Noise in neurons is message dependent

Guillermo A. Cecchi*†, Mariano Sigman*‡, José-Manuel Alonso‡, Luis Martínez‡, Dante R. Chialvo*, and Marcelo O. Magnasco*

*Center for Studies in Physics and Biology and ‡Laboratory of Neurobiology, The Rockefeller University, New York, NY 10021

Neuronal responses are conspicuously variable. We focus on one particular aspect of that variability: the precision of action potential timing. We show that for common models of noisy spike generation, elementary considerations imply that such variability is a function of the input, and can be made arbitrarily large or small by a suitable choice of inputs. Our considerations are expected to extend to virtually any mechanism of spike generation, and we illustrate them with data from the visual pathway. Thus, a simplification usually made in the application of information theory to neural processing is violated: noise is not independent of the message. However, we also show the existence of error-correcting topologies, which can achieve better timing reliability than their components.

neuronal reliability | temporal coding | information theory | error correcting

Brains represent signals by sequences of identical action potentials or spikes (1). Upon presentation of a stimulus, a given neuron will produce a certain pattern of spikes; when the stimulus is repeated, the pattern may repeat spike for spike in a highly reliable fashion, or may be similar only "rate-wise," or some spikes may be repeatable and others not. If individual spikes can be recognized and tracked across trials, then their timing accuracy can be ascertained unambiguously; this is not always the case. The existence of repeatable spike patterns and the reliability of their timing changes not only from neuron to neuron, but even for the same neuron under various circumstances. Bryant and Segundo (2) first noticed (in the mollusk *Aplysia*) that spike timing accuracy depends on the particulars of the input driving the neuron. This intriguing property has received renewed attention, notably among those searching for experimental evidence supporting various theories of neural coding (3, 4). In a study of the response of pyramidal neurons to injected current (3) the temporal pattern of firing was found to be unreliable when the injected current was constant, but highly reliable when the input current was "noisy" and contained high-frequency components. This study showed explicitly the difference between the irregularity of the spike pattern (which was irregular in both cases) as opposed to the *reliability* or accuracy of spike timing, and it also highlighted the connection that "natural" stimuli are noisy and contain sharp transitions. Similar results have been obtained from *in vivo* recordings of the H1 neuron in the visual system of the fly *Calliphora vicina* (4): constant stimuli produced unreliable or irreproducible spike firing patterns, but noisy input signals deemed to be more "natural" and representative of a fly in flight yielded much more reproducible spike patterns, in terms of both timing and counting precision. Although some aspects of the last study have been challenged recently (5), the phenomenon of different reliabilities for various inputs was reaffirmed: see, e.g., figure 1 in ref. 5. Thus one concludes that similar observations of the timing precision of spikes can be made in very different types of neurons under vastly different conditions, and that this must be a very universal, almost basic characteristic of neuronal function. We shall now show how these experimental observations follow from rather general and elementary considerations, and discuss the possible implications for brain function and architecture.

Stochastic Spiking Models. The essence of this phenomenon can be demonstrated easily in toy models. A simple example is the leaky integrate-and-fire (I-F) model (6), which assumes the neuron is a (leaky) capacitor driven by a current which simulates the actual synaptic inputs. We add membrane fluctuations representing several internal sources of noise (cluttering of ion channels, synaptic noise, etc.) to obtain a system described by the following Langevin (7) equation:

$$C\dot{V} = -gV + I(t) + \xi(t), \qquad [1]$$

where C is the cell membrane capacitance, V is the membrane potential, gV is the leakage term (g is a conductance), $I(t)$ is the input current, and $\xi(t)$ is Gaussian noise, $\langle\xi(t)\rangle = 0$ with autocorrelation $\langle\xi(t)\xi(t')\rangle = \sigma\delta(t - t')$; when the potential reaches the threshold V_0 an action potential is generated and the system returns to the equilibrium potential V_e, here set arbitrarily to zero.

