Receiver Operating Characteristic (ROC) Analysis for Characterizing Synaptic Efficacy

Frances S. Chance

J Neurophysiol 97:1799-1808, 2007. First published Dec 13, 2006; doi:10.1152/jn.00885.2006

J Neurophysiol 97: 1799–1808, 2007.
First published December 13, 2006; doi:10.1152/jn.00885.2006.

Receiver Operating Characteristic (ROC) Analysis for Characterizing Synaptic Efficacy

Frances S. Chance
Department of Neurobiology and Behavior, University of California, Irvine, California

Submitted 16 August 2006; accepted in final form 12 December 2006

Chance FS. Receiver operating characteristic (ROC) analysis for characterizing synaptic efficacy. *J Neurophysiol* 97: 1799–1808, 2007. First published December 13, 2006; doi:10.1152/jn.00885.2006. The role of background synaptic activity in cortical processing has recently received much attention. How do individual neurons extract information when embedded in a noisy background? When examining the impact of a synaptic input on postsynaptic firing, it is important to distinguish a change in overall firing probability from a true change in neuronal sensitivity to a particular input (synaptic efficacy) that corresponds to a change in detection performance. Here we study the impact of background synaptic input on neuronal sensitivity to individual synaptic inputs using receiver operating characteristic (ROC) analysis. We use the area under the ROC curve as a measure of synaptic efficacy, here defined as the ability of a postsynaptic action potential to identify a particular synaptic input event. An advantage of using ROC analysis to measure synaptic efficacy is that it provides a measure that is independent of postsynaptic firing rate. Furthermore, changes in mean excitation or inhibition, although affecting overall firing probability, do not modulate synaptic efficacy when measured in this way. Changes in overall conductance also affect firing probability but not this form of synaptic efficacy. Input noise, here defined as the variance of the input current, does modulate synaptic efficacy, however. This effect persists when the change in input variance is coupled with a change in conductance (as would result from changing background activity).

INTRODUCTION

Cortical neurons receive a continuous barrage of excitatory and inhibitory synaptic inputs that can significantly affect single-neuron processing in vivo (Destexhe and Paré 1999; Destexhe et al. 2001, 2003; Holt et al. 1996; Paré et al. 1998; Tiesinga et al. 2000). Much of this synaptic bombardment does not appear to be correlated with external stimuli (Heggelund and Albus 1978; Holt et al. 1996; Schiller et al. 1976; Softky and Koch 1993; Vogels et al. 1989), yet it has a significant impact on the membrane properties of in vivo neurons. Neurons have higher overall membrane conductances in vivo (Borg-Graham et al. 1998; Destexhe and Paré 1999; Hirsch et al. 1998; Paré et al. 1998; Steriade et al. 2001; Woody and Gruen 1978), presumably because of the high "background" of synaptic input. This background synaptic input also introduces a large amount of variability into responses of cortical neurons (Holt et al. 1996; Shadlen and Newsome 1998; Softky and Koch 1993; Stevens and Zador 1998).

In addition to increasing response variability, the high variance in the synaptic input current, referred to here as "input noise" (although this term is not meant to imply that such activity serves no useful function in cognition), enhances neuronal responses to subthreshold synaptic inputs (Anderson et al. 2000; Azouz 2005; Braun et al. 1994; Collins et al. 1995; Douglass et al. 1993; Hô and Destexhe 2000; Shu et al. 2003; Troyer and Miller 1997; Wenning et al. 2005). This enhancement of responses to otherwise subthreshold inputs is coupled with a decrease in the neuronal gain, defined in these studies by the relationship between output firing rate and input current (Chance et al. 2002; Fellous et al. 2003; Prescott and De Koninck 2003), excitatory postsynaptic potential (EPSP) rate (Mitchell and Silver 2003), or EPSP size (Hô and Destexhe 2000; Shu et al. 2003). Many excitatory inputs within this background barrage may be "balanced" or "cancelled out" by inhibitory inputs and not lead directly to depolarization, although they still affect neuronal responses to other inputs. How does a cortical neuron detect and transmit sensory information embedded within this background of seemingly random synaptic activity?

Figure 1*A* is a firing rate histogram of a model leaky integrate-and-fire neuron receiving a constant background of noisy random excitatory and inhibitory synaptic input (see METHODS). The baseline, or average, firing rate (indicated by the dashed line) was between 6 and 7 Hz. One particular excitatory synaptic input—referred to as the "input event"—arrived at a specific time in each trial (labeled as "0 ms" in the figure) and triggered a clear elevation in firing rate shortly after its arrival (Fig. 1*A*). We seek to measure the ability of the postsynaptic neuron to detect and report this input event by firing an action potential and define this quantity as its *synaptic efficacy*.

A common approach to describing synaptic efficacy is to measure the "response probability" or probability of an action potential occurring shortly after the input event (e.g., Hô and Destexhe 2000; Shu et al. 2003). Although this approach is useful for relating synaptic efficacy to the size of the input current (Hô and Destexhe 2000; Shu et al. 2003; also see London et al. 2002) or correlation between input events (Azouz 2005), it has the disadvantage that it requires that the baseline firing rate be specified. In Fig. 1*B*, we plot the probability that an action potential is fired within a 15-ms period after the arrival of the input event spike as a function of the size of the excitatory conductance change it produces. For each of the three different traces, current was injected to produce a different baseline firing rate. Changing the baseline firing rate significantly affected response probability, regardless of input

Address for reprint requests and other correspondence: F. Chance, University of California, Department of Neurobiology and Behavior, 2205 McGaugh Hall, Irvine, CA 92697-4550 (E-mail: fchance@uci.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked "*advertisement*" in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

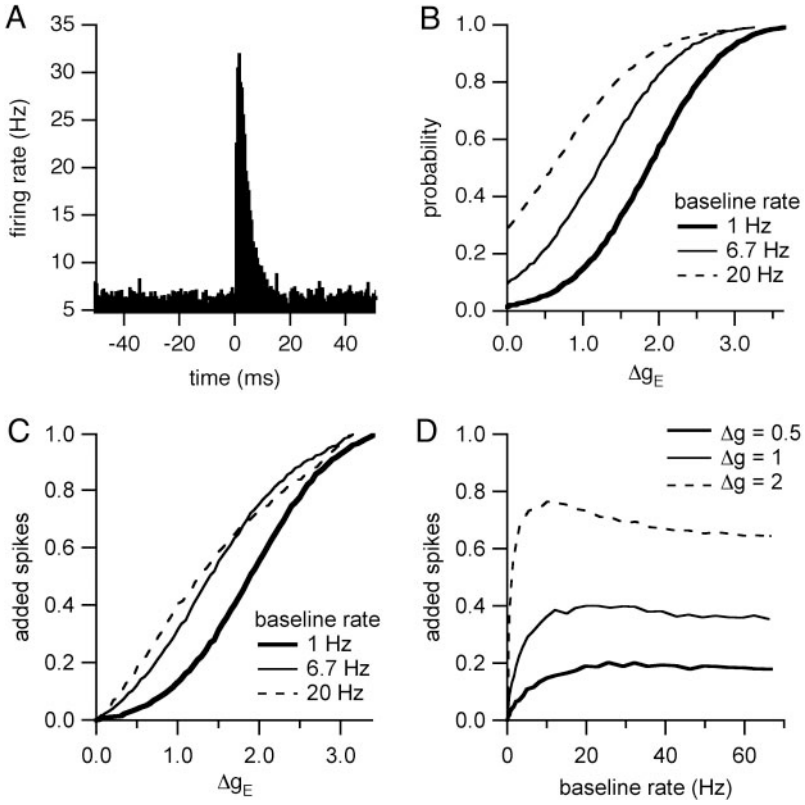

FIG. 1. *A*: firing rate histogram of the model neuron receiving noisy synaptic input (each bin is 1 ms). Synaptic inputs arrived randomly except at time t = 0 ms, when the input event occurred. Dashed line indicates the baseline firing rate. *B*: probability of a spike occurring within 15 ms after the arrival of the input event is plotted as a function of the maximum excitatory conductance change (given relative to g_0, the resting membrane conductance) triggered by the input event. For the thick solid line, the thin solid line, and the dashed line, the baseline firing rate was 1, 6.7, or 20 Hz, respectively. *C*: average number of spikes fired within a 15-ms window immediately after the input event minus the expected number based on the baseline firing rate. Baseline firing rates for the different traces are as in *B*. *D*: same quantity as in *C* plotted as a function of baseline firing rate. Thick solid line, thin solid line, and dashed line represent input event sizes of 0.5, 1.0, and 2.0, respectively.

