Dynamics of Neuronal Firing Correlation: Modulation of "Effective Connectivity"

A. M. H. J. AERTSEN, G. L. GERSTEIN, M. K. HABIB, AND G. PALM

(With the Collaboration of P. Gochin and J. Krüger)

Max-Planck-Institute for Biological Cybernetics, D-7400 Tübingen, Federal Republic of Germany; Department of Physiology, University of Pennsylvania, Philadelphia, Pennsylvania 19104-6085; Department of Bio-Statistics, University of North Carolina, Chapel Hill, North Carolina 27514; Vogt-Institute for Brain Research, University of Düsseldorf, D-4000 Düsseldorf, Federal Republic of Germany; Neurlogical Clinic, University of Freiburg, D-7800 Freiburg, Federal Republic of Germany.

SUMMARY AND CONCLUSIONS

1. We reexamine the possibilities for analyzing and interpreting the time course of correlation in spike trains simultaneously and separably recorded from two neurons.

2. We develop procedures to quantify and properly normalize the classical joint peristimulus time scatter diagram. These allow separation of the "raw" correlation into components caused by direct stimulus modulations of the single-neuron firing rates and those caused by various types of interaction between the two neurons.

3. A newly developed significance test ("surprise") is applied to evaluate such inferences.

4. Application of the new procedures to simulated spike trains allowed the recovery of the known circuitry. In particular, it proved possible to recover fast stimulus-locked modulations of "effective connectivity," even if they were masked by strong direct stimulus modulations of individual firing rates. These procedures thus present a clearly superior alternative to the commonly used "shift predictor."

5. Adopting a model-based approach, we generalize the classical measures for quantifying a direct interneuronal connection ("efficacy" and "contribution") to include possible stimuluslocked time variations.

6. Application of the new procedures to real spike trains from several different preparations showed that fast stimulus-locked modulations of "effective connectivity" also occur for real neurons.

INTRODUCTION

When examining the relation between activities of two neurons in a multi-neuron recording, a most useful tool has been the cross-correlation of the two spike trains (29; for reviews, see Refs. 12 and 20). This calculation measures the probability of firing of the "target" neuron at various times relative to the firing of the "reference" neuron, in which the probability is determined by averaging over many occurrences of the two spikes.

This procedure detects the (delayed) coincidence of neural firing among different neurons, which can be regarded in two ways. The first way is to regard such coincidences as a sign of a possible neural code being used by the working brain: a code based on the cooperative action of two or more impulse-carrying pathways (30). A second way is to use the near-coincidence to infer functional connectivity among the underlying elements. Such knowledge can be extracted by the experimenter, but it is not obvious how

that should be available to the brain under investigation. Results of this sort of analysis are traditionally interpreted as indicative of "effective connectivity" among the observed neurons. Appropriate control calculations are made to allow the separation of direct stimulus effects from other contributions to the "raw" correlation. The usual strategy has been to subtract the direct stimulus effects [estimated by the so-called "shift predictor" (10, 11, 29)] from the "raw" cross-correlogram, yielding a time-averaged "residual" correlation due to neural origin. Sources of the latter are further subdivided into "direct connections" and "shared input"; this subdivision is based on inspection of the correlogram shape. A review of topics related to crosscorrelation of spike trains together with many references to the original literature can be found in Glaser and Ruchkin (16).

Many of the concepts and measurements used with spike trains have been developed and calibrated with the use of pulse trains from simulated neuronal circuits. In that case, the measurements for a given circuit can be predicted: this is the "forward problem." We are fully aware that there is no unique solution to the "inverse problem," i.e., it is impossible to determine uniquely the underlying circuit by spike-train analysis (or any other method that avoids exhaustive enumeration of all states of all elements). Our goal is to define the minimum, simple neuronal model that would replicate the experimentally observed features of measurements made on simultaneously recorded spike trains. Thus, when we make the jump from observed, coincident spike events to a statement of "effective connectivity" between two neurons, this should be taken as an abbreviated description of an equivalent class of neuronal circuits.

Most applications of cross-correlation have assumed that the system is stationary at all time scales. This is, of course, rarely the case for real neurons. Variations may occur at three rough time scales: short (<1 s), medium (between 1 and 1,000 s), and longer. The variations on short and medium time scales may be tightly related to stimulus conditions; the long-term variations may be related to phenomena like learning. Tools more complex than ordinary cross-correlation, which integrates over time, are appropriate to deal with each of these nonstationarity time scales. For example, the gravitational clustering algorithm (3, 13, 14) and modifications of it (4) can demonstrate variations on the medium time scale.

In this paper we will examine the possibilities for dealing with short time scale, stimulus-locked variations of nearcoincident firing. Such variations of near-coincident firing are to be expected in most cases of effective stimulation, if only because the individual neuron spike rates will be modulated. After correction for such individual rate modulations, any residual variation of the near-coincident firing can be interpreted in terms of corresponding variation of the interactions between the observed neurons. We will demonstrate the detection, quantification, and interpretation of such residual stimulus-time-locked variations in near-coincident firing.

CRITICAL REVIEW OF PREVIOUS APPROACHES TO DYNAMIC CORRELATION

Most electrophysiological measurements on the nervous system involve some kind of time averaging. The reason for this is the apparent stochastic nature of the observed activity. Presumably, the stochastic behavior of the neurons consists partly of statistical fluctuations ("noise"), but may also include genuine (rapid) modulations of activity and/or connectivity. The latter would imply that correlation of firing is a time-varying function; obviously, the usual cross-correlogram would collapse such time variations and present the average values. What is really wanted, but presently unattained, is a new class of measurement that would be essentially instantaneous so as to deal with the modulations of activity, but yet that would be able to discard the statistical fluctuations

The joint peristimulus time scatter diagram

A partial solution to the above dilemma can be developed from the idea of the *peristimulus time (PST) histogram*: let us consider a time-dependent measure of correlation, averaged over many repetitions of the same stimulus and time-locked to the instant of stimulus presentation. This would allow detection of any time structure in the correlation that is related to the instant of stimulus presentation, and yet, being a special form of average, would cope with the statistical fluctuations. Obviously, any modulations (or long-term changes) that are not time-locked to the stimulus would not be detected by such an approach.

An early version of the required calculation was named the *joint peristimulus time (PST) scatter diagram* (10, 11). The basic idea is to create a two-dimensional scatter diagram of the firings of the two neurons relative to each stimulus onset. Technically, this diagram is related to the time-dependent cross-correlation function of the two spike trains. Each dot in the scatter diagram represents a (delayed) coincidence of the two spike trains during one stimulus period, as indicated in Fig. 1*A*, together with the underlying spike trains and their contributions to what will ultimately become PST histograms.

As this process is carried out for many repetitions of the same stimulus, we build up dot densities in the scatter diagram and along the PST axes (Fig. 1*B* shows a scatter diagram for data from a neuronal simulator circuit that will be described in more detail below; note that in such a scatter diagram we do not recognize possible superposition of dots). Regions of high or low density lying in bands parallel to the scatter diagram axes represent stimuluslocked activation or suppression of either neuron; the variation of dot density along such a band represents the stimulus-induced modulation of firing rate in the corresponding neuron. Regions of high or low density lying parallel to the principal diagonal represent excess or deficient amounts of delayed coincident firing, in which the time delay equals the distance from the high- (or low-) density band to the diagonal; the variation of dot density along such a band represents the stimulus-induced modulation of the nearcoincident firing at that particular delay. These various features are illustrated for various simulations of simple neuronal circuits in the catalog of joint-PST scatter diagrams, PST histograms, and cross-correlograms given by Gerstein and Perkel (11).