Fig. 1 shows the results of a numerical simulation of Eq. **1** in response to two different signals,§ in both cases in the presence of identical "internal" noise. Following refs. 3–5, Fig. 1 *Left* shows a time-independent input and the neuronal responses to repeated stimulation, whereas Fig. 1 *Right* ensues from a "noisy" input and its responses. When $I(t)$ contains high-frequency components (*Right*), spikes are clustered more tightly to the upward strokes of the input and overall there is less trial-to-trial variability than in the time-independent case. The phenomenology described in Fig. 1 captures the main feature that renewed interest in this subject: that apparently more "noisy" (in the sense of "severely fluctuating") inputs can generate less variable responses.

By concentrating on these two extreme stimuli, these studies have somewhat obscured an issue that is central to the variability phenomenon: the relationship of the output to the timescale or timescales over which the input itself varies. The constant stimulus can be thought to have no timescale, or at best, just a timescale since onset, whereas in principle the noise process has indefinitely fast timescales. Thus the two extreme cases under discussion have timescales outside the range of interaction with physiological processes.

A simple dimensional argument shows that timing variability depends on the characteristic timescale of the input. As a consequence of noise, the neuron's membrane potential will
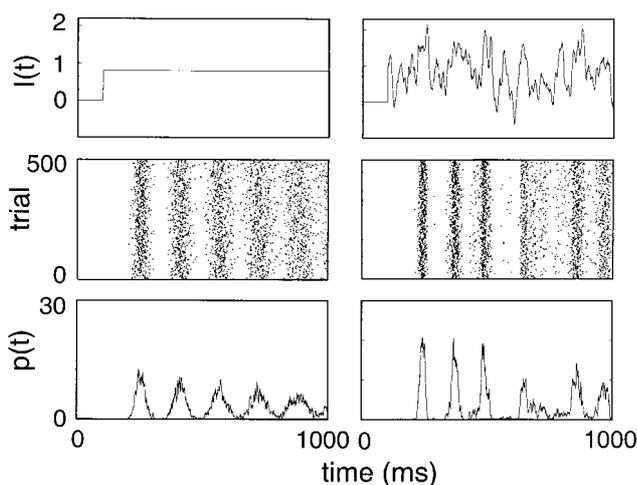
NEUROBIOLOGY

Fig. 1. Rapidly fluctuating (*Right*) inputs produce more reliable output spikes than do time-independent stimuli (*Left*). (*Top*) Input stimulus—*I*(*t*) in Eq. 1. (*Middle*) The 500 responses to the stimulus; the occurrence of each spike is denoted with a dot. (*Bottom*) The peristimulus time histogram describing the rate, *p*(*t*), at which spikes are generated in response to the stimulus *I*(*t*).

Fig. 2. Computed mean spike temporal jitter (and SEM) as a function of the estimated membrane potential's *dV/dt*. Estimates are plotted for two noise variances: $\sigma = 0.01$ (○) and $\sigma = 0.02$ (∗). The expected inverse-law fit (with exponent ~0.35) is depicted with dotted lines. The examples in the *Insets* illustrate two density distributions of spiking times resulting from relatively slow (*Inset A*) or relatively fast (*Inset B*) *dV/dt* threshold crossings. For each case, below each histogram, a few typical trajectories are also plotted, showing the membrane potential preceding the threshold crossing.

have fluctuations of a characteristic size (or root mean square) ΔV. To translate this ΔV into a timing noise Δt, we have to scale the potential noise by something having units of $\Delta t/\Delta V$; the most natural such quantity at hand is the inverse of the potential's time derivative.

While the distribution of the derivative of the potential is a reasonable measure of the degree of reliability of a particular input, one can see from a geometrical argument that the derivative of the potential at the particular time in which the potential crosses the threshold is a better estimator of stimulus reliability. The uncertainty in firing time is essentially the time during which the membrane potential is at a distance from the threshold smaller than the value of the voltage noise, and so is scaled by the derivative at this particular moment. Thus, the faster the voltage approaches the threshold (in the vicinity of the threshold), the more reliable the timing of the spike will be. The statement captures much of the phenomenology in refs. 2 and 4, and in our own data, as we will show.

