{\exp \left( -\theta_i[k]^2 / 2\sigma_i^2 \right)}{1 + \text{erf} \left( \theta_i[k] / \sqrt{2}\sigma_i \right)}}_{b_i[k]} \\ &= A \sqrt{\frac{2}{\pi}} \sum_{i=1}^n a_i[k] w_i[k] - b_i[k], \end{aligned} \tag{3.11}$$

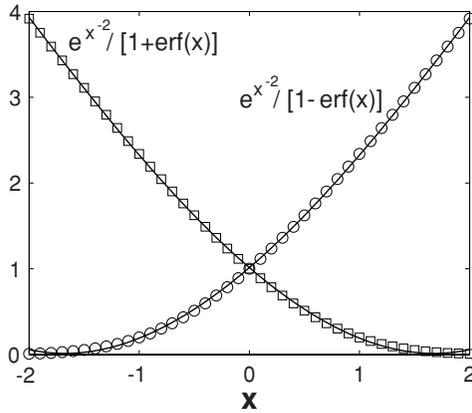


Figure 5: Plot of the function  $f(x) = e^{-x^2}/[1 - \text{erf}(x)]$  (circles) and its counterpart  $f(x) = e^{-x^2}/[1 + \text{erf}(x)]$  (squares), which appear in the dynamic synaptic weights and biases of the correlation-sensitive likelihood ratio in equation 3.11. Both functions are monotonic and well approximated by cubic polynomials (solid lines).

where  $A$  denotes the signal amplitude and  $\tilde{s}_i[k]$  is the normalized signal at receptor site  $i$ . The amplitude can thus be absorbed into the detector threshold  $\gamma$ . When integrating over time, the mean detection delay will be proportional to the amplitude if the signal amplitude is constant. Note that the synaptic weights and the bias terms are time dependent. Since they are functions of the firing threshold, they implicitly depend on the spike train history.

Although the analytical expressions for the weights and biases look complicated, they are smooth monotonic functions that can be well approximated by much simpler functions, such as the cubic polynomial shown in Figure 5. Thus, from the viewpoint of biological plausibility, the detector is not as computationally demanding as it may first appear.

**3.5 Threshold Prediction.** In order to construct a likelihood ratio detector based on equation 3.11, knowledge of the spike threshold  $\theta_i[k]$  is required at each time step. At first glance, this approach may seem infeasible since such information is not readily available. However, as we demonstrate in this section, it is possible to track the fluctuations of the firing threshold by feeding a model afferent spike train into a dynamical system that mimics the changes of the threshold in the spike-generating mechanism.

We introduce a predictor variable,  $\Theta$ , that is incremented or decremented, depending on whether a spike was received at the previous step, in the same

manner as in equation 3.4:

$$\Theta[k + 1] = \Theta[k] - \tilde{\beta}/\tilde{\alpha} + a[k] \tilde{\beta} \tilde{g}\{\Theta[k]\}. \quad (3.12)$$

Indices have been dropped for simplicity. The predictor variable is initialized with a random value,  $\Theta[0]$ , drawn from a gaussian distribution (though the choice of distribution is not critical). Equation 3.12 defines two alternative maps that describe how the subsequent threshold value is computed from its current value. The received spike train controls the choice of map. In effect,  $\Theta$  is a variable undergoing change in a dynamical system that randomly alternates between two deterministic components. If the parameters  $\tilde{\alpha}$ ,  $\tilde{\beta}$ , and  $\tilde{\eta}$  of the predictor system are identical to those of the spike generator and if this dynamical system has a stable orbit, the predictor sequence ( $\Theta[k]$ ) will converge toward the sequence of the thresholds ( $\theta[k]$ ) of the spike generator, regardless of the initial value  $\Theta[0]$ . Using equation 3.5,  $\Theta[k]$  can be transformed into the corresponding estimate of the firing probability,  $P_{\text{pred}}[k]$ . Figure 6A shows an example of such convergence of the predicted firing probability (dotted line) toward the actual firing probability used to generate the spike train (solid line). From approximately 20 time steps onward, the predictor sequence is tracking the actual firing probability accurately. The semilogarithmic plot of their difference, shown in Figure 6B, reveals an exponential convergence. It is this property that enables the predictor  $\Theta$  to track the actual firing threshold without knowledge of its initial value  $\theta[0]$ . Sensitivity to the initial condition would destroy this property so that the sequences could never synchronize.

Although the convergence in Figures 6A and 6B is demonstrated only for pure baseline activity, the predictor mechanism is not affected by the presence of a stimulus. A signal alters the membrane potential  $v$  in an additive fashion, thus changing the firing probability. However, the threshold predictor system does not require knowledge of the amount of change in signal amplitude  $s$ , since  $s$  does not appear explicitly in the transformation, equation 3.12. The predictor system implicitly receives information about the change of firing probability through the incoming spike train, thereby maintaining its ability to select the correct map to transform  $\Theta$  at any time instance.

**3.6 Synaptic Plasticity.** In order to illustrate the synaptic plasticity in our model detector neuron, we investigate the behavior of an individual input weight  $w_i[k]$  in equation 3.11 and the corresponding bias  $b_i[k]$  under stimulation with a test spike train, a sequence of bursts. This simulation resembles a typical neurophysiological test for short-term synaptic plasticity.

Figure 7 demonstrates the plasticity at an individual synapse under tetanic stimulation. Note that synaptic efficacy rather than a postsynaptic

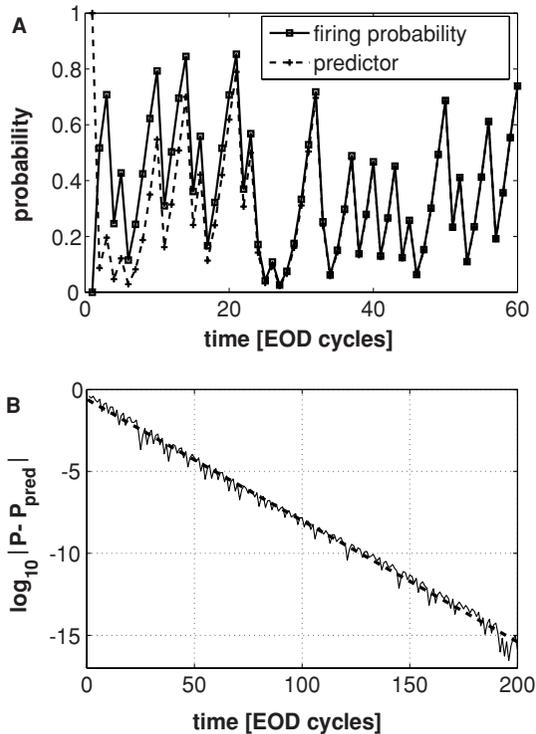


Figure 6: (A) Within about 20 time steps, the sequence of predicted firing probabilities (dashed line) has closely approached that of the spike-generating process (solid line). (B) The logarithm of the prediction error as a function of time reveals exponential convergence. The dashed line is a linear fit to the simulation data. The value of its slope closely approximates the Lyapunov exponent of the system (see appendix A) for the given set of parameters. A negative slope thus corresponds to a negative Lyapunov exponent, indicating convergence toward a stable orbit of the dynamical system.

current is plotted. Hence, this is a genuine nonlinear plasticity effect rather than a consequence of linear summation of overlapping excitatory postsynaptic potentials. The synaptic weight  $w_i$  exhibits facilitation (see Figure 7B), whereas the bias  $b_i$  shows rapid depression (see Figure 7C). The net contribution, the log-likelihood ratio for this afferent, is a combination of the two (see Figure 7D).