For ease of explanation, we assumed that the time axes of the Joint-PST scatter diagram start at stimulus onset and extend over an interval equal to the stimulus duration. Clearly, one may deviate from this simple scheme. By selecting an appropriate time window, possibly offset with respect to the stimulus trigger, one may "zoom in" on some arbitrary fraction of the stimulus period. One may also go beyond the time duration of the stimulus by incorporating suitably chosen pre- and/or poststimulus intervals, thus enabling explicit comparison of stimulus and nonstimulus conditions within a single scatter diagram. In the latter case, one should, however, make sure that in each of the trials the extrastimulus interval still represents the same experimental conditions; this is especially important when triggering on repeated occurrences of a particular stimulus, temporally embedded in a (pseudo-) randomly structured stimulus ensemble with possibly varying "neighboring" stimuli and/or interstimulus pauses.

We can estimate the density in a scatter diagram with various histograms. A Cartesian grid to define bins over which tallies are made (schematically indicated in Fig. 1C) leads to the joint-PST histogram (JPSTH) and the two ordinary PST histograms (Fig. 1F). (To allow easy timing comparisons in subsequent modifications of this figure, the PSTHs along the axes are preserved without change in all versions.) Tallies over the diagonally arranged set of bins indicated in Fig. 1D lead to a histogram of the time-locked, near-coincident firing shown in Fig. 1G. This measure shows the stimulus-locked modulation of near-coincident firing in the same way as the PSTH shows the stimuluslocked modulation of single-neuron firing. Where appropriate, we will call this diagonal histogram the PST coincidence histogram. We found it generally convenient to smooth the PST coincidence histogram to reduce statistical variation, especially when we choose a narrow band in which to make the tallies. This would be the case when we study the time distribution of the contributions to a narrow peak (or trough) in the cross-correlogram The smoothing procedure is equivalent to the assumption that variation along the PST coincidence histogram is slow compared to the binwidth. As smoothing function, we use a gaussian with sigma of four (sometimes two) bins; in each case, this





particular value as well as the location and width of the selected diagonal band are indicated in the figure (cf. Figs. 1, F and G).

Finally, tallies over the set of para-diagonal bins indicated in Fig. 1E, suitably normalized for the different length of each bin, estimate the ordinary *cross-correlation histogram* (Fig. 1G). Note that this operation is equivalent to an additional time averaging that inexorably washes out any possible stimulus-induced modulation. We will call this time-averaged measure the (ordinary) cross-correlogram.

Clearly, the statistical variance of counts in the crosscorrelogram bins will increase with the distance from the center bin, due to the decreasing length of the corresponding para-diagonal JPSTH bins; for this reason, we only show the central part of the cross-correlogram. An alternative would be to use para-diagonal bins of identical size, which inevitably would be the smallest of all values used here. This, however, would amount to an unnecessary sacrifice of statistical reliability in the center region of the cross-correlogram; in fact, that would be most unfortunate. because precisely this region is the more interesting part of the histogram in view of the latencies usually associated with (direct) interneuronal connections. If the spike train data would originate from an experiment that allows extending the time base of the JPSTH to encompass more than the duration of the stimulus itself, the para-diagonal JPSTH bins could be arranged such that they all have the extent of the stimulus duration, simply by going beyond the stimulus boundaries for off-diagonal bins. This would be equivalent to cutting a square (or rectangular) section from a temporally extended joint-PST histogram, the section being oriented parallel to the main diagonal and extending equally wide below and above it. The next logical step would be to take this section and rotate it clockwise over 45°; the result is an alternative temporal arrangement of the joint-PST histogram, with the x-axis representing running time and the v-axis corresponding to "difference time." [A more formal description of such operation has been given in the context of nonlinear signal and systems theory (1).] As noted earlier, however, such extension of the time window beyond the stimulus duration is only allowed if experimental conditions in the extrastimulus interval are the same in all trials. Because at this stage we do not want to impose any restrictions on the temporal experiment design, we adopted the above described procedure of compensating for variation in the extent of para-diagonal bins.

The actual display arrangement of the PST coincidence histogram and the cross-correlogram (along and perpendicular to the diagonal in Fig. 1G) was chosen to emphasize their logical relationship to the ordinary PST histograms and the joint-PSTH (Fig. 1F). The additional grid of horizontal and vertical broken lines in Fig. 1G should accommodate comparisons between the time course of the PST histograms and the diagonal histogram.

Prediction of direct stimulus effects

Just as in ordinary cross-correlation, the joint-PST histogram (JPSTH) will in general contain contributions from several sources of (delayed) coincident firing. Those contributions that come purely from stimulus-related modulations of single-neuron firing rates can be predicted. Comparison of the "raw" JPSTH with this predictor defines a "residual" that represents the intrinsic neuronal dependencies alone. Note that the predictor reflects the null hypothesis that spike events in the two neurons are statistically independent, only the firing probabilities of the two neurons are related to the stimulus.

In the past, as a practical matter, the prediction of purely stimulus-related effects in both joint-PST scatter diagrams and in ordinary cross-correlations has usually been calculated by the so-called "shift (or shuffle) predictor" (10, 11, 29). This predictor is based on the idea that most neuronal interactions occur on a much shorter time scale than the time elapsing between two successive stimuli. Thus (for the case of periodic stimuli) if we shift one of the two spike trains by one (or more) stimulus periods, the spike trains will still contain all direct stimulus-induced modulations at the single-neuron level; neural interaction effects will, however, be destroyed, because the time shift is much larger than the delays usually involved in neural interactions. Similar effects can be achieved by cutting one spike train

FIG. 1. Stimulus-locked dynamic correlation of neuronal firing. A: principle of generation of a 2-dimensional scatter diagram of the firings of 2 neurons relative to each stimulus onset (indicated by arrow on horizontal and vertical axis): each dot in the scatter diagram represents a (delayed) coincidence of the 2 spike trains during 1 stimulus period. The spike trains and their contributions to the single-unit dot displays are indicated along the x- and y-axes. B: as this process is carried out for many repetitions of the same stimulus, we build up the joint peristimulus time (PST) scatter diagram and, along the axes, the single-unit dot displays. Spike trains used for Fig. 1 were obtained from a neuronal simulator; the simulated neuronal circuit is depicted in Fig. 3.4, stimulus and constant connectivity in Fig. 2, A and B. C-E: different binning schemes for estimation of dot densities in the scatter diagram and the single-unit dot displays. C: Cartesian grid of bins over which tallies are made results in the joint-PST histogram (JPSTH) and the 2 ordinary PST histograms. D: tallies over a (para-)diagonal arrangement of bins lead to a histogram of the time-locked near-coincident firings: the PST coincidence histogram. E: tallies over the set of (long) para-diagonal bins, suitably normalized for the different length of each bin, estimate the time average of near-coincident firing as a function of relative delay: the ordinary cross-correlation histogram. F. joint-PST histogram (JPSTH matrix) and the 2 ordinary PST histograms along its x- and y-axis (binwidth: 4 ms). Values in the JPSTH matrix are displayed by using gray levels: the higher the value, the darker the gray, as indicated in the gray wedge with the associated values. The tic mark above the gray wedge corresponds to the value zero. All counts were divided by the number of stimulus presentations. G: PST coincidence histogram (running from lower left to upper right) and the ordinary cross-correlation histogram (running from upper left to lower right). The PST coincidence histogram was smoothed using a gaussian with σ of 4 bins; this particular value (gs4), as well as the location and width of the selected diagonal band are indicated in F and G. The position of true coincidence (zero delay) in the cross-correlogram coincides with the intersection point of the PST coincidence histogram and the cross-correlogram; it is indicated by a tic mark near the diagonal band marker above the correlogram.

on stimulus markers, shuffling the sections, and concatenating them. If we now recompute either the joint-PST scatter diagram or the ordinary cross-correlogram for such manipulated spike trains, we have a prediction for the purely stimulus-related effects. Note that a prediction calculated in this way will have exactly the same kind of statistical variation as the original joint-PST scatter diagram or cross-correlogram, so that comparison must be made between two equally "noisy" functions. Such comparisons, for the scatter diagram, are necessarily visual rather than quantitative.

