

Impact of Synaptic Unreliability on the Information Transmitted by Spiking Neurons

ANTHONY ZADOR

Salk Institute, La Jolla, California 92037

Zador, Anthony. Impact of synaptic unreliability on the information transmitted by spiking neurons. *J. Neurophysiol.* 79: 1219–1229, 1998. The spike generating mechanism of cortical neurons is highly reliable, able to produce spikes with a precision of a few milliseconds or less. The excitatory synapses driving these neurons are by contrast much less reliable, subject both to release failures and quantal fluctuations. This suggests that synapses represent the primary bottleneck limiting the faithful transmission of information through cortical circuitry. How does the capacity of a neuron to convey information depend on the properties of its synaptic drive? We address this question rigorously in an information theoretic framework. We consider a model in which a population of independent unreliable synapses provides the drive to an integrate-and-fire neuron. Within this model, the mutual information between the synaptic drive and the resulting output spike train can be computed exactly from distributions that depend only on a single variable, the interspike interval. The reduction of the calculation to dependence on only a single variable greatly reduces the amount of data required to obtain reliable information estimates. We consider two factors that govern the rate of information transfer: the synaptic reliability and the number of synapses connecting each presynaptic axon to its postsynaptic target (i.e., the connection redundancy, which constitutes a special form of input synchrony). The information rate is a smooth function of both mechanisms; no sharp transition is observed from an “unreliable” to a “reliable” mode. Increased connection redundancy can compensate for synaptic unreliability, but only under the assumption that the fine temporal structure of individual spikes carries information. If only the number of spikes in some relatively long-time window carries information (a “mean rate” code), an increase in the fidelity of synaptic transmission results in a seemingly paradoxical decrease in the information available in the spike train. This suggests that the fine temporal structure of spike trains can be used to maintain reliable transmission with unreliable synapses.

INTRODUCTION

A pyramidal neuron in the cortex receives excitatory synaptic inputs from 10^3 – 10^4 other neurons (Shepherd 1990). When an action potential invades the presynaptic terminal of one of these synapses, it sometimes triggers the release of a vesicle of glutamate, which causes current to flow into the postsynaptic dendrite. Some of this current then propagates, passively or actively, to the spike generator, where it may contribute to the triggering of an action potential.

The postsynaptic neuron can be viewed as an input-output element that converts the input spike trains from many presynaptic neurons into a single-output spike train. This input-output transformation is the basic computation performed by neurons. It is the foundation upon which cortical processing is based.

The computational strategies available to a neuronal circuit depend upon the fidelity of its components. For example, the computational power of a single integrate-and-fire neuron depends on the effective noise of the currents driving the spike generator (Zador and Pearlmuter 1996). In the cortex, the transformation of somatic current into an output spike train appears to be highly reliable (Mainen and Sejnowski 1995, see also Bryant and Segundo 1976), in marked contrast to the unreliability of synaptic transmission (Allen and Stevens 1994; Dobrunz and Stevens 1997; Stratford et al. 1996). In this paper, we use simple biophysical models of spike transduction and stochastic synaptic release to explore the implications of synaptic unreliability on information transmission and neural coding in the cortex. Our goal is to provide a quantitative answer to the question: How much information can the output spike train provide about the synaptic inputs? Our answer will be cast in an information-theoretic framework.

METHODS

Physiology

Standard slice recording methods were used to obtain Fig. 1. Briefly, patch-clamp recordings were obtained under visual guidance by using infrared optics from 400- μ m slices from Long Evans rats [postnatal day (P)14–P20]. Recordings were performed at 33–35°C. Slices were continuously perfused with a solution containing (in mM) 120 NaCl, 3.5 KCl, 2.6 CaCl₂, 1.3 MgCl₂, 1.25 NaH₂PO₄, 26 NaHCO₃, and 10 glucose, which was bubbled with 95% O₂-5% CO₂ and the pH of which had been adjusted to 7.35. All recordings were obtained in the presence of the α -amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid (AMPA) receptor antagonist 6-cyano-7-nitroquinoxaline-2,3-dione (CNQX, 50 μ M). Recording pipettes were filled with (in mM) 170 K gluconate, 10 *N*-2-hydroxyethylpiperazine-*N'*-2-ethanesulfonic acid (HEPES), 10 NaCl, 2 MgCl₂, 1.33 ethylene glycol-bis(β -aminoethyl ether)-*N,N,N',N'*-tetraacetic acid (EGTA), 0.133 CaCl₂, 3.5 MgATP, and 1.0 guanosine 5'-triphosphate (GTP), pH 7.2. Resistance to bath was 3–5 M Ω before seal formation.

Data were acquired by using a National Instruments (TX) AT-MIO-16-F-5 A/D card on a Pentium-based computer under the Window NT (Microsoft) operating system. Software written in Labview (National Instruments) with Dynamic Data Exchange links to Matlab (Mathworks) allowed convenient online synthesis and injection of arbitrary synthetic current waveforms.

Simulations

All simulations were performed using Matlab 4.2.

Model of spiking

We use an integrate-and-fire mechanism to model the transformation of synaptic inputs into spike trains in cortical neurons. Let

$i_{\text{syn}}(t)$ be the synaptic current driving a leaky integrator with a time constant τ and a threshold V_{thresh} . As long as the voltage is subthreshold, $v(t) < V_{\text{thresh}}$, the voltage is given by

$$\tau \frac{dv(t)}{dt} = -[v(t) - V_{\text{rest}}] + i_{\text{syn}}(t)R_n \quad (1)$$

where R_n is the input resistance and V_{rest} is the resting potential. At the instant the voltage reaches the threshold V_{thresh} , the neuron emits a spike, and resets to some level $V_{\text{reset}} < V_{\text{thresh}}$. The five parameters of this model, V_{thresh} , V_{reset} , V_{rest} , τ , and R_n , determine its response to a given input current.

The output of the model is a spike train, i.e. a sequence of times at which $v(t)$ exceeded threshold. If time is finely discretized into bins shorter than the shortest interspike interval, so that the number of spikes in each bin is either zero or one (but not greater than one), then the spike train can be represented as a binary string $z_o(t)$, with ones at times when the neuron fired and zeros at other times.

Model of synaptic drive

We assume that the synaptic current $i_{\text{syn}}(t)$ consists of the sum of very brief—essentially instantaneous—individual excitatory postsynaptic currents (EPSCs). This represents a reasonable simplification of the component of the excitatory input to cortical neurons mediated by fast AMPA receptors, which decay with a time constant of 2–3 ms (Bekkers and Stevens 1990), but not for the component mediated by the slower *N*-methyl-D-aspartate (NMDA) receptor-gated channels.

The synaptic current driving any neuron results from the spike trains of all the other neurons that make synapses onto it. The postsynaptic current depends both on the precise times at which each of the presynaptic neurons fired and on the response at each synapse to the arrival of a presynaptic action potential. If the response at each synapse is either unreliable or variable in amplitude, then even the arrival of precisely the same spike train at each terminal will fail to produce identical postsynaptic current. In what follows, the exact sequence of action potentials arriving at each of the presynaptic terminals is the “signal,” and any variability response to repeated trials on which precisely the same sequence is presented represents the “noise.”

After the basic quantal model of synaptic transmission (Katz 1966), we consider two sources of synaptic variability, or noise. The first is that the probability P_r that a glutamate-filled vesicle is released after presynaptic activation may be less than unity in the hippocampus (Allen and Stevens 1994; Hessler et al. 1993; Rosenmund et al. 1993) and the cortex (Castro-Alamancos and Connors 1997; Stratford et al. 1996). The second is that the postsynaptic current in response to a vesicle may vary even at single individual terminals (Bekkers and Stevens 1990). This quantal variability may arise, for example, from variable amounts of neurotransmitter filling each vesicle (Bekkers and Stevens 1990); but the results of the present study do not depend on the mechanism underlying this variability.