This example reveals how the detector interprets ISI sequences. Compared to the usual baseline activity with short ISIs followed by longer ones and vice versa, a tetanic burst is an unusual cluster of short ISIs, suggesting the presence of a stimulus. Hence, the likelihood ratio is increasingly

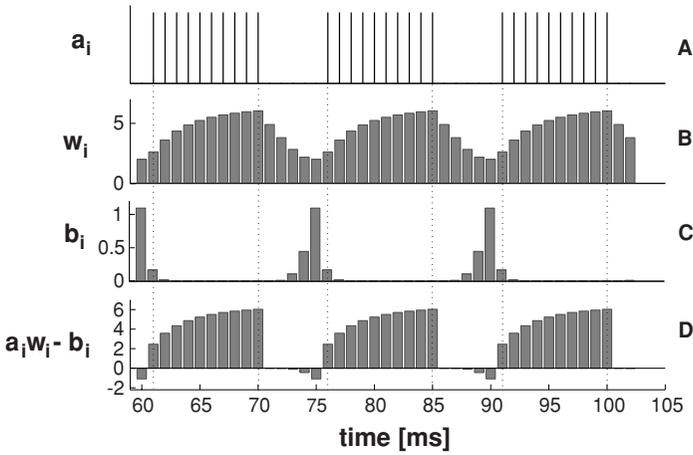


Figure 7: Behavior of one model synaptic weight and bias term under stimulation with tetanic bursts (A). The synaptic weight exhibits facilitation (B), whereas the bias shows rapid depression (C). The spike-controlled combination of the two is the log-likelihood ratio, the evidence contribution of the individual afferent (D). A positive value provides evidence for the presence of a stimulus, and a negative value suggests the opposite.

positive. In between bursts, there is an unusually long ISI, which is more likely to occur in the absence of a stimulus. Therefore, the likelihood ratio decreases into the negative range.

**3.7 Tracking the Firing Rate of Biological Spike Trains.** Given the close agreement in the ISI statistics of data and model in Figure 3, one might wonder whether it is possible to obtain convergence with natural spike trains. To investigate this, we fed our predictor system, equation 3.12, with baseline spike trains obtained by in vivo recording from electrosensory afferent fibers of a weakly electric fish (*Apteronotus leptorhynchus*). In this case, the only available information from the afferent neuron is the spike train. Since the recording was extracellular, the internal fluctuations of the threshold that determines the firing probability could not be observed. Therefore, without the “ground truth” sequence ( $\theta[k]$ ), one cannot directly verify convergence of firing probabilities as in Figure 6A.

To test whether a dynamical system, such as the one defined by equation 3.12, accurately tracks the firing probability of the received spike train, one can compare the predicted firing probability, denoted by  $P_{\text{pred}}$ , with the empirical spike probability of the received spike train,  $P(a = 1 | P_{\text{pred}})$ , for every value of  $P_{\text{pred}}$ . Since  $P_{\text{pred}}$  can assume any value from zero to one

(within the numerical precision), we divide the interval  $[0, 1]$  into  $N$  narrow bins of width  $\Delta P_{\text{bin}} = 1/N$ , into which the sequence  $P_{\text{pred}}[k]$  is sorted. For each bin, the empirical firing probability is approximated by a normalized spike count obtained through the following procedure:

- Let  $P_{\text{bin}}$  be the center of the considered bin.
- Define the set  $K$  of all time steps  $k$  at which  $P_{\text{pred}}$  is within the bin range:

$$K = \{k \mid P_{\text{pred}}[k] \in [P_{\text{bin}} - \Delta P_{\text{bin}}/2, P_{\text{bin}} + \Delta P_{\text{bin}}/2]\}.$$

Let  $n_{\text{bin}} = |K|$  be the number of instances for which  $P_{\text{pred}}[k]$  is within the bin.

- Count the spikes that occurred at the time steps recorded in  $K$ :

$$n_{\text{spikes}} = \sum_{k \in K} a[k].$$

- The empirical firing probability given the center value of the bin is then

$$P(a = 1 \mid P_{\text{bin}}) = \frac{n_{\text{spikes}}}{n_{\text{bin}}}.$$

If the spike train is sufficiently large, the bin width small, and the predictor correct,  $P_{\text{bin}} \approx P(a = 1 \mid P_{\text{bin}})$  for each bin. In other words, on average the midvalue of a bin matches the normalized actual spike count obtained at all instances  $k$  when  $P_{\text{pred}}[k] \approx P_{\text{bin}}$ . Equality would be reached in the limit  $\Delta P_{\text{bin}} \rightarrow 0$  and with an infinitely long spike train.

*3.7.1 Parameter Estimation.* Accurate tracking of the firing probability is possible only if the parameters of the predictor system match those of the spike generator. In order to fit our predictor system with the afferent spike train, we employed a gradient-ascent algorithm to determine the parameter set  $\tilde{\alpha}$ ,  $\tilde{\beta}$ ,  $\tilde{\eta}$ , and  $\tilde{\sigma}$  for which the recorded spike train becomes most likely.

The total spike train log likelihood is given by the sum of the conditional log likelihoods of individual spike states:

$$\mathcal{L}(\vec{a}) = \sum_{k=1}^m \left\{ a[k] \log P(a[k] = 1 \mid \Theta[k]) \right. \\ \left. + (1 - a[k]) \log P(a[k] = 0 \mid \Theta[k]) \right\},$$

where  $m$  is the length of the spike train (in time steps). The firing probability,  $P(a[k] = 1 \mid \Theta[k])$ , is calculated according to equation 3.5, and the

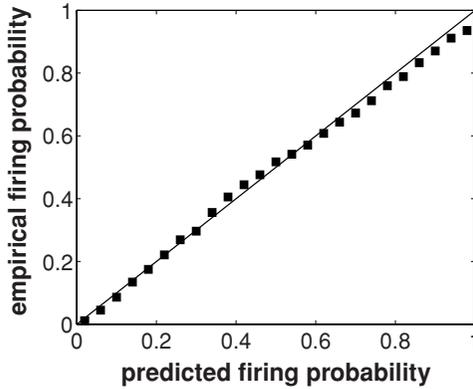


Figure 8: Scatter plot of predicted versus empirical firing probability of an afferent spike train (baseline activity) obtained from a weakly electric fish (*Apteronotus leptorhynchus*). The duration of the spike train was 275.2 s, and there were 95,027 spikes at a mean firing rate of 345 Hz ( $f_{EOD} = 970$  Hz). System parameters were obtained by maximizing the spike train likelihood using a gradient-ascent algorithm (see appendix C for details). The parameter values were  $\alpha = 2.6905$ ,  $\beta = 0.5062$ ,  $\eta = 1.54$ , and  $\sigma = 0.199$ . This fit yielded the parameter set used for the comparison in Figure 3.

probability of the complementary event  $a[k] = 0$  is simply

$$P(a[k] = 0 \mid \Theta[k]) = 1 - P(a[k] = 1 \mid \Theta[k]).$$

For the gradient ascent, one has to calculate the partial derivative of  $\mathcal{L}$  with respect to each parameter in order to obtain the increments of the update equations (details are provided in appendix C).

Figure 8 shows a scatter plot of the empirical firing probability versus its predicted value. Ideally, if the predictions were exact, all points would be on a diagonal. However, due to the finite bin size and length of the spike train, this is never the case. Interestingly, the tracking of the firing rate seems to work quite well provided that the parameters are chosen appropriately. There are only minor systematic errors in the data fit. Such good adaptation of the model system to a natural spike train is somewhat surprising, considering that the predictor system is not a biophysically detailed model and has but four free parameters. Additional confirmation that the obtained parameter set is indeed meaningful is provided by the closely matching ISI statistics of model and data, as shown in Figure 3. Note that the parameter optimization was not designed to fit the model's ISI histograms or serial correlation coefficient to the afferent spike train.