We proceed now to apply this method of estimation to the situation shown in Fig. 1. A rapidly fluctuating signal was constructed,[§] and the times at which spikes occurred were determined first for the deterministic condition (i.e., $\sigma = 0$), and subsequently for several hundred stochastic realizations [i.e., with identical *I*(*t*) but with different stochastic realizations of the noise term in Eq. **1**]. At the same time the voltage derivative preceding each spike was computed over the prior 50 time steps. The fluctuating character of *I*(*t*) (as illustrated in Fig. 1 *Top Right*) offers the opportunity to explore a wide range of *dV/dt* at threshold crossings. To estimate the precision of spike timing we define an index of "temporal jitter" as follows: it is the absolute difference between the time at which a spike is generated in the noise-free simulation and that of the corresponding nearest spike for the stochastic realization, averaged over realizations, and over all spikes that occur within a *dV/dt* interval of 0.1. This quantity is plotted in Fig. 2. The data points in Fig. 2 are scattered around the expected (dotted lines) inverse relationship between the temporal jitter of the spike and the speed at which the voltage crossed the threshold. The trajectories plotted in both *Insets* help to visualize the geometrical argument already discussed: inputs that rise more rapidly to the membrane potential threshold produce less variable spike timing.
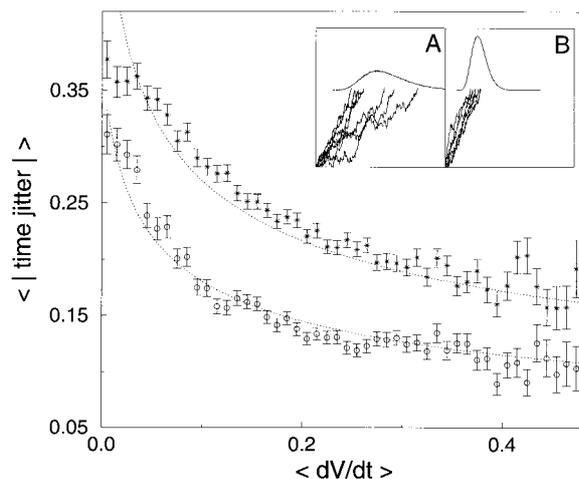
The same basic phenomenon will affect any model of spike generation in which the noise in the voltage-like variable is translated into jitter of its threshold crossings. To provide for a specific example, we choose a widely used model of excitable media with continuous dynamics, the FitzHugh–Nagumo model (FHN) with additive stochastic forcing (8), described by the following Langevin system:

$$\dot{V} = V - V^3/3 - W + I(t) + \xi(t) \qquad [2]$$

$$\dot{W} = \phi(V + a - bW). \qquad [3]$$

The variable *V* is the fast voltage-like variable, *W* is the *slow* recovery variable; *a*, *b*, and ϕ are constant parameters, and $\xi(t)$ is a zero-mean Gaussian white noise of intensity *D*. According to these equations, after a spike the recovery produces an absolute refractory period during which a second spike cannot occur, followed by a longer relative refractory period during which firing requires stronger perturbations.

As previously argued, it is biologically reasonable to assume that noise affects the voltage-like variable. Thus, according to our argument, a higher steady input will result in a faster approach to the threshold crossing and therefore will reduce the probability of an untimely crossing. This expectation is confirmed in Fig. 3, where the distributions of periods are plotted for two input amplitudes. Note the large change in variance (1.5 to 2.34, 56%) for comparably small change in the average period (36.4 to 39.1, 7%) in going from high to low input; also note that the low-input distribution has a supraexponential tail. For comparison, the distribution of interspike intervals corresponding to the integrate-and-fire model with constant input is shown. In both cases, the distribution for high input tends to a Gaussian (parabola on a semilog plot), whereas the low input has an exponential tail (linear on a semilog plot) as expected for Poisson-like statistics. Notice that even though this phenomenon will affect more strongly systems with zero-frequency quiescent-to-firing bifurcations, given the longer time to integrate fluctuations at low input (i.e., the system can evolve near a homoclinic orbit), the example presented here shows that it also affects
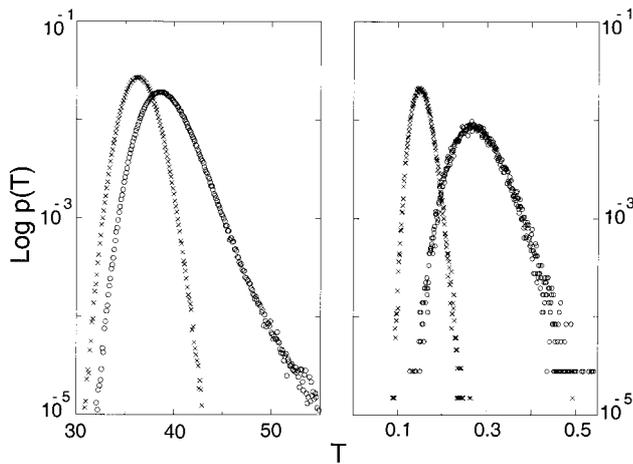
Cecchi *et al.*

