

Chapter 12

Correlated Neuronal Activity: High- and Low-Level Views

Emilio Salinas¹, and Terrence J. Sejnowski^{2,3}

¹Department of Neurobiology and Anatomy, Wake Forest University School of Medicine, Medical Center Boulevard, Winston-Salem, NC 27157-1010, U.S.,

²Computational Neurobiology Laboratory, Howard Hughes Medical Institute, The Salk Institute for Biological Studies, 10010 North Torrey Pines Road, La Jolla, CA 92037, U.S., ³Department of Biology, University of California at San Diego, La Jolla, CA 92093, U.S.

CONTENTS

12.1 Introduction: the timing game	341
12.2 Functional roles for spike timing	343
12.2.1 Stimulus representation	343
12.2.2 Information flow	345
12.3 Correlations arising from common input	347
12.4 Correlations arising from local network interactions	350
12.5 When are neurons sensitive to correlated input?	353
12.5.1 Coincidence detection	353
12.5.2 Fluctuations and integrator models	354
12.6 A simple, quantitative model	357
12.6.1 Parameterizing the input	357
12.6.2 A random walk in voltage	359
12.6.3 Quantitative relationships between input and output	362
12.7 Correlations and neuronal variability	364
12.8 Conclusion	365
12.9 Appendix	366
References	367

12.1 Introduction: the timing game

Correlated firing is a common expression used in Neuroscience. It refers to two or more neurons that tend to be activated at the same time. It is used so frequently in part because there are so many timescales at which one may analyze neural activity. In a

sense, correlation might appear as a trivial phenomenon. For instance, if one looks at day-long activity, practically the whole cerebral cortex fires in a correlated manner, because of the sleep-wake cycle. Similarly, whenever an object appears within the visual field, many neurons in visual cortex are expected to respond throughout the same time interval. Clearly, such correlations are to be expected. However, as the observation time window becomes smaller, explaining the presence of correlations becomes more difficult and, at the same time, potentially much more useful. Suppose the activity of two visual neurons is monitored during presentation of a visual stimulus, after its onset. Suppose also that within a short time window of, say, a few hundred milliseconds, spikes from the two neurons tend to appear at the same time. Why is this? Neither the sensory information nor the state of the subject are changing in an appreciable way, so the correlation must reflect something about the internal dynamics of the local circuitry or its connectivity. This is where correlations become interesting.

Thus, correlations at relatively short timescales become useful probes for understanding what neural circuits do, and how they do it. This is what this chapter is about. This analysis goes down to the one millisecond limit (or even further), where correlation changes name and becomes synchrony. Even at this point, the significance of correlated activity cannot be taken for granted. Some amount of synchrony is practically always to be expected simply because cortical neurons are highly interconnected [14, 91]. The question is not just whether there is any correlated activity at all, but whether timing is an issue and correlations make any difference. In other words, given the function of a particular microcircuit or cortical area, if the system were able to control the level and timescale of correlated activity, what would the optimal values be? For example, in a primary sensory area, stimulus representation is of paramount importance, so maybe measuring an excess of coincident spikes in this case is not an accident, but a consequence of the *algorithm* that local circuits use to encode stimulus features. This is just an example; the broader question is whether neurons exploit the precise coincidence of spikes for specific functions. There are several theoretical proposals that revolve around this concept; we discuss some of them below.

Asking about the functional implications of correlated activity is one way to attack the problem; this is a top-down approach. Another alternative is to take a bottom-up view and investigate the biophysical processes related to correlated firing. These come in two flavors, mechanisms by which correlations are generated, and mechanisms by which a postsynaptic neuron is sensitive to correlated input. In this case valuable information can be obtained about possible correlation patterns and timescales, and in general about the dynamics of correlated activity.

This approach is also important because it sheds some light on a fundamental question: how does a single cortical neuron respond under realistic stimulation conditions? The reason this is a problem is the interaction between spike-generating mechanisms, which are inherently nonlinear, and the input that drives the neuron, which typically has a complicated temporal structure. The major obstacle is not the accuracy of the single-neuron description; in fact, classic conductance-based models [24] like the Hodgkin-Huxley model [48] are, if anything, too detailed. The larger

problem is the complexity of the total driving input, which is mediated by thousands of synaptic contacts [14, 91]. What attributes of this input will the postsynaptic response be most sensitive to? Correlations between synaptic inputs are crucial here because they shape the total input and hence the postsynaptic response. Determining what exactly is their role is a key requisite for understanding how neurons interact dynamically and how the timing of their responses could be used for computational purposes.

This chapter reviews work related to both perspectives on correlated activity: the high-level approach at which function serves to guide the analysis, and the low-level approach that is bound to the biophysics of single neurons. Eventually (and ideally), the two should merge, but currently the gap between them is large. Nevertheless, comparing results side by side provides an interesting panorama that may suggest further clues as to how neurons and neural circuits perform their functions.

12.2 Functional roles for spike timing

There is little doubt that the correct timing of action potentials is critical for many functions in the central nervous system. The detection of inter-aural time differences in owls and the electrosensory capabilities of electric fish are two well-known examples [21]. In these cases it is not surprising that timing is important; it is ingrained in the nature of the sensory signals being detected. The issue of timing also arises naturally in the rodent somatosensory system [7]. To explore their surroundings, rats move their whiskers periodically. To locate an object, whisker deflections need to be interpreted relative to whisker position, which can be determined from the phase of the motor signal. Thus, the latencies of stimulus-evoked responses relative to such internal signal can be used to encode spatial information. This mechanism by which sensory-triggered activity is interpreted relative to an internal, reference signal may be applicable to other circuits and in a more general way [6, 57].

12.2.1 Stimulus representation

Spike timing, however, has been discussed in an even wider sense than implied by the above examples. No doubt, this is partly because oscillations at various frequencies and synchronous activity are so widespread [8–86]. One proposal that has received considerable attention is that the coordinated timing of action potentials may be exploited for stimulus representation [71–44]. Specifically, neurons that have different selectivities but fire synchronously may refer to the same object or concept, binding its features. The following experiment [51] illustrates this point. The receptive fields of two visual neurons were stimulated in two ways, by presenting a single object, and by presenting two objects. Care was taken so that in the two conditions practically the same firing rates were evoked. The synchrony between pairs of neurons varied

across conditions, even when the firing rates did not. Thus, correlations seemed to code whether one or two stimuli were shown [51].

In general, changes in firing rate pose a problem when interpreting variations in synchrony or correlations, first, because the latter can be caused by the former, and second, because the impact of a change in correlation upon downstream neurons becomes uncertain given a simultaneous change in firing rate. When neural activity is compared in two conditions involving different stimuli, it is likely that the evoked firing rates from the recorded neurons will change; even the populations that respond within a given area may be different. This is one of the main factors that muddles the interpretation of experiments in which correlations have been measured [72]. The most solid paradigms for investigating correlated activity are those in which variations in correlation are observed without variations in stimulation and without parallel changes in firing rate, but fulfilling all of these conditions requires clever experimental design and analysis.

There are many other studies in which correlations have been interpreted as additional coding dimensions for building internal representations. The following are cases in which the confounding factors just mentioned were minimized. Consider two neurons with overlapping receptive fields, and hence a considerable degree of synchrony. Analysis of the activity of such visual neurons in the lateral geniculate nucleus has shown [23] that significantly more information about the stimulus (around 20% more) can be extracted from their spike trains if the synchronous spikes are analyzed separately from the nonsynchronous ones. In a similar vein, recordings from primary auditory cortex indicate that, when a stimulus is turned on, neurons respond by changing their firing rates and their correlations [26]. In many cases the firing rate modulations are transient, so they may disappear if the sound is sustained. However, the evoked changes in correlation may persist [26]. Thus, the correlation structure can signal the presence of a stimulus in the absence of changes in firing rate.

Finally, the antennal lobe of insects is an interesting preparation in which this problem can be investigated. Spikes in this structure are typically synchronized by 20 Hz oscillations [90]. When these neurons are artificially desynchronized [55], the specificity of downstream responses is strongly degraded, selectivity for different odors decreases, and responses to new odors arise, even though this loss of information does not occur upstream. Apparently, what happens is that the downstream cells — Kenyon cells in the mushroom bodies — act as coincidence detectors that detect synchronized spikes from projection neurons in the antennal lobe. Kenyon cells have very low firing rates and are highly selective for odors, so in effect they sparsify the output of the antennal lobe [62]. In addition, disrupting synchrony in this system has a real impact on behavior: it impairs odor discrimination [79]. This preparation is also convenient for studying the biophysical mechanisms underlying such oscillatory processes [9, 10].

These examples show that the neural codes used to represent the physical world can be made more efficient by taking into account the pairwise interactions between neural responses. The degree to which this is actually a general strategy used by neurons is uncertain; the key observation is that, under this point of view, correlations

are stimulus-dependent, just like sensory-evoked firing rates. The studies discussed below suggest a different alternative in which correlations change rapidly as functions of internal events and may regulate the flow of neural information, rather than its meaning [68].

12.2.2 Information flow

The regulation of information flow is illustrated by the following result [70]. When intracortical microstimulation is applied during performance of a visual-motion discrimination task, the subject's response is artificially biased, but the bias depends strongly on the time at which the microinjected current is delivered relative to stimulus onset. Microstimulation has a robust effect if applied during presentation of the visual stimulus, but it has no effect if applied slightly earlier or slightly later than the natural stimulus [70]. This suggests that even a simple task is executed according to an internal schedule, such that the information provided by sensory neurons is effectively transmitted only during a certain time window. How does this internal schedule work? One possibility is that changes in correlations are involved [68]. This is suggested by a number of recent experiments in which correlations were seen to vary independently of stimulation conditions. To work around the usual problems with stimulus-linked correlations, investigators have studied correlated activity in paradigms where, across trials, stimulation conditions remain essentially constant and the most significant changes occur in the internal state of a subject.

Riehle and colleagues trained monkeys to perform a simple delayed-response task where two cues were presented sequentially [66]. The first cue indicated a target position and instructed the animal to get ready, while the second cue gave the go signal for the requested hand movement. Crucially, the go signal could appear 600, 900, 1200 or 1500 ms after the first cue, and this varied randomly from trial to trial. Neurons recorded in primary motor cortex increased their synchrony around the time of the actual sensory stimulus or around the time when the animal expected the go signal but it did not appear [66]. The latter case is the most striking, because there the firing rates did not change and neither did the stimulus; the synchronization depended exclusively on the internal state of the monkey.