**3.8 Synopsis and Biological Plausibility.** In summary, the correlation-sensitive detection algorithm involves the following steps:

1. Establishing a match between the synaptic dynamics of short-term plasticity (see equation 3.13) and the adaptation dynamics of the presynaptic spike-generating mechanism (see equation 3.4).
2. An initialization of a synaptic state variable  $\Theta$  and its subsequent temporal evolution (see equation 3.13)
3. A smooth, nonlinear transformation between the time-varying state variable  $\Theta$  and the time-varying synaptic weight  $w$  and bias  $b$  (see equation 3.12, and Figure 5)
4. A summation of time-varying bias terms ( $b_i$  terms in equation 3.12)
5. A summation of weighted spike inputs ( $w_i$  terms in equation 3.12)
6. Temporal integration of the weighted sum and comparison with a threshold level (equation 1.2)

Steps 1 to 3 apply to each individual synapse, and steps 4 to 6 apply to the sum over all synapses that occurs at the postsynaptic detector neuron. The biological plausibility of carrying out log-likelihood-type computations using the integrate-and-fire dynamics associated with steps 5 and 6 has been established previously (Gold & Shadlen, 2001). Here we discuss the plausibility of the initial steps, 1 to 4, which are associated with the proposed short-term plasticity mechanism.

**Step 1:** To the extent that the afferent population has homogeneous adaptation dynamics, a match between the synaptic dynamics of short-term plasticity and the adaptation dynamics of the input spike trains could be hard-wired into the system through an evolutionary process of variation and natural selection. In this scenario, the adaptation dynamics giving rise to ISI correlations in the input spike train would likely evolve first because even a detector with static weights can benefit passively from the reduced spike count variability (Ratnam & Nelson, 2000; Chacron et al., 2001; Goense & Ratnam, 2003). Further improvements in detection performance, and hence a selective advantage, would be afforded to individuals with genetically specified synaptic dynamics that were more closely matched to those of the spike-generating mechanism. An afferent population that exhibited heterogeneity in adaptation dynamics would presumably require some sort of online or developmental tuning of individual synapses in order to take advantage of the proposed mechanism.

**Step 2:** As discussed in section 3.5 and appendices A and B, the dynamical system associated with state prediction (see equation 3.13) robustly converges toward a trajectory that yields accurate estimates of firing probability, independent of state initialization or slight variations in parameter

values. The robust convergence properties should carry over into any biological implementation with similar underlying dynamics.

**Step 3:** The transformations between state estimates  $\Theta$  and synaptic weights  $w$  and biases  $b$  are analytically complex (see equation 3.12), but are smoothly varying, weakly nonlinear functions, as illustrated in Figure 5. Such nonlinear relationships could be established by a variety of biological mechanisms associated with calcium signaling and transmitter release at the synapse.

**Step 4:** In addition to a synaptic weight  $w_i$ , the application of statistical detection theory predicts an associated bias  $b_i$  for each synapse. When the input is a renewal spike train, these bias terms are constant (see equation 2.6) and can be absorbed into a redefinition of the threshold  $\gamma$  associated with the detection process (see equation 1.1). When the spike trains have ISI correlations, the individual synaptic biases  $b_i$  become time-dependent (see equation 3.12). If the detector neuron receives a large number of independent afferent inputs,  $n$ , the sum of the biases  $b(t) = \sum_{i=1}^n b_i(t)$  will have a variance that decreases linearly with  $n$ , according to the central limit theorem. Thus, when the degree of afferent convergence is large, the bias term is approximately constant and can once again be absorbed into a redefinition of the detection threshold  $\gamma$ , as was the case for renewal process inputs.

Thus, we see that all the elements of the proposed model are plausible under certain biologically relevant circumstances (e.g., multiple converging afferents with homogeneous adaptation dynamics). The model could also be applicable to more challenging circumstances (e.g., a small number of converging afferents with heterogeneous dynamics), but would require the specification of additional mechanisms for tuning the dynamics and accommodating time-varying threshold levels in the detection process.

**3.9 The Electrosensory Image Model.** We model a patch of fish skin as a segment of a cylindrical surface. Let  $n$  be the number of receptors contained in the skin patch and  $i$  be an index referencing individual receptors (afferent fibers). Let  $(r_0, \phi_i, x_i)$  be the position of the  $i$ th receptor organ in cylindrical coordinates and  $(r, \phi, x)$  be the (unknown) target position. The signal intensity at the  $i$ th receptor is modeled as a two-dimensional gaussian:

$$s_i = A \exp \left\{ \underbrace{-(x_i - x)^2 / 2\sigma_s^2}_{\text{main axis}} - \underbrace{r_0^2(\phi_i - \phi)^2 / 2\sigma_s^2}_{\text{polar angle}} \right\}, \quad (3.13)$$

where  $A$  is the signal amplitude,  $r_0$  the radius of curvature of the patch, and  $\sigma_s$  characterizes the width of the electrosensory image. Both the amplitude  $A$  and the stimulus width  $\sigma_s$  depend on the target distance  $r$ . Within a distance of a few centimeters, the amplitude follows a power law

(Rasnow, 1996):

$$A(r) = k_e r^{-\alpha}. \quad (3.14)$$

The factor  $k_e$  incorporates the fish's electric field strength, as well as the conductivity of the target object. For the exponent of the power law, a value of  $\alpha \approx 4$  has been observed in *Apteronotus albifrons* (Chen, House, Krahe, & Nelson, 2005). The steep power law limits the effective range for prey detection to a few centimeters (MacIver, Sharabash, & Nelson, 2001). For small targets, the width  $\sigma_s$  is approximately proportional to the distance of the target from the skin (Rasnow, 1996):

$$\sigma_s(r) \propto r - r_0.$$

The constant of proportionality is approximately unity, so we set

$$\sigma_s(r) = r - r_0. \quad (3.15)$$

Inserting equations 3.14 and 3.15 into 3.13 yields the intensity of the electrosensory stimulus at any receptor position  $(r_0, \phi_i, x_i)$  as a function of target coordinates  $(r, \phi, x)$ :

$$s_i = k_e r^{-\alpha} \exp \left\{ \frac{-r_0^2(\phi_i - \phi)^2 - (x_i - x)^2}{2(r - r_0)^2} \right\}. \quad (3.16)$$

**3.10 Performance Comparison.** In our simulation of electrosensory signal detection, we have restricted our analysis to a proof of principle using only stimuli with instantaneous onset and constant intensity. Optimal detection of time-varying signals would require a temporal receptive field matched to the expected time course of typical stimuli. In the electrosensory system of weakly electric fish, there is evidence that such expectations are relayed via feedback from higher brain areas (Maler & Berman, 1999; Bastian, 1999; Lewis & Maler, 2002). Such top-down information could be included in the likelihood ratio framework, but is beyond the scope of this article.

We performed computer simulations of a single detector "neuron" monitoring a  $15 \times 15$  array of receptors. The stimulus was centered on the receptor array and thus the receptive field of the detector. Using the natural time discretization provided by the periodicity of the fish's electric organ discharge (EOD) and assuming an EOD frequency of 1000 Hz resulted in a step size of  $\Delta t = 1/f_{\text{EOD}} = 1$  ms.

In order to assess the influence of ISI correlations on detection performance, we generated renewal spike trains using the probability encoder model described in section 2 and nonrenewal spike trains using the more

realistic nonlinear adaptive threshold model introduced in section 3.2. The parameters of both spike generators were adjusted, so that equal stimulus amplitudes resulted in the same increase in firing rate above equal baseline activities. We used a gain of  $\approx 250$  spikes/s/mV, in accordance with experimental observation (Nelson, Xu, & Payne, 1997).