event size (Azouz 2005; Reyes and Fetz 1993b; Shu et al. 2003; Türker and Powers 2002; Fig. 1*D*).

Compensating for the dependency of synaptic efficacy on baseline firing rate is not straightforward. In Fig. 1, *C* and *D*, synaptic efficacy is measured as the average number of spikes fired within 15 ms after the input event minus the average number of spikes fired within a 15-ms window without the input event. This quantity (referred to in the figure as "added spikes") is plotted as a function of input event size in Fig. 1*C* and is very similar to the "cumulative sum" technique for measuring synaptic efficacy (Ellaway 1977, 1978; Türker and Powers 2002). The differences between the three curves, representing three different baseline firing rates, demonstrate that simply subtracting baseline firing rate does not remove the dependency of synaptic efficacy on baseline firing rate.

We plot synaptic efficacy, measured as added spikes, as a function of baseline firing rate for three different input event sizes in Fig. 1*D*. At higher baseline firing rates, the curves in Fig. 1*D* tend to be flat, indicating that the added spike measure is relatively independent of baseline firing rate in this regime. The effect of baseline firing rate outside of this regime, however, is nonmonotonic, indicating that measures that normalize for baseline firing rate by dividing by the baseline firing rate (e.g., Fetz and Gustaffson 1983) or by subtracting the baseline firing rate (such as the "added spikes" or "cumulative sum" techniques) will not fully compensate for the effect of baseline firing rate. Also, the range of baseline firing rates for which the curves in Fig. 1*D* are flat varies depending on the size of the input event. This indicates that it may be difficult to know what the effect of baseline firing rate will be when measuring synaptic efficacy using these methods.

Clearly synaptic strength is not the only factor that determines response probability and the baseline firing rate can significantly affect its value. Because of this confound and the fact that in vivo cortical neuron firing rates rarely maintain a stable value, it would be valuable to have a measure of synaptic efficacy that is independent of firing rate. In situations where baseline firing rate varies significantly, such as when describing the effect of one input presented against a varying background (the example we will examine here), such a measure would be particularly useful for describing the efficacy of that input. Our proposed method of measuring synaptic efficacy provides a firing-rate–independent measure that should be particularly useful in such situations.

Receiver operating characteristic (ROC) analysis was used to determine the ability of neurons to signal changes in stimulus intensity in their firing rate responses (Britten et al. 1992; Cohn et al. 1975; Fitzhugh 1957; Goense et al. 2003; Shofner and Dye 1989; Tolhurst et al. 1983; for a review see Cohn 1977). These studies use ROC analysis to determine how effectively neurons can respond to weak stimuli, essentially determining a stimulus-detection threshold. ROC analysis was also previously used to examine how the response mode, bursting or tonic, of lateral geniculate nucleus (LGN) neurons affects their sensitivity to visual stimuli (Guido et al. 1995). Here, we use the area under an ROC curve as a measure of synaptic efficacy, describing the ability of a neuron to detect a particular input event within a "background" of many other inputs. One advantage to using this measurement is that it is independent of postsynaptic firing rate, a factor that affects all other measurements of synaptic efficacy of which we are aware. Although this measurement can be used to describe how synaptic strength or EPSP size affects synaptic efficacy, in this study we focus on the effects of changing levels of background activity as exist in vivo.

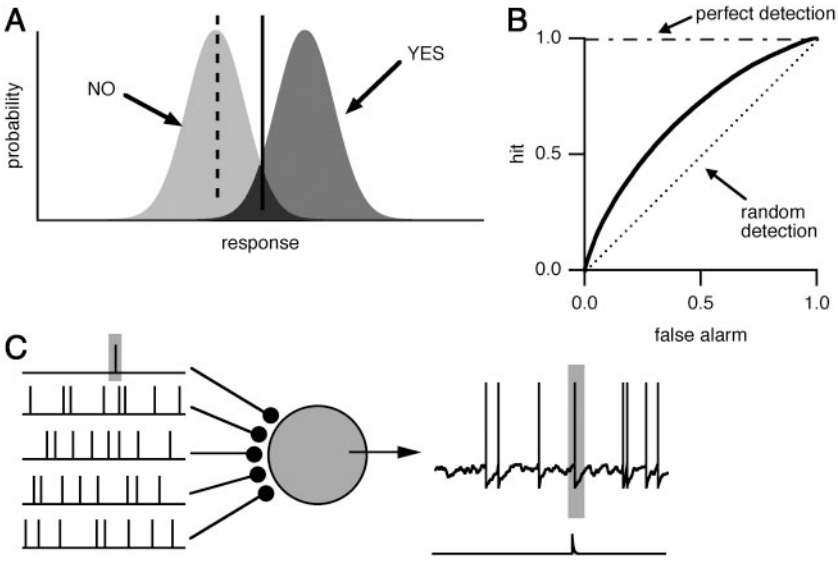

FIG. 2. *A*: example probability distributions for responses of a system to an event (dark gray) or in the absence of an event (light gray). Two different detection thresholds are illustrated by the solid and dashed lines. *B*: receiver operating characteristic (ROC) curve (solid line) is a plot of the probability that an event will be detected (hit) against the false-alarm rate for all possible detection thresholds. ROC curves corresponding to random responses and perfect detection are indicated by the dotted line and dash–dotted line, respectively. *C*: a neuron detecting an input presented against a background of excitatory and inhibitory synaptic inputs. Traces on the *right* are the membrane potential of the model neuron (*top*) and the excitatory postsynaptic current (EPSC) arising from the input event (*bottom*). Input event and the resulting action potential are highlighted.

Traditionally, ROC analysis has been used to distinguish between two distributions of responses (Fig. 2). The *index of discriminability* (d′) value describes how well these two response distributions are separated from each other (Green and Swets 1988; also see METHODS and Fig. 2*A*). For the purpose of measuring synaptic efficacy, we have found it useful to interpret the underlying distributions of neuronal membrane potential, with and without the input event, as these two response distributions, with the neuron acting as an "observer" of its own membrane potential and using an action potential to report the results of this observation.

Examining the effect of background activity on the membrane potential distributions of a model neuron is a useful way to study how synaptic efficacy is affected by background activity and d′ is an effective measure of synaptic efficacy. However, measuring membrane potential distributions and their d′ value is difficult, particularly in vivo. A more easily measured quantity that is related to the d′ value of the underlying response distributions is the area under the ROC curve (Fig. 2*B*). This provides an effective, firing-rate–independent measurement of neuronal sensitivity to synaptic inputs. We discuss how background activity affects the membrane potential distributions underlying action potential generation and the area under the resulting ROC curve. Through this analysis, we study how the different components of background activity—overall conductance, mean excitation, and high-input noise—modulate synaptic efficacy, defined in this way.

