Rate Coding and Signal Processing

Fabrizio Gabbiani

Introduction

In the peripheral and central nervous system, many neurons encode information and pass it on to other neurons by generating irregular sequences of short voltage pulses, typically less than 1 ms in duration, called action potentials. The shape of these action potentials, or spikes, is usually quite stereotyped over the course of time. The sequence of spike occurrence times generated by the cell, often called the *spike train*, is therefore thought to carry most of the information that a neuron communicates to its targets. When studying how sensory information might be encoded in neuronal spike trains, one is faced with the fact that spike trains are often quite variable under seemingly identical stimulation conditions (Figure 1). Is this variability simply noise, perhaps due to uncontrolled changes in the state of some internal variable, or does it carry information about the stimulus? Answering this question in a particular case would require a thorough knowledge of the biophysical mechanisms of spike generation and of stimulus coding—knowledge that is out of reach at present.

Although no universal definition exists, the term *rate coding* is applied in situations where the precise timing of spikes is not thought to play a significant role in carrying sensory information. Rate codes have been identified in many sensory systems and are probably the best understood means by which neurons encode information. In many cases, rate codes have been shown to play an important role in determining behavioral responses of animals.

Rate coding comes in two flavors: mean firing rate codes and instantaneous firing rate codes. The sensory information conveyed by these two types of codes can be studied rigorously by applying classical methods of statistical signal processing borrowed from the engineering literature. In the next two sections, we will show how these methods can be carried over to the analysis of neuronal spike trains. Before turning to more general examples in the third section, we will illustrate them in the case of electrical field amplitude–sensitive neurons of weakly electric fishes. These animals possess an unusual sense for the electrical properties of their environment that is favorable to computational investigations (see ELECTROLOCATION).

Rate coding is not the only mean by which neurons convey information. In weakly electric fishes and in other auditory-like sensory systems, the role played by spike timing information is well documented (see ELECTROLOCATION and ECHOLOCATION: COCHLEOTOPIC AND COMPUTATIONAL MAPS). Two articles address the issue of spike timing in cortical circuits (SYNFIRE CHAINS; SYNCHRONIZATION: BINDING AND EXPECTANCY). Finally, the role of rate coding in the context of neuronal populations is examined in POPULATION CODES and MOTOR CORTEX, CODING AND DECODING OF DIRECTIONAL OPERATIONS.

Mean Firing Rate Coding

An increase in firing rate is typically the most conspicuous change recorded from sensory neurons in response to external stimuli. It is therefore natural to ask how well the spike count observed in a single trial from such a cell can predict the presence of the stimulus. Let us take the example of a neuron having a mean spontaneous rate $\bar{\lambda}_0 = 30$ spk/s that fires at a rate of $\bar{\lambda}_s = 50$ spk/s under stimulus presentation (Figure 1*B*). We will first assume for simplicity that spikes are generated completely independently of each other (i.e., following a Poisson process) and thus do not carry any additional information beyond their mean rate of occurrence.

Figure 2*A* illustrates the distribution of spike counts observed during a 200 ms window in the baseline and stimulus condition for this model neuron. The overlap between these two distributions indicates that guessing the presence of the stimulus from the spike

count observed in a single trial will lead to a significant fraction of errors. A simple method to decide between the two alternatives "stimulus present" or "no stimulus" consists in choosing a threshold number of spikes, k_{thres}, and classifying the observed responses, n, as baseline activity or stimulus-induced activity according to whether the threshold is exceeded or not:

$$\begin{aligned} n < k_{\text{thres}} &\Rightarrow \text{baseline activity} \\ n > k_{\text{thres}} &\Rightarrow \text{stimulus present} \end{aligned} \quad (1)$$

This decision strategy leads to two types of errors. On the one hand, we might call for the stimulus to be present in a trial during which spontaneous activity was unusually high. This type of error is called a false alarm. On the other hand, we might confuse an unusually low stimulus response with spontaneous activity, an error called a false miss. Clearly, the proportion of false alarms to false misses depends on the choice of the threshold k_{thres}: high (low) threshold values correspond to low (high) probabilities of false alarms with higher (lower) fractions of false misses. It is customary to characterize the performance of this detection algorithm by varying the threshold from low to high values and plotting the probability of correct detection, p_D (i.e., 1 minus the probability of false misses), as a function of the probability of false alarms, p_{FA} (i.e., 1 minus the probability of correct rejections; Figure 2*B*).