To make meaningful comparisons, the mean false alarm rates of the two detectors were matched. The thresholds of both detectors were set to obtain a mean false alarm interval of 95 ms, which corresponds to a false alarm rate of approximately 10 Hz, similar to typical spontaneous firing rates of ELL neurons (Bastian & Nguyenkim, 2001).

At the beginning of each trial, the cumulative likelihood ratio was set to zero. After 50 time steps, a gaussian electrosensory image according to equation 3.15 was presented and the time counter started. The time interval before stimulus onset allowed for transients in the afferent spike generators to decay and ensured sufficient convergence of synaptic predictors. The value of 50 time steps was determined empirically (see Figure 6). The detection delay is then the time from stimulus onset to the first postsynaptic spike (detector decision in favor of hypothesis  $H_1$ ). The same procedure was performed with zero signal amplitude in order to test the detector's behavior for pure baseline input. The detection delay is then the time from the timer reset to the first false alarm.

*3.10.1 Distributions of Detection Delay.* To demonstrate the advantage of a dynamic detector, we analyze the distribution of detection delays. Figure 9 shows delay histograms for both detector types, with and without a stimulus. Under baseline conditions, the histograms for both detectors are virtually identical. However, in the presence of a stimulus, the dynamic detector, Figure 9C, exhibits a smaller coefficient of variation than the static detector, Figure 9D, and has more probability mass concentrated at shorter delay times.

There is no difference in mean detection delay between the two detector types, since false alarm rates are equal and the static detector is adjusted to match the mean afferent firing rates. Although the static detector cannot track the afferent firing probability, its estimate is correct on average.

*3.10.2 Detection Probability as a Function of Integration Time.* Integrating the delay distributions over time, that is, summing the counts of all histogram bins in Figures 9C and 9D up to a given delay time and normalizing by the total count, yields the "hit probability" as a function of integration time. This includes correct detections and false alarms. Integrating the baseline histograms, Figures 9A and 9B, over the same desired time window yields the probability of false alarm, which must be subtracted from the hit probability to obtain the probability of correct detection.

Figure 10 shows a plot of the probability of correct detection as a function of integration time for three situations: (1) dynamic detector with correlated

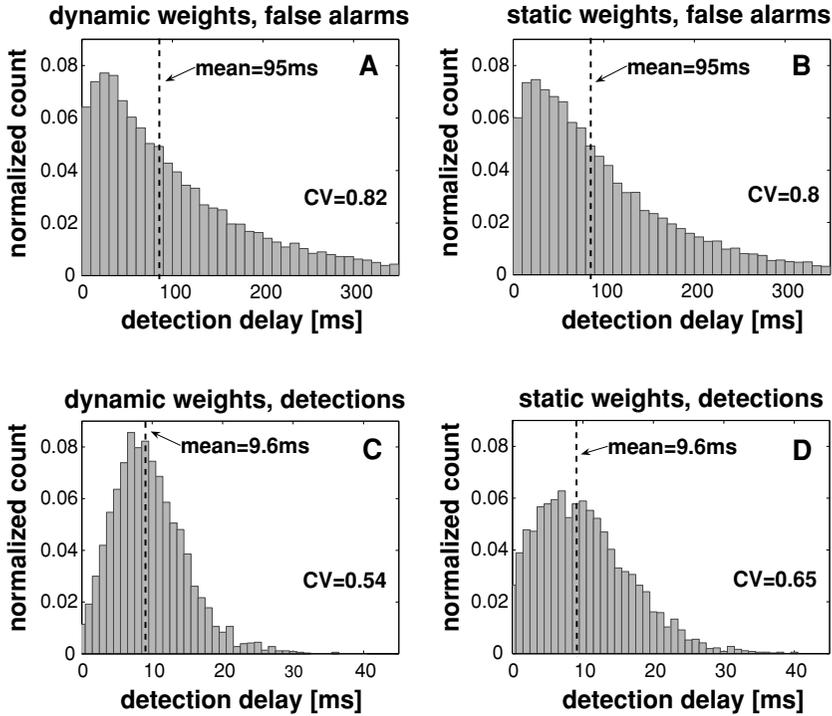


Figure 9: Distributions of the detection delay for both detectors under baseline and signal condition. Histograms were obtained for the static and the dynamic detector, both of which received correlated input. Thresholds were chosen so that false alarm rates were equal. Therefore, the mean detection delay (dashed lines) under baseline conditions is the same in both detectors (A and B). In the presence of a stimulus, the mean detection delay is still the same, but the dynamic detector (C) has more probability mass concentrated at shorter delay times and exhibits a smaller coefficient of variation (CV) than the static detector (D).

input, (2) static detector with correlated input, and (3) static detector with renewal input. (The histograms for the third case are not shown in Figure 9.)

In addition to the expected beneficial effect of reduced spike count variability in the correlated firing, the dynamic detector is able to exploit the higher degree of predictability of afferent spike events due to temporal correlations. Consequently, the same detection probability is reached within a shorter integration time. Interestingly, the optimal integration time for the dynamic detector, which is about 15 ms, matches typical values of membrane time constants observed in the ELL of weakly electric fish (Berman & Maler, 1998). Thus, the model suggests a biologically plausible integration timescale, even though it has no leak term.

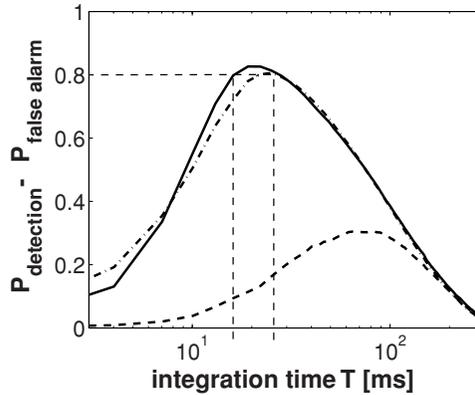


Figure 10: Detection probability as a function of integration time for a given, fixed signal intensity. The static detector with uncorrelated input spike train statistics (dashed curve) performs poorly, requiring at least 70 ms to reach a maximum of only 30%. Fed with a negatively correlated spike train, the detection performance improves significantly (dash-dotted curve). The dynamic detector (solid curve) requires a shorter integration time to reach the same detection probability. In all three simulations, the false alarm rate is kept at the same level. The signal intensity is equivalent to that caused by a small prey-like object at a typical detection distance of about 2 cm.

Moreover, the synaptic plasticity transforms the inherent redundancy in the presynaptic spike trains into a reduced variability of the postsynaptic spike output, indicated by a smaller coefficient of variation in the detection delay. Such a reduction in firing variability is a well-known consequence of short-term synaptic depression (Abbott & Regehr, 2004).

A population of such detector neurons would be more likely to fire within a small time interval than an equivalent population of static detectors. A neuron in a higher brain area, receiving input from a population of neurons with dynamic synaptic weights, could act as a coincidence detector and, due to the more precise firing of its input, the time window of coincidence could be tighter, resulting in a more efficient rejection of false alarms.

#### 4 Discussion

---

There is ample evidence of activity-dependent synaptic conductances varying on timescales comparable to the interspike interval of their presynaptic input (Zucker & Regehr, 2002; Xu-Friedman & Regehr, 2004). Hence, the question of the functional significance of such plasticity arises. Our primary focus has been the interplay between short-term synaptic plasticity and presynaptic input spike trains with correlated ISIs and its role in weak

signal detection. While it has been suggested that depressive synapses can achieve decorrelation (“whitening”) of positively correlated spike trains (Goldman, Maldonado, & Abbott, 2002), the same synaptic mechanism would have the opposite effect in the presence of negative correlations. Instead of decorrelating, short-term synaptic depression would preserve the negative correlations. The combination of facilitating and depressing plasticity in our dynamic detector model enables the postsynaptic neuron to differentiate between expected and unexpected spikes by exploiting the inherent redundancy in its correlated presynaptic input.