In the original papers on this subject (10, 11, 29), alternate predictors based on the PST histograms were in fact proposed. The PST histogram is appropriate for this purpose, because it is an estimate of the purely stimulus-induced modulation of the single-neuron firing. Such predictors are intrinsically "smooth" in comparison with the "raw" joint calculations, because the PST histograms are themselves already averages over many trials. Because such predictors are binned, it is mandatory to also transform the joint-PST scatter diagram into a histogram (see Fig. 1, C and F) to allow for *quantitative* rather than pictorial comparison. (The binning issue did not arise for the usual cross-correlogram, because historically it has always been presented as a binned histogram rather than as a plot of dots along a line.)

In the original papers, the preference for the shift predictor relative to the PST histogram-based predictor was explained by noting that "a two-dimensional histogram is difficult to compare with the original scatter diagram" (Ref. 11, p. 464). From this, in itself correct observation, the authors, however, chose to generate a predictor that was itself a scatter diagram rather than using a binned predictor and transforming the original scatter diagram into a histogram.

This choice implies a comparison between two scatter diagrams, which is, at best, a qualitative process. Humans perform the task by "squinting" at the two scatter plots, in effect smoothing them with a low-pass spatial filter. Obviously, the smoothing could be carried out quantitatively (at some cost in computer time) by passing an appropriate two-dimensional gaussian over the scatter diagram, thus attaining any desired degree of "spatial" resolution (31). A much more convenient and rapid way to attain similar results is by binning. This has the additional advantage that appropriate choice of bins transforms a spike train into the theoretically more attractive representation of a (0, 1)-process: per stimulus presentation at most one count is added to any histogram bin.

Once we have an original measurement that is a histogram, it is appropriate to choose a predictor that is also a histogram for comparison. This leads to the alternate, and superior, choice of predictors, i.e., those based on the PST histograms; a binned histogram of a shift predictor would be far noisier. For the ordinary cross-correlogram, the preferred predictor is the cross-correlogram of the two PST histograms. For the JPSTH, the predictor is the cross-product of the two PST histograms, taken bin by bin. All histograms should be normalized for the number of trials (stimulus presentations). It turns out that the shift predictor and the PST-based predictor are tightly related, and that the above discussion of relative merits of predictors holds only for the simple shift predictor (i.e., involving any one order of shift), which is, in fact, the one commonly used. It is possible to work with a binned compound shift predictor, i.e., an average over shifts of different order so as to arrive at a statistically smoother predictor. (For such a procedure in the context of ordinary cross-correlograms, see Ref. 8.) It is easy to show that an average over the set of all possible shift predictors (i.e., ALL different orders of shift, including 0) is entirely equivalent to the PST-based predictor for both ordinary and joint cross-correlation (27).

All the above, however, does not directly address a more fundamental problem, associated with the application of the ordinary cross-correlogram to stimulus-driven spike trains. In principle, the cross-correlation of two stochastic processes is a function of two time arguments, t_1 and t_2 . namely, the different observation times for each of the two processes. Only under special conditions concerning stationarity does this reduce to a function of a single time argument, i.e., the time shift $t_1 - t_2$, such as is done in the ordinary cross-correlogram. Processes that are called "jointly stationary in the wide sense" do fulfill these conditions (6, 28). Such processes are required, amongst others. to have an expected value that is independent of time. Because the very purpose of presenting "adequate" stimuli is to influence the neuron's firing behavior to a maximum extent (exemplified by large excursions in the PST histogram), the requirement of a constant firing rate throughout the entire stimulus presentation seems to impose a strong, if not contradictory, condition on the experimental data. Consequently, it is not a priori clear whether the current practice of applying ordinary cross-correlation to stimulusdriven spike trains is allowed at all, nor how the results should be evaluated. The issue of stimulus normalization. which only arises for stimulus-driven data to begin with, merely underlines this problematic aspect of the ordinary cross-correlogram of nonstationary spike trains.

COMPARISON OF ' RAW'' AND PREDICTED JOINT-PST HISTOGRAMS

For ease of explanation in what follows, we will assume that the stimulus and consequent single-neuron rate modulations are periodic. It is trivial to lift this restriction once the arguments and calculations are understood. We assume that all spike trains are "periodically stationary," i.e., that individual trials are statistically indistinguishable. This excludes long-term trends. Finally, we assume that binwidth is chosen such that in each trial we collect at most one spike per bin. We call these numbers of spikes per bin per sweep $n_i^{(k)}(u)$, where u is discrete time (bin number) relative to the kth, i.e., most recent occurrence of the stimulus marker. and *i* indicates the neuron. Our restriction on binwidth means that n can only take values of zero or one; these define the fundamental events with which we will be working. Note that at this stage we make no assumptions whatsoever about the way in which direct stimulus effects and neural effects combine.

A stimulus



FIG. 2. Time course of stimulus (A) and connectivity (B) used in simulations of neuronal circuits in Figs. 3A and 4, A and B. The broken line in B indicates the constant excitatory connection strength of 0.1 for the circuit in Fig. 3A; the continuous line corresponds to the stimulus-modulated connectivity (time average, 0.1) used for the circuits in Fig. 4, A and B.

To illustrate the various possibilities for comparison of "raw" and predicted joint-PST histograms (JPSTH), we will make use of spike trains produced by a neuronal simulator program. The design of this simple simulator has been described (2), the current version allows for stimulus modulation of both individual firing rates and of the strength of interneuronal connectivity. Let us for the moment continue to work with the data whose "raw" JPSTH is shown in Fig. 1F and repeated in Fig. 3C. The simulated circuit consisted of two neurons, both spontaneously active (some 4 spikes/s), both excited by a stimulus (schematically depicted in Fig. 2A), and with neuron 1 exciting neuron 2 with constant efficacy (i.e., probability of inducing a spike from the "postsynaptic" neuron) of 0.1 as indicated in Figs. 2B and 3A. Average spike rates were set to 7 spikes/s to mimic typical cortical recordings.

With the above given assumptions, we can develop some convenient notation. Let the stimulus-locked event density be $E_{stim}[n_i(u)]$; the k index is omitted because we are considering an expectation. The estimator of this quantity is the ordinary PST histogram

$$\langle n_i(u) \rangle = \frac{1}{K} \sum_{k=1}^K n_i^{(k)}(u) \tag{1}$$

Note that because of the normalization for stimulus number, values in the PST histogram are restricted to the range zero to one. With the same notation, the contribution to the joint-PST histogram by the events during the kth stimulus period is given by

$$n_{ii}^{(k)}(u, v) = n_{i}^{(k)}(u)n_{i}^{(k)}(v)$$
⁽²⁾

These joint events are also zeroes or ones. In analogy to the single-neuron situation, we can now introduce the stimulus-locked joint event density $E_{stim}[n_{ij}(u, v)]$ and its estimator, the JPSTH

$$\langle n_{ij}(u, v) \rangle = \frac{1}{K} \sum_{k=1}^{K} n_{ij}^{(k)}(u, v)$$
 (3)

Note that because of the normalization for stimulus number, also the JPSTH values are restricted to the range zero to one.

Under the null hypothesis of no interaction between the neurons, it can be formally shown that the cross-product matrix of the individual PST histograms (i.e., bin by bin) is an appropriate estimator of the expected JPSTH (27). This predictor is given by

$$\tilde{n}_{ij}(u,v) = \langle n_i(u) \rangle \langle n_j(v) \rangle \tag{4}$$

and is shown in Fig. 3B. Its values also lie between zero and one. As promised, this is a very smooth function in comparison with the "raw" JPSTH; it shows clear horizontal and vertical features, but no diagonal features. Nevertheless, its diagonal histogram, the predicted PST coincidence histogram, shows a clear, stimulus-locked modulation of near-coincidence firing. This is the amount and time course of near-coincident firing that would be expected purely on the basis of the single-neuron firing rates and their modulations by the stimulus.