This curve is called the receiver operating characteristic (ROC) curve of the detection algorithm (a term originating from early applications to radar observations). The dashed diagonal line $p_D = p_{FA}$ corresponds to chance performance (i.e., independent of the threshold, k_{thres}, the probability, p_D, of correctly detecting the stimulus is as good as the probability, p_{FA}, of incorrectly mistaking spontaneous activity with stimulus-induced activity). Thus, the higher the ROC curve lies above the diagonal, the better the performance of our detection algorithm and, in the limit of perfect performance, $p_D = 1$ over the entire interval $0 \leq p_{FA} \leq 1$.

From the ROC curve, one can obtain the values of p_{FA} and $p_D(p_{FA})$ that minimize the overall probability of stimulus detection error (comprising both false alarms and false misses). If the stimulus is presented on average in one-half of the trials, the error rate is given, for a fixed value of p_{FA}, by

$$\varepsilon(p_{FA}) = \frac{1}{2} p_{FA} + \frac{1}{2} (1 - p_D(p_{FA})) \quad (2)$$

The minimum of $\varepsilon(p_{FA})$ as a function of p_{FA} can be easily found by numerical methods (see Figure 2*C*). The corresponding threshold k_{thres} may then be obtained from $p_{FA}(k_{\text{thres}})$. Thus, in our example the minimal error rate $\varepsilon = 0.24$ is achieved for $p_{FA} = 0.26$, corresponding to a detection threshold k_{thres} of 8.5 spk/s.

One important question remains: Given the simplicity of this algorithm, could it be outperformed by a more sophisticated one? Remarkably, this is not the case: under fairly general assumptions, the threshold condition of Equation 1 is equivalent to a similar condition on the likelihood ratio, $l(n) = p_s(n)/p_0(n)$, where $p_s(\cdot)$ and $p_0(\cdot)$ are the probability distributions of the spike count in the presence and absence of the stimulus, respectively (Figure 2*A*). The likelihood ratio is a quantity central to signal detection theory, and this equivalence shows that, for a fixed value of p_{FA}, no other algorithm taking into account only the probability distributions of Figure 2*A* can outperform the threshold test. Thus, the ROC curve defines the performance of an ideal observer of the mean rate code, having complete access to $p_0(\cdot)$ and $p_s(\cdot)$. Whether neurons or neuronal networks in the brain adopt similar algorithms to read out information about the external world remains an open question.

The performance of the ideal observer algorithm will be affected by at least two additional factors, the first being the length of the time window over which spikes are registered. Longer windows typically lead to better performance by averaging out the noise component of the spike rate that causes deviations from the mean. In the case of the Poisson process discussed above, for a fixed mean firing rate $\bar{\lambda}$, the mean number of spikes observed in a time window T is given by $\bar{n} = \bar{\lambda}T$, whereas the standard deviation is $\sigma = \sqrt{\bar{\lambda}T}$. Thus, $\bar{n}/\sigma \propto \sqrt{T}$, and the signal grows as $\sqrt{T}$ with respect to noise over the course of time. Currently, the time interval that is relevant for behavioral responses often is only weakly constrained

by experimental observations.

The second factor is the regularity of the spike train or, in other words, the amount of noise that is present to start with. While many neurons in cortical areas have highly variable responses resembling those obtained from Poisson processes, other neurons can be much more regular. In the weakly electric fish *Apteronotus*, for example, the spike trains of primary sensory afferent neurons sensitive to amplitude modulations of the electrical field have a variability that is almost an order of magnitude smaller than that expected from a Poisson spike train on behaviorally relevant time scales (Ratnam and Nelson, 2000). These neurons are thought to encode information necessary for the detection of small prey, such as the water fleas on which the fish feeds. Computer simulations, behavioral observations, and electrophysiological recordings suggest that the firing rate of these cells will increase by only a few spikes per second during the 200 ms needed to detect the prey. An ROC analysis reveals that increases of 2–3 spk/s above baseline activity can be detected with greater than 90% accuracy even if the probability of a false alarm is very low, 0.1% (Figure 2*D*). Such low false alarm rates ($p_{FA} \leq 0.001$) are constrained from behavioral observations showing that fishes almost never strike a nonexistent prey. As illustrated in Figure 2*D*, three spike trains of models designed to reproduce the short-term variability of the experimental spike trains cannot reproduce these results. The first model (squares) reproduces only the mean firing rate of the afferents, while the second and third models also reproduce the interspike interval distribution (diamonds) or the joint statistical distribution of two successive interspike intervals (triangles), respectively. The regularity and statistical structure of the spike trains over at least three firing cycles is therefore responsible for this unusually low detection threshold.