The basic model for the postsynaptic current i_{syn} driving the neuron is as follows. We assume that the activity in the population of presynaptic neurons j is given by $z_j(t)$, where (by analogy with the output $z_o(t)$ above) $z_j(t)$ is a binary string whose entries are one if the neuron fired and zero otherwise. When an axon fires, the presynaptic terminal releases transmitter with a probability P_r . If transmitter is released at time t at synapse j , then the postsynaptic amplitude is given by $q_j(t)$, which is a random variable that represents the quantal variability. Thus the total postsynaptic current is given by

$$i_{\text{syn}}(t) = \sum_j z_j(t) \times f_j(t) \times q_j(t) \quad (2)$$

where the summation index j is over the input neurons, the random process $f_j(t)$ representing synaptic failures is a binary string that is one when transmitter is released and zero otherwise and $q_j(t)$ is a random variable that determines the quantal size of releases when they occur. The processes $i_{\text{syn}}(t_i)$, $z_j(t_i)$, $f_j(t_i)$, and $q_j(t_i)$ are discrete-time, but for notational convenience we will often suppress the time index i .

A single axon may sometimes make multiple synapses onto a postsynaptic target, or a single synapse might have multiple release sites. We use the term functional contact to describe both these situations. Equation 2 implicitly assumes that each axon has only a single “functional contact” onto the postsynaptic neuron. We also consider the case where each axon makes N_r multiple functional contacts. In this case, the current i_{syn} is given by

$$i_{\text{syn}}(t) = \sum_j z_j(t) \times \sum_k^{N_r} q_{jk}(t) \times f_{jk}(t) \quad (3)$$

where the summation index k is over functional contacts, each of which is driven by the same sequence of presynaptic action potentials $z_j(t)$. In this model, all the terminals k associated with a single presynaptic axon are activated synchronously, but release failures occur at each contact independently.

The Poisson rate S_{net} (impulses/second) at which EPSCs contribute to the postsynaptic current is given in this model by

$$S_{\text{net}} = A \times F_{\text{in}} \times N_r \times P_r \quad (4)$$

where A is the number of afferent axons, N_r is the number of functional contacts per axon (assumed to be the same for all axons), F_{in} is the Poisson rate at which each axon fires (assumed to be the same for all axons), and P_r is the release probability at each functional contact (assumed to be the same for all contacts). S_{net} determines the average postsynaptic current and thereby the output firing rate R .

In some of the simulations described below (Figs. 3–5), the parameters N_r and P_r were varied. To keep S_{net} fixed under these conditions, any decrease in these parameters was compensated for by a proportional increase in $A \times F_{\text{in}}$. For example, if the release probability P_r was reduced to 0.5 from 1, F_{in} was increased twofold.

Information rate of spike trains

A typical pyramidal neuron in the cortex receives synaptic input from 10^3 – 10^4 other neurons. We define the activity in each of these input neurons as the “signal,” and the variability due to the unreliability of synaptic transmission is the “noise.”

How much information does the output spike train $z_o(t)$ provide about the input spike trains $z_j(t)$? More formally, what is the mutual information $I(Z_{\text{in}}(t); Z_{\text{out}}(t))$ between the ensemble of input spike trains $Z_{\text{in}}(t) = \{z_1(t), \dots, z_j(t), \dots\}$ and the output spike train ensemble $Z_{\text{out}}(t)$? We assume that both $Z_{\text{in}}(t)$ and $Z_{\text{out}}(t)$ are completely specified by the activity (i.e., the precise list of spike times) in each spike train; that is, all the information in the spike trains can be represented by the list of spike times and there is no extra information contained in properties such as spike height or width. Characteristics of the spike train such as the mean or instantaneous rate can be derived from this representation; if such a derived property turns out to be the relevant one, then this formulation can be specialized appropriately.

The mutual information $I(Z_{\text{in}}(t); Z_{\text{out}}(t))$ is defined (Shannon and Weaver 1948) in terms of the entropy $H(Z_{\text{in}})$ of the ensemble of input spike trains, the entropy $H(Z_{\text{out}})$ of output spike trains, and their joint entropy $H(Z_{\text{in}}, Z_{\text{out}})$,

$$I(Z_{\text{in}}; Z_{\text{out}}) = H(Z_{\text{in}}) + H(Z_{\text{out}}) - H(Z_{\text{in}}, Z_{\text{out}}) \quad (5)$$

The entropies $H(Z_{in})$, $H(Z_{out})$ and $H(Z_{in}, Z_{out})$ depend only on the probability distributions $P(Z_{in})$, $P(Z_{out})$, and the joint distribution $P(Z_{in}, Z_{out})$, respectively.

Note that because the joint distribution is symmetric $P(Z_{in}, Z_{out}) = P(Z_{out}, Z_{in})$, the mutual information is also symmetric, $I(Z_{in}; Z_{out}) = I(Z_{out}; Z_{in})$. Note also that if the inputs $Z_{in}(t)$ and outputs $Z_{out}(t)$ are completely independent, then the mutual information is zero, because the joint entropy is just the sum of the individual entropies $H(Z_{in}, Z_{out}) = H(Z_{in}) + H(Z_{out})$. This is completely reasonable, because in this case the inputs provide no information about the outputs.

Methods for estimating spike train information rates

The expression given in Eq. 5 for the mutual information is in practice difficult to evaluate because estimating the distributions $P(Z_{in})$, $P(Z_{out})$, and $P(Z_{in}, Z_{out})$ may require very large amounts of data. For example, suppose that there are 1,000 input spike trains driving the output and that each spike train is divided into segments 100 ms in length and discretized into 1 ms bins. There are then 2^{100} possible output spike trains, $2^{100 \times 1,000}$ sets of input spike trains, and $2^{100 \times 1,000} \times 2^{100}$ possible combinations of input and output spike trains forming the space over which the joint distribution $P(Z_{in}, Z_{out})$ must be estimated. Although this naive calculation is in practice an overestimate (see Buracas et al. 1996 and de Ruyter van Steveninck et al. 1997 for methods that make use of the fact that most spike trains are very unlikely), it emphasizes the potential problems involved in estimating the mutual information. Below we describe two practical methods for computing information rates.

Reconstruction method

One approach to this dilemma (Bialek et al. 1991, 1993) is to compute a strict lower bound on the mutual information using the reconstruction method. The idea is to “decode” the output and use it to “reconstruct” the input that gave rise to it. The error between the reconstructed and actual outputs is then a measure of the fidelity of transmission and with a few testable assumptions can be related to the information. Formally, this method is based on an expression mathematically equivalent to Eq. 5 involving the conditional entropy $H(Z_{in}|Z_{out})$ of the signal given the spike train

$$I(Z_{in}, Z_{out}) = H(Z_{in}) - H(Z_{in}|Z_{out}) \quad (6)$$

In the present context, the quantity reconstructed is the sum of the input, $\sum_j z_j(t)$. The entropy $H(Z_{in})$ is just the entropy of the time series $\sum_j z_j(t)$ and can be evaluated directly from the Poisson synthesis equation (Eq. 3). Intuitively, Eq. 6 says that the information gained about the spike train by observing the stimulus is just the initial uncertainty about the synaptic drive (in the absence of knowledge of the spike train) minus the uncertainty that remains about the signal once the spike train is known. The reconstruction method estimates the input from the output and then bounds the errors of the outputs from above by assuming they are Gaussian. This method, which can provide a lower bound on the mutual information, has been used with much success in a variety of experimental preparations (Bialek et al. 1991; de Ruyter van Steveninck and Bialek 1988; de Ruyter Van Steveninck and Laughlin 1996; Rieke et al. 1997).