Fries and colleagues [35] used attention rather than expectation to investigate the synchrony of visual neurons in area V4. They used conditions under which firing rates varied minimally, taking advantage of the finding that, although attention may have a strong effect on the firing rates evoked by visual stimuli, this modulation is minimized at high contrast [65]. Monkeys were trained to fixate on a central spot and to attend to either of two stimuli presented simultaneously and at the same eccentricity. One of the stimuli fell inside the receptive field of a neuron whose activity was recorded. Thus the responses to the same stimulus could be compared in two conditions, with attention inside or outside the neuron's receptive field. At the same time, the local field potential (LFP) was recorded from a nearby electrode. The LFP measures the electric field caused by transmembrane currents flowing near the electrode, so it gives an indication of local average activity [38]. The correlation that was studied in these experiments [35] was that between the LFP and the recorded

neuron's spikes. The key quantity here is the spike-triggered average of the LFP, or STA. The STA is obtained by adding, for each spike recorded, a segment of the LFP centered on the time of the spike; the final sum is then divided by the total number of spikes. The result is the average LFP waveform that is observed around the time of a spike. STAs were computed for attention outside and inside the receptive field. They were similar, but not identical: rapid fluctuations were more pronounced when attention was directed inside the receptive field; in the Fourier decomposition, power in the low frequency band (0–17 Hz) decreased while power in the high frequency band (30–70 Hz) increased. Because the STA reflects the correlation between one neuron and the neighboring population, the interpretation is that, as attention shifts to the receptive fields of a cluster of neurons, these become more synchronized at high frequencies and less so at low frequencies. Although the changes in synchrony were modest — on average, low-frequency synchronization decreased by 23% and high-frequency synchronization increased by 19% — changes in firing rate were also small; these were enhanced by a median of 16% with attention inside the receptive field. Under these conditions the changes in synchrony could be significant in terms of their impact on the responses of downstream neurons.

The study just discussed [35] suggests that synchrony specifically in the gamma band (roughly 30–80 Hz) may enhance the processing of information in some way. But what exactly is the impact of such synchronization? Another recent study [34] suggests at least one measurable consequence: the latencies of synchronized neurons responding to a stimulus may shift in unison. In this case the paradigm was very simple: oriented bars of light were flashed and the responses of two or more neurons in primary visual cortex (V1) were recorded, along with LFPs. Neurons were activated by the stimuli, and the key quantity examined was the time that it took the neurons to respond — the latency — which was calculated on each trial. Latencies covaried fairly strongly from trial to trial (mean correlation coefficient of 0.34, with a range from 0.18 to 0.55), so pairs of neurons tended to fire early or late together. This tendency depended on the amount of gamma power in the LFPs right before the stimulus. When the LFPs from two electrodes both had a strong gamma component, the latency covariation between the two recorded neurons from the same pair of electrodes was high. Note that the spectral composition of the LFPs was only weakly related to changes in firing rate, so short latencies were probably not due to changes in excitability. This means that, if neurons get synchronized around 40 Hz right before a stimulus is presented, they will respond at about the same time [34]. In other words, while the mean firing rates are mostly insensitive to shifts in oscillation frequencies, the time spread in the evoked spikes from multiple neurons is much smaller when the gamma oscillations are enhanced. This could certainly have an impact on a downstream population driven by these neurons [18–67]. Thus, the modulation of latency covariations [34] is a concrete example of how the synchrony of a local circuit may be used to control the strength of a neural signal.

Finally, we want to mention two other studies [36, 37] that also investigated the synchronization of V1 neurons, this time using an interocular rivalry paradigm. In rivalry experiments, different images are shown to the two eyes but only one image is perceived at any given moment [52]. The perception flips from one image

to the other randomly, with a characteristic timescale that depends on the experimental setup. The studies in question [36, 37] were done in awake strabismic cats, a preparation with two advantages: V1 neurons are dominated by a single eye, so their firing rates essentially depend on what their dominant eye sees regardless of the other one, and it is relatively easy to know which of the two images is perceived (at equal contrasts for the two images, one eye always suppresses the other, and this can be measured by tracking the cat's eye movements in response to conflicting moving stimuli). The two conditions compared were: a single image presented to the eye driving the recorded neurons, or the same stimulus shown to the driving eye plus a conflicting image presented to the other eye. The firing rates in these two conditions should be the same, because strabismus makes most neurons monocular; indeed, the rates did not change very much across conditions and did not depend on which image was perceived. However, synchrony within the 40 Hz band did change across conditions [36, 37]. When neurons were driven by the eye providing the percept, synchrony was much stronger in the rivalrous condition than in the monocular one. In contrast, when neurons were driven by the eye whose image became suppressed, synchrony was much lower in the rivalrous condition than in the monocular one. In other words, when conflicting images were presented, neurons responding to the image being perceived were always more synchronized. In this case, stronger synchronization in the high frequency band (30–70 Hz) is suggested to be a neural correlate of stimulus selection [36, 37].

In summary, it is possible that correlations between neurons can be controlled independently of firing rate. Two ideas that have been put forth are: that this may serve to generate more efficient neural codes [71, 44], which follows from theoretical arguments and experiments in which correlations vary in a stimulus-dependent way; or to regulate the flow of information [68], which follows from experiments in which correlations have been linked to expectation, attention, sensory latencies and rivalry — all processes that regulate the strength but not the content of sensory-derived neural signals. Other alternatives may become apparent in the future.

Next we discuss some common types of correlated activity patterns. In part, the goal is to describe them mathematically, at least to a first-order approximation.

12.3 Correlations arising from common input

As mentioned above, oscillations and synchronous responses are commonly observed throughout the nervous system [8–86]. This is not particularly surprising; in fact, correlations are to be expected simply because neurons in the brain are extensively interconnected [14, 91]. Now we will discuss two major mechanisms that give rise to correlated activity, common input and recurrent connectivity. The distinction between them is somewhat artificial, but it is useful in portraying the range of correlation patterns that may arise. Although they will not be included, it should

be kept in mind that intrinsic oscillatory properties of neurons are also important in determining global rhythmic activity [53–54].

An important analytical tool used to study the joint activity of neurons is the cross-correlation histogram or cross-correlogram [63–15], which is constructed from pairs of spike trains. This function shows the probability (or some quantity proportional to it) that neuron B fires a spike t milliseconds before or after a spike from neuron A , where t is called the time shift or time lag. When the two spike trains are independent, the cross-correlogram is flat; when they covary in some way, one or more peaks appear [15]. A peak at zero time shift means that the two neurons tend to fire at the same time more often than expected just by chance. Interpreting a cross-correlogram constructed from experimental data can be quite difficult because any covariation during data collection will show up as a peak [15]. Two neurons, for example, may respond at the same time to changes in stimulation conditions even if they are independent; this will produce a peak that has nothing to do with the functional connectivity of the circuit, which is what one is usually interested in. Another problem with this technique is that it requires large amounts of data. These disadvantages, however, have much lesser importance with simulated spike trains because they can be very long and their statistics can be constant.

Figure 12.1 shows synthetic, computer-generated spike trains from neurons that share some of their driving inputs but are otherwise disconnected. Responses from 20 neurons are displayed in each panel. Continuous traces superimposed on the spike rasters show the mean spike density or instantaneous firing rate, averaged over all neurons; this quantity is proportional to the probability of observing a spike from any of the neurons at any given time. Cross-correlograms are shown below. As mentioned above, the y-axis indicates the probability of observing a pair of spikes separated in time by the amount on the x-axis. The normalization is such that the probability expected by chance is equal to 1. The spikes shown were produced by integrate-and-fire model neurons [24, 67, 82], each driven by two time-varying signals, $g_E(t)$ and $g_I(t)$, representing the total excitatory and inhibitory conductances generated by large numbers of synaptic inputs. Details of the model are given in the Appendix. To generate synchronous activity between postsynaptic responses, the conductances $g_E(t)$ and $g_I(t)$ were correlated across neurons. This is exactly what would happen if pairs of postsynaptic neurons shared some fraction of all presynaptic spike trains driving them. In Figure 12.1a the mean correlation between conductances was 0.2. This means that, for any pair of neurons i and k , the correlation coefficient between excitatory conductances,

$$\rho_E^{ik} = \frac{\langle (g_E^i - \langle g_E^i \rangle) (g_E^k - \langle g_E^k \rangle) \rangle}{\sqrt{\langle (g_E^i - \langle g_E^i \rangle)^2 \rangle} \sqrt{\langle (g_E^k - \langle g_E^k \rangle)^2 \rangle}}, \quad (12.1)$$

was approximately 0.2. In this expression the angle brackets $\langle \rangle$ indicate an average over time, and neurons are indexed by a superscript. Inhibitory conductances also had a correlation of 0.2 across neurons, but all excitatory and inhibitory conductances were independent of each other. The nonzero correlation between conduc-

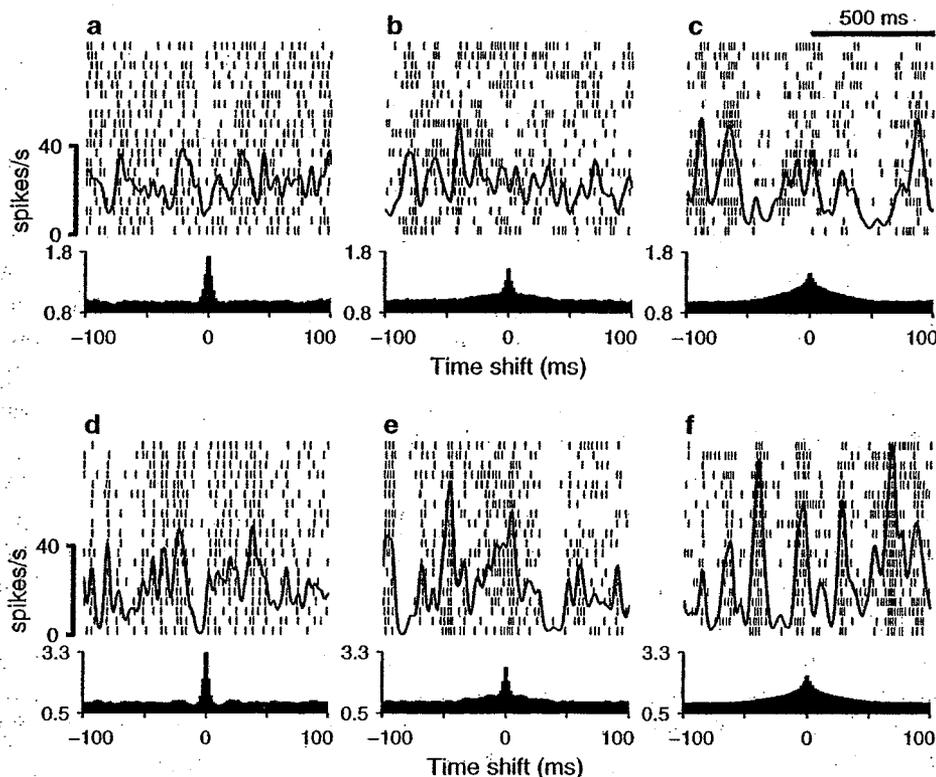


Figure 12.1

Spike trains correlated by common input. Each panel includes 20 computer-generated spike trains. Each row represents one neuron and each small, vertical line one spike. Neurons were modeled as leaky integrate-and-fire units disconnected from each other but driven by synaptic conductances that co-fluctuated across neurons. Continuous traces superimposed on the rasters are firing rates, averaged over all neurons, obtained by smoothing the spike trains with a Gaussian function with $\sigma=10$ ms. Plots below the rasters are cross-correlation histograms averaged over multiple distinct pairs of units. These were based on longer spike trains that included the segments shown.

tances gives rise to the sharp peak in the histogram of Figure 12.1a.