A number of ways can be envisioned to compare the raw and predicted JPSTH. Because we want to test the null hypothesis of no interaction, we are particularly interested in the dissimilarity between the two. The simplest measure for this dissimilarity is the algebraic difference

$$D_{ij}(u, v) = \langle n_{ij}(u, v) \rangle - \tilde{n}_{ij}(u, v)$$
$$= \langle n_{ij}(u, v) \rangle - \langle n_i(u) \rangle \langle n_j(v) \rangle$$
(5)

The values of this function lie between minus and plus one. Equation 5 is equivalent to the usual definition of the cross-covariance function (6, 28); we shall refer to it accordingly. This function is shown in Fig. 3D. We note that the bands of high counts parallel to the x- and y-axes were present both in the "raw" and predicted JPSTH; the operation of taking an algebraic difference in the cross-covariance removes these features of net positive values. The diagonal feature was present only in the "raw" JPSTH; therefore, it is not surprising that it persists in the cross-covariance. There is still a modulation along the diagonal; the time course of this modulation roughly follows that in the two ordinary PST histograms shown along the axes. We will return to this modulation later on. Note that the average of the background in both the cross-covariance and, consequently, in the differential correlogram (at upper *right*) has been shifted to zero by the subtraction operation. The variance of the background may not be uniform, however, and will depend on the detailed firing rates as shown by the PST histograms. This means that different portions of the cross-covariance may show different "texture" (notice the horizontal and vertical bands of texture around the zero value).

The range of values comprising both the statistical variance and the "real" features in the algebraic difference of the cross-covariance will clearly depend on the details of the data. To compare the importance of various features in different pieces of data (e.g., corresponding to the presentation of different stimuli to the same two neurons), we need an appropriate normalization. In other words, we need a data-dependent measuring stick to arrive at a data-*in*dependent measure of departure from the null hypothesis. To give qualitative meaning to the extent of such departure, let





us in addition require from an appropriate normalization procedure that the resulting values in diagonal features bear a monotonic relation to the "effective connectivity" in the underlying neuronal network: the stronger the connectivity, the larger the departure, and vice versa. This additional and relatively mild requirement allows us to interpret temporal variations of diagonal features in a properly normalized JPSTH in terms of modulations of "effective connectivity."

One attempt to achieve this normalization goal is to measure the departure from predicted value relative to that predicted value. This means we want to divide the cross-covariance by the cross-product of the PST histograms

$$R_{ij}(u, v) = \frac{D_{ij}(u, v)}{\tilde{n}_{ij}(u, v)}$$
$$= \frac{\langle n_{ij}(u, v) \rangle}{\langle n_i(u) \rangle \langle n_j(v) \rangle} - 1$$
(6)

Note that the values of this function are not a priori limited to any particular range. This quantity has been called "scaled cross-correlation density" by Kuznetsov and Stratonovich (21; see also Ref. 7). We show the results of this calculation in Fig. 3F; it is clearly unsatisfactory. It has reintroduced vertical and horizontal features in the off-diagonal regions: a clear modulation of variance (note the "blur" in Fig. 3F, coinciding with the original horizontal and vertical features from Fig. 3D) where there should be only uniformity. In addition, the diagonal histogram is not uniform along its length, but rather has clearly lower values in the early part of its time course, i.e., where the PST histograms have high values. Such variation along the diagonal is hard to reconcile with the known constant strength of connection in the simulated neuronal circuit. We have also applied this attempted normalization to other simulated spike train data (with known underlying circuitry), and in all cases the results of this procedure were similarly inadequate. In summary, this attempted stimulus normalization procedure seems to be decompressing too strongly in regions corresponding to low firing rates of either or both neurons and, therefore, will not be shown again.

A second possible normalization of the cross-covariance is inspired by standard statistical procedures. Here we divide the cross-covariance by the standard deviation of the predictor. Under the null hypothesis this standard deviation is simply the cross-product of the standard deviations of the PST histograms $\tilde{s}_{ii}(u, v) = s_i(u)s_i(v)$

$$= \left\{ \frac{1}{K} \sum_{k=1}^{K} (n_i^{(k)}(u) - \langle n_i(u) \rangle)^2 \cdot \frac{1}{K} \sum_{k=1}^{K} (n_j^{(k)}(v) - \langle n_j(v) \rangle)^2 \right\}^{\gamma_i}$$
(7a)

As can be observed from comparison of Eqs. 2-5 and 7a, an equivalent formulation of this quantity is the square root of the product of the two single-neuron autocovariances, each matrix considered only along its diagonal

$$\tilde{s}_{ij}(u, v) = \{ D_{ii}(u, u) D_{jj}(v, v) \}^{1/2}$$
(7b)

It should be noted that there is a simple relation between the predictor in Eq. 4 and the standard deviation in Eq. 7a. With the assumptions stated at the outset, the expectations $E[n_i(u)]$ and $E[n_i^2(u)]$ are equal for any u because we are dealing with a "zero-one" process. From this it is easily shown that the standard deviation matrix is

$$\tilde{s}_{ij} = s_i s_j = \{ \langle n_i \rangle (1 - \langle n_i \rangle) \cdot \langle n_j \rangle (1 - \langle n_j \rangle) \}^{\nu_1}$$
(8)

where we omit the time bin reference (u, v) for brevity. For $\langle n_i \rangle$ and $\langle n_j \rangle$ small compared to one (the usual case), say both equal to ϵ , it follows that the predictor (Eq. 4) behaves as ϵ^2 , whereas the standard deviation (Eq. 8) behaves as ϵ^2 , whereas the standard deviation is the square root of the predictor. Furthermore, the ratio of maximum to minimum values of expression (Eq. 8) will be the square root of the corresponding ratio in the predictor of expression (Eq. 4). The standard deviation thus has a smaller relative dynamic range than the predictor; it is a very smooth function, and we omit explicitly showing it here. Such smoothness is an obvious requirement for any division process in order not to increase the "noise" level.

The normalization of the cross-covariance using the standard deviation of the predictor as scaling matrix is, therefore, given by

$$C_{ij}(u, v) = \frac{D_{ij}(u, v)}{\tilde{s}_{ij}(u, v)} = \frac{D_{ij}(u, v)}{\{D_{ii}(u, u)D_{ji}(v, v)\}^{v_i}}$$
(9)

Note that this function, unlike the relative scaling operation in Eq. 6, is a true normalization, i.e., it delivers numbers between plus and minus one. Expression 9 is referred to in the literature as the "normalized cross-covariance function" or "correlation function coefficient" (4, 6); it has also been called "cross-correlation surface" (18). Finding the above terminology unsatisfactory for various reasons,

FIG. 3. Various possibilities for comparison of "raw" and predicted joint-PST histograms (JPSTH). A: the simulated neuronal circuit, giving rise to the spike train data used in Figs. 1 and 3, consisted of 2 neurons, both excited by a stimulus, and with neuron 1 exciting neuron 2 with constant efficacy (see also Fig. 2, A and B). Average spike rates were set to 7/s to mimic typical cortical recordings. Total numbers of events: 2,774 (neuron 1, x-axis), 2,951 (neuron 2, y-axis), 2,000 (stimulus triggers). B: estimator of expected JPSTH under the null hypothesis of noninteracting neurons: the cross-product matrix of the individual PST histograms (Eq. 4). Display format of this and all following figures as in Fig. 1, F and G. C: "raw" JPSTH (Eq. 3) as shown in Fig. 1, F and G, and repeated here for ease of reference. D: cross-covariance function (Eq. 5): the algebraic difference of "raw" (C) and predicted JPSTH (B). E: normalized JPSTH (Eq. 9): normalization of the cross-covariance (D) using the standard deviation of the predictor (cross-product of the standard deviations of the PST histograms) as scaling matrix. F: scaled cross-correlation density (Eq. 6): ratio of cross-covariance (D) and predicted JPSTH (B). G: significance of correlation: algebraic difference between the surprise for "excitation" (i.e., for an excess joint count compared with the individual counts), and the surprise for "inhibition" (i.e., for a deficit of joint counts). The surprise of a value is defined as the negative natural logarithm of the probability for finding that value or a more deviant one (Eqs. A1.1 and A1.2). Surprise matrices were scaled for display between -2.996 (corresponding to a significance level of 5% for "inhibition") and +4.605 (similarly of 1% for "excitation") as indicated along the gray wedge.

we propose to call the result of Eq. 9 the normalized JPSTH. We show the result of the calculation in Fig. 3E. The result is clearly better than that of Fig. 3F. No obvious horizontal or vertical features reappear the off-diagonal "texture" is completely uniform. Furthermore, the diagonal feature is approximately constant along its length, as it should be for this neuronal circuit of constant connectivity. This is also shown by the normalized PST coincidence histogram in Fig. 3E.