METHODS

Model

We study a standard leaky integrate-and-fire neuron model with excitatory and inhibitory synaptic inputs. In this model, the membrane potential V is determined by

$$\tau \frac{dV}{dt} = V_0 - V + g_K(E_K - V) + g_E(E_E - V) + g_I(E_I - V) + \frac{I}{g_0}$$

In the preceding equation and in the rest of this paper, all conductances are expressed relative to g_0, the resting membrane conductance in the absence of any synaptic input. The three variables g_E, g_I, and g_K are, respectively, the excitatory synaptic conductance, the inhibitory synaptic conductance, and a potassium conductance. The reversal potentials of these conductances are $E_E = 0$ mV for the excitatory synaptic conductance and $E_I = E_K = -80$ mV for the inhibitory synaptic conductance and the potassium conductance. In the absence of any additional input, the resting membrane time constant τ is 20 ms and the resting membrane potential V_0 is -57.8 mV. When the neuronal membrane potential depolarizes above -52 mV, an action potential is fired and the membrane potential is immediately set to -70 mV. A short refractory period is generated immediately after each action potential by increasing g_K to 5.0. Between action potentials, g_K exponentially decays to zero with a time constant of 5 ms.

Excitatory and inhibitory synaptic inputs are generated as if arising from presynaptic populations of randomly firing neurons, simulated using Poisson spike trains with underlying rates of 1,500 Hz for excitatory inputs and 2,600 Hz for inhibitory inputs. With each presynaptic excitatory input, g_E is increased by 0.16 and, with each inhibitory input, the inhibitory synaptic conductance g_I is increased by 0.24. The equivalent reversal potential of the synaptic conductances with these parameters is V_0. Between synaptic events, both g_E and g_I decay exponentially to zero with a time constant of 5 ms. To create the higher level of synaptic activity (Fig. 6), the rates of excitatory and inhibitory inputs were tripled. To create the higher level of membrane conductance (without the corresponding increase in input noise), the membrane leak conductance was set to 9.65, the equivalent conductance in the presence of this higher level of synaptic activity, rather than increasing the background input rate. To create the higher level of noise, the size of each excitatory and inhibitory input was tripled and input rates were divided by three (500 and 867 Hz for excitatory and inhibitory inputs, respectively). This tripled the variance of the synaptic current without significantly changing the overall conductance.

At a specified time in each trial, a large excitatory synaptic input is triggered. This "input event" increases the excitatory synaptic conductance by 0.5. In the absence of any additional synaptic activity, the input event evoked an EPSP with a peak voltage deflection of 4.3 mV. For Figs. 2–6, the model neuron was driven at a particular firing rate by injecting additional constant current I. In Fig. 7 the model neuron was driven at different firing rates by varying the rate of inhibitory input.

ROC analysis

We use ROC curves to describe the sensitivity of a neuron to input events. This measure is commonly used in signal detection theory to

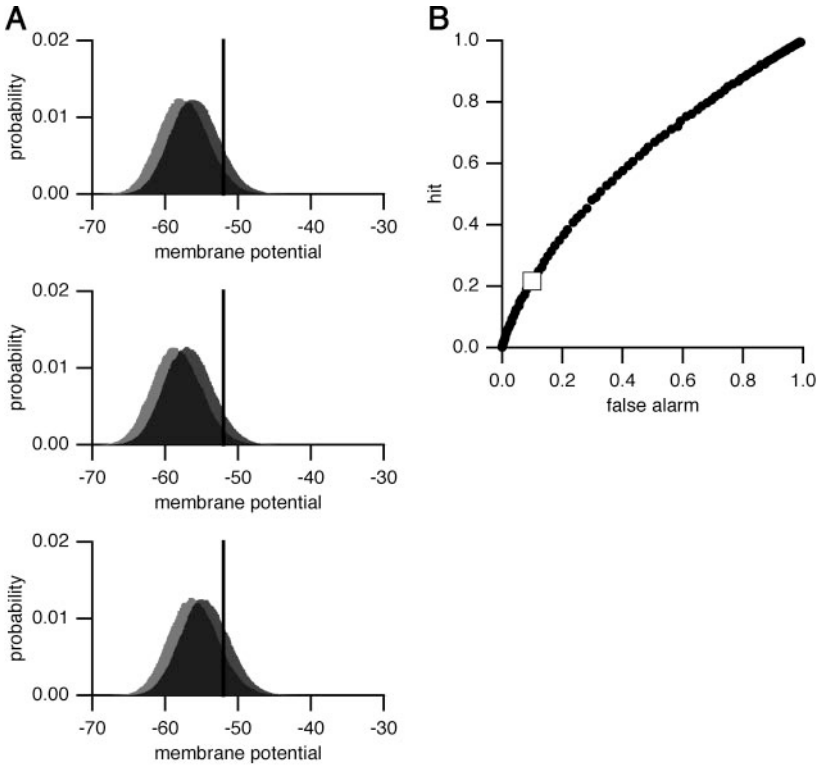

FIG. 3. *A*: membrane potential distributions of the model neuron without the input event (*left*, referred to as "no") and within 5 ms after the input event (*right*, referred to as "yes"). Solid vertical line indicates the action potential threshold. For the *top panel*, the "no" distribution has a mean of −57.9 mV and a SD of 2.134; the "yes" distribution has a mean of −56.2 mV and a SD of 2.136. For the *middle panel*, the neuron was hyperpolarized with constant injected current. Means of the "no" and "yes" distributions are −58.8 and −57.16, and the SDs are 2.129 and 2.118, respectively. For the *bottom panel*, the neuron was depolarized with constant injected current and the means (and SDs) for the "no" and "yes" distributions are −56.5 (2.132) and −54.8 (2.133). *B*: probability of an action potential within 15 ms after the input event (hit) plotted against the probability of an action potential occurring within a 15-ms window that does not contain the input event (false alarm). Open square represents the responses of the neuron for zero injected current (*A*, *top*). To generate other points on the curve, the neuron was hyperpolarized or depolarized by different levels of constant injected current.

describe the performance of an observer. We consider the case where the observer must decide whether a stimulus was present based on the responses of the system. Figure 2*A* illustrates an example of the response probability distributions of such a system. The probability of a response if a stimulus occurred is plotted in dark gray and labeled "yes"; the probability of a response in the absence of the stimulus is plotted in light gray and labeled "no." If these distributions of yes and no responses do not overlap, the task of identifying the stimulus is trivial. However, if the distributions overlap (darkest area in Fig. 2), the observer must choose a threshold (solid line) to determine whether the event did or did not take place. In the example of Fig. 2*A* the observer would thus report the event if the response was to the right of the threshold and its absence if the response was to the left of the threshold. The cost of reporting the event when it did not occur (false alarm, overlap area to the *right* of the threshold in Fig. 2) and the cost of not reporting the event when it did occur (miss, overlap area to the *left* of the threshold) determine the optimal value for the threshold. For example, if there is a high cost for missing an event but a relatively low cost for reporting a false alarm, the threshold should be placed to the *left* of the overlap area (dashed line).