Figure 12.1b was generated using the same correlation values, but the excitatory signals $g_E(t)$ varied more slowly (in addition, their magnitude was adjusted so that similar output rates were produced). The characteristic time at which $g_E(t)$ varies is its correlation time. Below we describe this quantity more accurately; for the moment the crucial point is that in Figure 12.1 the correlation time of $g_E(t)$ corresponds to the time constant of excitatory synapses, τ_E . This is essentially the duration of a unitary synaptic event. In Figure 12.1a the synaptic time constants for excitation and inhibition were both equal to 2 ms. In Figure 12.1b τ_I stayed the same but τ_E was increased to 20 ms. As can be seen in the raster, this changed the postsynaptic

responses: the spike trains are more irregular; the spikes of a single neuron tend to appear in clusters. The two timescales show up in the cross-correlogram as a sharp peak superimposed on a wider one. Figure 12.1c shows what happens when both synaptic time constants are set to 20 ms. Now the clustering of spikes in individual spike trains is even more apparent and the cross-correlogram shows a single, wide peak.

The correlations between conductances parameterize the degree of synchrony among output responses. When the correlations are 0 the responses are independent and the cross-correlogram is flat; when the correlations are equal to 1 all neurons are driven by the exact same signals and thus produce the same spike train — this is perfect synchrony. Figures 12.1a–12.1c were generated with correlations of 0.2, whereas Figures 12.1d–12.1f were generated with correlations of 0.5. Notice that the shapes of the histograms in the top and bottom rows are the same, but the y-axis scales in the latter are much larger. Larger correlations always produce more synchrony and larger fluctuations in instantaneous firing rates (continuous traces). In addition, they may also alter the postsynaptic firing rates, but this effect was intentionally eliminated in Figure 12.1 so that different synchrony patterns could be compared at approximately equal firing rates.

These examples show that there are at least two important factors determining the synchronous responses caused by common input: the amount of common input, which corresponds to the magnitudes of the correlations between conductances, and the timescales of the input signals, which in this case are determined by synaptic parameters. Analogously, there are two aspects of the cross-correlation function that are important, the height of the peak and its width.

12.4 Correlations arising from local network interactions

Networks of recurrently interconnected neurons may naturally give rise to oscillatory and synchronous activity at various frequencies; this is a well documented finding [94–17]. The type of activity generated depends on the network's architecture, on its inputs, and on single-cell parameters. Here we illustrate this phenomenon with a highly simplified network with the following properties. (1) Model neurons, excitatory and inhibitory, are of the integrate-and-fire type, without any intrinsic oscillatory mechanisms. (2) Synaptic connections between them are all-to-all and random, with strengths drawn from a uniform distribution between 0 and a maximum value g_{max} ; this is both for excitatory and inhibitory contacts. (3) All neurons receive an external input drive implemented through fluctuating conductances $g_E(t)$ and $g_I(t)$, which are uncorrelated across neurons.

Figure 12.2 illustrates some of the firing patterns produced by such a network. For Figure 12.2a the recurrent connections were weak, i.e., g_{max} was small. The

peak in the cross-correlogram is also small, indicating that the neurons fired nearly independently. The peak is narrow because all synaptic time constants were set to

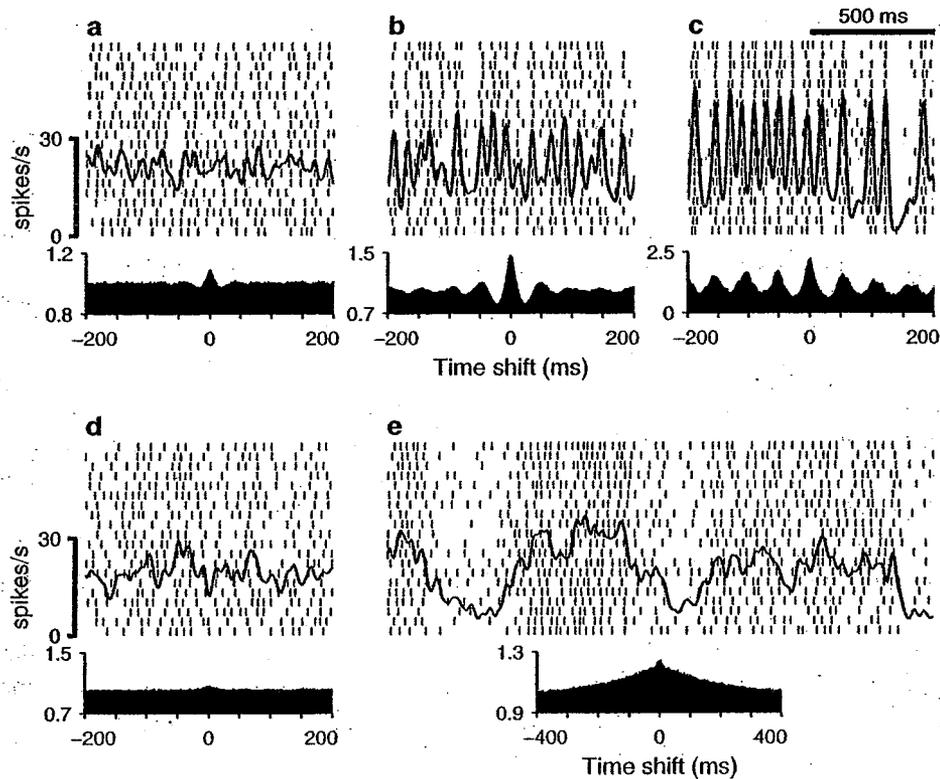


Figure 12.2

Spike trains correlated by local network interactions. Same format as in Figure 12.1. Neurons were modeled as integrate-and-fire units receiving two types of inputs: a background synaptic drive that was independent across neurons, and recurrent synaptic input from other units in the network. The full network consisted of 100 excitatory and 25 inhibitory model neurons. Synaptic connections were all-to-all, with conductances chosen randomly (uniformly) between 0 and a maximum value g_{max} . Note the different y-axes for cross-correlation histograms.

3 ms. Figures 12.2b and 12.2c show what happens when the connections are made progressively stronger. The central peak becomes much taller (notice the different y-axes), and secondary peaks, indicating oscillatory activity, become apparent. In contrast to the rest of the cross-correlograms, the one in Figure 12.2c was generated from a short segment of data. This enhanced the secondary peaks, which practically disappeared when longer stretches of data were used (not shown). This is because the frequency of the oscillatory activity is not constant so, over a long time, many

phases are averaged out, making the correlation flat everywhere except in the central region. As in Figure 12.1, compensatory adjustments were made so that average firing rates remained approximately the same; in this case the external excitatory drive was slightly decreased as the connection strengths increased.

Figures 12.2d and 12.2e show that even such a simplified network may have quite complex dynamics. Parameters in Figure 12.2d were identical to those of Figure 12.2b, except for two manipulations. First, for recurrent excitatory synapses only, the synaptic time constant was increased from 3 to 10 ms; and second, to compensate for this, the synaptic conductances were multiplied by 3/10. This generated approximately the same firing rates and also preserved the average recurrent conductance level. However, as a consequence of these changes the correlations between postsynaptic spikes almost disappeared. Thus, the tendency to fire in phase is much larger when the characteristic timescales for excitatory and inhibitory synaptic events are the same. This is reminiscent of resonance.

Figure 12.2e illustrates another interesting phenomenon. In this case the timescales of both excitatory and inhibitory recurrent synapses were set to 10 ms, while the characteristic time of all external input signals stayed at 3 ms. The mean conductance levels, averaged over time, were the same as in Figure 12.2c, so the connections were relatively strong. Now the peak in the cross-correlation histogram (Figure 12.2e) is much wider than expected, with a timescale on the order of hundreds of milliseconds. Such long-term variations are also apparent in the spike raster and in the firing rate trace. This is quite surprising: firing fluctuations in this network occur with a characteristic time that is at least an order of magnitude longer than any intrinsic cellular or synaptic timescale. Discussion of the underlying mechanism is beyond the scope of this chapter, but in essence it appears that the network makes transitions between two pseudo steady-state firing levels, and that the times between transitions depend not only on cellular parameters but also on how separated the two firing levels are.

In any case, a key point to highlight is that, in all examples we have presented, the cross-correlation histograms show a common feature: a central peak with a shape that resembles a double-exponential. That is, the correlation function can be described as

$$C(t) = C_{max} \exp\left(-\frac{|t|}{\tau_{corr}}\right), \quad (12.2)$$

where t is the time lag and τ_{corr} , which determines its width, is the correlation time. This function is related to many types of random processes [87, 42]. Indeed, below we will use it to characterize the total input that drives a typical cortical neuron.

So far, we have looked at a variety of correlation patterns that a network may display. Next, we take the point of view of a single downstream neuron that is driven by this network. From this perspective we return to an important question posed in the introduction: how does the response of a postsynaptic neuron depend on the full set of correlated spike trains that typically impinge on it? Equation 12.2 will be used as a rough characterization of those correlations. The answer will be presented in two parts. First we will discuss some of the main factors determining whether input correlations have an impact on the postsynaptic response, and roughly to what

degree. Then we will present a mathematical model that is somewhat abstract but that can be solved analytically and can provide some quantitative insight into the problem.

12.5 When are neurons sensitive to correlated input?

The goal of this section is to identify some of the main factors that determine the sensitivity of a postsynaptic cortical neuron to the presence of correlations in its inputs.

Synapses generate discrete events that are localized in time. Hence the basic intuition suggesting that timing is important: if action potentials from two excitatory neurons arrive simultaneously or within a short time window of each other to the same postsynaptic neuron, the two synaptic events may add up, producing a larger conductance change. Roughly, depending on the time interval between their arrivals, two presynaptic action potentials may act as two separate events of unit amplitude and duration, as one event of unit amplitude but lasting twice as long, or as one event of double amplitude and unit duration. If excitatory spikes have a tendency to arrive simultaneously more often than expected by chance, the target neuron might respond more vigorously.