Integration along the diagonal (tallies in the para-diagonal bins as in Fig. 1E) results in a histogram that we call the normalized cross-correlogram. It should be noted that this normalization in general can not be achieved by any operation on the usual cross-correlogram: the proper procedure fundamentally involves the temporal decomposition of near-coincident firing in the JPSTH, its subsequent dynamic correction (subtraction and scaling), and reintegration along the diagonal. The order of these various operations can not be interchanged. The usual procedure for correcting the ordinary cross-correlogram for stimulus effects has been to subtract the shift predictor correlogram or PST predictor correlogram (with or without additional scaling for numbers of events). It is clear that the normalized cross-correlogram, i.e., a calculation from the normalized JPSTH, will generally give a different and more correct result.

SIGNIFICANCE TESTING: SURPRISE

In the normalized JPSTH (Eq. 9), we have a calculation that is appropriate for the intercomparison of different pieces of data with the null hypothesis However, we now need a measure of significance to interpret the deviations from the null hypothesis. The usual approach to significance testing consists of comparing the actual outcome of an experiment with the postulated distribution of values under the null hypothesis. In our case, we need to identify this postulated distribution for the JPSTH value in a given bin.

The point of departure is to calculate the distribution of possible JPSTH values for a particular bin (u, v), given 1) the null hypothesis of independence, 2) the values at the corresponding bins in the two PST histograms, and 3) the number of trials (stimulus presentations) over which these PST histograms were gathered It is possible to give this distribution in closed form, expressed in terms of only these algebraic quantities (for details, see Ref. 27). By substituting the experimental values for these quantities, we obtain an explicit distribution appropriate for this particular piece of data; it generally turns out to be different from Gaussian, especially for realistic, relatively low numbers of stimulus presentations. We can now test the significance of our actual experimental JPSTH value against this theoretically derived, but experimentally particularized, distribution: the significance of a value is defined as the probability for finding that value or a more deviant one. One can view this significance test in two different ways: 1) we test the significance of deviations of the "raw" JPSTH from the null hypothesis (i.e., the predictor), or equivalently 2) we test the significance of deviations of the normalized JPSTH from the (predicted) value of zero.

An informative display of the results of the significance test (27) can be made by using the information-theoretical concept of "surprise" (22, 26). The surprise of a value in this context is the negative natural logarithm of the probability for finding that value or a more deviant one (formal expressions are given in Appendix 1). It is, therefore, directly related to the usual statistical notion of "level of significance." For instance, a 5% level of significance corresponds to a surprise of 2.996 (= $-\ln 0.05$); a 1% level corresponds to a surprise of 4.605. The logarithmic transformation serves to expand the scale in the interesting region in which the probability density has low values; this is somewhat comparable to the use of a decibel scale. Moreover, introduction of the logarithmic scale allows a more sensible comparison of different values of significance; the same measure is assigned to equal ratios of probabilities rather than to equal algebraic differences.

We can define a surprise measure for "excitation," S(E), i.e., for an excess joint count compared to the individual counts, and comparably, a surprise for "inhibition," S(I), i.e., for a deficit of joint counts. We could show these measures separately. However, if for a particular bin one of the measures is high, the other is necessarily close to zero. Thus, in case of a significant "excitation" (or "inhibition"), the algebraic difference S(E) - S(I) practically equals S(E)(or -S(I), respectively). For reasons of display economy we have combined the two measures accordingly and show the algebraic difference S(E) - S(I) for all (u, v)-bins in the form of a single matrix (Fig. 3G). Note that "excitation" leads to positive values of this combined measure, whereas "inhibition" leads to negative values; the larger the (absolute) value in the surprise matrix, the more statistically significant the normalized JPSTH in the corresponding bin. (In Appendix 1 we give explicit relations between the surprise values for significant "excitation" or "inhibition" and the corresponding values of the normalized JPSTH.)

We will scale all surprise matrices so that the positive range is set to 4.605 (corresponding to a significance level of 1% for "excitation") and the negative range to -2.996 (similarly, 5% for "inhibition"). The reason for this asymmetry is the difficulty in detecting inhibition as compared with excitation for the usual small counts (2, 25, 27).

All normalizations introduced above, as well as the surprise measure for significance are bin-by-bin measures. Yet in the above exposition, we have regularly referred to "features" in the matrix. In other words, we are implicitly searching for regions of approximate "spatial" coherence to assign meaning; isolated bins that reach extreme values are considered "noise." This implicit additional requirement for "spatial" coherence should somehow be incorporated into the formal bin-by-bin measure of significance. Just as in the normalizations, we have introduced this regional aspect into the surprise significance measure by using a moving average (gaussian) along the corresponding diagonal histogram of surprise values. The underlying hypothesis for this smoothing process is that the values in individual bins are independent¹; recall that the surprise

¹ This (null) hypothesis is based on the assumption that the time interval between the bins considered is large enough such that the single neuron's spike-generating mechanism will not produce appreciable correlations

calculation is bin by bin. In such a situation, the joint probability for having any particular combination of bin values is simply the product of the individual probabilities. Here we encounter another advantage of the logarithmic surprise measure of significance: the joint surprise for any particular bin constellation is found simply by adding the corresponding individual surprise values. In our specific case. we require a large proportion of high values over some limited region of the surprise matrix; the addition of surprise values in such a region translates into the smoothing operation. Note that the added surprise values, however, cannot directly be interpreted as a measure for statistical significance. This is quite obvious, because the addition would lead to amazingly low significance probabilities. The correct statistical procedure for significance analysis of combined surprise values is the subject of current investigation.

We can now interpret the surprise measure in Fig. 3G. The surprise histogram shows no coherent structure except along the diagonal, where many points surpass the 1% "excitation" criterion. The corresponding diagonal histogram is relatively flat along its length. We take this result to indicate that the diagonal feature seen in the normalized JPSTH (Fig. 3E) is indeed a significant departure from the null hypothesis; the two neurons are not to be considered independent, and their interaction (which is "excitatory") is not modulated by the stimulus. This means that our calculations are capable of recovering the detailed structure that had been built into the simulated circuit. Although there are occasional single bins in the remainder of the surprise histogram that reach extreme values (note extrema indicated in Fig. 3G), we do not consider these to be significant according to our extended criterion which incorporates the additional requirement of "spatial" coherence in the surprise matrix.

CALIBRATION BY SIMULATED NEURONAL CIRCUITS

As we continue to calibrate the sensitivity of the normalized JPSTH and the surprise significance, we turn to several additional simulated neuronal circuits, as shown in Fig. 4, A and B. Both circuits have stimulus-locked modulation of the connectivity between the neurons. The objective is to investigate whether the normalized JPSTH allows detection of such stimulus-modulated connectivity, even when this modulation is masked by direct stimulus influences on each of the two neurons.