Fig. 3. First passage time distributions for the FitzHugh–Nagumo (*Left*) and integrate-and-fire (*Right*) models, forced with constant input. In both panels, × corresponds to high input and ○ to low.

non-zero-frequency transitions like the Hopf bifurcation present in the FitzHugh–Nagumo model.

Experimental Results. Thus, in neural data, one should expect to find that the output noise generically depends on the structure of the input. Experiments discussed next demonstrate the relevance of this observation for brain function. We recall that an assumption usually made in the application of information theory to neurons is that the noise introduced by the communication channel is independent of the message being transmitted, and now we explore to what extent this assumption is valid in real neural data.

The experimental methods were as described by Alonso and Martinez (9). Briefly, cats (2.5–3 kg) were anesthetized with ketamine (10 mg/kg) and thiopental sodium (20 mg/kg; maintenance 2 mg/kg per hr, i.v.) and paralyzed with Norcuron (Organon; 0.2 mg/kg per hr, i.v.). Temperature (37.5–38°C), EKG, EEG, and expired CO_2 were monitored throughout the experiment. Pupils were dilated with 1% atropine sulfate and the nictitating membranes were retracted with 10% phenylephrine. A multi-electrode matrix was used for the cortical recordings. Recorded signals were collected and analyzed by a computer running Datawave Systems software (Broomfield, CO).

The result is presented in Fig. 4, where different visual stimuli (i.e., the four messages) were presented to an anesthetized cat while electrophysiological recordings were made in the lateral geniculate nucleus, which is the second stage of processing in the visual pathway.

The raster plot in Fig. 4*A* shows the response to a moving bar with high contrast and high speed, and Fig. 4*B* shows the responses to a low-contrast bar moving at low speed; both results are shown in a window of 200 ms centered at the peak of the ON response. The responses in *A* clearly display a higher temporal precision than those for the condition in *B*. Two further conditions were also presented, for high contrast and low speed, and for low contrast and high speed. In the *Inset* of Fig. 4 we quantify the noise for the four messages as the jitter of the response onset, which we define as the average of the absolute value of the position with respect to the mean of the first spike in each trial. The messages include two conditions for the velocity and two for the contrast. For each velocity condition, the high contrast is more reliable than the low one: 8.4 ms vs. 22.1 ms, and 16.8 ms vs. 23.8 ms. Note, however, that the high-contrast/high-speed and low-contrast/low-speed stimuli, although they have a very similar spike response (4.17 vs. 3.86) have the largest difference
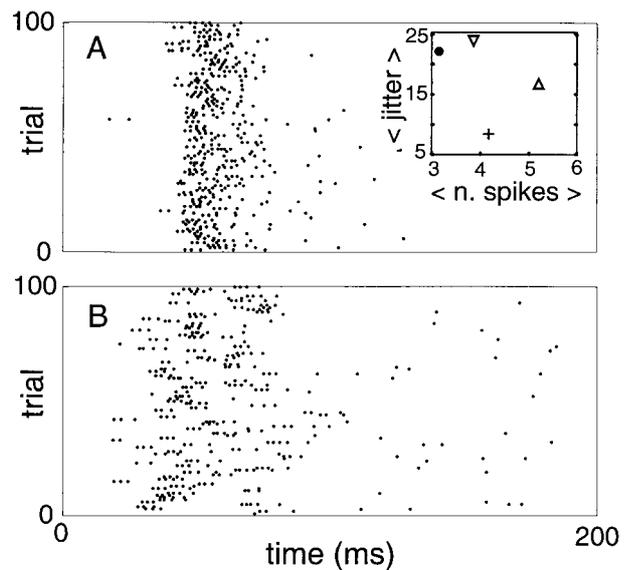