This idea has been confirmed through simulation studies [13, 60]. Compared to independent spike trains, synchronous spikes may evoke stronger responses, but only up to a point, after which further synchronization actually decreases the response [13, 60]. This decrease occurs for two reasons. First, only a certain number of simultaneous excitatory synaptic events are required to trigger an action potential, so, once this number is reached, other simultaneous spikes cannot enhance the response. Second, excitatory spikes that arrive while the postsynaptic cell is in its refractory period are wasted. Thus, there is a tradeoff between two effects: on one hand, grouping excitatory spikes in time so that synaptic events summate; on the other, spreading them so that refractory effects are avoided.

This line of argument, however, has serious limitations. Refractory effects become important only when the output neuron is firing near its maximum rate, which is rarely the case. And more importantly, inhibition is not considered. Inhibition alters the scenario in three ways. (1) It may affect the sensitivity of the postsynaptic neuron to synchronous excitatory spikes. (2) Synchrony affects not only the average response of the cell but also the variability of the output spike train, and this too may depend on the level of inhibition. (3) Additional questions arise about the effects of synchrony between pairs of inhibitory input spikes and between excitatory-inhibitory pairs as well. In short, the situation gets considerably more complicated. The spectrum of possible firing modes of a neuron is often split between two extreme cases, integration and coincidence detection.

12.5.1 Coincidence detection

The classic mechanism underlying a neuron's sensitivity to temporal patterns is coincidence detection [2–50]. Neurons can certainly be sensitive to the arrival of spikes from two or more inputs within a short time window; the most notable examples are from the auditory system [5, 21]. The question is, however, whether this mechanism is commonly used throughout the cortex.

In the traditional view, coincidence detection is based on a very short membrane time constant [2–50]. However, it may be greatly enhanced by the spatial arrangement of synapses and by nonlinear processes. For instance, nearby synapses may interact strongly, forming clusters in which synaptic responses to simultaneous activation are much stronger than the sum of individual, asynchronous responses [58]. A neuron could operate with many such clusters which, if located on electrotonically distant parts of the dendritic tree, could act independently of each other. Voltage-dependent channels in the dendrites may mediate or boost such nonlinear interactions between synapses [58–64]. These nonlinearities could in principle increase the capacity for coincidence detection to the point of making the neuron selective for specific temporal sequences of input spikes, and the very idea of characterizing those inputs statistically would be questionable. However, the degree to which the cortex exploits such nonlinearities is uncertain.

The coincidence detection problem can also be posed in terms of the capacity of a network to preserve the identity of a volley of spikes fired by multiple neurons within a short time window [18, 31]. Suppose a neuron receives a volley of input spikes; what is the likelihood of evoking a response (reliability), and what will its timing be relative to the center of mass of the input volley (precision)? Theoretical studies suggest that the temporal precision of the response spikes is not limited by the membrane time constant, but rather by the up-slope of excitatory synaptic events. Thus, under the right conditions a volley of synchronized action potentials may propagate in a stable way through many layers [31]. Whether areas of the cortex actually exchange information in this way is still unclear, and other modes of information transmission are possible [88].

12.5.2 Fluctuations and integrator models

The flip side of coincidence detection is integration. Neurons may also sum or average many inputs to generate an action potential [2, 50, 73]. Earlier theoretical arguments suggested that neurons acting as integrators would not be sensitive to temporal correlations [74], or that these would only matter at high firing rates, where refractory effects become important [13, 60]. However, later results [67, 69] show that neurons may still be highly sensitive to weak correlations in their inputs even if there is no spatial segregation along the dendritic tree and no synaptic interactions beyond the expected temporal summation of postsynaptic currents.

A key quantity in this case is the balance of the neuron, which refers to the relative strength between inhibitory and excitatory inputs [67, 82, 73]. When the neuron is not balanced, excitation is on average stronger than inhibition, such that the net

synaptic current is depolarizing and the mean steady-state voltage is near or above threshold. In this case the main driving force is the drift toward steady state, and input fluctuations have a small effect on the rate of output spikes [67, 33]. On the other hand, when the neuron is balanced, both excitation and inhibition are strong, the mean input current is zero or very small, and the mean steady-state voltage remains below threshold. However, the neuron may still fire because there are large voltage fluctuations that lead to random threshold crossings. In this mode, any factor that enhances the fluctuations will produce more intense firing [67, 32].

There is a subtle but important distinction between mechanisms that may alter input fluctuations. Higher rates should be seen in a balanced neuron if fluctuations increase without affecting the mean synaptic conductances, as when only the correlations change [67]. But if stronger fluctuations are accompanied by increases in total conductance, as when both excitatory and inhibitory inputs fire more intensely, the firing rate may actually decrease [32–22]. In a complex network these effects may be hard to disentangle.

Figure 12.3 compares the responses of balanced (upper traces) and unbalanced (lower traces) model neurons [67]. These were driven by excitatory and inhibitory input spike trains similar to those illustrated in Figure 12.1. For the balanced neuron both excitatory and inhibitory synaptic conductances were strong, and the combined current they generated near threshold was zero. In contrast, for the unbalanced unit both conductances were weak, but their combined current near threshold was excitatory. The four panels correspond to different correlation patterns in the inputs. In Figure 12.3a all inputs are independent, so all cross-correlograms are flat. The voltage traces reveal a typical difference between balanced and unbalanced modes: although the output rate is approximately the same, the subthreshold voltage of the balanced neuron is noisier and its interspike intervals are more variable [67, 82]. Figure 12.3b shows what happens when the excitatory inputs fire somewhat synchronously due to common input. The firing rate of the balanced neuron always increases relative to the response to independent inputs, whereas the rate of the unbalanced neuron may show either a smaller (although still substantial) increase or a decrease [13, 60]. Another effect of synchrony is to increase the variability of the output spike trains, both for balanced and unbalanced configurations [67, 69, 78, 80]; this can be seen by comparing Figures 12.3b and 3d with Figure 12.3a. Correlations between inhibitory inputs can also produce stronger responses. When the inhibitory drive oscillates sinusoidally, as in Figure 12.3c, the balanced neuron practically doubles its firing rate compared to no oscillations; in contrast, the unbalanced does not change.

The balance of a neuron is important in determining its sensitivity to correlations, but there is another key factor [67]. There are three correlation terms: correlations between pairs of excitatory neurons, between pairs of inhibitory neurons, and between excitatory-inhibitory pairs. The first two terms increase the voltage fluctuations but the last one acts in the opposite direction, decreasing them. The total effect on the postsynaptic neuron is a function of the three terms. In Figure 12.3 d, all inputs to the model neurons are equally correlated, but the balanced model shows no change in firing rate. Thus, it is possible to have strong correlations between all

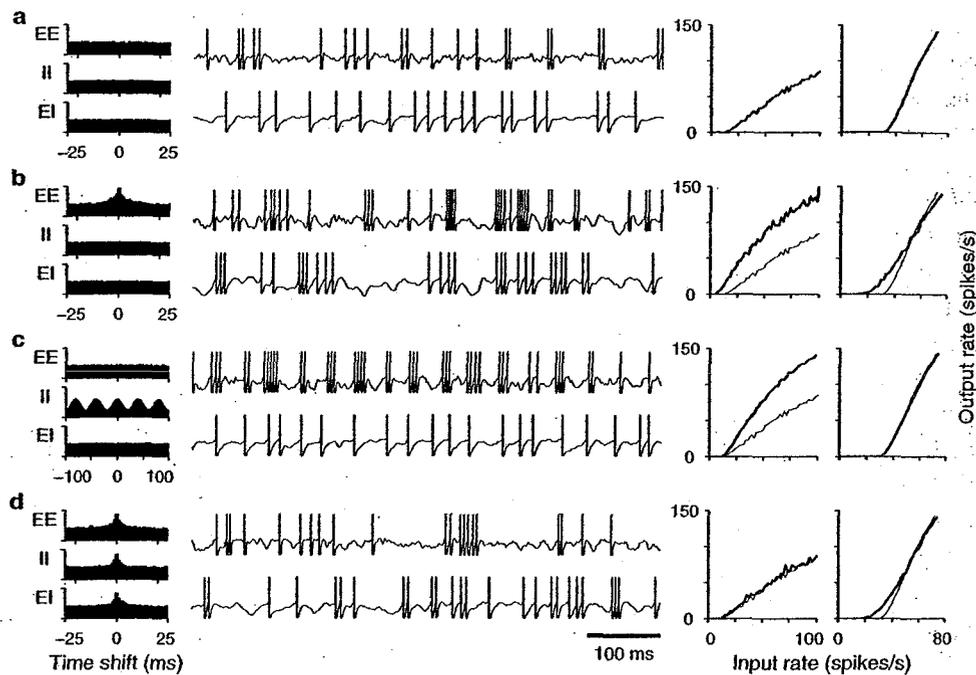


Figure 12.3

Responses of two model neurons to four input correlation patterns. Histograms on the left show average cross-correlations between pairs of excitatory input spike trains (EE), between inhibitory pairs (II), and between excitatory-inhibitory pairs (EI). Y-axes in the correlograms go from 0.7 to 1.4. Upper and lower traces in each panel show the responses of balanced and unbalanced neurons, respectively. The rate of inhibitory inputs was always equal to 1.7 times the excitatory rate. For all middle traces the excitatory input rate was 42 spikes/s. The plots on the right show the firing rates of the two (postsynaptic) model neurons versus the mean firing rate of the (presynaptic) excitatory inputs. Thin black lines are the curves obtained with independent inputs (top panel). The two output neurons were leaky integrate-and-fire units with identical parameters; they differed in the relative strength of their excitatory and inhibitory inputs. (Adapted from [67] and [68].)

inputs but still not see a change in the firing rate of the postsynaptic neuron relative to the case of independent inputs.

In summary, a balanced-neuron is much more sensitive to input correlations than an unbalanced one because correlations affect the fluctuations in synaptic drive, which cause the balanced neuron to fire. However, the postsynaptic response depends on the relative values of the three correlation terms, which may cancel out. The key point here is that even when neurons act as integrators they can, in a statistical sense, be highly sensitive to the temporal patterns of their input spikes.