Direct method

In this paper we will use a direct method (DeWeese 1995, 1996; de Ruyter van Steveninck et al. 1997; Stevens and Zador 1996) to estimate the mutual information. Direct methods use another form of the expression Eq. 5 for mutual information

$$I(Z_{in}, Z_{out}) = H(Z_{out}) - H(Z_{out}|Z_{in}) \quad (7)$$

The first term $H(Z_{out})$ is the entropy of the output spike train itself, whereas the second term $H(Z_{out}|Z_{in})$ is the conditional entropy of the output given the inputs. The first term measures of the variability of the spike train in response to the ensemble of different inputs, whereas the second measures the reliability of the response to repeated presentations of the same inputs. The second term depends on the reliability of the synapses and spike generating mechanism: to the extent that the same inputs produce the same outputs, this term approaches zero.

The direct method has two advantages over the reconstruction method in the present context. First, it does not require the construction of a “reconstructor” for estimating the input from the output. Although the optimal linear reconstructor is straightforward to estimate, the construction of more sophisticated (i.e., nonlinear) reconstructors can be a delicate art. Second, it provides an estimate of information that is limited only by the errors in the estimation of $H(Z_{out})$ and $H(Z_{out}|Z_{in})$; the reconstruction method by contrast provides only a lower bound on the mutual information that is limited by the quality of the reconstructor.

As noted above, the estimation of $H(Z_{out})$ and $H(Z_{out}|Z_{in})$ can require vast amounts of data. If, however, interspike intervals (ISIs) in the output spike train were independent, then the entropies could be simply expressed in terms of the entropy of the associated ISI distributions. The information per spike $I(Z_{in}, T)$ is then given simply by

$$I(Z_{in}, T) = H(T) - H(T|Z_{in}) \quad (8)$$

where $H(T)$ and $H(T|Z_{in})$ are total and conditional entropies, respectively, of the ISI distribution. The information rate (units: bits/second) is then just the information per spike (units: bits/spike) times the firing rate R (units: spikes/second)

$$I(Z_{in}, Z_{out}) = R \times I(Z_{in}, T) \quad (9)$$

The representation of the output spike train as a sequence of firing times $\{t_o, \dots, t_n\}$ is entirely equivalent (except for edge effects) to the representation as a sequence of ISIs $\{T_o, \dots, T_n\}$, where $T_i = t_{i+1} - t_i$. The advantage of using ISIs rather than spike times is that $H(T)$ depends only on the ISI distribution $P(T)$, which is a univariate distribution. This dramatically reduces the amount of data required.

In the sequel we assume that spike times are discretized at a finite time resolution Δt . The assumption of finite precision keeps the potential information finite. If this assumption is not made, each spike has potentially infinite information capacity; for example, a message of arbitrary length could be encoded in the decimal expansion of a single ISI.

Equation 8 represents the information per spike as the difference between two entropies. The first term is the total entropy per spike

$$H(T) = - \sum_{i=0}^{\infty} p(T_i) \log_2 P(T_i) \quad (10)$$

where $P(T_i)$ is the probability that the length of the ISI was between T_i and T_{i+1} . The distribution of ISIs can be obtained from a single long (ideally, infinite) sequence of spike times.

The second term is the conditional entropy per spike. The conditional entropy is just the entropy of the ISI distribution in response to a particular set m of input spikes $[Z_{in}(t)]_m$, averaged over all possible sets of inputs spikes

$$H(T|Z_{in}) = - \left\langle \sum_{j=1}^{\infty} P(T_j|[Z_{in}(t)]_m) \log_2 P(T_j|[Z_{in}(t)]_m) \right\rangle_{|[Z_{in}(t)]_m} \quad (11)$$

where $\langle \cdot \rangle$ represents average. Here $P(T_j|[Z_{in}(t)]_m)$ is the probability of obtaining an ISI of length T_j in response to a particular set of input spikes $[Z_{in}(t)]_m$.

We used the following algorithm for estimating the conditional entropy:

1) *Generate ensemble of input spikes.* Some particular ensemble of input spikes (corresponding, for example, to $m = 17$) is generated, $[Z_{in}(t)]_{17} = [z_1^{17}(t), \dots, z_j^{17}(t), \dots]$, where $[z_1^{17}(t), \dots, z_j^{17}(t), \dots]$ are independent homogenous Poisson processes (for convenience we assume they have the same rate, but this is not essential).

2) *Compute conditional ISI distribution.* The conditional distribution $P(T|[Z_{in}(t)]_{17})$ of ISIs of the model neuron is obtained by measuring the ISIs on a large (ideally, an infinite) number of trials in which a synaptic current is generated from $[Z_{in}(t)]_{17}$ by using the synaptic noise equations Eq. 2 or Eq. 3. If the noise is nonzero, then each realization of the synaptic current $i_{syn}^{17}(t)$ is slightly different, leading to variability in the output ISI.

3) *Compute conditional entropy for this input ensemble.* From the conditional distribution, the conditional entropy in response to this particular input ensemble is computed as $H(T|[Z_{in}(t)]_{17}) = -\sum_T P(T|[Z_{in}(t)]_{17}) \log_2 P(T|[Z_{in}(t)]_{17})$. This ISI distribution depends on the amount of synaptic noise assumed; if there is no noise, the output distribution assumes only a single value and the conditional entropy is zero.

4) *Repeat and average over conditional entropies for other ensemble.* The average conditional entropy per spike is calculated by repeating this procedure for a large (ideally, infinite) number of input patterns $[Z_{in}(t)]_1, [Z_{in}(t)]_2, \dots$ and averaging over the resulting conditional entropies.

In summary, we have described the three steps required to compute the information rate in our model. First, the total entropy per spike is computed from Eq. 10 and the conditional entropy per spike is computed from Eq. 11. Next, the information per spike is computed from Eq. 8. Finally, the information rate (information per time) is computed from Eq. 9.

Model assumptions

We have assumed a model of neuronal dynamics in which ISIs are independent. This assumption simplifies the estimation of the information rate, because it reduces the estimation of the multidimensional distribution of spike times to the estimation of the one dimensional ISI distributions ($P(T)$ and $P(T|Z_{in}(t))$), from which the mutual information can be calculated exactly. Under what conditions will ISIs be independent? Because correlated ISIs can arise either from the spike generating mechanism itself or the input signal, we consider the validity of our assumptions about each in turn.

The first assumption is that the spike-generating mechanism does not induce correlations between ISIs. We have used a standard ‘‘memoryless’’ integrate-and-fire model in which the length of one ISI has no influence on subsequent ISIs. At least in cortical neurons, this assumption is not strictly valid for at least two reasons. First, on long time scales, adaptation (i.e., a change in the firing rate that depends on the firing rate itself) becomes important. Second, low-pass filtering by dendrites may induce temporal correlations in the effective synaptic current reaching the spike generator, even if they did not exist in the input ensemble. Correlations between ISIs may either increase or decrease the information rate.

The second assumption is that the correlations do not arise from the synaptic drive. This assumption may be inadequate for at least three reasons. First, it requires that EPSCs be much shorter than typical ISIs. Correlations in the synaptic drive are unlikely to arise from the fast AMPA component, but might well arise from the NMDA component, which decays much more slowly. A second potential source of correlations in the synaptic drive is correlations in the spike trains of each of the input neurons. To the extent that each input spike train is not a homogenous Poisson spike train, the model must be reevaluated. Finally, correlations might arise

through the history-dependence of efficacy at individual synapses (Abbott et al. 1997; Dobrunz and Stevens 1997; Markram and Tsodyks 1996; Varela et al. 1997; Zador and Dobrunz 1997). We have made no attempt to explore the potentially important consequences of such use-dependent effects.