Instantaneous Rate Coding

Stimuli that vary on a fast time scale—comparable to the 200 ms observation window introduced in the last section—cannot be encoded by the mean spike count alone. Such stimuli are ubiquitous in the sensory environment of many animals. Motion of an object or self-motion, for example, result in rapid changes in light intensity across the visual field. Sound stimuli used for communication or localization correspond to rapidly varying changes in air pressure. In the case of weakly electric fishes considered in the last section, time-varying electrical field amplitude modulations occur as the fish moves through an electrically dense environment in water.

Such time-varying stimuli could be encoded by time-varying changes in the instantaneous firing rate of a neuron, even if the precise timing of spikes does not play an essential role in the process (Gabbiani and Koch, 1999). Consider, for example, the Poisson spike train model of the previous section with a spontaneous rate $\bar{\lambda} = 30$ spk/s. We assume that changes in the instantaneous firing rate from its mean value, $\bar{\lambda}$, are caused by changes of the stimulus, $s(t)$, from its mean value, $\bar{s}$,

$$\lambda(t) = \alpha(s(t) - \bar{s}) + \bar{\lambda} \tag{3}$$

The constant α is a conversion factor between stimulus and firing rate, and $\lambda(t)$ is assumed to be positive. In the following discussion the stimulus will usually be assumed to have zero mean, i.e., $\bar{s} = 0$.

How much information does such a spike train convey about the stimulus? Using an approach analogous to that introduced in the last section, this question can be addressed by presenting a random stimulus $s(t)$ and estimating it from the spike train (Figure 3*A*). Because $s(t)$ varies randomly in time, the estimate $s_{\text{est}}(t)$ will also have to vary in time to track $s(t)$. Thus, this estimation problem is more complex than the detection problem considered in the last section. It is customary to minimize the root mean square error between the stimulus and its estimate,

$$\varepsilon = \langle (s(t) - s_{\text{est}}(t))^2 \rangle^{1/2} \tag{4}$$

where the average is taken over the stimulus presentation interval (Figure 3*A*). It is much more difficult to find an optimal estimate $s_{\text{est}}(t)$ from the spike train than it is to find an optimal classification strategy based on the spike count. A simplification is therefore

made by looking only at estimates obtained from linear superpositions of a waveform, $h(t)$, around each spike. If $r(t) = \Sigma_i \delta(t - t_i)$ is a sum of delta functions representing the sequence of spikes at times $\{t_i\}$, the estimated stimulus is assumed to be of the form

$$s_{\text{est}}(h, t) = \int h(t - t_0) r(t_0) dt_0 - \bar{r} \int h(t_0) dt_0 = \sum_i h(t - t_i) - \bar{r} \int h(t_0) dt_0 \quad (5)$$

where $\bar{r}$ is the mean firing rate of the cell and the second term ensures that $s_{\text{est}}(t)$ has zero mean value, as was assumed for $s(t)$ (Figure 3*B*). Under this assumption, the optimal waveform, $h(t)$, minimizing the root mean square error in Equation 4 can be obtained by standard statistical and Fourier transform techniques. The minimum value obtained for the root mean square error, ε, is usually normalized by the stimulus standard deviation, σ, which corresponds to chance guessing (i.e., to the error obtained in Equation 4 when $s_{\text{est}}(t) = \bar{s} = 0$). Figure 3*C* illustrates the result of this estimation procedure using the spike train of an amplitude-sensitive afferent obtained in response to a random electrical field amplitude modulation in a second species of weakly electric fishes, *Eigenmannia*. From the spike train, the amplitude modulation could be estimated with an error $\varepsilon/\sigma = 0.17$. Equivalently, our ideal linear observer could reproduce 83% of the stimulus standard deviation using a single spike train.