Interestingly, at least in some pyramidal neurons, distal dendrites seem to act much more like coincidence detectors than proximal dendrites [93], so real neurons may,

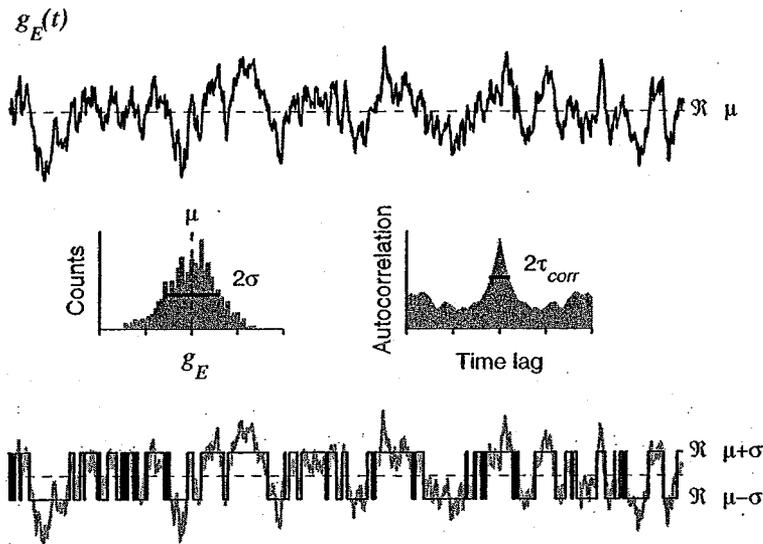


Figure 12.4

Parameterization of a continuous conductance trace. The top graph represents the total synaptic conductance generated by excitatory spikes driving a postsynaptic neuron. This conductance can be characterized statistically by its mean μ , standard deviation σ , and correlation time τ_{corr} . The left histogram is the distribution of conductance values of the top trace. The histogram on the right is its autocorrelation function. The width of the peak is parameterized by τ_{corr} . The bottom graph shows a binary variable that approximates the continuous trace. The binary function has the same mean, standard deviation and correlation time as the original function [69].

to some extent, combine both types of firing modes.

12.6 A simple, quantitative model

Now we discuss a simple model for which the responses to correlated input can be calculated analytically [69]. The first step is to describe its input.

12.6.1 Parameterizing the input

The input to a neuron consists of two sets of spike trains, ones that are excitatory and others that are inhibitory. What are the total synaptic conductances generated by these spikes? How can they be characterized? An example generated through a computer simulation is shown in Figure 12.4. The top trace represents the total excitatory conductance $g_E(t)$ produced by the constant bombardment of excitatory

synapses onto a model neuron. When plotted versus time, the time course looks noisy, random. Because $g_E(t)$ is the result of thousands of individual synaptic events, the distribution of conductance values should be approximately Gaussian, with mean and standard deviation

$$\begin{aligned}\mu &= \langle g_E \rangle \\ \sigma &= \sqrt{\langle (g_E - \mu)^2 \rangle},\end{aligned}\tag{12.3}$$

where the angle brackets $\langle \rangle$ indicate an average over time. The histogram on the left in Figure 12.4 shows the distribution of g_E values for the top trace. Indeed, it is close to a Gaussian, even though the trace is relatively short (1 s long sampled at 1 ms intervals). The mean and standard deviation, however, are not enough to characterize the conductance trace because it fluctuates with a typical timescale that has to be determined independently. The histogram on the right shows the autocorrelation function of the trace. This is akin to the cross-correlation functions discussed earlier, except that the correlation is between a continuous function and itself. Now a peak centered at zero time lag indicates that $g_E(t)$ and $g_E(t + \Delta t)$ tend to be similar to each other, and the width of the peak tells how fast this tendency decreases. A flat autocorrelation means that all values of g_E were drawn randomly and independently of each other. Thus, an autocorrelation function that is everywhere flat except for a peak centered at zero is the signature of a stochastic function that varies relatively smoothly over short timescales but whose values appear entirely independent when sampled using longer intervals. The autocorrelation function can be computed analytically for a variety of noise models, and it is typically a double exponential, as in Equation 12.2, with $C_{max} = \sigma^2$. Identical considerations apply to the conductance generated by inhibitory synapses.

From Figures 12.1 and 12.2 and from these observations, it appears that a reasonable framework to describe the total excitatory and inhibitory conductances that drive a cortical neuron is to model them using two random signals with given means, standard deviations and correlation times. Indeed, this approach has been tested experimentally, with highly positive results [29, 30]. This is also what was done to generate the spikes in Figure 12.1 (see Appendix). As explained below, this approximation is very good; for the leaky integrate-and-fire model the responses obtained using this method versus actual spike trains are virtually identical within a large parameter range (not shown).

In general, calculating the three parameters for $g_E(t)$ or $g_I(t)$ from the quantities that parameterize the corresponding input spike trains is difficult. However, this can be done under the following simplifying assumptions. Suppose there are N_E excitatory spike trains that are independent, each with Poisson statistics and a mean rate r_E . Also suppose that the synapses operate like this: whenever an input spike arrives, $g_E(t)$ increases instantaneously by an amount G_E ; otherwise, $g_E(t)$ decreases exponentially toward zero with a time constant τ_E (see ref. [24]). For this simple

scheme it can be shown [81] that

$$\begin{aligned}\mu &= G_E N_E r_E \tau_E \\ \sigma^2 &= \frac{G_E^2 N_E r_E \tau_E}{2} \\ \tau_{corr} &= \tau_E;\end{aligned}\tag{12.4}$$

with the correlation function having a double-exponential shape. Thus, for this situation, a model neuron in a simulation can be driven by two methods. First, by generating independent Poisson spikes and increasing the conductance every time one such spike arrives, exactly as described above. In this case parameters G_E , r_E , N_E , τ_E , and the corresponding quantities for inhibitory inputs need to be specified. The second method is to generate the fluctuating signals $g_E(t)$ and $g_I(t)$ directly by combining random numbers, in which case only the respective μ , σ and τ_{corr} are strictly required. Neuronal responses evoked using this type of model can match experimental data quite well [29, 30].

When the assumptions of the case just discussed are violated, for instance, when the spikes driving a neuron are not independent, determining μ , σ and τ_{corr} analytically becomes much more difficult. However, in general one should expect correlations to increase σ , and the correlation time should be equal to the synaptic time constant, although it may increase further if the input spikes are correlated over longer timescales.

Next, we ask how each of the three key parameters, μ , σ and τ_{corr} , affects the response of a postsynaptic neuron.

12.6.2 A random walk in voltage

The model neuron we consider is the non-leaky, integrate-and-fire neuron [67, 69, 40], whose dynamics resemble those of random walk models used to study diffusion in physical systems [87, 42, 41, 12]. The voltage V of this unit changes according to the input $I(t)$ that impinges on it, such that

$$\tau \frac{dV}{dt} = I(t),\tag{12.5}$$

where τ is its integration time constant. In this model an action potential is produced when V exceeds a threshold V_θ . After this, V is reset to an initial value V_{reset} and the integration process continues evolving according to the equation above. This model is related to the leaky integrate-and-fire model [24, 67, 82] but it lacks the term proportional to $-V$ in the right-hand side of the differential equation. An additional and crucial constraint is that V cannot fall below a preset value, which acts as a barrier. For convenience the barrier is set at $V=0$, so only positive values of V are allowed. This choice, however, makes no difference in the model's dynamics. Except for the barrier and the spike-generating mechanism, this model neuron acts as an ideal integrator, with an integration time constant τ .

Here the input $I(t)$ is the total current, including excitatory and inhibitory components. To simplify things even further, we will consider $I(t)$ to be a noisy function with mean μ , standard deviation σ , and correlation time τ_{corr} . Note, however, that these quantities now refer to the total current, not to the conductances, as before; this is just to simplify the notation. In this scheme it is not clear how exactly $I(t)$ is related to $g_E(t)$ and $g_I(t)$, which in principle are the measurable parameters of real neurons. However, evidently μ should depend on the means of the conductances, σ should depend on their standard deviations, and τ_{corr} should depend on their correlation times. This qualitative relationship is good enough to proceed because the model is somewhat abstract anyway.

The quantity that we are interested in is T , the time that it takes for V to go from reset to threshold. T is known as the first passage time or the interspike interval; what we want to know are its statistics. The key for this [69] is to rewrite Equation 12.5 as follows

$$\tau \frac{dV}{dt} = \mu + \sigma Z(t), \quad (12.6)$$

where $Z(t)$ is a binary variable that can only be either +1 or -1 and whose correlation function is a double exponential with correlation time τ_{corr} . Thus $I(t)$ has been replaced by a stochastic binary function that indicates whether $I(t)$ is above or below its average. This approximation is illustrated in Figure 12.4 (bottom trace). The binary function has the same mean, standard deviation and correlation time as $I(t)$. This substitution allows us to solve Equation 12.6 analytically [69]. Notice also that the neuron's time constant τ simply acts as a scale factor on the input. Hereafter it will be considered equal to 1.

Figure 12.5 shows examples of spike trains produced by the model when driven by the binary, temporally correlated input. In this figure μ was negative, so on average the voltage tended to drift away from threshold, toward the barrier. In this case the spikes are triggered exclusively by the random fluctuations, as measured by σ ; without them the neuron would never reach threshold. In Figures 12.5a–12.5c the correlation time is $\tau_{corr}=1$ ms. For a binary variable like Z , which switches between +1 and -1, the correlation time corresponds to the average time one needs to wait to observe a change in sign. In other words, the correlation time is equal to half the average time between sign changes. Thus, the input in Figure 12.5a (lower trace) flips state approximately every 2 ms. Figure 12.5c shows that, under these conditions, the neuron fires at a relatively low rate and irregularly; the times between spikes or interspike intervals are quite variable, which can also be seen from the interspike-interval distribution in Figure 12.5b.

When τ_{corr} is increased to 5 ms, as in Figures 12.5d–12.5f, the changes in input state occur approximately every 10 ms (Figure 12.5d, lower trace). This produces a large increase in mean firing rate and, to a lesser extent, an increase in variability. This can be seen by comparing the spike trains from Figures 12.5c and 12.5f. The respective mean rates are 10 and 37 spikes/s. Notice that there is a short time interval that appears very frequently. The short interval results when the input stays positive for a relatively long time, as is the case with the pair of spikes in Figure 12.5d. This interval is equal to $(V_\theta - V_{reset})/(\mu + \sigma)$, which is the minimum separation between

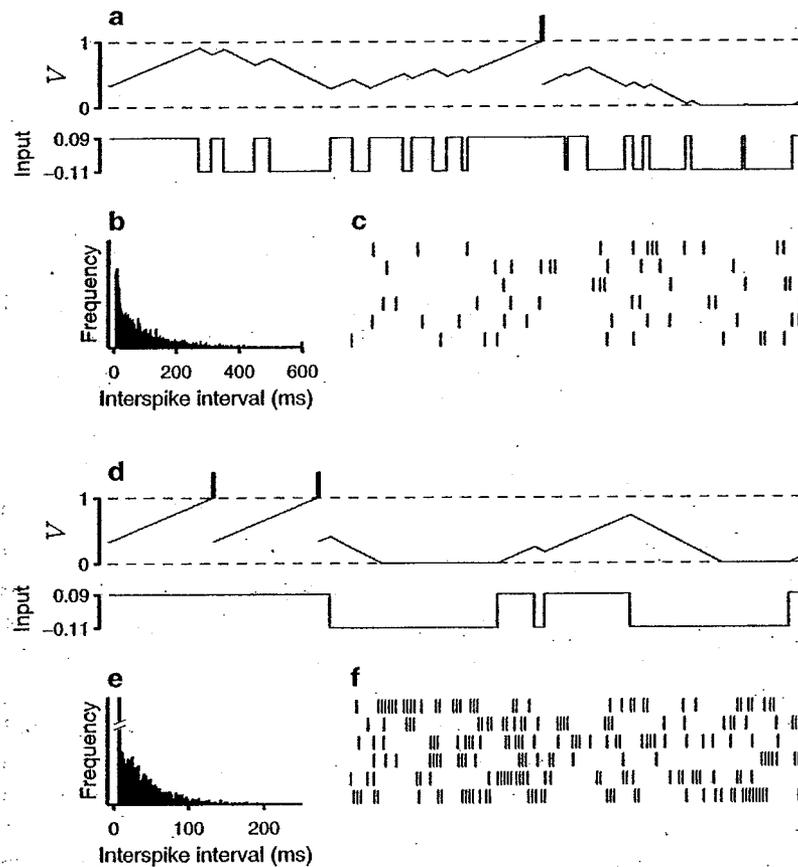


Figure 12.5

Responses of the nonleaky integrate-and-fire model driven by correlated, binary noise. The input switches states randomly, but on average the same state is maintained for $2\tau_{corr}$ ms. Sample voltage and input time courses are 50 ms long. Raster plots show 6 seconds of continuous simulation time. For the three top panels the correlation time τ_{corr} was 1 ms; for the lower panels it was 5 ms. (Adapted from [69].)