The ROC curve is generated by plotting the probability of reporting the event when it was present (hit rate) against the probability of a false alarm (Fig. 2*B*). The ROC curve of a perfect detector (dash–dotted line) would travel vertically up the *y*-axis and then horizontally (hit rate of 1) across all possible false-alarm rates. The area under such an ROC curve is one. If detection is random, the hit rate equals the false-alarm rate and the curve lies on the diagonal (dotted line). In this case, the area under the ROC curve is 1/2. The area under the ROC curve may thus be used as a measure of detection performance. In this study, we define *hit rate* as the probability that an action potential is fired within 15 ms (the same window of time used for Fig. 1*B*) after the input event. If more than one spike is fired, it is counted as only one hit. The *false-alarm rate* is the probability of an action potential within a 15-ms time window when the input event was not present.

Drawing the ROC curve is equivalent to plotting the area of the yes-distribution to the right of the threshold against the area of the no-distribution to the right of the threshold for each possible threshold value. Intuitively, the greater the overlap between the two distributions, the closer the ROC curve will be to the diagonal (random performance). The index of discriminability (d′) value (Green and

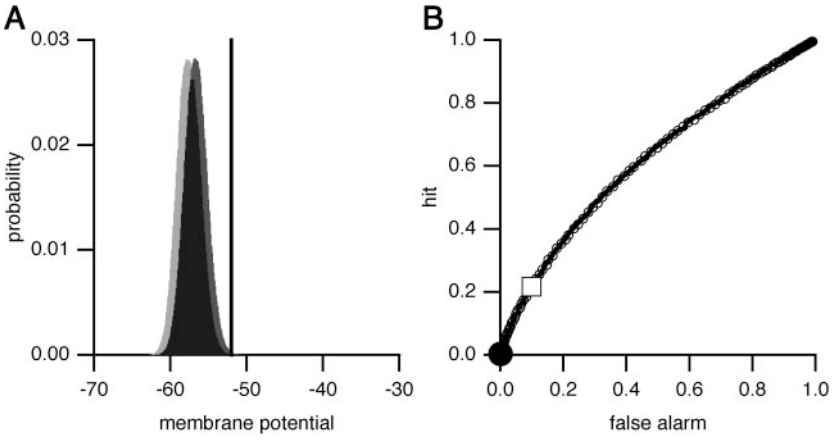

FIG. 4. *A*: membrane potential distributions of the model neuron with increased leak conductance. Means of the distributions are −57.89 for the "no" distribution and −56.88 for the "yes" distribution. SD of both distributions is 1.415. Solid vertical line indicates action potential threshold. *B*: solid curve is the ROC curve from Fig. 3, with the open square indicating the point on the curve for zero additional current. Open symbols demonstrate the ROC curve generated with increased leak conductance. Filled circle represents the point on the ROC curve for increased conductance, with zero additional current.

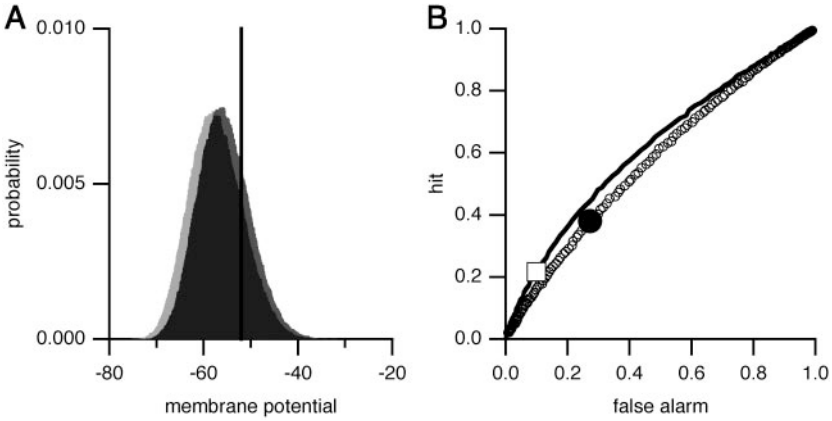

FIG. 5. *A*: membrane potential distributions in the presence of increased input noise and zero additional injected current. Light histogram represents the membrane potential probability in the absence of the input event ("no") and the dark histogram represents the membrane potential probability in the presence of the input event ("yes"): the "no" distribution has a mean of −58.14 and SD of 2.80; the "yes" distribution has a mean of −56.33 mV and SD of 2.77. Solid vertical line indicates the threshold for an action potential. *B*: solid line and open square are the ROC curve and point for zero additional current from Fig. 3. For the open circles, the input noise was increased without changing conductance. To increase input noise, the unitary size of each excitatory and inhibitory input was tripled, whereas the input rates were divided by 3. Filled circle shows the response of the neuron with the higher level of noise but without any additional injected current, as in Fig. 3.

Swets 1988) is the difference between the means of the two distributions divided by their standard deviation (SD) and describes how easily an observer can distinguish between the two response distributions. Herein, we calculated the d′ value by measuring the mean and SD of the best Gaussian fit to each membrane potential distribution. Although there is a slight correlation between SD and mean in our simulations, the SD of the membrane potential distribution was virtually unaffected by the input event, so the two distributions we compare, with and without the input event, have the same SD. The exception was Fig. 5, where the SDs of the two distributions differed by 0.03. In this case, to calculate d′ we used the average of the two SD measurements.

As an intuitive example for the meaning of the d′ value, consider the effect of increasing the peak conductance change evoked by the input event. Increasing the size of the input event will increase the hit probability without affecting the false-alarm probability. Because hit rate is now increased relative to false-alarm rate, the area under the ROC curve is increased. The effect of increasing the input event size on the SD of the membrane potential is negligible because this measure is dominated by the background activity. However, increasing the size of the input event will significantly increase the distance between the means of the membrane potential distributions with and without the input event. Thus increasing input event size will lead to an increase in the d′ value. The larger the d′ value, the more easily the two distributions may be separated from each other and the larger the area under the ROC curve.

RESULTS

We study a single-compartment neuronal model receiving a barrage of excitatory and inhibitory synaptic input (referred to as the "background" input). An "input event" occurs at a specific point in time for each trial. In our simulations, the input event was modeled as a large unitary excitatory synaptic conductance, but it may also be thought of as several smaller synchronous excitatory inputs (see METHODS). We now examine the performance of the model neuron at detecting the input event in the presence of noisy random background input. *Detection* is defined as firing an action potential, as illustrated in Fig. 2*C*, where the presynaptic action potential that represents the input event and the action potential evoked in the postsynaptic neuron are both highlighted. We use both d′ and ROC area values to describe and understand how background activity modulates synaptic efficacy, defined as the ability of a synaptic input to influence postsynaptic firing. Because of practical considerations, we propose the area under the ROC curve as an effective measure of synaptic efficacy.

The *top panel* in Fig. 3*A* shows the membrane potential distributions of the model neuron without the input event (light histogram) and measured within 5 ms after the input event (dark histogram). The action potential generator was turned off for these simulations. The d′ value for the distributions with (the "yes" distribution in Fig. 2*A*) and without (the "no" distribution in Fig. 2*A*) the input event is 0.79. Although measuring the d′ value of the membrane potential of an in vivo neuron is impractical, drawing this analogy is useful for developing intuition about how changes in conductance and input noise affect input event detection. In this study, we provide the d′ values as a quantitative basis for this intuition.