In contrast, estimation of the same stimulus using a spike train generated using a Poisson process and Equation 3 is considerably less accurate (only 25% of the stimulus standard deviation is recovered; Figure 3*D*), because a Poisson process is more variable than the spiking of *Eigenmannia* afferents (Kreiman et al., 2000). Thus, as in the signal detection case, spike train variability plays an important role in stimulus estimation. Other factors that play a significant role in the encoding capacity of the instantaneous firing rate are the contrast of the stimulus (or its standard deviation σ; typically, higher contrasts result in larger firing rate modulations and thus better encoding), the mean firing rate of the cell (higher mean firing rates lead to a finer temporal sampling of the stimulus), and the cutoff frequency of the stimulus (accurate encoding is possible only when temporal stimulus frequencies are well below the mean firing rate of the cell).

The assumption relating spike train and stimulus estimate by a linear transform embodied in Equation 5 works very well in practice when the encoding of the stimulus by the spike train can be described by equations analogous to Equation 3. This result can be justified theoretically (Gabbiani and Koch, 1999). In contrast, no systematic studies have been carried out on the effect of nonlinear relations between stimulus and firing rate; only a few scattered examples have been examined (Gabbiani and Koch, 1999; Haag and Borst, 1998).

Rate Coding in Neural Systems

Instantaneous and mean firing rate codes have been extensively characterized in a variety of different neuronal systems. In the following discussion, we will highlight a few directions of investigation and some open questions relevant to the subject.

Starting in the late 1940s, signal detection methods have been applied to characterize the information conveyed by neuronal spike trains, along the lines prescribed earlier in this article (see Parker and Newsome, 1998). The investigation of neuronal signals carried by optic nerve fibers of the horseshoe crab *Limulus* was one of the earliest examples of work on this topic (see Parker and Newsome, 1998). Over the next 30 years, signal detection methods were extended to the activity of sensory neurons in the early auditory, somatosensory, and visual pathways of vertebrates. The variability of retinal ganglion cell spike trains, for example, has been extensively investigated in an attempt to explain its impact on the encoding reliability of visual signals (Parker and Newsome, 1998).

More recently, signal detection methods have been applied to neurons in cerebral cortical areas (visual and somatosensory, for instance) of monkeys trained to perform discrimination tasks. In some cases, the reliability of neuronal firing could be compared to

the behavioral accuracy of the animal performing the task. These results, together with analyses of variability and correlation across cells, have led to neural models of signal encoding that can account for the animal's behavior (see Parker and Newsome, 1998). The neural mechanisms underlying behavioral selection in those discrimination tasks, however, remain difficult to test experimentally.

In the cockroach, directional escape responses to wind stimuli are thought to rely on the mean firing rate of 14 giant interneurons (GIs). Several models that could in principle explain escape behaviors on the basis of the mean firing rate of GIs have been tested by directly manipulating them through current injections (Levi and Camhi, 2000). The results of these experiments suggest that a directional average of the GIs' mean firing rate is the most accurate description of the behavior. Mean firing rate codes across population of neurons have also been shown to play similar roles in vertebrate neurons, in the generation of visual saccades in the superior colliculus of monkeys, and in the generation of limb movements in motor cortical areas (Sparks, Kristan, and Shaw, 1997).

Given that in the engineering literature signal estimation is usually considered a close relative of signal detection (Poor, 1994), it is perhaps surprising that it has been applied to neural spike trains only within the past 10 years. Estimation of time-varying stimuli along the lines developed earlier in this article has shown that single spike trains of sensory neurons can accurately convey information about time-varying stimuli, although the results are usually less spectacular than those shown in Figure 3*B* (Borst and Theunissen, 1999). At present, these methods have been applied mainly to invertebrate and lower vertebrate preparations. Mechanisms of encoding across multiple neurons and their relation to behavior have received little attention so far (but see Stanley, Li, and Dan, 1999).

In contrast, instantaneous firing rate codes have been extensively studied by characterizing how stimulus attributes are encoded in the instantaneous firing rate of neurons through generalizations of Equation 3. Such models are particularly well developed for the early visual pathways of mammals, from the retina to early visual cortical areas (Dayan and Abbott, 2001).