The predictive synaptic plasticity introduced in this article is a novel mechanism that goes beyond previous models of weak signal detection in spike trains with correlated ISIs (Ratnam & Nelson, 2000; Chacron et al., 2001; Goense & Ratnam, 2003). In all of these approaches the emphasis is on the long-term regularization (i.e., on a timescale of multiple ISIs) rather than short-term predictability of spike trains. While the framework of statistical detection theory lends itself very well to weak signal detection in correlated spike trains, an explicit representation of conditional spiking probabilities would seem biologically implausible. As a history-dependent process, short-term synaptic plasticity offers a way to address the computational challenge of spike forecasting by implicitly representing conditional firing probabilities. Though most of the increase in detection performance compared to a renewal input stems from long-term regularization, there is also a modest, but significant, benefit in exploiting the statistical dependence of ISIs via synaptic plasticity.

The apparent similarity between neuron models with spike-driven threshold adaptation and short-term synaptic depression has been noted by Chacron and colleagues (Chacron et al., 2003), though the authors caution that for electrosensory afferents, the required time constant of neurotransmitter recovery would have to be significantly shorter than those typically found in cortical neurons. However, the improvement of detection performance in our dynamic detector model raises the question of whether a neural correlate of such a detector exists. It would be surprising if electrosensory systems, and perhaps sensory neurons in other modalities, did not in some way actively exploit the correlated nature of afferent spike trains. We therefore speculate that a rapid form of short-term plasticity in excitatory synapses of postsynaptic (ELL) neurons might be able to mimic afferent threshold fluctuations, thus enabling the synapses to track the firing probabilities associated with presynaptic spikes. The unusually high firing rates of electrosensory afferents may be matched by unusually small synaptic time constants unique to sensory neurons in the ELL.

From a bioengineering point of view, the described form of predictive synaptic plasticity may also have implications for the design of neuromorphic systems or bioelectronic interface technology in sensory prostheses, where it may be advantageous to precisely match the ISI statistics of a

sensory device output to the synaptic properties of the cells to which the device is connected.

### Appendix A: The Lyapunov Exponent of the Afferent Model \_\_\_\_\_

The stability of threshold orbits is determined by the Lyapunov exponent of the map that transforms threshold  $\theta[k]$  into its successor  $\theta[k + 1]$  (see equation 3.4):

$$f_k : \mathbb{R} \longrightarrow \mathbb{R}, \quad \theta \longmapsto \theta - \beta/\alpha + a[k] \beta \frac{2e^{-\eta\theta}}{1 + e^{-\eta\theta}}. \quad (\text{A.1})$$

The subscript  $k$  denotes that the map  $f_k$  is time dependent. The Lyapunov exponent of this map is given by (Strogatz, 1994),

$$\lambda = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} \ln |f'_k(\theta[k])|, \quad (\text{A.2})$$

where the prime indicates the derivative with respect to  $\theta$ . If the limit in equation A.2 exists, the absolute difference between actual and predicted threshold changes exponentially:

$$|\Theta[k] - \theta[k]| = |\Theta[0] - \theta[0]| e^{\lambda k}.$$

Consequently, the Lyapunov exponent must be negative (see Figure 6B) in order to produce stable threshold orbits that the predictor can converge toward. The time constant of this convergence is

$$\tau = \frac{1}{|\lambda|}.$$

For a positive Lyapunov exponent, even the slightest difference between  $\Theta[0]$  and  $\theta[0]$  would be substantially magnified within a small number of time steps. Such sensitivity to initial conditions is a characteristic feature of deterministic chaos. Therefore, it is important to investigate the behavior of the Lyapunov exponent for different parameter settings in order to avoid chaotic regimes if they exist.

Apart from special cases, the Lyapunov exponent must be calculated numerically since no general analytical expression exists. However, it is relatively easy to assess the dependence of  $\lambda$  on the boost parameter  $\eta$ . The

derivative of the map  $f_k$  with respect to  $\theta$  is given by

$$f'_k(\theta) = 1 - 2a[k]\beta \frac{\eta e^{-\eta\theta[k]}}{[1 + e^{-\eta\theta[k]}]^2}. \quad (\text{A.3})$$

In the special case of constant threshold boost, that is, for  $\eta = 0$  (see equation A.1), the expression reduces to

$$f'_k \equiv 1, \quad \text{for all } k \text{ and all } \theta.$$

By inserting this result into equation A.2, one finds that the Lyapunov exponent is zero, implying an infinite convergence time constant. Obviously, threshold saturation plays a crucial role in system stability. For instance, a predictor system based on the simpler linear adaptive threshold model (Brandman & Nelson, 2002), which lacks a threshold saturation term, does not possess the convergence property and would require precise knowledge of the initial value of  $\theta[0]$  in order to set  $\Theta[0] = \theta[0]$ .

The remainder of the analysis is based on numerical calculation of the Lyapunov exponent, approximating equation A.2 by summing over a large number of time steps (c. 10,000). Since the derivative of the map in equation A.3 does not depend on  $\alpha$ , the Lyapunov exponent is determined by only two parameters,  $\beta$  and  $\eta$ . This considerably simplifies numerical stability analysis and enables us to visualize the Lyapunov exponent in the relevant parameter space. As Figure 11A shows, stable threshold orbits are guaranteed over a wide parameter range.

## Appendix B: Robustness of Threshold Tracking

To test the predictor robustness, we generated artificial spike trains using the values for the parameters  $\alpha$  and  $\beta$  obtained from the gradient ascent fit with afferent data (see section 3.7.1) and chose slightly deviant values for  $\tilde{\alpha}$  and  $\tilde{\beta}$  in the predictor system. Keeping these parameters fixed, we varied only  $\eta$  and  $\tilde{\eta}$ , setting  $\tilde{\eta} = \eta$ . The overall discrepancy between prediction  $P_{\text{pred}}$  and actual firing probability  $P$  was measured in terms of the root mean square (RMS) deviation,

$$\langle P - P_{\text{pred}} \rangle_{\text{rms}} = \sqrt{\langle (P - P_{\text{pred}})^2 \rangle} = \sqrt{\frac{1}{m} \sum_{k=1}^m (P[k] - P_{\text{pred}}[k])^2},$$

where the angular brackets denote the temporal average. Figure 12A shows the RMS prediction error as a function of  $\eta$ . From this plot, one might conclude that the fitted parameter value  $\eta^*$  is quite far from the minimum and would therefore be a suboptimal choice.