Fig. 4. Experimental example showing that different "messages" elicit different noise (i.e., spike temporal jitter). Results were gathered from recordings in cat lateral geniculate nucleus, where neuronal action potentials are recorded in response to the presentation of a moving bar with different contrast and speed. *Inset* in *A* shows, for the conditions tested, the temporal jitter versus average number of spikes (collected during the window of 200 ms); ● corresponds to low contrast/high speed; +, to high contrast/high speed; ▽, to low contrast/low speed; and △, to high contrast/low speed. Notice the data from moving bars of different contrast originating very different jitter. The raw data for two of the experimental conditions are presented in the raster plots of *A* and *B*. Responses obtained with a moving bar of low contrast moving at low speed are plotted in *B*. The (less variable) responses obtained with a moving bar of high contrast moving at high speed are depicted in *A*.

in jitter (8.4 ms vs. 23.8 ms). Alternatively, we can measure the entropy difference between the corresponding first spike probability distributions (1-ms bins), which varies from 2.0 bits for low contrast/high speed vs. high contrast/high speed, to 0.6 bit for high contrast/low speed vs. high contrast/high speed. Similar results showing contrast-dependent temporal precision in the lateral geniculate nucleus have been reported previously in ref. 10, where it is also proposed that the variability is due to intrinsic noise. We also show that messages with indistinguishable mean rate responses can show large differences in their variability, in agreement with the hypothesis of the threshold-crossing derivative as the relevant parameter, which can be independent of the firing rate for particular inputs.

Network Architecture and Reliability

The observations discussed in the previous section are restricted to the behavior of individual neurons. How can they be extended to networks of spiking elements? In principle, we can assume that input-dependent noise will be present in early stages of processing, given the relative simplicity of the neural architecture. The relevance of this phenomenon in higher-order areas remains to be explored; nevertheless, we may hypothesize that unless dedicated architectures are implemented to eliminate it, significant input-dependent noise will be ubiquitous in neural processing.

Indeed, a recurring question in brain theory and computer theory is whether there are reliable ways to compute that use unreliable components. We shall now show explicitly the existence of a network topology for which the overall time reliability
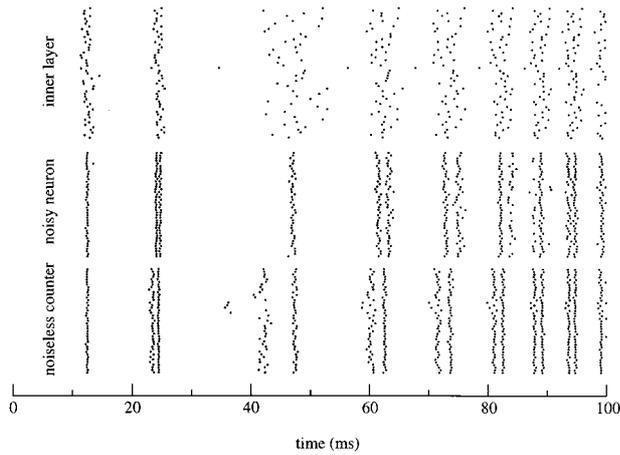
Fig. 5. Reliability of a fan-out/fan-in structure. An input $I(t)$ is fanned out in parallel to N identical neurons; one extra (identical) neuron fans in: it receives input from all N, with the weights adjusted so as to fire once per $N/2$ input spikes. In this example, $N = 51$ and $I(t)$ is a constant plus an Ornstein–Uhlenbeck process. (*Top*) Rasters for *one* instance and all 51 neurons, to illustrate the reliability of the middle layer. (*Middle*) Raster for the output neuron, 40 instances. We emphasize that the output neuron has *identical* parameters to the 51 neurons in the middle layer, including identical levels of internal noise. (*Bottom*) A noiseless counter firing twice per 51 spikes will implement a median measurement. Notice that both *Middle* and *Bottom* show *two* spikes per each spike of the middle layer neurons, and one of them (the closer to the median) is the more reliable.

can be made arbitrarily better than that found in its individual neurons.