The simpler case is shown in the *left* column of Fig. 4. (The layout of each column in Fig. 4 matches the *right*-

hand column of Fig. 3, as well as the columns of Fig. 5.) For this simulation, the average firing rates are $\sim 4/s$; this is a lower rate than in Fig. 3, because there are no direct stimulus effects on the neurons. The connectivity was a triangular function of time as shown in Fig. 2B, with the peak after the first quarter of the stimulus cycle. The numbers were specifically chosen to give an average connectivity of 0.1. the same value as the constant connectivity in Fig. 3. Fig. 4, C, E, and G, in that order, shows the "raw" JPSTH, the normalized JPSTH, and the surprise measure. The PST histogram along the x-axis, corresponding to neuron 1 (the "driver"), is flat, as expected from Fig. 4A. The connection between neurons 1 and 2 is sufficiently weak (although physiologically realistic) to make the PST histogram along the y-axis (neuron 2, the "driven") also appear flat. The only feature in the "raw" JPSTH (Fig. 4C) is diagonal, and its intensity, shown in the PST coincidence histogram, matches what we know of the modulation of connectivity in the circuit. Given that the PST histograms are both flat, the normalized JPSTH (Fig. 4E) simply reflects a changed scale and base line. Finally, the surprise measure shows a diagonal feature that has many bins exceeding the 1% criterion level. The corresponding diagonal histogram also follows the known connectivity function. To interpret these smoothed values, we also calculated a comparable histogram for the diagonal stripe that is two bins laterally displaced. That histogram does not rise above one-tenth of the peak values here. From this we conclude that the diagonal feature in the normalized JPSTH is highly significant; we have recovered the structure of the simulated circuit.

The right-hand column of Fig. 4 presents a combination of direct stimulus modulation of the two firing rates and the stimulus-modulated connectivity that we have just analyzed. The circuit is shown in Fig. 4B. The "raw" JPSTH of Fig. 4D shows the expected vertical, horizontal, and diagonal features. The PST histograms show the direct stimulus effects, but again there is no obvious signature of the modulated connectivity in the driven PST histogram (along the y-axis). The corresponding PST coincidence histogram shows a clear stimulus modulation of the nearcoincident events; presumably this represents both the modulation of the firing rates and of the connectivity.

The direct stimulus effects are removed in the normalized JPSTH of Fig. 4F; only the diagonal feature persists. Note that Fig. 4F looks generally like Fig. 4E, its counterpart for the circuit that only has stimulus-modulated connectivity: in both cases, the normalized PST coincidence histogram closely follows the known modulation profile of the connectivity. Finally, the significance measure in Fig. 4H indicates that the diagonal feature in the normalized JPSTH represents a highly significant departure from the null hypothesis of no interaction. Therefore, even in this case of mixed direct and indirect stimulus effects, we have consistently recovered the simplest network that is sufficient to explain these data; this network appears to be identical to the (known) simulated circuit.

QUANTITATIVE ASSESSMENT OF "EFFECTIVE CONNECTIVITY"

It should be emphasized that in all of the above development of normalization, no assumptions whatsoever were

among these bins. A possible source for such correlations would be the neuron's refractory mechanism: the firing in one bin would be negatively correlated with the firing in adjacent bins (especially because we assume our bins small enough to have at most one spike per bin per trial). Such negative correlations would manifest themselves as para-diagonal features of decreased coincidence counts in the auto-JPSTH, which is obtained by taking the same single-neuron spike train along both the x- and y-axis of the JPSTH. What we are, in fact, assuming here is "sparse firing": firing rates are so low that, even in stimulus-driven activity, the intervals between adjacent spikes are (considerably) larger than the neuron's refractory period. With this assumption, which does not appear implausible for cortical neurons with typical, low firing rates, we can extend the null hypothesis to the effect that also the spikes in different bins for the same neuron are independent.



made regarding the structure of the underlying neuronal network: the approach is essentially *model free*. Under these conditions, the proper strategy is to test the data against the null hypothesis of no interaction; the possible outcome of this test is basically "yes" or "no" to some significance level.

Once a significant departure from the null hypothesis of independence has been established, we can move to a second, model-related stage of interpretation. So far, we have only implicitly and qualitatively assigned meaning to the magnitude of the departure from the null hypothesis: the larger the departure, the stronger the "connectivity." This notion can be made more explicit and quantitative by making specific assumptions about the underlying network structure. This means a *model-based* approach and aims at evaluation of the model's parameters. From an excess (or deficit) of correlated events, we can make the conceptual transition to a statement of "effective connectivity." The latter is expressed as a minimal neuronal circuit model that would reproduce such correlation measurements.

An important example of such a model would be the simple circuit in which neuron a (x-coordinate in the JPSTH) drives neuron b (y-coordinate) with a specifically defined (linear or nonlinear) synaptic transfer function. The normalized JPSTH in this case will show an above-diagonal feature that is displaced from the principal diagonal by an amount equal to the latency of the connection. In fact, we used just this model for the simulations that produced the data analyzed in Figs. 1, 3, and 4.

To quantitatively assess the strength of direct interaction among neurons, Levick et al. (23) introduced the notions of "efficacy" and "contribution." Efficacy was defined as the fraction of spikes from the driver that are time related to spikes from the driven neuron, contribution as the fraction of spikes from the driven neuron that are time related to spikes from the driver. These two quantities remain the principal quantitative measures for "connectivity" developed so far; consequently, they are widely used in the experimental multi-unit literature (references can be found in Refs. 12 and 20). Adopting a model-based approach, we may generalize the quantities of efficacy and contribution to include possible stimulus-locked time variations. Such manipulations allow the assignment of specific numerical values for this type of dynamic "effective connectivity." A discussion along these lines is given in Appendix 2; here we will only list the final results.

The dynamic generalization of efficacy and contribution can be expressed in terms of the experimentally observable JPSTH and the single-neuron PST histograms

$$e(u, v) = \frac{\langle n_{ab}(u, v) \rangle - \langle n_a(u) \rangle \langle n_b(v) \rangle}{\langle n_a(u) \rangle (1 - \langle n_a(u) \rangle)}$$
(A2.11)

$$c(u, v) = \frac{\langle n_{ab}(u, v) \rangle - \langle n_a(u) \rangle \langle n_b(v) \rangle}{\langle n_b(v) \rangle (1 - \langle n_b(v) \rangle)}$$
(42.12)

These expressions for the dynamic efficacy and contribution can be related in an interesting way to the normalization of the JPSTH. When we compare Eqs. A2.11 and A2.12 with Eqs. 8 and 9, we observe that the normalized JPSTH takes an "intermediate" position between the complementary pair of efficacy and contribution. In contrast to the efficacy and contribution, the expression for the normalized JPSTH is symmetric with respect to both neurons, as one should expect for a model free descriptor. The "intermediate" position of the normalized JPSTH between efficacy and contribution takes the form of a geometrical average

$$(normalized JPSTH)^2 = Efficacy \times Contribution$$
(10)

MULTI-NEURON RECORDINGS FROM REAL NEURONS

In the above, we have described a new, extended quantification of the old joint-PST Scatter Diagram (10, 11) in terms of the "raw" JPSTH, the normalized JPSTH, and the surprise measure of significance. These quantities were calibrated by the use of spike trains derived from simulations of particular neuronal circuits. Analysis of these simulated spike trains showed that stimulus-time-locked variations of connectivity can indeed be untangled from the direct stimulus modulations of individual firing rates. In this section, we present illustrative examples of the "inverse problem": we analyze the spike trains from real neurons to infer an as yet unknown "effective connectivity."