spikes in the model given μ and σ . The number of spikes separated by this interval grows as the correlation time increases. At the same time, however, longer correlation times also give rise to long interspike intervals, which occur because the input can stay in the low state for longer stretches of time. This is why correlation time increases variability: it produces both short and long interspike intervals. The quantity that is most often used to measure the regularity of a spike train is the coefficient of variation, or CV_{ISI} , which is equal to the standard deviation of the interspike intervals divided by their mean. The CV_{ISI} in Figure 12.5c is equal to 1, as for a Poisson process; in Figure 12.5f it is equal to 1.18, which reflects the higher variability. Note that μ and σ are the same for all panels. This demonstrates that the input correlation

time may have a very strong impact on the response of a postsynaptic neuron [69]. This is an interesting observation because little is known about the dynamic role of this parameter.

12.6.3 Quantitative relationships between input and output

The solution to the non-leaky model of Equation 12.6 consists in the moments of T , $\langle T \rangle$, $\langle T^2 \rangle$ and so forth. For each of these moments there are three sets of analytic expressions, because details of the solutions depend on the relative values of μ and σ . Here we only discuss the expressions for the average interspike interval $\langle T \rangle$, which is the inverse of the mean firing rate, but $\langle T^2 \rangle$ and therefore the CV_{ISI} can also be obtained in closed form [69].

When $\mu > \sigma$,

$$\langle T \rangle = \frac{V_\theta - V_{reset}}{\mu}. \quad (12.7)$$

In this case there is a strong positive drift toward threshold. Even when Z is equal to -1 the total input is positive; in other words, the voltage gets closer to threshold in every time step, whether the fluctuating component is positive or negative. The mean firing rate behaves as if the input were constant and there were no fluctuations. This can be seen in Figure 12.6, which plots the mean firing rate and the CV_{ISI} of the model neuron as a function of σ for various combinations of the other two input parameters. The values of μ are indicated in each column, and the three curves in each plot correspond to τ_{corr} equal to 1, 3 and 10 ms, with higher correlation values always producing stronger responses and higher variability. Continuous lines and dots correspond to analytic solutions and simulation results, respectively. Notice how, when $\mu=0.02$, the firing rate stays constant for σ below 0.02, although the variability increases most sharply precisely within this range.

When $\mu=0$,

$$\langle T \rangle = \frac{2(V_\theta - V_{reset})}{\sigma} + \frac{V_\theta^2 - V_{reset}^2}{2\tau_{corr}\sigma^2}. \quad (12.8)$$

Clearly, the average interspike interval decreases with both τ_{corr} and σ . In this case there is no drift, no net displacement; the voltage advances toward threshold when $Z=+1$ and retreats toward the barrier when $Z=-1$. Under these conditions the neuron is driven exclusively by fluctuations. The middle column of Figure 12.6 corresponds to this regime. As can be seen, the variability of the neuron also increases monotonically with σ and τ_{corr} .

Finally, when $\mu \leq \sigma$,

$$\langle T \rangle = \frac{V_\theta - V_{reset}}{\mu} + \tau_{corr}(c-1)^2(\exp(-\alpha V_\theta) - \exp(-\alpha V_{reset})), \quad (12.9)$$

where we have defined

$$c \equiv \frac{\sigma}{\mu} \\ \alpha \equiv \frac{1}{\mu\tau_{corr}(c^2-1)}. \quad (12.10)$$

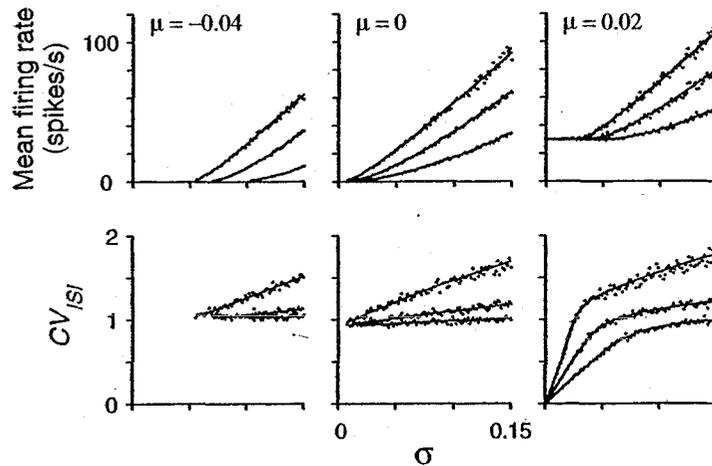


Figure 12.6

Mean firing rate and coefficient of variation for the nonleaky integrate-and-fire neuron driven by correlated binary noise. Continuous lines are analytic expressions and dots are results from computer simulations. Each simulation data point was based on spike trains containing 2000 spikes. The three curves in each graph are for different values of τ_{corr} : 1 ms (lower curves), 3 ms (middle curves), and 10 ms (upper curves). (Adapted from [69].)

As with the above equations, Figure 12.6 reveals the excellent agreement between this expression and computer simulations. An interesting special case is obtained when $\sigma = \mu$, or $c = 1$. Then the total input is zero every time that Z equals -1 , so half the time V does not change and half the time V increases by 2μ in each time step. Therefore, the average time to threshold should be equal to $(V_\theta - V_{reset})/\mu$, which is precisely the result from Equation 12.9. This quantity does not depend on the correlation time, but the CV_{ISI} does. The analytic expression for the CV_{ISI} is particularly simple in this case:

$$\sqrt{\frac{2\mu\tau_{corr}}{V_\theta - V_{reset}}} \quad (12.11)$$

Thus, the variability of the output spike train diverges as τ_{corr} increases, but the mean rate does not.

This last observation is valid in a more general sense, and is an important result regarding the effects of correlations. In most cases, the limit behaviors of the firing rate and the CV_{ISI} as the correlation time increases are quite different: the rate tends to saturate, whereas the variability typically diverges. This is illustrated in Figure 12.7. The one condition in which the variability saturates as the correlation time tends to infinity is when μ is larger than σ (thickest line on right column). The asymptotic value of the CV_{ISI} in this case is $c/\sqrt{1-c^2}$. In this parameter regime the drift is strong, so it usually produces high firing rates as well.

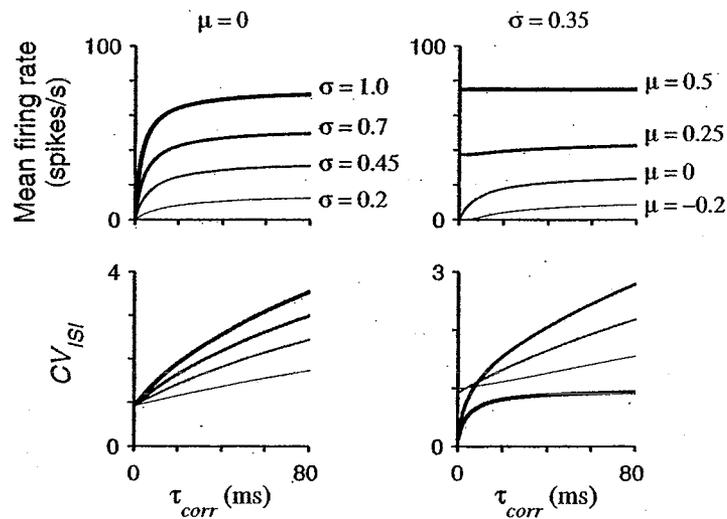


Figure 12.7

Responses of the nonleaky integrate-and-fire neuron as functions of input correlation time τ_{corr} . Only analytic results are shown. As the correlation time increases, the firing rate always tends to an asymptotic value. In contrast, the CV_{ISI} diverges always, except when $\mu > \sigma$; this case corresponds to the thickest line in the plots on the right. (Adapted from [69].)

The key to obtain all the analytic expressions was the use of a binary input. One may wonder, however, whether the results are valid with a more realistic input signal. It turns out that, for this model, the mean firing rate and the CV_{ISI} obtained using correlated Gaussian noise are very similar to those obtained with binary noise. This is not entirely unexpected, first, because the neuron essentially adds its inputs, and second, because Gaussian noise can be properly approximated as the sum of multiple binary random samples, as a consequence of the central limit theorem. This is strictly true when all binary and Gaussian samples are independent, that is, when the autocorrelation functions are everywhere flat, but the approximation works quite well even when there is a correlation time. For example, the rate and the CV_{ISI} still increase as functions of correlation time, and the same asymptotic behaviors are seen [69].

12.7 Correlations and neuronal variability

The spike trains of neurons recorded in awake animals are highly variable [25, 75–78]. However, spike generation mechanisms themselves seem to be highly reliable [20, 49, 56]. The contrast between these two observations stirred a fair amount of

discussion, especially after the work of Softky and Koch [76], who pointed out that although the CV_{ISI} of typical cortical neurons is close to 1, this number should be much lower for an integrator that adds up many small contributions in order to fire, especially at high output rates. However, their arguments applied in the absence of inhibition, and later work [82, 73] showed that including incoming inhibitory spikes produces higher CV_{ISI} values even in integrator models without any built-in coincidence detection mechanisms [2–50] or similar nonlinearities [58–64], a result that is consistent with early stochastic models [41, 84]. So-called ‘balanced’ models, in which inhibition is relatively strong, typically bring the CV_{ISI} to the range between 0.5 and 1 [73, 82], which is still lower than reported from recorded data [25, 75–78]. Other intrinsic factors have also been identified as important in determining spike train variability; for instance, combining the proper types of conductances [11], tuning the cellular parameters determining membrane excitability [47, 82], and bistability [92].