Assume a single input source $I(t)$, which "fans out" to a middle layer of N noisy neurons X_i; these in turn "fan in" or input into a single (also noisy) neuron Y. Let us imagine an $I(t)$ constructed so that the X_i fire once each at times t_i. If the size of the connections $X \rightarrow Y$ is set so that Y fires after seeing half of the expected number of spikes coming from the Xs, then clearly neuron Y will be fire near the *median* time of the spike times t_i. The median, as a descriptor of the t_i, has a number of exceedingly nice properties in this context: (*i*) its variance decreases as $1/\sqrt{N}$, and thus the timing accuracy of Y can be made better than the individual Xs simply by choosing an appropriately large N; (*ii*) the median, unlike the mean, is exceedingly robust against outliers and heavy-tailed noise; (*iii*) the median does not require that it see the entire set t_i but only half of it; and (*iv*) the median is expected to lie at the time of highest concentration of the t_i, thus maximizing the timing accuracy of Y. This property can be fully appreciated in Fig. 5. Thus, we see that different topologies will propagate spike timing accuracy in different fashions, and that one probably should expect architectural correlates in place in the brain.

To gain further insight into this issue and to study the implications of our observations on the synchronization of a neuronal network, we now concentrate on random but fixed time delays across pathways, rather than on individual timing accuracy. We modify the previous model: all neurons of the middle layer are now noise-free and in principle identical, and so are their spike trains; the neurons are integrate-and-fire units as in Eq. **1,** with the further addition of a refractory period. We next model a dispersion of the transmission delay times by rigidly shifting the spike train of each individual neuron by a random value uniformly distributed between 0 and δ. The total charge that the integrating neuron receives does not depend on this shift, but when all neurons are synchronized, the derivative of the charge will be high, and then we expect more reliable responses. Fig. 6 shows how the timing reliability (open circles), measured
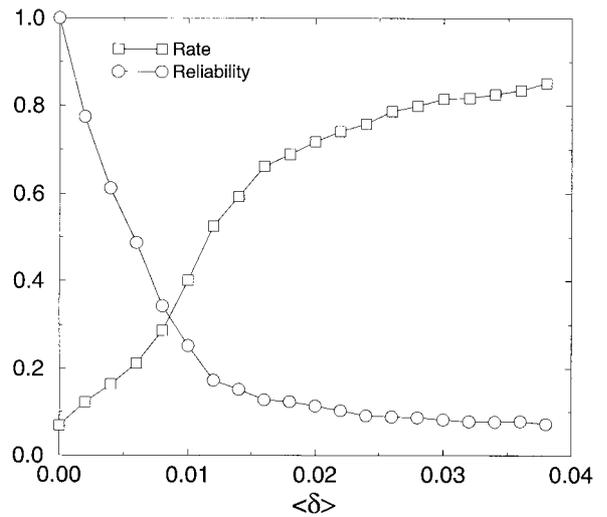

Fig. 6. Simple model to study the effects of synchronization on the reliability of the neuronal responses. Seventy-five neurons synapse on a unique neuron modeled as a leaky integrate-and-fire as in Eq. **1**, with the addition of a refractory period; $g = 0.005/s$, $\sigma = 0.07$, refractory period = 10 ms. Mean rate (□) and reliability (○) are shown for a neuron integrating spikes over an input layer of 75 neurons. A pattern of spikes was randomly generated (3.75 Hz), and each neuron of the input layer feeds the same pattern of spikes with a delay δ. The values of the other parameters are as in Fig. 5. When the input layer is synchronized, $\dot{V}$ of the integrating neuron is big and the response is very reliable. In this case, part of the charge from the input layer is lost because there is no summation during the refractory period. When δ increases, the reliability decreases but the mean firing rate is increased. This result suggests that a function of synchronization, which is known to occur in many different neuronal networks, might be to increase the reliability on the system at the expense of the intensity of the signal.

by the method of spike metrics (11), decreases with the decorrelation of the input layer. Interestingly, when the derivatives are high the response is more reliable, but by the same token, the current supplied during the refractory period will be lost. This means that smooth distributions will result in an increased effective charge and thus the rate of response of the integrating neuron is bigger when the input layer is decorrelated, as shown in Fig. 6 (open squares). [There is a second regime (not shown), if the decorrelation is large compared with the decay time of the neuron, in which rate decreases with increasing δ because of current losses.]