Informative upper bound

The assumption that successive ISIs are independent (i.e., that the spike train is a renewal process) leads to an exact expression (rather than the upper bound provided by the reconstruction method) for the mutual information, subject only to error in the estimation of the ISI distribution. Here we review the well known result that a Poisson process (the special case where the ISI distribution is exponential) leads to the maximum entropy spike train, and give the simple closed-form expression for the entropy in this case.

The upper bound on the possible information transmitted in this model is straightforward to calculate (MacKay and McCulloch 1952). The output is a binary string—we have disallowed the possibility of multiple spikes per bin. If the conditional entropy is zero (i.e., if there is no noise whatsoever), then all the entropy is information, and the upper bound on the entropy is equal to the upper bound I_{ub} on the information.

The probability of observing a spike in a bin of length Δt depends on the firing rate R as $P_1 = R \times \Delta t$ and the probability of not observing a spike is $P_0 = 1 - R \times \Delta t$. If spikes are independent—that is, if the probability of observing a spike in one bin does not depend on whether there was a spike in any neighboring bin, so that the spike train is a Poisson process—then the entropy per bin is $\sum_i P_i \log_2 \frac{1}{P_i} = P_0 \log_2 \frac{1}{P_0} + P_1 \log_2 \frac{1}{P_1}$. At low firing rates, $P_0 \rightarrow 1$, and $P_0 \log_2 \frac{1}{P_0} \rightarrow 0$, so the entropy per bin is approxi-

mately $P_1 \log_2 \frac{1}{P_1} = R \times \Delta t \log_2 \frac{1}{R \times \Delta t}$. The entropy rate (entropy per time) is then the entropy per bin divided by the time per bin Δt or

$$I_{ub} = R \log_2 \left(\frac{1}{R \times \Delta t} \right) \quad (12)$$

This upper bound on the information encoded in a discretized spike train is achieved if 1) there is no noise, 2) spikes are independent, and 3) the spike rate is low compared to the inverse bin size, $R \ll 1/\Delta t$. It shows that the information rate increases almost linearly with the firing rate, but the information per spike $\log_2 \left(\frac{1}{R \times \Delta t} \right)$ decreases logarithmically.

RESULTS

Synaptic variability is the dominant source of output variability

Mainen and Sejnowski (1995) have previously shown that the timing of spikes produced by cortical neurons in response to somatic current injection can be highly reliable. The currents they injected were obtained by passing a Gaussian signal through a low-pass filter representing the time course of an EPSC and adding a constant offset. Although such a Gaussian current is obtained in the limit as the number of inputs becomes large, Mainen and Sejnowski (1995) did not explicitly relate the current they injected to the underlying synaptic drive.

Figure 1A shows a typical experiment in which the same current was injected into the soma of a pyramidal neuron in layer II/III of a slice of rat neocortex and the response on 20 consecutive trials was recorded. In the experiment shown, a single 1,024 ms waveform was generated according to Eq. 2 and then stored; this precise waveform was injected on 20 trials. Figure 1A shows that most of the spikes are aligned with a “jitter” of ≤ 1 ms, although a few “stray” or “displaced” spikes are also seen. In agreement with the observations of Mainen and Sejnowski (1995), these results show that cortical neurons can generate precisely repeated outputs in response to precisely repeated inputs, even when the driving current corresponds to a synthetic synaptic current generated by an ensemble of independent inputs. The small remaining output variability seen in Fig. 1A is due to some combination of experimental instability and the intrinsic imprecision of spike generator. Experiments in which precisely the same current is injected establish a limit on the output precision of which these neurons are capable. The output variability increases as other sources of variability, such as synaptic noise are considered.

Synaptic failures occurring at even a relatively low rate dramatically increase the output variability. Figure 1B shows the response of the same neuron to injected current (as in

Fig. 1A, generated according to the synaptic model described in Eq. 2), but assuming that synapses failed to elicit a postsynaptic response on average 3 of every 10 spikes ($P_r = 0.7$). The response to 20 consecutive trials was recorded. Thus in contrast to Fig. 1A—in which precisely the same current was injected on each trial—for this experiment a somewhat different waveform, corresponding to the random removal of 3/10 spikes from the input ensemble, was injected on every trial. Figure 1B shows that spikes are no longer well aligned, indicating that under these conditions synaptic failures are the dominant source of output variability.

Information rate depends on firing rate

Experiments like those shown in Fig. 1 suggest that synaptic noise represents an important source of output variability. Such experiments can be used to estimate information rates in cortical neurons by using techniques developed elsewhere (Buracas et al. 1996; de Ruyter van Steveninck et al. 1997). In an experimental setting, however, information estimates can be distorted by nonstationarity, finite data sizes, variability between neurons, and a number of other factors. Although it is possible to correct for such factors (subject to certain




FIG. 1. Synaptic variability is dominant source of output variability. *Top*: spike generator is reliable. Response of a neuron from layer II/III of a slice of rat neocortex to 20 consecutive even-numbered trials in which precisely same synthetic synaptic current was injected through a somatic electrode (see METHODS). Most of spikes are aligned to a precision of ~ 1 ms, although a few “stray” or “displaced” spikes are also seen. This experiment places a lower bound on precision with which spikes can be generated in response to identically repeated stimuli; remaining variability is due to some combination of experimental noise and intrinsic variability of spike generator. *Bottom*: noisy synapses introduce output variability. Response to 20 consecutive odd-numbered trials (interleaved with even-numbered trials presented in *top*) is shown. In this experiment, synthetic currents were generated from same ensemble as in *top*, using a fixed pattern of presynaptic spikes drawn from a Poisson ensemble, but assuming that, because of synaptic failures, 3/10 spikes failed to elicit an EPSC ($P_r = 0.7$). [Current repeatedly injected in *top* is equivalent to assumption that precisely the same 3/10 spikes failed to elicit an excitatory postsynaptic current (EPSC) on every trial.] Under these conditions, effective output reliability is markedly decreased, as seen by poor alignment of spikes giving a haphazard appearance to raster. For this experiment, quantal fluctuations, which would tend to further decrease output reliability, were suppressed ($CV = 0$). Parameters for synthetic synaptic currents: quantal size (mean), 30 pA; quantal size (coefficient of variation): 0 $P_r = 0.7$, $N_r = 1$.

reasonable assumptions), here we focus on the results from a model neuron in which all assumptions are explicit; this permits us to focus specifically on the role of synaptic variability in governing transmitted information.

In what follows, we consider a model in which the spike generating mechanism is completely deterministic, known, and stationary. Thus variability in the output spike train is due solely to variability among the stochastic inputs. In this section we begin with the limiting case in which the only source of variability among the inputs is the quantal variability of the synapses, i.e., to the variation in the postsynaptic response that occurs even when only a single functional contact is successfully activated. Thus in this section we assume not only that 1) the spike generating mechanism is completely deterministic, but also that 2) synapses release transmitter reliably when an action potential invades the presynaptic terminal ($P_r = 1$). Here as elsewhere, the exact sequence of action potentials arriving at each of the presynaptic terminals is the signal and any variability response to repeated trials on which precisely the same sequence is presented represents the noise.

The information per spike is defined (Eq. 8) as the difference between the total and conditional entropies per spike. Figure 2A shows how these quantities depend on the firing rate for the integrate-and-fire spike generation model given by Eq. 1. The dashed curve represents the total entropy, which quantifies the total output variability of the spike train. The dotted line represents the conditional entropy, which quantifies the variability that remains when the signal (i.e. the precise firing times of each of the inputs) is held constant. The solid line is the mutual information between the input and the output and is the difference between these quantities. If there were no quantal variability, the conditional entropy would be zero and all the entropy would be information. Figure 2A shows that even when the only source of synaptic variability is quantal, only about 3/4 of the spike entropy (6 bits/spike conditional entropy vs. 8 bits/spike total en-

trophy at 4 Hz) is information. As seen in the next section, additional sources of synaptic variability reduce this fraction further.