To produce the *middle* and *bottom panels* of Fig. 3*A*, we injected either hyperpolarizing (*middle*) or depolarizing (*bot-*

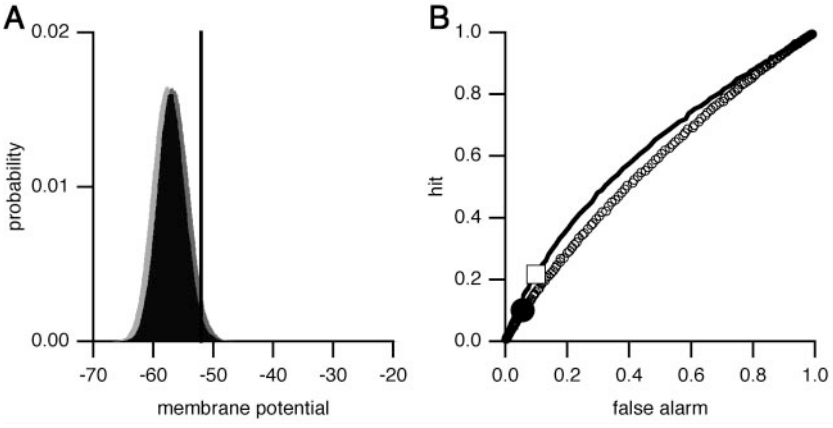

FIG. 6. *A*: membrane potential distributions in the presence of increased background activity, with ("yes," darker histogram) and without ("no," lighter histogram) the input event. Mean of the "no" distribution is −57.86 mV and mean of the "yes" distribution is −56.88. SD is 1.86 for both distributions. Solid vertical line is action potential threshold. *B*: solid line and open square are, respectively, the ROC curve and the point representing zero additional current (as in Fig. 3*B*). Open circles are the ROC curve of the model neuron when the rates of excitatory and inhibitory inputs were tripled. Filled circle demonstrates the response of the neuron, with the higher level of background input and zero additional injected current, as in Fig. 3.

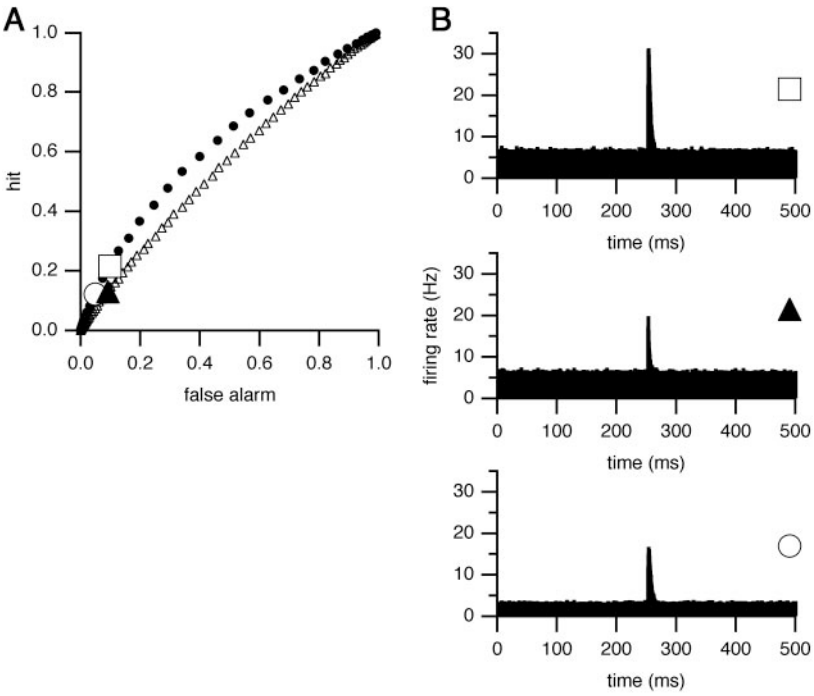

FIG. 7. *A*: to draw these ROC curves, the rates of inhibitory synaptic input were varied (as opposed to Figs. 2–5, where injected current was varied). For the filled circles, the excitatory input rate r_E was set to 1,500 Hz and the inhibitory input rate r_I varied between 300 and 4,500 Hz. To draw the ROC curve indicated by the open triangles, r_E was set to 9,000 Hz and r_I varied between 8,900 and 18,500 Hz. *B*: firing rate histograms illustrating the neuron responses to the input event (at time t = 250 ms). For the *top histogram*, corresponding to the open square in *A*, r_E = 1,500 Hz and r_I = 2,600 Hz. For the *middle histogram* (the filled triangle in *A*), r_E = 9,000 Hz and r_I = 14,300 Hz. For the *bottom histogram* (the open circle in *A*), r_E = 1,500 Hz and r_I = 2,850 Hz.

tom) constant current. As in the *top panel*, the light histogram represents the membrane potential distribution without the input event and the dark histogram the membrane potential distribution within 5 ms after the input event. The dominant effect of current injection was to shift the means of the membrane potential distributions without significantly changing their shape or width. The d′ values of the distributions are effectively unchanged, indicating that injecting constant current is equivalent to holding the membrane potential distributions fixed and moving the action potential threshold (indicated in each panel by the vertical line at −52 mV).

Intuitively, one may think of the action potential threshold as the detection threshold for these distributions. In Fig. 3*B*, we trace the ROC curve of the neuron by plotting the hit rate (the probability of an action potential within 15 ms after the input event) as a function of the false-alarm rate of the neuron (the probability of an action potential within a 15-ms time window in the absence of the input event). The open square in this figure (and also Figs. 4–7) indicates the point on the ROC curve that corresponds to the *top panel* of Fig. 3*A*, where no additional current was injected. To generate each additional point (filled symbols) on the ROC curve, we injected different levels of constant injected current. We used constant current rather than additional synaptic input to complete the ROC curves in Figs. 3–6 to avoid the possibility that changing the levels of excitation and inhibition might introduce conductance and noise changes in the background activity. (However, see Fig. 7 for an example where we generated the ROC curves by varying the rates of inhibitory synaptic input. Using either method did not significantly affect the results presented here.)

A fundamental consequence of background activity is an increase in membrane conductance. Figure 4*A* demonstrates the two membrane potential distributions, with (darker) and without (lighter) the input event, when the membrane conductance of the neuron was increased (see METHODS). As a result of shunting, the additional conductance decreases the effects of the input noise arising from the background synaptic input, reducing the standard deviation (SD) of the membrane potential distribution. The effect of the input event is also decreased, reducing the difference between the means of the two membrane potential distributions. Because the effects on the numerator and the denominator in the expression for d′ are similar, significantly increasing the conductance only slightly lowered the d′ value from 0.79 (Fig. 3) to 0.71 (Fig. 4).

Because of the negligible effect on the d′ value, one might expect that changing the conductance of the neuron would have a similarly slight effect on the ROC-defined synaptic efficacy. Increasing the neuronal membrane conductance caused the performance of the neuron without additional injected current to move from the open square in Fig. 4*B* (identical to the open square in Fig. 3*B*) to the filled circle. The increased conductance decreases both the hit rate and false-alarm rate, all the way to the point where the neuron stops firing action potentials. Completing the ROC curve and computing the area beneath it demonstrates that increasing membrane conductance does not affect synaptic efficacy (Fig. 4*B*). Although increasing overall membrane conductance significantly inhibits neuronal firing, it decreased the probability of response for conditions both with and without the input event. This does not result in a change in the shape of the ROC curve for the neuron. Indeed, the ROC curves with (open symbols) and without (solid curve) the additional membrane conductance are identical. Because the neuron is still operating with the same ROC curve, increasing membrane conductance does not have a significant effect on synaptic efficacy as we define it. However, this manipulation does change the ratio of hit rate to false-alarm rate and thus represents a modification in the detection strategy rather than efficacy.