However, several lines of evidence point to correlations in the conductances (or currents) that drive a neuron as a primary source of variability. First, correlated firing is ubiquitous. This has been verified through a variety of techniques, including *in vivo* experiments in which pairs of neurons are recorded simultaneously. The widths of the corresponding cross-correlograms may go from a few to several hundred milliseconds [43–77], so they may be much longer than the timescales of common AMPA and GABA-A synapses [27]. Second, *in vitro* experiments in which neurons are driven by injected electrical current suggest that input correlations are necessary to reproduce the firing statistics observed *in vivo* [28, 29, 30, 78]. This is in line with the suggestion that fluctuations in eye position are responsible for a large fraction of the variability observed in primary visual neurons, because they provide a common, correlating signal [46]. Third, this also agrees with theoretical studies [33, 67]; in particular with results for the non-leaky integrate-and-fire model showing that the CV_{ISI} depends strongly on the correlation time of the input [69]. In addition, similar analyses applied to the traditional leaky integrate-and-fire model reveal the same qualitative dependencies [69]. This, in fact, can be seen in Figure 12.1, where the model with leak was used: increases in the synaptic time constants give rise to longer correlation times and to higher CV_{ISI} values (compare Figures 12.1a and 12.1c), an effect that has nothing to do with the synchronization between output spike trains. Finally, high variability is also observed in simulation studies in which network interactions produce synchronized recurrent input [85–89], as in Figure 12.2.

12.8 Conclusion

The activity of a local cortical microcircuit can be analyzed in terms of at least two dimensions, its intensity, which is typically measured by the mean firing rates of

the neurons, and its coherence across neurons, which is often described in terms of synchrony or cross-correlations between pairs of units. These correlations serve as probes for the organization and dynamics of neural networks. There is strong evidence, both theoretical and experimental, indicating that correlations may be important dynamic components of cortical microcircuits. Here we have discussed two general hypotheses, the encoding of stimulus features and the gating of information from one structure to another. Although quite different, both are based on the premise that correlations have a specific functional role. Interestingly, there is an interpretation that is entirely opposite to the sensory-coding hypothesis, which suggests that correlations between cortical neurons limit the accuracy with which neural populations may encode stimulus features [95] (see also [1]). These top-down ideas have generated considerable debate, but the crucial question remains unresolved as to whether correlations have a specific, separate functional role, or whether they simply participate in all functions, just as firing rates do. It is also conceivable that there is no generic strategy, and that the meaning and impact of correlations vary from one local microcircuit to another.

A second, critical question is whether correlations can be controlled independently of firing rates. That is, a group of neurons M may affect another group A in two, not necessarily exclusive, ways: by changing the firing rates of A or the correlations between local neurons in A . There are two knobs that can be turned, and the question is whether these can be turned independently of each other. Here we reviewed some studies that begin to address this issue by taking the point of view of a single neuron: what intrinsic properties make it sensitive to correlations? How do correlations affect its response? Can changes in input correlations and input firing rates be distinguished? The hope is that this bottom-up perspective will eventually help clarify the top-down ideas by identifying and constraining the role of correlations in local circuit dynamics. A good example of this is the above section on neuronal variability. The highly variable discharge of cortical neurons is observed and characterized in recordings from awake, behaving preparations; and experiments *in vitro*, as well as computational and theoretical studies, identify a variety of biophysical mechanisms responsible for the observation. In this particular case, input correlations seem to play a major role because they can generate highly variable output spike trains in the absence of any additional intrinsic mechanisms [28, 46, 69, 78].

In conclusion, the two questions just pondered may represent high- and low-level interpretations of the same phenomenon, but a conceptual framework providing a unified view of this problem is still lacking. Establishing such framework, however, may serve as a guidelight for future investigations.

12.9 Appendix

Here we describe the leaky integrate-and-fire model [24, 67, 82, 84] driven by conductance changes that was used to generate Figure 12.1. In this model, the membrane potential V evolves according to

$$\tau_m \frac{dV}{dt} = -V - g_E(t)(V - V_E) - g_I(t)(V - V_I), \quad (12.12)$$

where the resting potential has been set to 0 mV. The spike-generating currents are substituted by a simple rule: whenever V exceeds a threshold (20 mV), a spike is emitted and V is clamped to a reset value (10 mV) for a refractory period (1.8 ms). After that, V continues evolving according to the above equation. The excitatory and inhibitory conductances, $g_E(t)$ and $g_I(t)$, were generated by combining Gaussian random numbers [69], so that the resulting traces would have the desired mean, standard deviation and correlation time. These parameters were related to input rates and model synaptic conductances through Equations 12.4. For Figures 12.1a and 1d, $N_{E/E} = 27.5$ spikes/ms, $G_E = 0.02$, $\tau_E = 2$ ms, $N_{I/I} = 12.15$ spikes/ms, $G_I = 0.06$, and $\tau_I = 2$ ms, with G_E and G_I in units of the leak conductance (i.e., where the leak conductance equals 1). For Figures 12.1b and 12.1e, $\tau_E = 20$ ms. For Figures 12.1c and 12.1f, $\tau_I = \tau_E = 20$ ms. Correlations between the conductances of different neurons were generated by drawing correlated Gaussian samples during generation of the $g_E(t)$ and $g_I(t)$ traces for different neurons. The correlation coefficient for a pair of conductances, Equation 12.1, is equal to the correlation coefficient between the corresponding Gaussian samples. Other parameters were: $\tau_m = 20$ ms, $V_E = 74$ mV, $V_I = -10$ mV, $\Delta t = 0.1$ ms.

Acknowledgments Research was supported by the Howard Hughes Medical Institute and by startup funds from Wake Forest University School of Medicine. Parts of this chapter were adapted from [68].

References

- [1] Abbott L.F., and Dayan P. (1999) The effect of correlated activity on the accuracy of a population code. *Neural Comput.* 11: 91–101.
- [2] Abeles M. (1982) Role of the cortical neuron: Integrator or coincidence detector? *Israel J Med Sci* 18: 83–92.
- [3] Aertsen A., and Arndt M. (1993) Response synchronization in the visual cortex. *Curr. Opin. Neurobiol.* 3: 586–594.

- [4] Aertsen A.M.H.J., Gerstein G.L., Habib M.K., and Palm G. (1989) Dynamics of neuronal firing correlation: modulation of "effective connectivity". *J Neurophysiol* **61**: 900–917.
- [5] Agmon-Snir H., Carr C.E., and Rinzel J. (1998) The role of dendrites in auditory coincidence detection *Science* **393**: 268–272.
- [6] Ahissar E., and Arieli A. (2001) Figuring space by time. *Neuron* **32**: 185–201.
- [7] Ahissar E., Sosnik R., and Haidarliu S. (2000) Transformation from temporal to rate coding in a somatosensory thalamocortical pathway. *Nature* **406**: 302–306.
- [8] Barlow J.S. (1993) *The Electroencephalogram: Its Patterns and Origins*. Cambridge, MA: MIT Press.
- [9] Bazhenov M., Stopfer M., Rabinovich M., Abarbanel H.D.I., Sejnowski T.J., and Laurent G. (2001) Model of cellular and network mechanisms for odor-evoked temporal patterning in the locust antennal lobe. *Neuron* **30**: 569–581.
- [10] Bazhenov M., Stopfer M., Rabinovich M., Huerta R., Abarbanel H.D.I., Sejnowski T.J., and Laurent G. (2001) Model of transient oscillatory synchronization in the locust antennal lobe. *Neuron* **30**: 553–567.
- [11] Bell A.J., Mainen Z.F., Tsodyks M., and Sejnowski T.J. (1995) 'Balancing' of conductances may explain irregular cortical spiking. Technical Report INC-9502, Institute for Neural Computation, UCSD, San Diego, CA, 92093–0523.
- [12] Berg H.C. (1993) *Random Walks in Biology*. Princeton, NJ: Princeton University Press.
- [13] Bernander Ö., Koch C., and Usher M. (1994) The effects of synchronized inputs at the single neuron level. *Neural Comput.* **6**: 622–641.
- [14] Braitenberg V., and Schüz A. (1997) *Cortex: Statistics and Geometry of Neuronal Connectivity*. Berlin: Springer-Verlag.
- [15] Brody C.D. (1999) Correlations without synchrony. *Neural Comput.* **11**: 1537–1551.
- [16] Brosch M., and Schreiner C.E. (1999) Correlations between neural discharges are related to receptive field properties in cat primary auditory cortex. *Eur. J. Neurosci.* **11**: 3517–3530.
- [17] Brunel N., and Hakim V. (1999) Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Comput.* **11**: 1621–1671.
- [18] Burkitt A.N., and Clark G.M. (1999) Analysis of integrate-and-fire neurons: synchronization of synaptic input and spike output. *Neural Comput.* **11**: 871–901.
- [19] Bush P., and Sejnowski T.J. (1996) Inhibition synchronizes sparsely connected cortical neurons within and between columns in realistic network models. *J.*

Comput. Neurosci. **3**: 91–110.

- [20] Calvin W.H., and Stevens C.F. (1968) Synaptic noise and other sources of randomness in motoneuron interspike intervals. *J. Neurophysiol.* **31**: 574–587.
- [21] Carr C.E., and Friedman M.A. (1999) Evolution of time coding systems. *Neural Comput.* **11**: 1–20.
- [22] Chance F.S., Abbott L.F., and Reyes A.D. (2002) Gain modulation from background synaptic input. *Neuron* **35**: 773–782.
- [23] Dan Y., Alonso J.M., Usrey W.M., and Reid R.C. (1998). Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nature Neurosci.* **1**: 501–507.
- [24] Dayan P., and Abbott L.F. (2001) *Theoretical Neuroscience*. Cambridge, MA: MIT Press.
- [25] Dean A. (1981) The variability of discharge of simple cells in the cat striate cortex. *Exp. Brain Res.* **44**: 437–440.
- [26] DeCharms R.C., and Merzenich M.M. (1995) Primary cortical representation of sounds by the coordination of action potential timing. *Nature* **381**: 610–613.
- [27] Destexhe A., Mainen Z.F., and Sejnowski T.J. (1998) Kinetic models of synaptic transmission. In: *Methods in Neuronal Modeling* (second edition), C. Koch, I. Segev, eds., pp. 1–25, Cambridge, MA: MIT Press.
- [28] Destexhe A., and Paré D. (1999) Impact of network activity on the integrative properties of neocortical pyramidal neurons *in vivo*. *J. Neurophysiol.* **81**: 1531–1547.
- [29] Destexhe A., and Paré D. (2000) A combined computational and intracellular study of correlated synaptic bombardment in neocortical pyramidal neurons *in vivo*. *Neurocomputing* **32-33**: 113–119.
- [30] Destexhe A., Rudolph M., Fellous J.M., and Sejnowski T.J. (2001) Fluctuating synaptic conductances recreate *in vivo*-like activity in neocortical neurons. *Neuroscience* **107**: 13–24.
- [31] Diesmann M., Gewaltig M.-O., and Aertsen A. (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* **402**: 529–533.
- [32] Doiron B., Longtin A., Berman N., and Maler L. (2001) Subtractive and divisive inhibition: effect of voltage-dependent inhibitory conductances and noise. *Neural Comput.* **13**: 227–248.
- [33] Feng J., and Brown D. (2000) Impact of correlated inputs on the output of the integrate-and-fire model. *Neural Comput.* **12**: 671–692.
- [34] Fries P., Neuenschwander S., Engel A.K., Goebel R., and Singer W. (2001) Rapid feature selective neuronal synchronization through correlated latency