The information and entropies per spike decrease monotonically with firing rate. These quantities diverge logarithmically to infinity as the firing rate goes to zero and in fact the entropy rates were calculated for firing rates only as low as ~ 4 Hz. The behavior of the total entropy per spike at low firing rates can be understood in terms of the results for the limiting case of Poisson model outlined (see Eq. 12).

In contrast to the entropy and information per spike, the entropy and information per second increase with increasing firing rate. The reason is that the entropy and information per time depend only logarithmically on firing rate, so the overall dependence, $I \propto R \log_2 1/R$, is increasing (see Eq. 12). Figure 2B illustrates the entropy and information rates (units: bits/second) corresponding to the curves shown in Fig. 2A. Because of our assumption that time is discretized into bins of length Δt , each containing only at most one spike, the information declines back to zero at very high firing rates (not shown).

The information rate is a nearly linear function of the firing rate (Fig. 2B). This is precisely the behavior that would be expected from the maximum entropy Poisson process (Eq. 12). Although the output of the integrate-and-fire model is not a Poisson process, the dependence on firing rate is qualitatively similar: an increased firing rate compensates for a logarithmic decrease in the entropy per spike.

Information rate depends on release probability

The invasion of a synaptic terminal by an action potential often fails to induce a postsynaptic response both in the hippocampus (Allen and Stevens 1994; Dobrunz and Stevens 1997) and in the cortex (Stratford et al. 1996). Although the release probability P_r varies across synapses onto the same neuron (Castro-Alamancos and Connors 1997; Hessler et al. 1993;

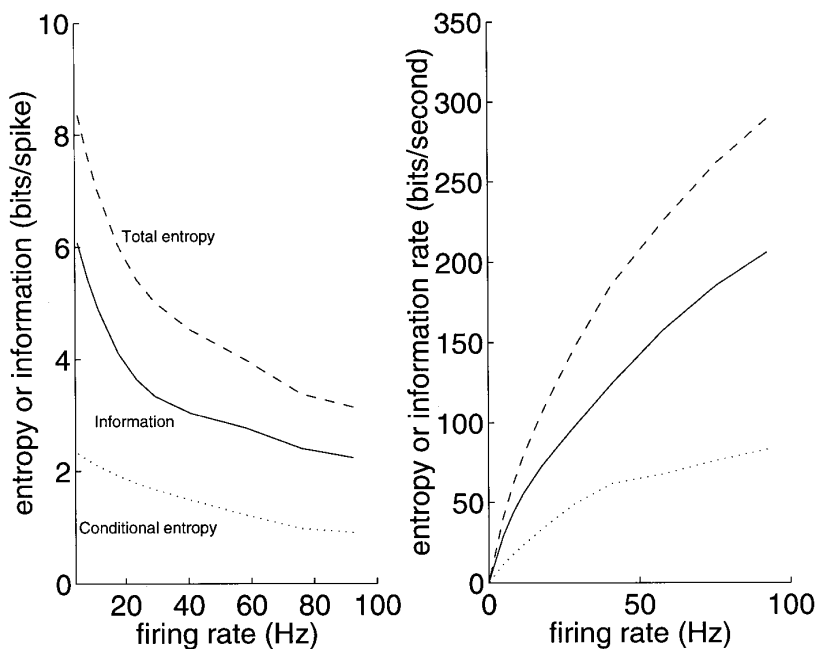


FIG. 2. Dependence of entropy and information firing rate in a model neuron. *Left*: entropy and information per spike are plotted as a function of firing rate in a model integrate-and-fire neuron. (---): total entropy, which quantifies total output variability of spike train. (···): conditional entropy, which quantifies variability that remains when signal is held constant. (—): mutual information between input and output and is difference between these quantities. *Right*: corresponding entropy and information rates in bits/ms are shown. Parameters: $V_{\text{thresh}} = -40$ mV; $R_n = 150$ M Ω ; $\tau = 50$ ms; $V_{\text{reset}} = -50$ mV; $V_{\text{rest}} = -60$ mV; quantal size (mean): 30 pA; quantal size (coefficient of variation): 0.2; $P_r = 1$; and $N_r = 1$. Spike rate was varied by increasing presynaptic Poisson input rate. Smooth curves shown represent fit of a high-order polynomial to values computed at a large number of firing rates. In this and all other simulations presented, a binsize of 1 ms was used.

Rosenmund et al. 1993) and as a function of history of use (Abbott et al. 1997; Dobrunz and Stevens 1997; Markram and Tsodyks 1996; Varela et al. 1997), for simplicity we make the assumption here that the release probability P_r is the same at all terminals.

Figure 3A shows the dependence of information on firing rate for several values of P_r . The top curve shows $P_r = 1$ and is the same as the solid curve in Fig. 2B. The lower three curves show that as P_r is decreased (to 0.9, and 0.6, and 0.3), the form of the dependence is largely preserved, but the curves are shifted down. Thus as expected, synaptic unreliability lowers the information rate. In these simulations, the input Poisson rate S_{net} was held constant as described in *Model of synaptic drive*, so the decrease in the information rate was due solely to an increase in the conditional entropy per spike and not to a change in firing rate. Fig. 3B illustrates the dependence on P_r in more detail. For this curve, the firing rate was held constant at 40 Hz and P_r was varied from zero to one. The information is a monotonically increasing function of P_r . This is reasonable, because as the synaptic reliability increases, so should the reliability with which information is transmitted. No sharp transition is observed from an unreliable to a reliable mode.

Information rate depends on the number of functional contacts per axon

A single axon may sometimes make multiple synapses onto a postsynaptic target, or a single synapse (such as the neuromuscular junction) might have multiple release sites. To avoid ambiguity, we use functional contact to refer to any release site from a presynaptic axon to a postsynaptic target, whether it involves multiple synapses per axon or multiple release sites per bouton. At the neuromuscular junction, functional contacts are counted by the thousands (Katz 1966). At excitatory synapses in the cortex, the number of functional contacts is much smaller, but still sometimes

greater than one (Markram and Tsodyks 1996; Sorra and Harris 1993). We have therefore explored the consequences of multiple functional contacts on the information rate.

Figure 4A shows the dependence of information rate on release probability for three different values of N_r , the number of functional contacts per axon. As in the previous simulations, the input Poisson rate S_{net} was held constant as described in *Model of synaptic drive*. The bottom curve is the same as that shown in Fig. 3B. As the number of functional contacts is increased, the information available in the output spike train is increased as well. Because multiple functional contacts can be seen as a form of redundancy, the increase in transmitted information is not unexpected. In these simulations, the input Poisson rate S_{net} was held constant as described in *Model of synaptic drive*. Although the firing rate was slightly ($<10\%$) higher for large N_r , the increase in the information rate was due primarily to both an increase in the total entropy per spike and a decrease in the conditional entropy per spike and not to the increased firing rate.

Figure 4B illustrates the dependence on the number of functional contacts N_r in more detail. For this curve, release probability was held constant at $P_r = 0.5$ and N_r was varied from 1 to 25. The information saturates at high N_r , but no sharp transition is seen from a low to a high reliability mode.

Reliability of mean rate coding

It may seem obvious that because multiple functional contacts increase the fidelity with which a presynaptic signal is propagated, it can overcome the noise induced by synaptic failures and quantal fluctuations and thereby increase the fidelity of neuronal signaling. In the previous section we quantified this intuition under the hypothesis that the precise timing of spikes carries information. To what extent does this conclusion depend on the particular assumptions we are making about the neural code?