A different effect is observed if the input noise is changed. Figure 5*A* shows the membrane potential distributions, with (dark histogram) and without (light histogram) the input event for a higher level of input noise than that in Fig. 3 but the same

overall membrane conductance. The increased input noise, the variance of the input current, results in an increase in the SD of the membrane potential distributions. However, because the increase in input noise does not shunt the current arising from the input event, the distance between the means of the distributions remains the same. Increasing the input noise therefore significantly decreases the d′ value (from 0.79 in Fig. 3 to 0.65 in Fig. 5).

Increasing input noise without affecting conductance increases the firing rate of the neuron. For zero additional injected current, the neuron moves from the open square (identical to the open square in Fig. 3*B*) to the filled circle (the responses for high-input noise but zero injected current). Comparing the open square and filled circle in Fig. 5*B* demonstrates the net excitatory effect of increasing input noise. The solid line in Fig. 5*B* is the ROC curve from Fig. 3*B*. The open symbols are the ROC curve in the presence of high-input noise. Increasing input noise flattens the ROC curve toward the diagonal, decreasing the area under the curve and thus decreasing the synaptic efficacy of the input event.

The effects of increasing conductance (Fig. 4) and increasing input noise (Fig. 5) are combined when background synaptic activity is increased (Fig. 6). We chose the values of the change in conductance and input noise in the previous figures to match those induced by tripling the synaptic input rates. Increasing the membrane conductance has two effects on the membrane potential distributions: the distance between the means of the distributions is decreased and also the widths of the distributions are reduced. The latter effect, however, is counteracted by the increased input noise. The net effect is a decrease in distance between the means of the distributions, resulting in a decreased d′ value of 0.53 (Fig. 6*A*) relative to 0.79 for Fig. 3.

With increased background activity, the neuron moves from the open square to the filled circle (three times the background synaptic input rates without additional injected current). As with increasing membrane conductance, increasing background activity decreases the firing rate of the neuron. However, as with increasing input noise, increasing background activity flattens the ROC curve toward the diagonal. This is seen by comparing the solid line in Fig. 3*B* (redrawn in Fig. 6*B*) with the ROC curve of the neuron in the presence of increased background activity (open circles in Fig. 6*B*). The synaptic efficacy of the input event is decreased by an increase in background activity.

With a change in synaptic efficacy, the hit rate and false-alarm rate can vary independently, suggesting that a neural circuit can modulate a neuron's response to particular synaptic inputs without changing the overall level of activity. Such an effect is possible only with a change in the shape of the ROC curve that describes the neuron's responses. We demonstrate the distinction between changing synaptic efficacy and simply changing response probability in Fig. 7. Here, the open square in Fig. 7*A* represents the same conditions as the open square in Fig. 3*B*; the firing rate histogram corresponding to this point on the ROC curve is the *top panel* of Fig. 7*B*. For this figure, we completed each ROC curve by varying the inhibitory background synaptic input rate rather than injecting current. We made this change to demonstrate that tracing the ROC curve by varying the level of excitation arising from synaptic input produces identical effects to tracing the ROC curve by varying the level of injected current.

To produce the ROC curve indicated by the open triangles (Fig. 7), we multiplied the background excitatory input rate by six. As in Fig. 6, the increased background synaptic activity decreased synaptic efficacy, flattening the ROC curve. To illustrate this effect, we chose the point on the ROC curve (indicated by the filled triangle) where, relative to the open square, the hit rate is lowered but the false-alarm rate is unchanged. Moving from the open square to the filled triangle therefore illustrates the effect of decreasing synaptic efficacy. The firing-rate histogram for this point is shown in Fig. 7*B* (*middle panel*).

The open circle in Fig. 7*A* was chosen to illustrate the effect of decreasing response probability. This point was chosen because the hit rate was the same as that for the filled triangle. However, because in this case the neuron is moving along the ROC curve rather than off it, both the hit rate and false-alarm rate decrease together. This is made apparent by comparing the *bottom histogram* in Fig. 7*B* with the *middle histogram*. Although the elevation in firing rate resulting from the input event is similar in both panels, it is accompanied by a decrease in baseline firing rate in the *bottom panel* only. A measurement of synaptic efficacy that is confounded by baseline firing rate will not distinguish between these two effects. Only with a firing-rate–independent measure can a change in neuronal sensitivity (*middle panel*) be distinguished from a general decrease in excitability (*bottom panel*).

DISCUSSION

We propose using ROC analysis for measuring synaptic efficacy, defined here as the impact of a synaptic input on the firing rate of the postsynaptic neuron. Although there are several ways of measuring synaptic efficacy, this is the only one (to our knowledge) that is independent of the baseline firing rate (the firing rate of the postsynaptic neuron). We have demonstrated that the baseline firing rate can be changed either by varying the mean input (or mean conductance level) or by varying the level of input noise. Only one of these manipulations leads to a true change in neuronal sensitivity. When trying to describe the efficacy of an input under conditions in which the baseline firing rate varies, it will prove extremely useful to have a measurement that is firing-rate independent. For an intuitive understanding of this measure and what factors modulate it, we have drawn an analogy in which the postsynaptic neuron acts as an observer of its own membrane potential and functions as a detector of a particular input event. In this analogy, the neuron reports detection of the input event by firing an action potential.

Intuitively, synaptic efficacy is a measure of how easily the membrane potential probability distribution in the presence of the input event can be discriminated from the membrane potential distribution in its absence. An often-used measurement for determining how easily two distributions (in this case the membrane potential distributions with and without the input event) may be discriminated from each other is the d′ value. However, because measuring the membrane potential distributions is difficult and because the exact transformation between the area of membrane potential distribution above action-potential threshold into firing rate is not straightforward,

the d′ value cannot practically be used as a measure of synaptic efficacy in many situations. Herein, the d′ values are provided as an intuitive explanation for why some mechanisms affect synaptic efficacy whereas others do not. We propose a related quantity—the area under the ROC curve—instead.

The ROC curve is a plot of the hit rate, the probability that a neuron will fire an action potential within a certain time window after an input event, against the false-alarm rate, the probability of an action potential in the same time window but in the absence of the input event. The area under the ROC curve is a firing-rate–independent measure of detection performance or, in this situation, synaptic efficacy. We have focused on examining how a noisy "background" of excitatory and inhibitory synaptic input, much like what exists in vivo, modulates synaptic efficacy. Increasing neuronal excitation or inhibition, whether by current injection or changing conductance, moves the neuron along its ROC curve. This movement is a change in the probability of spiking in which the hit rate and false-alarm rate are changing according to a specified relationship (the ROC curve). However, the shape of the ROC curve is unaffected by these manipulations, demonstrating that the synaptic efficacy is unchanged. Rather than changing synaptic efficacy or, equivalently, detection sensitivity, these modifications represent a change in detection strategy as reflected in the trade-off between missed events versus false alarms. Increasing the variance of the input current, here called the input noise, changes the shape of the ROC curve, flattening it toward the diagonal and indicating a decrease in synaptic efficacy. For the simulations presented here, when the input noise alone is increased, the decrease in synaptic efficacy is accompanied by an excitatory effect. If the input noise increase is accompanied by an increase in the membrane leak conductance, then the synaptic efficacy decrease is accompanied by an inhibitory effect. For larger inputs, however, it is also possible for increases in input noise to have an inhibitory effect (Shu et al. 2003). Nevertheless, increasing input noise will result in a decrease of synaptic efficacy.