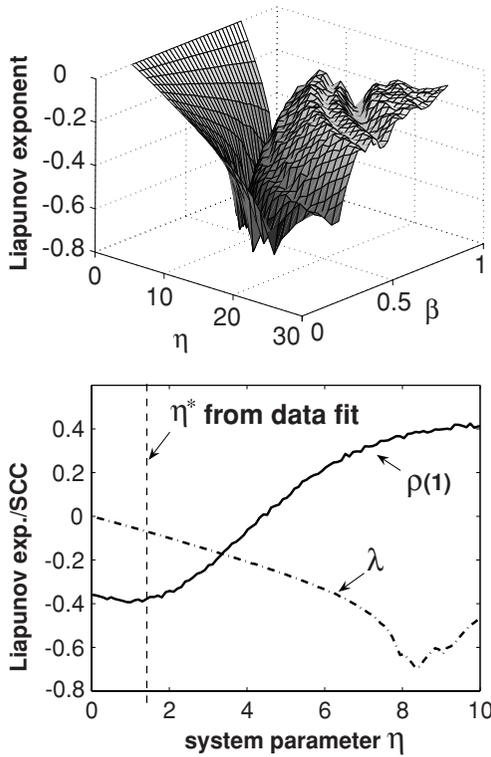


Figure 11: (Top) The Lyapunov exponent of the map  $f_k$  in equation A.1 as a function of the parameters  $\eta$  and  $\beta$ . Over the entire plotted parameter range, the Lyapunov exponent is negative, indicating stable orbits and the absence of chaos. (Bottom) Double plot of the Lyapunov exponent (dashed-dotted curve) and the correlation coefficient of adjacent ISIs (solid curve) as a function of threshold saturation parameter  $\eta$ . Parameters  $\alpha$  and  $\beta$  remain fixed. The dashed vertical line marks the value  $\eta^*$ , obtained from fitting the model to an electrosensory afferent spike train (see Figure 8). Increasing  $\eta$  lowers the convergence time constant (i.e., the Lyapunov exponent becomes more negative), while reducing the degree of negative correlativeness. Thus, the two desirable properties, negative correlativeness of subsequent ISIs and short convergence time constant, cannot be optimized simultaneously. Interestingly, when  $\eta$  assumes the value  $\eta^*$ , obtained from fitting neural data, it realizes a trade-off between the two quantities.

Intuitively, one expects the robustness against parameter perturbations to depend on the Lyapunov exponent, which determines the convergence time constant. If the unperturbed system has a negative Lyapunov exponent, the predictor variable partially self-corrects so that the asymptotic prediction error is bounded, and thus the sequences  $(\theta[k])$  and  $(\Theta[k])$  do

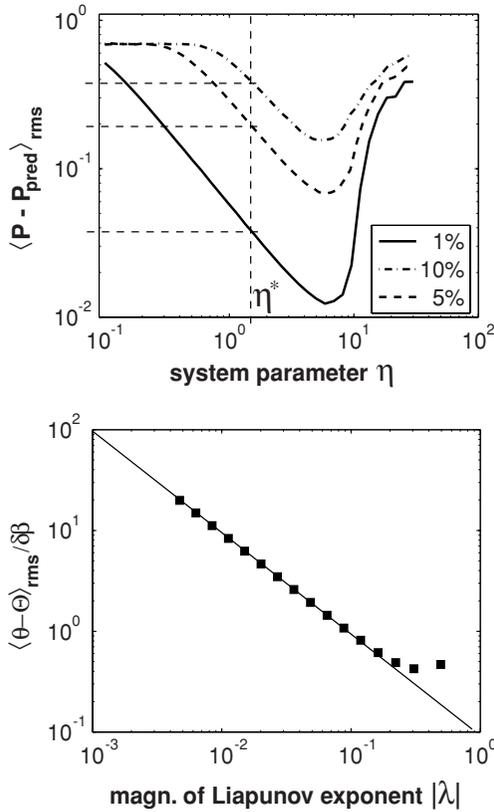


Figure 12: (Top) The saturation parameter  $\eta$  controls the robustness of the predictor system against small inaccuracies of its parameter values. The difference between actual and predicted firing probability, measured in terms of the root mean square error of the sequences  $P[k]$  and  $P_{\text{pred}}[k]$ , is plotted as a function of system parameter  $\eta$ . The error stems from the fact that predictor parameters  $\tilde{\alpha}$  and  $\tilde{\beta}$  deviated from their counterparts  $\alpha$  and  $\beta$  in the spike generator by 1%, 5%, and 10%, respectively. Both systems had the same fixed noise variance. The spike generator values were  $\alpha = 2.6905$  and  $\beta = 0.5062$ . For the fitted value,  $\eta = \eta^*$ , the error is close to 4% for a 1% parameter deviation. For larger deviations, the error increases significantly, but the system remains operational. This graceful degradation is important for parameter learning and supports the biological plausibility of the model. (Bottom) The root mean square error of threshold prediction (for 1% parameter deviation) as a function of the magnitude of the Liapunov exponent,  $|\lambda|$ . As predicted in equation B.1, the error scales with  $|\lambda|^{-1}$ , except for larger values of  $|\lambda|$ , for which the integral approximation becomes invalid.

not diverge. The shorter the time constant, the more quickly the perturbed predictor sequence is drawn back to the original attractor orbit, resulting in a smaller RMS prediction error.

As shown in Figure 11 bottom, the Lyapunov exponent, and thus the perturbation robustness, is controlled by the parameter  $\eta$  (dashed curve). Also plotted is the serial correlation coefficient of adjacent ISIs ( $\text{lag}=1$ ),  $\rho(1)$ , defined as

$$\rho(1) = \frac{\langle (I_{k+1} - \bar{I})(I_k - \bar{I}) \rangle}{\sigma^2\{I\}},$$

where  $\bar{I}$  and  $\sigma^2\{I\}$  are the mean and variance of the ISI. Obviously, increasing the robustness (magnitude of negative Lyapunov exponent) reduces the degree of negative ISI correlation, which would affect detection performance. Interestingly, the parameter value  $\eta^*$  appears to realize a trade-off between these two desirable properties that cannot be maximized simultaneously.

**B.1 Prediction Error and Lyapunov Exponent.** In order to quantify how the RMS prediction error depends on the Lyapunov exponent of the threshold map, we consider a small perturbation in the parameter  $\beta$ , denoted by  $\delta\beta$ . Consequently, the predictor variable  $\Theta$  is transformed via the perturbed map:

$$\Theta[k] = \Theta[k-1] - \frac{(\beta + \delta\beta)}{\alpha} + 2a[k-1](\beta + \delta\beta)g(\Theta[k-1], \eta).$$

Thus, at each time step, the resultant perturbation in  $\Theta$  is

$$\delta\Theta[k] = -\frac{\delta\beta}{\alpha} + 2a[k-1]\delta\beta g(\Theta[k-1], \eta).$$

In order to obtain a relation between  $\Theta[k]$  and  $\theta[k]$ , one must consider the perturbations from previous time steps. All past perturbations have decayed exponentially with time constant  $1/|\lambda|$ . Consequently, the difference between threshold and predictor at time  $k$  is the sum of all the decayed previous perturbations:

$$\Theta[k] - \theta[k] = \sum_{j=0}^k \delta\Theta[k-j] e^{-|\lambda|j}.$$

For the RMS error, one obtains

$$\sqrt{\langle (\Theta[k] - \theta[k])^2 \rangle} = \sqrt{\langle (\delta\Theta)^2 \rangle} \sum_{j=0}^{\infty} e^{-|\lambda|j}.$$

In the asymptotic limit  $k \rightarrow \infty$ , the summation over exponentials can be approximated by an integral that is easy to evaluate:

$$\begin{aligned} \langle \Theta[k] - \theta[k] \rangle_{\text{rms}} &= \sqrt{\langle (\delta\Theta)^2 \rangle} \sum_{j=0}^{\infty} e^{-|\lambda|j} \\ &\approx \sqrt{\langle (\delta\Theta)^2 \rangle} \int_0^{\infty} e^{-|\lambda|t} dt = \frac{\langle \delta\Theta \rangle_{\text{rms}}}{|\lambda|}. \end{aligned} \quad (\text{B.1})$$

Hence, the RMS prediction error is expected to scale inverse proportionally to the Lyapunov exponent. As Figure 12 bottom shows, the theoretical result is in good agreement with the numerical simulation. Only for larger magnitudes of the Lyapunov exponent is the power law scaling incorrect. In this range of  $\lambda$ , the convergence time constant  $\tau = 1/|\lambda|$  is of the same order of magnitude as the length of the time step between iterations, and thus the rapidly decaying integrand is a poor approximation of the piecewise constant entries in the summation in equation B.1.