- shifting. *Nature Neurosci.* **4**: 194–200.
- [35] Fries P., Reynolds J.H., Rorie A.E., and Desimone R. (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* **291**: 1560–1563.
- [36] Fries P., Roelfsema P.R., Engel A.K., König P., and Singer W. (1997) Synchronization of oscillatory responses in visual cortex correlates with perception in interocular rivalry. *Proc. Natl. Acad. Sci. U.S.A.* **94**: 12699–12704.
- [37] Fries P., Schröder J.-H., Roelfsema P.R., and Singer W., Engel AK (2002) Oscillatory neural synchronization in primary visual cortex as a correlate of stimulus selection. *J. Neurosci.* **22**: 3739–3754.
- [38] Frost Jr. J.D. (1967) An averaging technique for detection of EEG-intracellular potential relationships. *Electroenceph. Clin. Neurophysiol.* **23**: 179–181.
- [39] Fuentes U., Ritz R., Gerstner W., and Van Hemmen J.L. (1996) Vertical signal flow and oscillations in a three-layer model of the cortex. *J. Comput. Neurosci.* **3**: 125–136.
- [40] Fusi S., and Mattia M. (1999) Collective behavior of networks with linear (VLSI) Integrate and Fire Neurons *Neural Comput.* **11**: 633–652.
- [41] Gerstein G.L., and Mandelbrot B. (1964) Random walk models for the spike activity of a single neuron. *Biophys. J.* **4**: 41–68.
- [42] Gardiner C.W. (1985) *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*. Berlin: Springer-Verlag.
- [43] Gochin P.M., Miller E.K., Gross C.G., and Gerstein G.L. (1991) Functional interactions among neurons in inferior temporal cortex of the awake macaque. *Exp. Brain Res.* **84**: 505–516.
- [44] Gray C.M. (1999) The temporal correlation hypothesis of visual feature integration: still alive and well. *Neuron* **24**: 31–47.
- [45] Gray C.M., and McCormick D.A. (1996) Chattering cells: superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. *Science* **274**: 109–113.
- [46] Gur M., Beylin A., and Snodderly D.M. (1997) Response variability of neurons in primary visual cortex (V1) of alert monkeys. *J. Neurosci.* **17**: 2914–2920.
- [47] Gutkin B.S., and Ermentrout G.B. (1998) dynamics of membrane excitability determine interspike interval variability: a link between spike generation mechanisms and cortical spike train statistics. *Neural Comput.* **10**: 1047–1065.
- [48] Hodgkin A.L., and Huxley A.F. (1952) A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* **117**: 500–544.
- [49] Holt G.R., Softky W.R., Koch C., and Douglas R.J. (1996) Comparison of

- discharge variability *in vitro* and *in vivo* in cat visual cortex neurons. *J. Neurophysiol.* **75**: 1806–1814.
- [50] König P., Engel A.K., and Singer W. (1996) Integrator or coincidence detector? The role of the cortical neuron revisited. *Trends Neurosci.* **19**: 130–137.
- [51] Kreiter A.K., and Singer W. (1996) Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *J. Neurosci.* **16**: 2381–2396.
- [52] Leopold D.A., and Logothetis N.K. (1999) Multistable phenomena: changing views in perception. *Trends Cogn. Sci.* **3**: 254–264.
- [53] Llinás R.R. (1988) The intrinsic electrophysiological properties of mammalian neurons: insights into central nervous system function. *Science* **242**: 1654–1664.
- [54] Lüti A., and McCormick D.A. (1998) H-current: properties of a neuronal and network pacemaker. *Neuron* **21**: 9–12.
- [55] MacLeod K., Bäcker A., and Laurent G. (1998) Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature* **395**: 693–698.
- [56] Mainen Z.F., and Sejnowski T.J. (1995) Reliability of spike timing in neocortical neurons. *Science* **268**: 1503–1506.
- [57] Mehta M.R., Lee A.K., and Wilson M.A. (2002) Role of experience and oscillations in transforming a rate code into a temporal code *Nature* **417**: 741–746.
- [58] Mel B.W. (1993) Synaptic integration in an excitable dendritic tree. *J. Neurophysiol.* **70**: 1086–1101.
- [59] Mel B.W. (1999) Why have dendrites? A computational perspective. In: *Dendrites*, Stuart G, Spruston N, Häusser M, eds., pp. 271–289. Oxford: Oxford University Press.
- [60] Murthy V.N., and Fetz E.E. (1994) Effects of input synchrony on the firing rate of a three-conductance cortical neuron model. *Neural Comput.* **6**: 1111–1126.
- [61] Nelson J.I., Salin P.A., Munk M.H.-J., Arzi M., and Bullier J. (1992) Spatial and temporal coherence in cortico-cortical connections: a cross-correlation study in areas 17 and 18 in the cat. *Vis. Neurosci.* **9**: 21–37.
- [62] Perez-Orive J., Mazor O., Turner G.C., Cassenaer S., Wilson R.I., and Laurent G. (2002) Oscillations and sparsening of odor representations in the mushroom body. *Science* **297**: 359–65.
- [63] Perkel D.H., Gerstein G.L., and Moore G.P. (1967) Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophys. J.* **7**: 419–440.
- [64] Poirazi P., and Mel B.W. (2001) Impact of active dendrites and structural plas-

- ticity on the memory capacity of neural tissue. *Neuron* **29**: 779–796.
- [65] Reynolds J.H., Pasternak T., and Desimone R. (2000) Attention increases sensitivity of V4 neurons. *Neuron* **26**: 703–714.
- [66] Riehle A., Grün S., Diesmann M., and Aertsen A. (1997) Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* **278**: 1950–1953.
- [67] Salinas E., and Sejnowski T.J. (2000) Impact of correlated synaptic input on output firing rate and variability in simple neuronal models. *J. Neurosci.* **20**: 6193–6209.
- [68] Salinas E., and Sejnowski T.J. (2001) Correlated neuronal activity and the flow of neural information. *Nat. Rev. Neurosci.* **2**: 539–550.
- [69] Salinas E., and Sejnowski T.J. (2002) Integrate-and-fire neurons driven by correlated stochastic input. *Neural Comput.* **14**: 2111–2155.
- [70] Seidemann E., Zohary U., and Newsome W.T. (1998) Temporal gating of neural signals during performance of a visual discrimination task. *Nature* **394**: 72–75.
- [71] Singer W., and Gray C.M. (1995) Visual feature integration and the temporal correlation hypothesis. *Annu. Rev. Neurosci.* **18**: 555–586.
- [72] Shadlen M.N., and Movshon J.A. (1999) Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* **24**: 67–77.
- [73] Shadlen M.N., and Newsome W.T. (1994) Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* **4**: 569–579.
- [74] Shadlen M.N., and Newsome W.T. (1998) The variable discharge of cortical neurons: implications for connectivity, computation and information coding. *J. Neurosci.* **18**: 3870–3896.
- [75] Shinomoto S., Sakai Y., and Funahashi S. (1999) The Ornstein-Uhlenbeck process does not reproduce spiking statistics of neurons in prefrontal cortex. *Neural Comput.* **11**: 935–951.
- [76] Softky W.R., and Koch C. (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.* **13**: 334–350.
- [77] Steinmetz P.N., Roy A., Fitzgerald P.J., Hsiao S.S., Johnson K.O., and Niebur E. (2000) Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature* **404**: 187–190.
- [78] Stevens C.F., and Zador A.M. (1998) Input synchrony and the irregular firing of cortical neurons. *Nature Neurosci.* **1**: 210–217.
- [79] Stopfer M., Bhagavan S., Smith B.H., and Laurent G. (1997) Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nat-*

ture **390**: 70–74.

- [80] Svirskis G., and Rinzel J. (2000) Influence of temporal correlation of synaptic input on the rate and variability of firing in neurons. *Biophys. J.* **79**: 629–637.
- [81] Tiesinga P.H.E., José J.V., and Sejnowski T.J. (2000) Comparison of current-driven and conductance-driven neocortical model neuron with Hodgkin-Huxley voltage-gated channels. *Phys. Rev. E* **62**: 8413–8419.
- [82] Troyer T.W., and Miller K.D. (1997) Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Comput.* **9**: 971–983.
- [83] Tsodyks M.V., and Sejnowski T.J. (1995) Rapid state switching in balanced cortical network models. *Network* **6**: 111–124.
- [84] Tuckwell H.C. (1988) *Introduction to Theoretical Neurobiology*, Volumes 1 and 2. New York: Cambridge University Press.
- [85] Usher M., Stemmler M., Koch C., and Olami Z. (1994) Network amplification of local fluctuations causes high spike rate variability, fractal firing patterns and oscillatory local field potentials. *Neural Comput.* **6**: 795–836.
- [86] Usrey W.M., and Reid R.C. (1999) Synchronous activity in the nervous system. *Annu. Rev. Neurosci.* **61**: 435–456.
- [87] Van Kampen N.G. (1981) *Stochastic Processes in Physics and Chemistry*. Amsterdam: North Holland.
- [88] van Rossum M.C.W., Turrigiano G.G., and Nelson S.B. (2002) Fast propagation of firing rates through layered networks of noisy neurons. *J. Neurosci.* **22**: 1956–1966.
- [89] Van Vreeswijk C., and Sompolinsky H. (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**: 1724–1726.
- [90] Wehr M., and Laurent G. (1999) Relationship between afferent and central temporal patterns in the locust olfactory system. *J. Neurosci.* **19**: 381–390.
- [91] White E.L. (1989) *Cortical Circuits*. Boston: Birkhäuser.
- [92] Wilbur W.J., and Rinzel J. (1983) A theoretical basis for large coefficient of variation and bimodality in neuronal interspike interval distribution. *J. Theo. Biol.* **105**: 345–368.
- [93] Williams S.R., and Stuart G.J. (2002) Dependence of EPSP efficacy on synapse location in neocortical pyramidal neurons. *Science* **295**: 1907–1910.
- [94] Wilson M., and Bower J.M. (1992) Cortical oscillations and network interactions in a computer simulation of piriform cortex. *J. Neurophysiol.* **67**: 981–995.

- [95] Zohary E., Shadlen M.N., and Newsome W.T. (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* **370**: 140–143.