According to the ‘‘mean-rate’’ hypothesis for the neural




FIG. 3. Information depends on synaptic release probability. *Top*: left information rate is plotted as a function of firing rate for 4 values of release probability $P_r = 1, 0.9, 0.6, 0.3$, in a model integrate-and-fire neuron (top to bottom). Top curve is same as the middle curve shown in Fig. 2, *bottom*. *Bottom*: right information rate is plotted as a function of release probability P_r at $F = 40$ Hz. In each simulation, P_r was same at all synapses. To maintain Poisson input rate S_{net} constant, Poisson rate at each synapse was increased to compensate for decrease in Poisson rate because of synaptic failures. Thus for all curves, EPSCs arrived at a net rate of 2.4/ms (see *Model of synaptic drive* for details). Except as indicated, parameters are same as in Fig. 2, *bottom*.

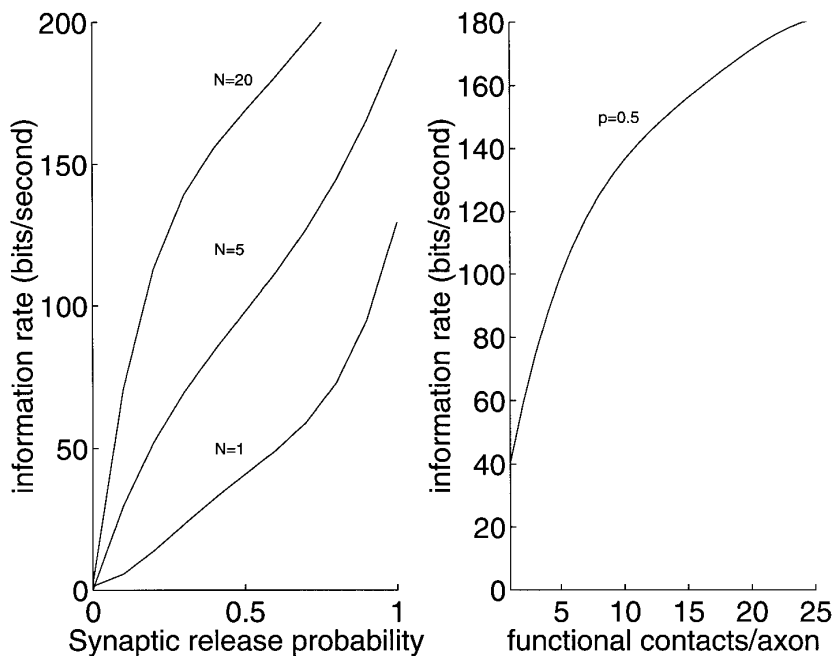


FIG. 4. Information rate depends on number of functional contacts. *Left*: information rate is plotted as a function of release probability P_r for 3 values of number of functional contacts $N_r = 1, 5,$ and 20 (bottom to top) in a model integrate-and-fire neuron. Bottom curve is same as that shown in Fig. 3, *bottom*. *Right*: information rate is plotted as a function of number of functional contacts for $P_r = 0.5, F = 40$ Hz. To maintain Poisson input rate S_{net} constant, Poisson rate at each synapse was increased to compensate for changes in S_{net} because of synaptic failures or number of functional contacts; thus for all curves, EPSCs arrived at a net rate of $2.4/\text{ms}$ (see *Model of synaptic drive* for details). Except as indicated, parameters are same as in Fig. 2.

code, the signal is carried not by the times at which spikes occur, but instead by the number of output spikes generated in some relatively long window. Under this hypothesis, multiple functional contacts can actually have the seemingly paradoxical effect of decreasing the transmitted information.

We use the Fano factor (Fano 1947) to assess the reliability of coding under the mean rate hypothesis. The Fano factor is defined as the variance σ_N^2 divided by the mean μ_N of the spike count N in some time window W . The Fano factor can be viewed as a kind of “noise-to-signal” ratio; it is a measure of the reliability with which the spike count could be estimated from a time window that on average contains several spikes. In fact, for a renewal process like the neuronal spike generator considered here, the distribution $P_N(N, W)$ of spike counts can be shown (Feller, 1971) by the central limit theorem to be normally distributed (asymptotically, as the number of trials becomes large), with $\mu_N = W/\mu_{\text{isi}}$ and $\sigma_N = W\sigma_{\text{isi}}^2/\mu_{\text{isi}}^3$, where σ_{isi} and μ_{isi} are, respectively, the mean and the standard deviation of the ISI distribution $P(T_i)$. Thus the Fano factor F is related to the coefficient of variation $C_v = \sigma_{\text{isi}}/\mu_{\text{isi}}$ of the associated ISI distribution by $C_v = \sqrt{F}$.

Figure 5 shows the Fano factor as a function of the number of functional contacts. The spike trains are the same as those analyzed in Fig. 4B. The Fano factor increases monotonically with the firing rate. Because the reliability with which the spike count can be estimated is inversely related to its variability, an increase in the number of functional contacts results in a decrease in the effective signal-to-noise ratio. This suggests that if a mean-rate coding scheme is used, an increase in the number of functional contacts could actually decrease the coding fidelity. This behavior stands in marked contrast to that observed in the previous section, where the increase in functional contacts produced the expected increase in information rate.

How can we account for this seemingly paradoxical decrease in signal-to-noise with increased redundancy? The resolution rests in the normalization used to increase the

connection redundancy. In these simulations, the net input Poisson rate S_{net} was held constant (as described in *Model of synaptic drive*) to keep the mean postsynaptic current $i_{\text{syn}}(t)$, and therefore the firing rate, constant. A large N_r leads to a redistribution of the presynaptic spikes into a small number of highly synchronous events, surrounded by longer




FIG. 5. Information is inversely proportional to number of functional contacts in a mean rate code. In these simulations, input Poisson rate S_{net} was held constant. Fano factor (variance divided by mean of spike count) during a 250-ms window is plotted as a function of number of functional contacts. This measure can be thought of as an effective “noise-to-signal” ratio for a mean rate code, because it reflects how well spike count can be estimated. A larger ratio indicates that spike count is harder to estimate. Curve illustrates that an increase in number of functional contacts leads to an increase in variance of synaptic current driving neuron and thereby an increase in Fano factor. To maintain rate of Poisson input S_{net} constant, Poisson rate at each synapse was increased to compensate for changes in S_{net} because of synaptic failures or number of functional contacts; thus for all curves, EPSCs arrived at a net rate of $2.4/\text{ms}$ (see *Model of synaptic drive* for details). Except as indicated, parameters are same as in Fig. 2.

periods during which no spike occurred. Normalizations that do not increase the effective synchrony might have given a different result.

These synchronous events have two effects. First, they tend to trigger postsynaptic action potentials at precise times. This increased timing precision decreases the conditional entropy (and thereby increases the total information by Eq. 8) under the coding assumptions analysed in this section, but has no effect on the available information under the mean rate hypothesis. Second, the increased synchrony increased the variance of the postsynaptic input current, which in turn leads to an increase in the output variance (as assessed by the Fano factor). This increases the total entropy and hence the total information under the coding assumptions analyzed in this section, but actually decreases the effective signal-to-noise ratio under the mean rate hypothesis. Thus the increased connection redundancy has diametrically opposed effects on the available information, depending on how the spike trains are decoded.

DISCUSSION

We have estimated the mutual information between the synaptic drive and the resulting output spike train in a model neuron. We have adopted a framework in which the time at which individual spikes occur carries information about the input. In this formulation, the exact sequence of action potentials arriving at each of the presynaptic terminals is the “signal,” and the “noise” is any variability in the response to repeated trials on which precisely the same sequence is presented. We found that the information was a smooth function of both synaptic reliability and connection redundancy: no sharp transition was observed from an “unreliable” to a “reliable” mode. However, connection redundancy can only compensate for synaptic unreliability under the assumption that the fine temporal structure of individual spikes carries information. If only the number of spikes in some relatively long time window carries information (a “mean rate” code), an increase in the fidelity of synaptic transmission results in a seemingly paradoxical decrease in the information available in the spike train.