The scaling law for the robustness shows that the predictor parameters do not have to perfectly match those of the spike generator in order to achieve a reasonably accurate threshold prediction. To reduce the prediction error, one would merely have to increase the Lyapunov exponent. However, due to the trade-off between Lyapunov exponent and serial correlation coefficient, any increase in robustness would be at the expense of a lower degree of ISI correlation (see Figure 11 bottom).

### Appendix C: Maximum Likelihood Estimation of Parameters by Gradient Ascent

---

In this section we describe the gradient-ascent procedure used to estimate the model parameters  $\tilde{\alpha}$ ,  $\tilde{\beta}$ ,  $\tilde{\eta}$ , and  $\tilde{\sigma}$  that best match a given afferent spike train. Let the spike train be given as binary vector  $\vec{a} \in \{0, 1\}^m$  of length  $m$ . This format is obtained by re-sampling the original afferent recording at the EOD frequency, since the afferents never produce more than one spike per cycle of the oscillating field. The spike train log likelihood is the sum of the

conditional log likelihoods of the individual spike events:

$$\mathcal{L}(\vec{a}) = \sum_{k=1}^m \left\{ a[k] \log P(a[k] = 1 \mid \Theta[k]) \right. \\ \left. + (1 - a[k]) \log P(a[k] = 0 \mid \Theta[k]) \right\}.$$

To make equations more compact, we introduce the following shorthand notation for the conditional firing probability:

$$P[k] := P(a[k] = 1 \mid \Theta[k]).$$

Thus, the spike train likelihood function is written as

$$\mathcal{L}(\vec{a}) = \sum_{k=1}^m \{a[k] \log P[k] + (1 - a[k]) \log(1 - P[k])\}. \quad (\text{C.1})$$

Note that  $P[k]$  contains an account of the entire spike train history, since it is a function of  $\Theta[k]$  and hence indirectly a function of all previous threshold values  $\Theta[k - 1]$ ,  $\Theta[k - 2]$ ,  $\dots$ ,  $\Theta[0]$ . By virtue of equation 3.5, the firing probability of a baseline spike train is obtained by setting the stimulus intensity to zero ( $s[k] \equiv 0$ ). Hence,

$$P[k] = \frac{1}{2} \operatorname{erfc} \left( \frac{\Theta[k]}{\sqrt{2}\sigma} \right). \quad (\text{C.2})$$

In order to find the parameter values that maximize the likelihood function, we employ a gradient-ascent algorithm. At each iteration, the current parameters are updated by adding a fraction of the gradient of the likelihood function:

$$(\tilde{\alpha}, \tilde{\beta}, \tilde{\eta}, \tilde{\sigma})_{n+1} = (\tilde{\alpha}, \tilde{\beta}, \tilde{\eta}, \tilde{\sigma})_n + l \nabla_{\tilde{\alpha}, \tilde{\beta}, \tilde{\eta}, \tilde{\sigma}} \mathcal{L},$$

where  $l$  is the (heuristically determined) learning rate and the gradient is

$$\nabla_{\tilde{\alpha}, \tilde{\beta}, \tilde{\eta}, \tilde{\sigma}} = \left( \frac{\partial}{\partial \tilde{\alpha}}, \frac{\partial}{\partial \tilde{\beta}}, \frac{\partial}{\partial \tilde{\eta}}, \frac{\partial}{\partial \tilde{\sigma}} \right). \quad (\text{C.3})$$

Using equation C.1, we obtain the gradient of the likelihood function:

$$\nabla_{\tilde{\alpha}, \tilde{\beta}, \tilde{\eta}, \tilde{\sigma}} \mathcal{L} = \sum_{k=1}^m \left\{ a[k] \frac{\nabla_{\tilde{\alpha}, \tilde{\beta}, \tilde{\eta}, \tilde{\sigma}} P[k]}{P[k]} - (1 - a[k]) \frac{\nabla_{\tilde{\alpha}, \tilde{\beta}, \tilde{\eta}, \tilde{\sigma}} P[k]}{1 - P[k]} \right\} \\ = \sum_{k=1}^m \left[ \frac{a[k]}{P[k]} - \frac{1 - a[k]}{1 - P[k]} \right] \nabla_{\tilde{\alpha}, \tilde{\beta}, \tilde{\eta}, \tilde{\sigma}} P[k]. \quad (\text{C.4})$$

Hence, one has to calculate the partial derivatives of the firing probability, equation C.2. For  $\tilde{\sigma}$ , this is straightforward:

$$\frac{\partial P[k]}{\partial \tilde{\sigma}} = \frac{\Theta[k]}{\sqrt{2\pi}\sigma^2} \exp\left(\frac{-\Theta[k]^2}{2\sigma^2}\right). \tag{C.5}$$

However, with the exception of  $\tilde{\sigma}$ ,  $P[k]$  is not an explicit function of the parameters. Therefore, the partial derivatives with respect to  $\tilde{\alpha}$ ,  $\tilde{\beta}$ , and  $\tilde{\eta}$  have to be obtained via the chain rule:

$$\nabla_{\tilde{\alpha}, \tilde{\beta}, \tilde{\eta}} P[k] = \frac{dP[k]}{d\Theta[k]} \nabla_{\tilde{\alpha}, \tilde{\beta}, \tilde{\eta}} \Theta[k]. \tag{C.6}$$

Differentiating equation C.2 yields

$$\frac{dP[k]}{d\Theta[k]} = -\frac{1}{\sigma\sqrt{2\pi}} \exp\left(\frac{-\Theta[k]^2}{2\sigma^2}\right). \tag{C.7}$$

To evaluate the gradient of  $\Theta[k]$ , one must take into account that  $\Theta[k]$  is the result of  $k$ -fold iteration of the map (see appendix A):

$$f_k : \mathbb{R} \rightarrow \mathbb{R}, \quad \Theta \mapsto \Theta - \beta/\alpha + a[k] \beta \frac{2e^{-\eta\Theta}}{1 + e^{-\eta\Theta}}.$$

Hence, one can write

$$\frac{\partial \Theta[k]}{\partial \tilde{\alpha}} = \frac{\partial}{\partial \tilde{\alpha}} f^k(\Theta[0]) = \prod_{j=0}^{k-1} \frac{\partial}{\partial \tilde{\alpha}} f(\Theta, \tilde{\alpha}) \Big|_{\Theta=\Theta[j]}.$$

This yields

$$\frac{\partial}{\partial \tilde{\alpha}} f(\Theta, \tilde{\alpha}) = \frac{\tilde{\beta}}{\tilde{\alpha}^2} \Rightarrow \frac{\partial \Theta[k]}{\partial \tilde{\alpha}} = \left(\frac{\tilde{\beta}}{\tilde{\alpha}^2}\right)^k.$$

Together with equation C.7, one obtains

$$\frac{\partial P[k]}{\partial \tilde{\alpha}} = -\frac{1}{\sigma\sqrt{2\pi}} \left(\frac{\tilde{\beta}}{\tilde{\alpha}^2}\right)^k \exp\left(\frac{-\Theta[k]^2}{2\sigma^2}\right). \tag{C.8}$$