Discussion

Mean and instantaneous firing rate codes are undoubtedly the best documented and best understood way by which neurons transmit information. Several other codes have also been studied, among them the mechanisms of coincidence detection in auditory processing (Pena and Konishi, 2001). More elaborate coding schemes are likely to be found, particularly across populations of neurons, although the highly sophisticated codes at the heart of information theory seem unlikely to find a place in describing the signaling repertoire of sensory and motor neurons.

One question that has long intriguied neuroscientists is whether the spike train variability usually observed in neurons using rate coding also carries further sensory information (Bullock, 1970). This question is difficult to answer rigorously. In the case of the cockroach, the pattern of spikes in GIs does not appear to play a role in determining escape behaviors (Liebenthal, Uhlmann, and Camhi, 1994). On the other hand, it has been suggested that in electric fishes, coincidence detection could be used to integrate the information conveyed by the amplitude-sensitive receptors described in this article and in Berman and Maler (1999). Thus, neurons might use a combination of different codes simultaneously at different levels of a neuronal circuit.

Roadmap: Neural Coding
Related Reading: Population Codes; Spansory Coding and Information Transmission

References

Borst, A., and Theunissen, F. E., 1999, Information theory and neural coding, *Nature Neurosci.*, 2:947–957. ◆

Berman, N. J., and Maler, L., 1999, Neural architecture of the electrosensory lateral line lobe: Adaptations for coincidence detection, a sensory searchlight and frequency-dependent adaptive filtering, *J. Exp. Biol.*, 202:1243–1253.

Comp: asterisk becomes diamond symbol at end of entryBullock, T. H., 1970, The reliability of neurons, *J. Gen. Physiol.*, 55:565–584. ◆

Dayan, P., and Abbott, L. F., 2001, *Theoretical Neuroscience*, Cambridge, MA: MIT Press. ◆

Gabbiani, F., and Koch, C., 1999, Principles of spike train analysis, in *Methods in Neuronal Modeling: From Synapses to Networks* (C. Koch and I. Segev, Eds.), 2nd ed. Cambridge, MA: MIT Press, pp. 313–360. ◆

Haag, J., and Borst, A., 1998, Active membrane properties and signal encoding in graded potential neurons, *J. Neurosci.*, 18:7972–7986.

Kreiman, G., Krahe, R., Metzner, W., Koch, C., and Gabbiani, F., 2000, Robustness and variability of neuronal coding by amplitude-sensitive afferents in the weakly electric fish *Eigenmannia, J. Neurophysiol.*, 84:189–204.

Levi, R., and Camhi, J. M., 2000, Population vector coding by the giant interneurons of the cockroach, *J. Neurosci.*, 20:3822–3829.

Liebenthal E., Uhlmann, O., and Camhi, J. M., 1994, Critical parameters of the spike trains in a cell assembly: Coding of turn direction by the giant interneurons of the cockroach, *J. Comp. Physiol. A*, 174:281–296.

Parker, A. J., and Newsome, W. T., 1998, Sense and the single neuron: Probing the physiology of perception, *Annu. Rev. Neurosci.*, 21:227–277. ◆

Pena, J. L., and Konishi, M., 2001, *Science*, 292:249–252.

Poor, H. V., 1994, *An Introduction to Signal Detection and Estimation*, New York: Springer-Verlag. ◆

Ratnam, R., and Nelson, M. E., 2000, Nonrenewal statistics of electrosensory spike trains: Implications for the detection of weak sensory signals, *J. Neurosci.*, 20:6672–6683.

Sparks, D. L., Kristan, W. B., and Shaw, B. K., 1997, The role of population coding in the control of movement, in *Neurons, Networks, and Motor Behavior* (P. S. G. Stein, S. Grillner, A. I. Selverston, and D. G. Stuart, Eds.), Cambridge, MA: MIT Press, pp. 21–32. ◆

Stanley, G. B., Li, F. F., and Dan, Y., 1999, Reconstruction of natural scenes from ensemble responses in the lateral geniculate nucleus, *J. Neurosci.*, 19:8036–8042.