Related work

Information rates for sensory neurons in a wide variety of experimental systems have now been measured for both static (Golomb et al. 1997; Optican and Richmond 1987; Richmond and Optican 1990; Tovee et al. 1993) and time-varying (Bair et al. 1997; Bialek et al. 1991; Buracas et al. 1996; Dan et al. 1996; de Ruyter van Steveninck and Bialek 1988; Gabbiani and Koch 1996; Gabbiani et al. 1996; Rieke et al. 1997; Warland et al. 1997) stimuli. Most of the work on time-varying stimuli used reconstruction methods to obtain a lower bound on the transmitted information; typical values were in the range of 1–3 bits/spike. De Ruyter van Steveninck and Laughlin (1996) applied similar techniques to estimate information rates across graded synapses in the blowfly.

The present model is a direct extension of that considered in Stevens and Zador (1996) and closely related to that in DeWeese (1996). Both used a direct rather than a recon-

struction method to estimate the information in a spiking neuron model. In Stevens and Zador (1996), the key assumption was that ISIs were independent, whereas in DeWeese (1996), the key assumption was that spikes were independent.

Neural code

Although it is generally agreed that the spike train output by a neuron encodes information about the inputs to that neuron, the code by which the information is transmitted remains unclear (see Ferster and Spruston 1995; Stevens and Zador 1995) for recent discussions. One idea (the conventional view in systems physiology) is that it is the mean firing rate alone that encodes the signal and that variability about this mean is noise (Shadlen and Newsome 1994, 1995). An alternative view that has recently gained increasing support is that it is the variability itself that encodes the signal, i.e. that the information is encoded in the precise times at which spikes occur (Abeles et al. 1994; Bialek et al., 1991; Rieke et al. 1997; Softky 1995).

Our results make no assumptions about the neuronal code. Rather, they provide an exact expression for the maximum information that could possibly be transmitted, given the stimuli and the neuronal parameters. The precise timing of spikes is used to achieve this maximum; how much of this available information is actually used by “downstream” neurons is a separate question.

The importance of spike timing in encoding time-varying signals is now well-established in some systems, such as the motion-sensitive H1 neuron of the fly (Bialek et al. 1991). A comparable role for spike timing in mammalian cortex has been more controversial. It has been suggested that motion-sensitive neurons in area MT of awake monkeys encode only fractions of a bit per second and that all of the encoded information is available in the spike count over a relatively long time window (Britten et al. 1992). However, more recent experiments (Bair et al. 1997; Buracas et al. 1996) suggest that these neurons encode information at rates (1–2 bits/spike) comparable with those of the H1 neuron of the fly, when presented with visual stimuli that have appropriately rich temporal structure. Thus it may be wrong to speak of *the* neural code: it may well turn out that some components of the input stimulus (e.g., those that are changing rapidly) are encoded by precise firing times, whereas others are not.

We have shown that an increase in the number of functional contacts per axon can lead to an increase in transmitted information if the timing of spikes encodes the signal, but not if a mean rate code is used. Such an increase can be seen as a special case of neuronal synchrony in which all synapses from a single axon are stimulated at precisely the same instant. This seemingly paradoxical observation is a consequence of the manner in which synchrony affects firing patterns: it increases timing precision, but also increases the trial-to-trial variability in the spike count. It is not clear how synaptic unreliability could be compensated for in a mean rate scenario.

Information and synaptic unreliability

The present paper is the first to interpret information rates in single cortical neurons in terms of the underlying biophys-

ical sources of the signal and noise. Here signal is the set of firing times over the ensemble of presynaptic neurons, whereas noise is synaptic variability that leads to variability in the firing times of the postsynaptic neuron.

The present study was centrally motivated by the hypothesis that the nervous system is under selective evolutionary pressure to preserve as much information as possible during processing. In the limit this is trivially true: a retina that transmits no information whatsoever about the visual input is no better than no retina at all! Less trivially, computational power in some models increases as the precision of the underlying components increases (Zador and Pearlmuter 1996). If such principles apply to cortical computation, then the cortex may have evolved strategies to compensate for synaptic unreliability, given other constraints.

The most obvious strategy would be simply to increase the synaptic release probability. Indeed, there are synapses (e.g., in the fly retina) (de Ruyter Van Steveninck and Laughlin 1996) where the number of release sites N_r per terminal is large enough to guarantee a high-fidelity connection under normal conditions. But such multirelease synapses are large and the cortex may be under an additional constraint to minimize size.

It is reasonable to wonder why the more direct approach, setting the release probability P_r to unity, does not appear to be common. It is well-known that the release probability P_r changes in a history-dependent manner during short-term plasticity (e.g., paired-pulse facilitation and depression, posttetanic potentiation, etc) (Abbott et al. 1997; Dobrunz and Stevens 1997; Fisher et al. 1997; Magleby 1987; Markram and Tsodyks 1996; Tsodyks and Markram 1997; Varela et al. 1997; Zador and Dobrunz 1997; Zucker 1989). We speculate that a dynamic P_r is essential to cortical computation. A dynamic P_r could function as a form of gain control (Abbott et al. 1997; Tsodyks and Markram 1997; Varela et al. 1997). More generally, it could be used to permit efficient computation on time-varying signals (Maass and Zador 1998). Thus we propose that the “reason” that P_r does not simply approach unity may be that cortical computation requires that a P_r retain a large dynamic range.

The cortex appears to adopt the “redundant connection” approach, albeit on a smaller scale. Figure 4B shows that even a modest increase in the connection redundancy from one to five can double the information rate, from one to two bits/spike. Although a direct comparison is difficult, it is interesting to note that information rates in both anesthetized (Bair et al. 1997) and alert (Buracas et al. 1996) primate visual cortex are in the same range.

In our formulation, the fraction of the signal entropy transmitted by the spike train is small, even when the signal is not corrupted by noise. This follows immediately when we consider that to drive the model neuron to fire at, for example, 40 Hz, impulses must arrive at 2,400 Hz, which is equivalent to 60 input neurons each firing at 40 Hz, with each input axon presumably carrying comparable (and, by assumption, independent) information. This captures what may be an essential feature of the cortex: each pyramidal neuron must in some sense “summarize” with a single spike train the spike trains from 10^4 other neurons. It is this “summary” that represents the “computation” that a neuron performs. Understanding the fidelity with which this computation can

occur is a necessary step toward understanding the computation.

This work was supported by The Sloan Center for Theoretical Neurobiology at the Salk Institute and by a grant to Charles F. Stevens from the Howard Hughes Medical Institute.

Received 12 September 1997; accepted in final form 19 November 1997.