In a similar manner, one can calculate the partial derivatives of  $P[k]$  with respect to  $\tilde{\beta}$  and  $\tilde{\eta}$ . Inserting the expressions for the partial derivatives of  $P[k]$  into equation C.4 yields the components of the gradient of the likelihood

function:

$$\frac{\partial \mathcal{L}}{\partial \tilde{\alpha}} = \sum_{k=1}^m Q[k] \left( \frac{\tilde{\beta}}{\tilde{\alpha}^2} \right)^k, \quad (\text{C.9})$$

$$\frac{\partial \mathcal{L}}{\partial \tilde{\beta}} = \sum_{k=1}^m Q[k] \prod_{j=0}^{k-1} \left( -\frac{1}{\tilde{\alpha}} + a[j] \left[ 1 - \frac{1}{1 + \exp(-\eta \Theta[j])} \right] \right) \quad (\text{C.10})$$

$$\frac{\partial \mathcal{L}}{\partial \tilde{\eta}} = - \sum_{k=1}^m Q[k] \tilde{\beta}^k \prod_{j=0}^{k-1} \frac{a[j] \Theta[j] \exp(-\eta \Theta[j])}{[1 + \exp(-\eta \Theta[j])]^2}, \quad (\text{C.11})$$

$$\frac{\partial \mathcal{L}}{\partial \tilde{\sigma}} = - \frac{1}{\sigma} \sum_{k=1}^m Q[k] \Theta[k], \quad (\text{C.12})$$

where

$$Q[k] = - \frac{1}{\sigma \sqrt{2\pi}} \exp \left( \frac{-\Theta[k]^2}{2\sigma^2} \right) \left[ \frac{a[k]}{P[k]} - \frac{1 - a[k]}{1 - P[k]} \right]. \quad (\text{C.13})$$

## Acknowledgments

---

This work has been supported by a grant from the National Institutes of Health. We thank the reviewer and Mark Mahler for valuable comments on the manuscript.

## References

---

- Abbott, L., & Regehr, W. (2004). Synaptic computation. *Nature*, 431, 796–803.
- Bahar, S., Kantelhardt, J., Neiman, A., Rego, H., Russel, D., Wilkens, L., Bunde, A., & Moss, F. (2001). Long-range temporal anti-correlations in paddlefish electroreceptors. *Europhysics Letters*, 56(3), 454–460.
- Bastian, J. (1999). Plasticity of feedback inputs in the apteronotid electrosensory system. *Journal of Experimental Biology*, 202, 1327–1337.
- Bastian, J., & Nguyenkim, J. (2001). Dendritic modulation of burst-like firing in sensory neurons. *Journal of Neurophysiology*, 85, 10–22.
- Berman, N., & Maler, L. (1998). Inhibition evoked from primary afferents in the electrosensory lateral line lobe of weakly electric fish. *Journal of Neurophysiology*, 80, 3173–3196.
- Blitz, D., Foster, K., & Regehr, W. (2004). Short-term synaptic plasticity: A comparison of two synapses. *Nature Reviews Neuroscience*, 5(8), 630–640.
- Brandman, R., & Nelson, M. (2002). A simple model of long-term spike train regularization. *Neural Computation*, 14, 1575–1597.

- Chacron, M., Longtin, A., St. Hilaire, M., & Maler, L. (2000). Suprathreshold stochastic firing dynamics with memory in p-type electroreceptors. *Physical Review Letters*, *85*, 1576–1579.
- Chacron, M., Maler, L., & Longtin, A. (2001). Negative interspike interval correlations increase the neuronal capacity for encoding time dependent stimuli. *Journal of Neuroscience*, *21*, 5328–5343.
- Chacron, M., Pakdaman, K., & Longtin, A. (2003). Interspike interval correlations, memory, adaptation, and refractoriness in a leaky integrate-and-fire model with threshold fatigue. *Neural Computation*, *15*, 253–278.
- Chen, L., House, J., Krahe, R., & Nelson, M. E. (2005). Modeling signal and background components of electrosensory scenes. *Journal of Comparative Physiology A*, *191*(4), 331–345.
- Goense, J., & Ratnam, R. (2003). Continuous detection of weak sensory signals in afferent spike trains: The role of anticorrelated interspike intervals in detection performance. *Journal of Comparative Physiology*, *189*, 741–759.
- Gold, J., & Shadlen, M. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences*, *5*(1), 10–16.
- Gold, J., & Shadlen, M. (2002). Banburismus and the brain: Decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, *36*, 299–308.
- Goldman, M., Maldonado, P., & Abbott, L. (2002). Redundancy reduction and sustained firing with stochastic depressing synapses. *Journal of Neuroscience*, *22*(2), 584–591.
- Kay, S. (1998). *Fundamentals of statistical signal processing*. Upper Saddle River, NJ: Prentice Hall.
- Lewis, J., & Maler, L. (2002). Dynamics of electrosensory feedback: Short-term plasticity and inhibition in a parallel fiber pathway. *Journal of Neurophysiology*, *88*, 1695–1706.
- Lowen, S., & Teich, M. (1992). Auditory-nerve action-potentials form a nonrenewal point process over short as well as long-time scales. *Journal of the Acoustical Society of America*, *92*(2), 803–806.
- Maass, W., & Zador, A. (1999). Dynamic stochastic synapses as computational units. *Neural Computation*, *11*, 903–917.
- MacIver, M., Sharabash, N., & Nelson, M. (2001). Prey-capture behavior in gymnotid electric fish: Motion analysis and effects of water conductivity. *Journal of Experimental Biology*, *204*, 543–557.
- Maler, L., & Berman, N. (1999). Neural architecture of the electrosensory lateral line lobe: Adaptations for coincidence detection, a sensory searchlight and frequency-dependent adaptive filtering. *Journal of Experimental Biology*, *202*, 1243–1253.
- Markram, H. (2003). Elementary principles of nonlinear synaptic transmission. In R. Hecht-Nielsen & T. McKenna (Eds.), *Computational models for neuroscience* (pp. 126–169). Berlin: Springer-Verlag.
- Moustakides, G. (1986). Optimal stopping times for detecting changes in distributions. *Annals of Statistics*, *14*, 1379–1387.
- Nelson, M., & MacIver, M. (1999). Prey capture in the weakly electric fish *Apteronotus albifrons*: Sensory acquisition strategies and electrosensory consequences. *Journal of Experimental Biology*, *202*, 1195–1203.

- Nelson, M., Xu, Z., & Payne, J. (1997). Characterization and modeling of p-type electrosensory afferent responses to amplitude modulations in a wave-type electric fish. *Journal of Comparative Physiology*, *181*, 532–544.
- Page, E. (1954). Continuous inspection schemes. *Biometrika*, *41*, 100–114.
- Raciot, D., & Longtin, A. (1997). Spike train patterning and forecastability. *BioSystems*, *40*, 111–118.
- Rasnow, B. (1996). The effects of simple objects on the electric field of *Apteronotus*. *Journal of Comparative Physiology A*, *178*, 397–411.
- Ratnam, R., & Nelson, M. (2000). Nonrenewal statistics of electrosensory afferent spike trains: Implications for the detection of weak sensory signals. *Journal of Neuroscience*, *20*(17), 6672–6683.
- Strogatz, S. (1994). *Nonlinear dynamics and chaos*. Boulder, CO: Westview Press.
- Teich, M., Turcott, R., & Siegel, R. (1996). Temporal correlation in cat striate-cortex neural spike trains. *IEEE Transactions on Engineering in Medicine and Biology*, *15*, 79–87.
- Tsodyks, M., & Markram, H. (1997). The neural code between neocortical pyramidal cells depends on neurotransmitter release. *Proceedings of the National Academy of Sciences*, *94*, 719–723.
- Wald, A. (1948). *Sequential analysis*. New York: Wiley.
- Xu-Friedman, M., & Regehr, W. (2004). Structural contributions to short-term synaptic plasticity. *Physiological Reviews*, *84*, 69–85.
- Zucker, R., & Regehr, W. (2002). Short-term synaptic plasticity. *Annual Review of Physiology*, *64*, 355–405.