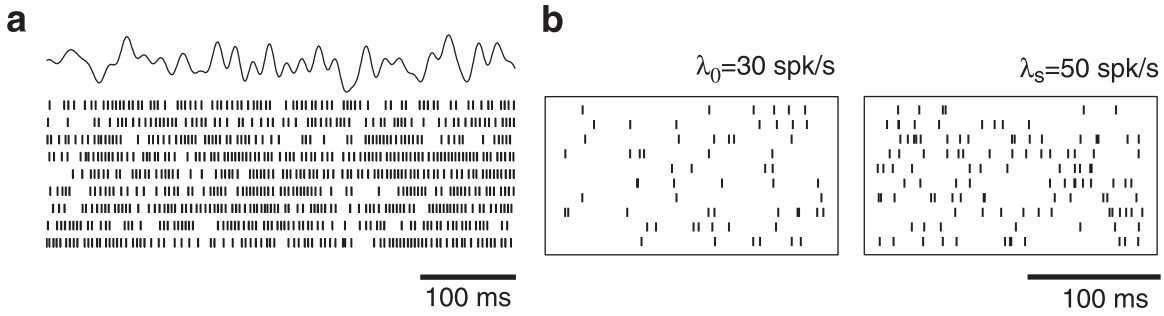

Figure 1. *A*, Nine spike trains recorded from an amplitude-sensitive afferent in the weakly electric fish *Eigenmannia* in response to repeated presentations of the same random electrical field amplitude modulation (shown on top). (Adapted from Kreiman et al., 2001.) *B*, Ten spike trains (200 ms long) obtained from a Poisson process with mean firing rate $\lambda_0 = 30$ spk/s (spontaneous rate, left) and $\lambda_s = 50$ spk/s (stimulus-induced rate, right).

391

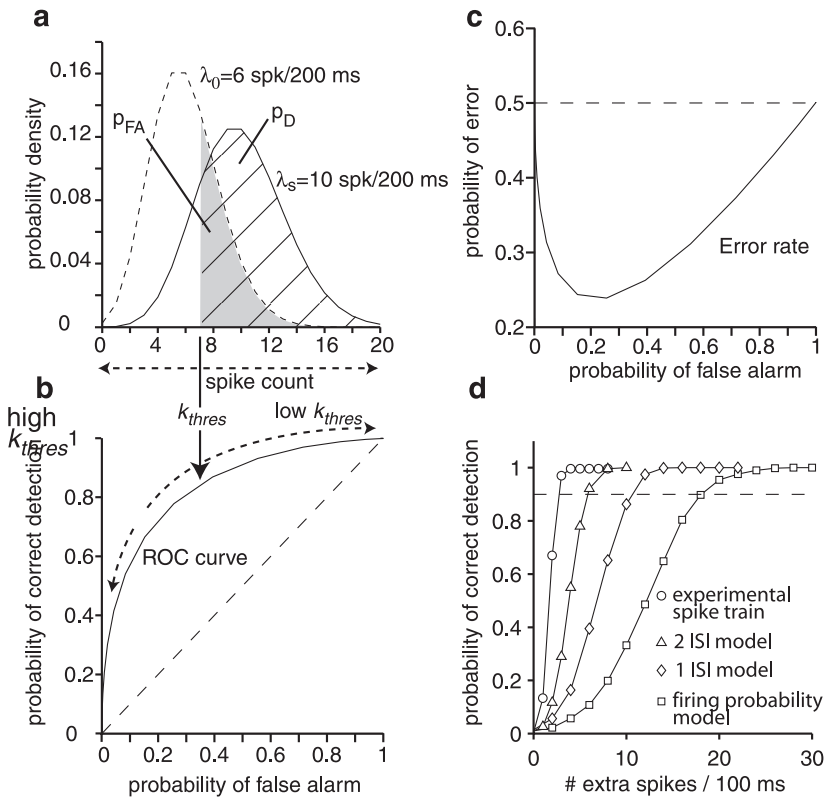

Figure 2. *A*, Probability distributions of the spike count observed in a 200 ms window for the two Poisson processes illustrated in Figure 1*B*. Choosing a threshold number of spikes (k_{thres}) to discriminate between the presence or absence of the stimulus leads to errors because of the overlap of the distributions. The probability of false alarm (p_{FA}) and of correct detection (p_D) are illustrated by the gray and hatched areas, respectively. *B*, Plot of p_D versus p_{FA}, called an ROC curve. Different thresholds will correspond to different values of p_D and p_{FA} (dashed double arrows in *A* and *B*). *C*, Overall probability of error (see Equation 2), computed from the ROC curve in *B*. *D*, Probability of correct detection obtained from the spike trains of an amplitude-sensitive afferent neuron as a function of the number of spikes above spontaneous activity generated by the cell (circles). Note that the cell can discriminate with more than 90% accuracy increases of three spikes or more, corresponding to a 1% increase in firing rate. Different models that take into account only the mean firing rate (squares), the mean firing rate and the interspike interval distribution (diamonds), or in addition the joint properties of successive intervals (triangles) are unable to match the experimental performance. (Adapted from Ratnam and Nelson, 2000.)