We have taken data from two preparations: cochlear nucleus of the rat and visual cortex of the cat, both lightly anesthetized. These data were available on tape, and we did not perform any exhaustive or systematic searches. The only criterion for selection of data to be examined was an "adequate" number of firings over an "adequate" number of stimulus presentations (for obvious statistical reasons). For more detail of the experimental arrangements, original intent, and results on other issues than dealt with here, see Refs. 17, 19, and 20, respectively. Cochlear nucleus experiments were done with a repeated long sequence of 50-ms tone pips with different frequencies and amplitudes. Data corresponding to the repetitions of a particular tone pip near best frequency, to which the two observed neurons responded with an on-burst and an off-burst (see PST his-

FIG. 4. Evaluation of dynamic correlation and its significance for spike trains from 2 additional simulated neuronal circuits. A and B: schematic representation of the 2 simulated circuits. In both circuits, neuron 1 excites neuron 2 with identical stimulus-locked modulation of interneuronal connectivity (cf. Fig. 2B). In addition, the circuit in B has direct stimulus modulation of the single-unit firing rates (cf. Fig. 2A); this circuit thus presents a combination of the 2 circuits in Figs. 3A and 4A. The layout of each column in Fig. 4 matches the *right-hand* column of Fig. 3. Furthermore, to enable comparison across results for different circuits, scaling within the lower 2 horizontal rows of Fig. 4 was set to identical values (scaling between rows is different though). Binwidth in all cases was 4 ms. Total numbers of events were as follows: left: 1,578 (neuron 1, x-axis), 1,911 (neuron 2, y-axis), 2,000 (stimulus triggers). C, E, and G (left): the "raw" JPSTH, the normalized JPSTH, and the significance of correlation for the spike trains generated by the circuit in A. D, F. H (right): similarly for the spike trains generated by the circuit in A.



tograms in *left-hand* column of Fig. 5), was sorted out for analysis by the JPSTH. Cortical data involved repetitions of long sequences of moving bar stimuli (different orientations; one cycle of moving back and forth lasted 5 s), and was similarly sorted to isolate responses to the repetitions of a particularly effective stimulus (in this case optimal bar orientation).

We show results of the JPSTH analysis in Fig. 5. The *left* column is for a pair of rat cochlear nucleus neurons, the *right* column for a pair of cat visual cortex neurons. The layout is similar to that in the preceding figures, with the "raw" JPSTH on top, followed by the normalized JPSTH and the surprise measure for significance.

Both these sets of data show in the "raw" JPSTH all features that we have encountered above for the simulated spike trains: features parallel to the axis, corresponding to direct stimulus effects on the single-neuron firing rates, and diagonal features corresponding to the correlation of firing. The normalized JPSTHs in both cases remove most of the horizontal and vertical features, and leave only the diagonal feature. In both cases, when we examine the diagonal feature with the normalized PST coincidence histogram (tallied over the band indicated in the figure), we observe that the normalized correlation is clearly time varying along the diagonal. The significance of these effects is shown in the surprise matrices, which clearly confirm significance of the diagonal features.

To "explain" the diagonal features, we make the conceptual transition to the notion of "effective connectivity": the minimal neuronal circuit model that will reproduce these features. In such a minimal model there is a monotonic relation between diagonal feature strength and connectivity. Thus time variations along the diagonal features of the normalized JPSTHs in Fig. 5, C and D, are interpreted as indicators of time modulated "effective connectivity."

The notion of "effective connectivity" usually can be split into two types of circuitry: direct interaction and shared input. This difference is reflected in characteristics of the diagonal feature such as width and (a)symmetry around the principal diagonal. The inferred "equivalent" circuits based on these criteria are shown in Fig. 5, G and H, bottom. The narrow, slightly above-diagonal feature in the normalized JPSTH for the cochlear nucleus pair (Fig. 5C) is interpreted as the signature of a direct, excitatory

synaptic connection from neuron 1 to neuron 2, its time course is taken to reflect stimulus modulation of the connection; we thus arrive at the model circuit in Fig. 5G, similar to our simulated circuit in Fig. 4B. Such a direct connection is not visible in the cortical data (Fig. 5, D and F): the diagonal feature is symmetrical around time shift 0, suggesting shared neuronal input. The strong time variations along the diagonal are taken to reflect stimulus modulation of the (unobserved) source(s) of shared input and/ or of the connections from source(s) to neurons 1 and/or 2 (cf. Fig. 5H).

In these two bodies of data, without systematic or exhaustive search, we have encountered other neuron pairs with similar stimulus-modulated "effective connectivity." We have also encountered neuron pairs in which the normalized PST coincidence histogram is relatively flat, indicating that the "effective connectivity" is not modulated in a stimulus-locked fashion. Obviously, we have also encountered neuron pairs that showed no diagonal features at all; these presumably were noninteracting neuron pairs.

DISCUSSION

In this paper we have developed a quantification procedure for the study of stimulus-locked, time-dependent correlation of firing between two neurons. This procedure starts from the ancient joint-PST scatter diagram and makes it possible to unravel and quantitatively describe direct and indirect stimulus effects on correlated firing of two neurons. The direct stimulus effects are shown by the predictor for the JPSTH: this predictor is the cross-product of the two single-neuron PST histograms and represents the null hypothesis of independent firing. The indirect stimulus effects (residual correlation) are described by the normalized JPSTH (Eq. 9); this is the algebraic difference between the "raw" JPSTH and the predictor, with subsequent dynamic scaling by the cross-product of the standard deviations in the single-neuron PST histograms. Finally, we have demonstrated the use of a new measure for significance of features in the normalized JPSTH: surprise.

We have investigated several other possible normalization procedures (one of which is explicitly considered in this paper; see also Ref. 27). In comparison, these alternatives proved to be unsatisfactory in performance. The normalization of the JPSTH that we have selected is basically a test statistic for departure from the null hypothesis. It starts

FIG. 5. Evaluation of dynamic correlation and its significance for physiological spike trains recorded from 2 neurons simultaneously. The left column is for a pair of rat cochlear nucleus neurons, the right column for a pair of cat visual cortex (area 17) neurons. In each of the experiments, spike trains were recorded during repeated presentation of an effective stimulus. Total numbers of events were as follows: left: 511 (neuron 1, x-axis), 481 (neuron 2, y-axis), 242 (stimulus triggers); right: 756 (neuron 1, x-axis), 2,131 (neuron 2 y-axis), 75 (stimulus triggers). The layout is similar to that in Figs. 3 and 4: the "raw" JPSTH on top (A and B), followed by the normalized JPSTH (C and D) and the significance of correlation (E and F). At the bottom, we show the inferred "effective" circuitry (G and H). The U symbol in H denotes unobserved other neuron(s). Scaling is done independently for each of the figure elements. Note the different time resolutions for these 2 sets of data: binwidth was 1.5 ms (left) and 50 ms (right), sweep duration 75 ms (left) and 5,000 ms (right). The lack of counts in the band along the main diagonal of the "raw" JPSTH for the cochlear nucleus neuron pair (A) and, hence, the negative values in the cross-correlograms of C and E, are an artifact: the multi-unit spike trains in this case were recorded on a single microelectrode and subsequently sorted on the basis of spike waveforms. The spike sorter used in these experiments has a dead time that discards temporally overlapping spikes. This dead time, admittedly, is a most unfortunate artifact of all spike sorter algorithms developed so far. This holds in particular, because it affects the central features in the cross-correlogram, which presumably are the signs of direct interactions. The alternative of using data from different microelectrodes, with consequently larger interneuron distance, however, has the disadvantage of only rarely showing direct interaction among recorded neurons (such as shown here).

from the deviation of actual stimulus-locked correlation from its expectation, but weighs such deviation, bin by bin, for the a priori expected statistical spread as expressed by the single-neuron standard deviations. This normalization is intrinsically model free; accordingly, it is symmetric with respect to the two neurons.

Assumptions

Throughout the development of the normalization procedure, we have made two major assumptions. First, we have assumed that spike trains during successive stimulus presentations are "periodically stationary," i.e., that individual trials are statistically indistinguishable. This assumption in itself is not particularly new: any analysis of (single- or multi-unit) spike trains involving averaging over stimulus sweeps (as in, e.g., the common PST histogram or any measure derived from it) is based on the same assumption. The second assumption involves the choice of binwidth in the analysis window: bins should be taken small enough such that per trial we collect at most one spike per bin. This assumption is essentially connected to the probabilistic nature of our approach: it allows the interpretation of trial averages of various counts as probabilities, i.e., with values between zero and one. This implies that also the second assumption is not restricted to the present issue: any analysis procedure of spike trains that treats a trial average (e.g., a PST histogram) as an event density in the probability sense must necessarily make the same assumption to guarantee interpretable values between zero and one.