The effects of background activity and its components, input noise and conductance changes, on the slope of the firing rate curve (plotted against injected current) were previously studied. These studies reported that the slope of the firing-rate curve is unaffected by overall conductance, but decreases when input noise is increased (Chance et al. 2002; Doiron et al. 2000; Holt and Koch 1997; Prescott and De Koninck 2003). These reports are consistent with the conclusions drawn here. The slope of the firing rate curve is essentially a measure of the sensitivity of the neuron to small changes in input current. Similar studies (Hô and Destexhe 2000; Shu et al. 2003) focused on the effects of background activity by plotting hit rate as a function of synaptic input size (EPSC size or unitary conductance change). Synaptic efficacy, as defined here, is a measure of how the neuron will respond to a small change in input and therefore is related to the slope of these curves. In the studies presented here, the input event consists of a brief increase in excitatory synaptic conductance. We repeated these studies using EPSCs and observed the same results (not shown). Combined, these results suggest that mechanisms of gain modulation of steady-state neuronal responses (Chance et al. 2002; Fellous et al. 2003; Hô and Destexhe 2000; Mitchell and Silver 2003; Prescott and De Koninck 2003; Shu et al. 2003; although see Murphy and Miller 2003) extend to a change in synaptic efficacy as well. This result is consistent with a previous study by Powers and Binder (1996) demonstrating that the change in firing rate evoked by a pulse of current approximates the product of the additional current reaching the soma multiplied by the slope of the steady-state *f–I* curve.

Previous studies demonstrated that changing conductance can have an effect on the precision of spike timing in response to an excitatory input (London et al. 2002; Zsiros and Hestrin 2005). This effect can be observed in Fig. 7*B* (the elevation in firing rate in the *middle histogram* is slightly narrower than that for either of the other histograms because of the increased conductance). Many works investigated the transformation of synaptic input current into postsynaptic firing rate, focusing on the time course of individual synaptic inputs and the timing of synaptic inputs (Fetz and Gustafsson 1983; Hermann and Gerstner 2001, 2002; London et al. 2002; Manwani and Koch 1999; Reyes and Fetz 1993a,b; Türker and Powers 1999, 2005). In addition, noisy synaptic input can affect input detection by modifying the propagation properties of the dendritic cable (Manwani and Koch 1999) or by activating voltage-gated conductances (Svirskis et al. 2002). Because of the simplicity of the ROC measure, some information about the underlying membrane potential is necessarily lost. It is possible that ROC analysis could be used to explore changes in spike timing by reanalyzing the same recordings with different time windows for hit and false-alarm probability. Our focus here, however, is to examine the effect of conductance and other aspects of in vivo synaptic activity on input detection rather than on response timing and provide that information in a single measurement. London et al. (2002) suggested an information-theoretic measure of synaptic efficacy. The ROC measure introduced here is similar in spirit to this approach, but it has the useful feature of being firing-rate independent.

Short-term synaptic plasticity, such as synaptic depression or facilitation, or any other effect that changes the strength of a synapse, will affect synaptic efficacy. For the results presented here, we assume that synaptic strength remains constant from trial to trial. A major objective of this study was to examine the effects of background activity on synaptic efficacy. Because other factors, such as baseline firing rate, are perhaps more likely to confound our measure of synaptic efficacy when changing the level of background synaptic input received by a neuron, we sought to find a method of quantifying synaptic efficacy that was independent of these factors.

Our results show that there are two types of manipulations that increase the probability that a neuron will fire an action potential shortly after the arrival of an input event. When current is injected or an excitatory conductance is tonically activated, the overall response probability of the neuron is increased. This causes the "hit" rate, or the probability of the neuron firing an action potential shortly after the input event, to increase, but this increase is accompanied by an increase in baseline firing rate, or "false alarm" rate. Because both hit rate and false-alarm rate are linked in a prescribed fashion that is not altered by this manipulation, there is no change in synaptic efficacy. A true change in neuronal sensitivity to synaptic input arises when the relationship between hit rate and false-alarm rate is altered and this is what we refer to as a change in synaptic efficacy. For the manipulations studied here, which did not involve directly changing synaptic conductances, this

arises only when there is a change in input noise, such as when the level of background activity increases or becomes more correlated.

In this study we have focused on the neuron's ability to detect excitatory inputs and what conditions change the neuron's sensitivity to those inputs. Inhibitory inputs result in a transient decrease in firing probability, although the exact transformation between the shape of the transient inhibitory current and resulting depression in firing rate appears to be more complex (Fetz and Gustafsson 1983; Türker and Powers 1999) than for excitatory inputs. With an appropriately chosen time window, ROC analysis should provide a similarly effective, firing-rate–independent measurement of the "efficacy" of an inhibitory event. In this case, detection of an inhibitory input decreases the hit rate, and greater sensitivity to inhibitory inputs will result in a smaller area under the ROC curve.

Measuring synaptic efficacy by this method in vitro simply requires driving the postsynaptic neuron at different firing rates and recording the probability of input event detection. Repeating such an analysis in vivo may be slightly more difficult because control over the baseline firing rate is limited. In this situation we envision that false-alarm rates could be estimated by measuring the firing rate over a slightly longer period of time (compared with the time window used in this study) before the input event. The amount of data required would depend on how quickly the baseline firing rate of the system under study varies.

ACKNOWLEDGMENTS

A portion of this research was conducted in the laboratory of A. D. Reyes. The author is grateful to both A. D. Reyes and L. F. Abbott for helpful comments and discussion.

GRANTS

This research was supported by National Science Foundation Grant NSF-IOB-0446129, funds provided by the University of California, and an Alfred P. Sloan Foundation Research Fellowship to F. S. Chance.

REFERENCES

Anderson JS, Lampl I, Gillespie DC, Ferster D. The contribution of noise to contrast invariance of orientation tuning in cat visual cortex. *Science* 290: 1968–1972, 2000.

Azouz R. Dynamic spatiotemporal synaptic integration in cortical neurons: neuronal gain, revisited. *J Neurophysiol* 94: 2785–2796, 2005.

Borg-Graham LJ, Monier C, Frégnac Y. Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature* 393: 369–373, 1998.

Braun HA, Wissing H, Schäfer K, Hirsch MC. Oscillation and noise determine signal transduction in shark multimodal sensory cells. *Nature* 367: 270–273, 1994.

Britten KH, Shadlen MN, Newsome WT, Movshon JA. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12: 4745–4765, 1992.

Chance FS, Abbott LF, Reyes AD. Gain modulation from background synaptic input. *Neuron* 35: 773–782, 2002.

Cohn TE. Receiver operating characteristic analysis of sensitivity in neural systems. *Proc IEEE* 65: 781–786, 1977.

Cohn TE, Green DG, Tanner WP. Receiver operating characteristic analysis. Application to the study of quantum fluctuation effects in optic nerve of *Rana pipiens*. *J Gen Physiol* 66: 583–616, 1975.

Collins JJ, Chow CC, Imhoff TT. Stochastic resonance without tuning. *Nature* 376: 236–238, 1995.

Destexhe A, Paré D. Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. *J Neurophysiol* 81: 1531–1547, 1999.

Destexhe A, Rudolph M, Fellous J-M, Sejnowski TJ. Fluctuating synaptic conductances recreate *in vivo*-like activity in neocortical neurons. *Neuroscience* 107: 13–24, 2001.

Destexhe A, Rudolph M, Paré D. The high-conductance state of neocortical neurons *in vivo*. *Nat Rev Neurosci* 4: 739–751, 2003.