392

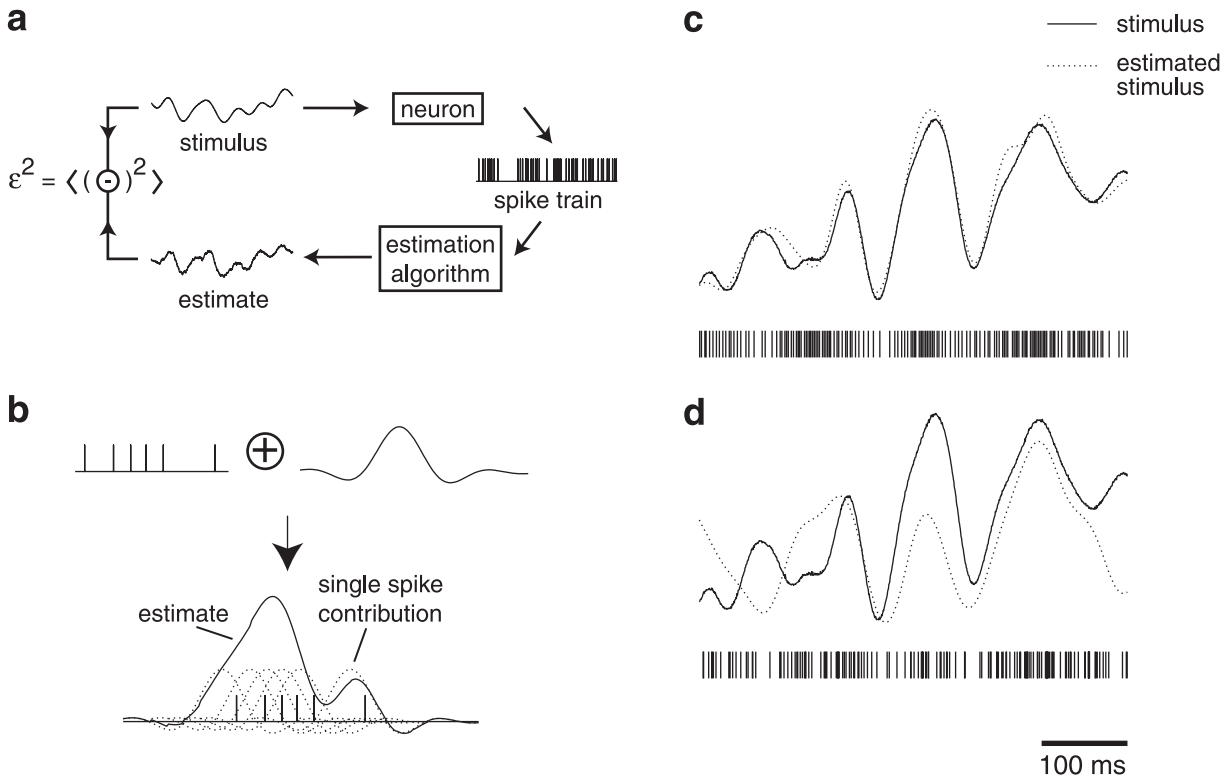

Figure 3. *A*, Stimulus estimation is performed using a linear algorithm (see *B*) that is based on a comparison of the stimulus and its estimate aimed at minimizing the mean square error between the two (Equation 4). *B*, The linear algorithm consists in taking a spike train (left) and placing a waveform (right) around each spike. Linear superposition of these waveforms (bottom) yields the estimate. The waveform is chosen to minimize the mean square error between stimulus and estimate (see *A*). *C*, Estimation of a random amplitude modulation from the spike train of an amplitude sensitive afferent neuron in *Eigenmannia* (mean firing rate 314 spk/s). *D*, Same stimulus estimated from a Poisson spike train encoding the stimulus according to Equation 3 at the same mean firing rate as in *C*.