We think that this very simple concept might be of physiological relevance, in that it implies a tradeoff between temporal reliability and the dynamic range of the signal: when the dynamic range is large, it is possible to afford the cost of losing part of the charge to achieve reliability, whereas in a situation in which the dynamic range is limited, reliability has to be killed to preserve and propagate the signal. This idea suggests the possibility of integration pathways that multiplex the signal, as found in retinal adaptation.

Conclusions

The relevance of the message-dependent nature of noise in spiking elements is, to our understanding, twofold. One aspect is its consequences for the use of Shannon's Information Theory as a framework to measure the information content present in the output of a neuron. This is the subject of much work recently (12). In this regard, a basic assumption commonly made is that the noise introduced by the communication channel is independent of the message being transmitted, which allows the modeling of a neuron as a Gaussian channel (13, 14). This assumption is indeed desirable, given that departures from the Gaussian

channel have proved to be rather cumbersome; for instance, no general theory for the computation of channel capacity exists (15). Our results render average measures of information per spike less meaningful than usually thought, and speak for the necessity of concentrating on particular individual spikes. This concentration is especially critical when the message consists of only a few spikes (for instance, in the auditory pathway), or in the event of fast conductance modulations (16), which can dramatically affect the reliability. Thus, perhaps the most fundamental consequence of Mainen and Sejnowski's result (3) is the implicit demonstration that cortical neurons are not always classical Gaussian channels.

A second relevant aspect is the possible anatomical correlate of this phenomenon. In this regard, we have shown that unless dedicated architecture is implemented to reduce the multipli-cation of noise along the processing pathways, the encoding of behaviorally relevant (not only "artificial" stimuli) can be highly degraded. Finally, we demonstrated that it is possible to design network topologies with arbitrarily large temporal reliability, and consequently one should expect an evolutionary pressure to implement them in specific areas of the brain where time accuracy is essential.

1. Adrian, E. D. (1928) *The Basis of Sensation* (Norton, New York).
2. Bryant, H. L. & Segundo, J. P. (1976) *J. Physiol. (London)* **260,** 279–314.
3. Mainen, Z. F. & Sejnowski, T. J. (1995) *Science* **268,** 1503–1506.
4. de Ruyter van Steveninck, R. R., Lewen, G. D., Strong, S. P., Koberle, R. & Bialek, W. (1997) *Science* **275,** 1805–1808.
5. Warzecha, A.-K. & Egelhaaf, M. (1999) *Science* **283,** 1927–1930.
6. Knight, B. W. (1972) *J. Gen. Physiol.* **59**, 734–766.
7. Risken, H. (1989) *The Fokker-Planck Equation: Methods of Solution and Applications* (Springer, Berlin).
8. Longtin, A. & Chialvo, D. R. (1998) *Phys. Rev. Lett.* **81,** 4012–4014.
9. Alonso, J. M. & Martinez, L. M. (1998) *Nat. Neurosci.* **1,** 395–403.
10. Reich, D. S., Victor, J. D., Knight, B. W., Ozaki, T. & Kaplan, E. (1997) *J. Neurophysiol.* **77,** 2836–2841.
11. Victor, J. D. & Purpura, K. P. (1997) *Network* **8,** 127–164.
12. Rieke, F., Warland, D., de Ruyter van Steveninck, R. & Bialek, W. (1997) *Spikes: Exploring the Neural Code* (MIT Press, Cambridge, MA).
13. Shannon, C. & Weaver, W. (1949) *The Mathematical Theory of Communication* (Univ. of Illinois Press, Urbana).
14. Borst, A. & Theunissen, F. E. (1999) *Nat. Neurosci.* **2,** 947–957.
15. Ash, R. B. (1990) *Information Theory* (Dover, New York).
16. Borg-Graham, L. J., Monier, C. & Fregnac, Y. (1998) *Nature (London)* **393,** 369–373.

NEUROBIOLOGY