Doiron B, Longtin A, Berman N, Maler L. Subtractive and divisive inhibition: effect of voltage-dependent inhibitory conductances and noise. *Neural Comput* 13: 227–248, 2000.

Douglass JK, Wilkens L, Pantazelou E, Moss F. Noise enhancement of information transfer in crayfish mechanoreceptors by stochastic resonance. *Nature* 365: 337–340, 1993.

Ellaway PH. An application of cumulative sum technique (cusums) to neurophysiology. *J Physiol* 265: 1P–2P, 1977.

Ellaway PH. Cumulative sum technique and its application to the analysis of peristimulus time histograms. *Electroencephalogr Clin Neurophysiol* 45: 302–304, 1978.

Fellous J-M, Rudolph M, Destexhe A, Sejnowski TJ. Synaptic background noise controls the input/output characteristics of single cells in an *in vitro* model of *in vivo* activity. *Neuroscience* 122: 811–829, 2003.

Fetz EE, Gustafsson B. Relation between shapes of post-synaptic potentials and changes in firing probability of cat motoneurones. *J Physiol* 341: 387–410, 1983.

Fitzhugh R. The statistical detection of threshold signals in the retina. *J Gen Physiol* 40: 925–947, 1957.

Goense JBM, Ratnam R. Continuous detection of weak sensory signals in afferent spike trains: the role of anti-correlated interspike intervals in detection performance. *J Comp Physiol A Sens Neural Behav Physiol* 189: 741–759, 2003.

Green DM, Swets JA. *Signal Detection Theory and Psychophysics.* Los Altos Hills, CA: Peninsula Publishing, 1988, p. 60.

Guido W, Lu S-M, Vaughan JW, Godwin DW, Sherman SM. Receiver operating characteristic (ROC) analysis of neurons in the cat's lateral geniculate nucleus during tonic and burst response mode. *Vis Neurosci* 12: 723–741, 1995.

Heggelund P, Albus K. Response variability and orientation discrimination of single cells in striate cortex of cat. *Exp Brain Res* 32: 197–211, 1978.

Hermann A, Gerstner W. Noise and the PSTH response to current transients: I. General theory and application to the integrate-and-fire neuron. *J Comput Neurosci* 11: 135–151, 2001.

Hermann A, Gerstner W. Noise and the PSTH response to current transients: II. Integrate-and-fire model with slow recovery and application to motoneuron data. *J Comput Neurosci* 12: 83–95, 2002.

Hirsch JA, Alonso J-M, Reid RC, Martinez LM. Synaptic integration in striate cortical simple cells. *J Neurosci* 18: 9517–9528, 1998.

Hô N, Destexhe A. Synaptic background activity enhances the responsiveness of neocortical pyramidal neurons. *J Neurophysiol* 84: 1488–1496, 2000.

Holt GR, Koch C. Shunting inhibition does not have a divisive effect on firing rates. *Neural Comput* 9: 1001–1013, 1997.

Holt GR, Softky WR, Koch C, Douglas RJ. Comparison of discharge variability *in vitro* and *in vivo* in cat visual cortex neurons. *J Neurophysiol* 75: 1806–1814, 1996.

London M, Schreibman A, Häusser M, Larkum M, Segev I. The information efficacy of a synapse. *Nat Neurosci* 5: 332–340, 2002.

Manwani A, Koch C. Detecting and estimating signals in noisy cable structures. II: Information theoretical analysis. *Neural Comput* 11: 1831–1873, 1999.

Mitchell SJ, Silver RA. Shunting inhibition modulates neuronal gain during synaptic excitation. *Neuron* 38: 433–445, 2003.

Murphy BK, Miller KD. Multiplicative gain changes are induced by excitation or inhibition alone. *J Neurosci* 23: 10040–10051, 2003.

Paré D, Shink E, Gadreau H, Destexhe A, Lang EJ. Impact of spontaneous synaptic activity on the resting properties of cat neocortical pyramidal neurons *in vivo*. *J Neurophysiol* 79: 1450–1460, 1998.

Powers RK, Binder MD. Experimental evaluation of input–output models of motoneuron discharge. *J Neurophysiol* 75: 367–379, 1996.

Prescott SA, De Koninck Y. Gain control of firing rate by shunting inhibition: roles of synaptic noise and dendritic saturation. *Proc Natl Acad Sci USA* 100: 2076–2081, 2003.

Reyes AD, Fetz EE. Two modes of interspike interval shortening by brief transient depolarizations in cat neocortical neurons. *J Neurophysiol* 69: 1661–1672, 1993a.

Reyes AD, Fetz EE. Effects of transient depolarizing potentials on the firing rate of cat neocortical neurons. *J Neurophysiol* 69: 1673–1683, 1993b.

Schiller PH, Finlay BL, Volman SF. Short-term response variability of monkey striate neurons. *Brain Res* 105: 347–349, 1976.

Shadlen MN, Newsome WT. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18: 3870–3896, 1998.

Shofner WP, Dye RH. Statistical and receiver operating characteristic analysis of empirical spike-count distributions: quantifying the ability of cochlear nucleus units to signal intensity changes. *J Acoust Soc Am* 86: 2172–2184, 1989.

Shu Y, Hasenstaub A, Badoual M, Bal T, McCormick DA. Barrages of synaptic activity control the gain and sensitivity of cortical neurons. *J Neurosci* 23: 10388–10401, 2003.

Softky WR, Koch C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13: 334–350, 1993.

Steriade M, Timofeev I, Grenier F. Natural waking and sleep states: a view from inside neocortical neurons. *J Neurophysiol* 85: 1969–1985, 2001.

Stevens CF, Zador AM. Input synchrony and the irregular firing of cortical neurons. *Nat Neurosci* 1: 210–217, 1998.

Svirskis G, Kotak V, Sanes DH, Rinzel J. Enhancement of signal-to-noise ratio and phase locking for small inputs by a low-threshold outward current in auditory neurons. *J Neurosci* 22: 11019–11025, 2002.

Tiesinga PHE, José JV, Sejnowski TJ. Comparison of current-driven and conductance-driven neocortical model neurons with Hodgkin–Huxley voltage-gated channels. *Phys Rev E* 62: 8413–8419, 2000.

Tolhurst DJ, Movshon JA, Dean AF. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 23: 775–785, 1983.

Troyer TW, Miller KD. Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Comput* 9: 971–983, 1997.

Türker KS, Powers RK. Effects of large excitatory and inhibitory inputs on motoneuron discharge rate and probability. *J Neurophysiol* 82: 829–840, 1999.

Türker KS, Powers RK. The effects of common input characteristics and discharge rate on synchronization in rat hypoglossal motoneurones. *J Physiol* 245: 245–260, 2002.

Türker KS, Powers RK. Black box revisited: a technique for estimating postsynaptic potentials in neurons. *Trends Neurosci* 28: 379–385, 2005.

Vogels R, Spileers W, Orban GA. The response variability of striate cortical neurons in the behaving monkey. *Exp Brain Res* 77: 432–436, 1989.

Wenning G, Hoch T, Obermayer K. Detection of pulses in a colored noise setting. *Phys Rev E* 71: e021902, 2005.

Woody CD, Gruen E. Characterization of electrophysiological properties of intracellularly recorded neurons in the neocortex of awake cats: a comparison of the response to injected current in spike overshoot and undershoot neurons. *Brain Res* 158: 343–357, 1978.

Zsiros V, Hestrin S. Background synaptic conductance and precision of EPSP-spike coupling at pyramidal cells. *J Neurophysiol* 93: 3248–3256, 2005.