REFERENCES

- ABBOTT, L., VARELA, J., SEN, K., AND N.S.B. Synaptic depression and cortical gain control. *Science* 275: 220–224, 1997.
- ABELES, M., PRUT, Y., BERGMAN, H., AND VAADIA, E. Synchronization in neuronal transmission and its importance for information processing. *Progress in Brain Research* 102: 395–404, 1994.
- ALLEN, C. AND STEVENS, C. An evaluation of causes for unreliability of synaptic transmission. *Proc. Natl. Acad. Sci. USA* 91: 10380–10383, 1994.
- BAIR, W., CAVANAUGH, J., AND MOVSHON, J. Reconstructing stimulus velocity from neuronal responses in area MT. *Advances Neural Information Processing Systems* 10: 34–40, 1997.
- BEKKERS, J. AND STEVENS, C. Origin of variability in quantal size in cultured hippocampal neurons and hippocampal slices. *Proc. Natl. Acad. Sci. USA* 87: 5359–5362, 1990.
- BIALEK, W., DEWEESE, F., RIEKE, F., AND WARLAND, D. Bits and brains: Information flow in the nervous system. *Physica A* 200: 581–593, 1993.
- BIALEK, W., RIEKE, F., DE RUYTER VAN STEVENINCK, R., AND WARLAND, D. Reading a neural code. *Science* 252: 1854–1857, 1991.
- BRITTEN, K., SHADLEN, M., NEWSOME, W., AND MOVSHON, J. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* 12: 4745–4765, 1992.
- BRYANT, H. AND SEGUNDO, J. Spike initiation by transmembrane current: a white-noise analysis. *J. Physiol. (Lond.)* 260: 279–314, 1976.
- BURACAS, G., ZADOR, A., DEWEESE, M., AND ALBRIGHT, T. Measurements of information rates in monkey MT neurons in response to time-varying stimuli. *Soc. Neurosci. Abstr.* 22: 717, 1996.
- CASTRO-ALAMANOS, M. AND CONNORS, B. Distinct forms of short-term plasticity at excitatory synapses of hippocampus and neocortex. *Proc. Natl. Acad. Sci. USA* 94: 4161–4166, 1997.
- DAN, Y., ATICK, J., AND REID, R. Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J. Neurosci.* 15: 3351–3362, 1996.
- DE RUYTER VAN STEVENINCK, R. AND BIALEK, W. Real-time performance of a movement-sensitive neuron in the blowfly visual system: coding and information transmission in short spike sequences. *Proc. R. Soc. Lond. B Biol. Sci.* 234: 379–414, 1988.
- DE RUYTER VAN STEVENINCK, R., LEWEN, G. D., STRONG, S. P., AND KOBBERLE, R. Reproducibility and variability in neural spike trains. *Science* 275: 1805–1807, 1997.
- DE RUYTER VAN STEVENINCK, R. R. AND LAUGHLIN, S. B. The rate of information transfer at graded-potential synapses. *Nature* 379: 642–645, 1996.
- DEWEESE, M. *Optimization principles for the neural code*. (PhD thesis). Princeton, NJ: Princeton Univ, 1995.
- DEWEESE, M. Optimization principles for the neural code. *News Physiol. Sci.* 8: 281–287, 1996.
- DOBUNZ, L. AND STEVENS, C. Heterogeneity of release probability, facilitation and depletion at central synapses. *Neuron* 18: 995–1008, 1997.
- FANO, U. Ionization yield of rations. II. The fluctuations of the number of ions. *Physiol. Rev.* 72: 26–29, 1947.
- FELLER, W. An introduction to probability theory and its applications (2nd Ed.). New York: Wiley, 1971, vol. 2.
- FERSTER, D. AND SPRUSTON, N. Cracking the neuronal code. *Science* 270: 756–757, 1995.
- FISHER, S., FISCHER, T., AND CAREW, T. Multiple overlapping processes underlying short-term synaptic enhancement. *Trends Neurosci.* 20: 170–177, 1997.
- GABBIANI, F. AND KOCH, C. Coding of time-varying signals in spike trains. *Neural Comput.* 8: 44–66, 1996.
- GABBIANI, F., METZNER, W., WESSEL, R., AND K. C. Coding of time-varying signals in spike trains. *Nature* 384: 564–567, 1996.
- GOLOMB, D., HERTZ, J., PANZERI, S., TREVES, A., AND R. B. How well can

- we estimate the information carried in neuronal responses from limited samples? *Neural Comput.* 3: 649–665, 1997.
- HESSLER, N., SHIRKE, A., AND M. R. The probability of transmitter release at a mammalian central synapse. *Nature* 366: 569–572, 1993.
- KATZ, B. *Nerve, Muscle, and Synapse*. New York: McGraw-Hill, 1966.
- MAASS, W. AND ZADOR, A. Dynamic stochastic synapses as computational units. *Advances Neural Information Processing Systems*. In press.
- MACKEY, D. AND MCCULLOCH, W. The limiting information capacity of a neuronal link. *Bull. Math. Biophys.* 14: 127–135, 1952.
- MAGLEBY, K. Short-term synaptic plasticity. In: *Synaptic Function*, edited by G. M. Edelman, W. E. Gall, and W. M. Cowan. New York: Wiley, 1987.
- MAINEN, Z. AND SEJNOWSKI, T. Reliability of spike timing in neocortical neurons. *Science* 268: 1503–1505, 1995.
- MARKRAM, H. AND TSODYKS, M. Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature* 382: 807–810, 1996.
- OPTICAN, L. AND RICHMOND, B. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *J. Neurophysiol.* 57: 162–178, 1987.
- RICHMOND, B. AND OPTICAN, L. Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. II. Information transmission. *J. Neurophysiol.* 64: 370–380, 1990.
- RIEKE, F., WARLAND, D., DE RUYTER VAN STEVENINCK, R., AND BIALEK, W. *Spikes: Exploring the Neural Code*. Cambridge, MA: MIT press, 1997.
- ROSENMUND, C., CLEMENTS, J., AND WESTBROOK, G. Nonuniform probability of glutamate release at a hippocampal synapse. *Science* 262: 754–757, 1993.
- SHADLEN, M. AND NEWSOME, W. Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* 4: 569–579, 1994.
- SHADLEN, M. AND NEWSOME, W. Is there a signal in the noise (comment)? *Curr. Opin. Neurobiol.* 5: 248–250, 1995.
- SHANNON, C. AND WEAVER, W. *A Mathematical theory of communication*. Urbana, IL: Univ. Illinois Press, 1948.
- SHEPHERD, G. *The Synaptic Organization of the Brain (3rd edition)*. Oxford, UK: Oxford Univ. Press, 1990.
- SOFTKY, W. Simple codes versus efficient codes. *Curr. Opin. Neurobiol.* 5: 239–247, 1995.
- SORRA, K. AND HARRIS, K. Occurrence and three-dimensional structure of multiple synapses between individual radiatum axons and their target pyramidal cells in hippocampal area CA1. *J. Neurosci.* 13: 3736–3748, 1993.
- STEVENS, C. AND ZADOR, A. Neural coding: the enigma of the brain. *Curr. Biol.* 12: 1370–1371, 1995.
- STEVENS, C. AND ZADOR, A. Information through a spiking neuron. *Advances in Neural Information Processing Systems* 8: 75–81, 1996.
- STRATFORD, K., TRACZY-HORNOCH, K., MARTIN, K., BANNISTER, N., AND J.J.J. Excitatory synaptic inputs to spiny stellate cells in cat visual cortex. *Nature* 382: 258–261, 1996.
- TOVEE, M., ROLLS, E., TREVES, A., AND B.R.P. Information encoding and the responses of single neurons in the primate temporal visual cortex. *J. Neurophysiol.* 70: 640–654, 1993.
- TSODYKS, M. AND MARKRAM, H. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Acad. Sci. USA* 94: 719–723, 1997.
- VARELA, J. A., SEN, K., GIBSON, J., FOST, J., ABBOTT, L. F., AND NELSON, S. B. A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex. *J. Neurosci.* 17: 7926–7940, 1997.
- WARLAND, D., REINAGEL, P., AND MEISTER, M. (1997). Decoding visual information from a population of retinal ganglion cells. *J. Neurophysiol.* 78: 2336–2350, 1997.
- ZADOR, A. AND DOBRUNZ, L. Dynamic synapses in the cortex. *Neuron* 19: 1–4, 1997.
- ZADOR, A. M. AND PEARLMUTTER, B. A. VC dimension of an integrate and fire neuron model. *Neural Comput.* 8: 611–624, 1996.
- ZUCKER, R. Short-term synaptic plasticity. *Annu. Rev. Neurosci.* 12: 13–31, 1989.