The assumption of "periodic stationarity" excludes long-term trends or other sources of variability on medium to large time scales. Clearly, if such sources should be present, our predictors for the JPSTH (Eq. 4) and its statistical variation (Eq. 7) would be incorrect and, consequently, the normalized JPSTH would yield erroneous results. The way in which the predictors and, consequently, the normalized JPSTH are affected is highly dependent on the nature and form of the nonstationarities. A simple example may serve to illustrate this point. In the case of a trend in overall firing rate of only one of the two neurons, while the other one remains constant, the PST-based predictor (Eq. 4) still is a correct estimator of the expected coincidence counts under the null hypothesis [the variance estimator (Eq. 7), however, in general is no longer correct]. For a linear trend in both neuron firing rates, however, the PST-predictor (Eq. 4) can either be an underestimate or an overestimate, depending on whether the trends in the two neurons go in the same or in obverse directions. The result thus is an incompletely normalized JPSTH: remaining horizontal/vertical features and/or possibly delusive suggestions of modulation of "effective connectivity" (as we have sometimes encountered). Careful screening of the spike train data for such nonstationarities, therefore, is mandatory (such screening was performed for all results presented here). Especially in data from the cortex, where medium- to long-term neuronal variability, not under the control of the experimenter and highly correlated within groups of neurons, has been demonstrated (5), and in spike data from awake, behaving animals, the possible effects of nonstationarities across stimulus presentations on any kind of average measure should seriously be considered. If necessary (and possible), one may use the procedure of "slicing" the spike trains into different sections, each one periodically stationary, and compare the results across sections. A more general analysis of the effects of nonstationarities and the development of appropriate control calculations is currently in progress.

Implementation

Software realizations of the various calculations we have described were implemented in Fortran 77; a wide range of graphic display devices is trivially accomodated. Computation time on any current computer is negligible: reading the spike data and calculation (without specific attempts at optimization) of the "raw" JPSTH plus its normalized version, the surprise matrix and the ordinary PST histograms, as well as the diagonal histograms and cross-correlograms for each of the three matrices for spike trains involving some 3,000 spikes over 2,000 sweeps (Fig. 4, D, F, and G) on the Institute's Vax-750 took ~35 CPU-seconds; comparable timing should be accomplished on any modern microcomputer. Both display and publication printing, however, are more difficult. We have used a Vectrix-384 color display device (672×480 pixels, 512-color palette) attached serially to our computers. This type of device with its serial connection requires a considerable time to build the kind of pictures used in this paper. Currently available display devices that are directly attached to the computer bus eliminate this time problem. We have found the use of color in these displays quite useful in the laboratory and for making slides. Publication costs of such color material, unfortunately, remain prohibitive, so that we had to use gray scales here. The results on the screen are quite satisfactory; however, we have found the technical process of getting from the grav screen to the printed page to be quite difficult and disproportionately time consuming.

"Effective connectivity"

The normalized JPSTH erases all vertical and horizontal features (signatures of direct stimulus effects) and leaves only the diagonal features as it should. These diagonal features represent correlation of firing that goes beyond stimulus-induced modulations of single-neuron firing rates. To "explain" such residual correlation, we invoked the notion of "effective connectivity": the minimal neuronal circuit model that will reproduce the diagonal features. Note that "effective connectivity" should not be understood as a unique statement of the actual anatomic connectivity, because more than one arrangement (involving, e.g., extra interneurons) could provide the same overall behavior.

The usual approach when interpreting neuronal correlations in terms of "effective connectivity" is to separate two types of circuitry: direct interaction and shared input. In the present context, we can base this separation on the characteristics of the diagonal features. In Appendix 2 we have elaborated some of the quantitative measures that are appropriate for the case of direct neuronal interaction. We generalized the commonly used measures of "efficacy" and "contribution" to incorporate stimulus-time-locked variations. It turns out that the normalized JPSTH is the geometrical mean of these two measures. We note that the variance in the denominator of the expression for the normalized JPSTH (Eq. 9) was originally introduced for statistical reasons. It is interesting (and reassuring) that the variance now also appears as the result of probabilistic reasoning underlying dynamic generalization of efficacy and contribution.

One of the benefits of the normalized JPSTH procedure is an improved way to normalize the ordinary cross-correlogram to remove the direct effects of stimulus. We have shown that tallies over long bins parallel to the diagonal of the normalized JPSTH provide a better estimate of such a normalized cross-correlogram than possible with any procedure applied directly to the "raw" cross-correlogram. Of course, this procedure loses all stimulus-locked time structure of the correlation—it is a time average. The effects of this were clearly demonstrated in Figs. 3, 4, and 5.

Modulation of "effective connectivity"

Application of the normalized JPSTH to simulated spike trains in which we knew the underlying circuit allowed recovery of the known "effective connectivity." In particular, we showed that the normalized JPSTH has the sensitivity to detect stimulus-locked modulation of "effective connectivity," even when strongly masked by direct stimulus rate modulations. We then applied this procedure to several sets of spike trains from physiological recordings in different preparations. Without exhaustive search, we found examples of both rapidly modulated and constant "effective connectivity"; these findings occurred both for direct interaction and for shared input. We observed time constants for these stimulus-locked modulations as low as tens of milliseconds for cochlear nucleus; in cortex the time constants seem to be somewhat longer. Similar observations have recently been made in multi-neuron recordings from a number of different laboratories and preparations; the phenomenon seems to be robust across animals, brain areas, stimulus modalities, and detailed recording procedures (15 and Gerstein and Aertsen, in preparation).

We suggest that it is appropriate to incorporate rapid modulation of "effective connectivity" into brain theory and modeling. Note that there are various physiologically reasonable mechanisms for achieving such rapid modulation of "effective connectivity," in terms of both direct connections and shared input. One possible type of mechanism would involve rapidly modulated synaptic strength (24), but there are many alternative possibilities that will work with constant synaptic strength (e.g., Ref. 9). One possibility is that the observed time variations in "effective connectivity" among two neurons are due to shared input from a large network of neurons in which the stimulus induces a reverberation. The underlying system-theoretical question in choosing the proper model-to represent residual correlation is whether one regards a description in terms of a small network with varying connectivity or in terms of a large network with constant connectivity as the simpler one. Both types of description are always logically possible for data such as presented in this paper. Further theoretical and experimental work is needed to select among the various possibilities.

APPENDIX 1

Surprise

The surprise for "excitation" S(E), given a count z of m coincidences in a particular bin is defined as (27)

$$S(E) = -\ln (\operatorname{prob} [z \ge m]) \qquad (A1.1)$$

Similarly, for "inhibition" it is defined as

$$S(E) = -\ln (\operatorname{prob} [z \le m]) \qquad (A1.2)$$

Under the assumption that the coincidence counts z are approximately normally distributed, we can derive an explicit relation between the surprise values for significant "excitation" (or "inhibition") and the corresponding large positive (or negative) values of the normalized JPSTH. The assumption of normality is more realistic, the larger the number of stimulus presentations; see Ref. 27 for a discussion on this issue. In that case, we get for significant "excitation"

$$\operatorname{prob}\left[z \ge m\right] = 1 - \operatorname{erf}\left(m'\right) \tag{A1.3}$$

where erf denotes the error function, and m' is the standardized coincidence count, i.e., after subtraction of the expected value and scaling by standard deviation. For large m' this can be approximated by

$$\operatorname{prob}\left[z \ge m\right] \approx a \exp(-bm^{\prime 2}) \tag{A1.4}$$

for some appropriate constants a and b. Thus we obtain for the surprise

$$S(E) \approx bm^{\prime 2} + c \tag{A1.5}$$

for some appropriate c. Analogous reasoning for the case of significant "inhibition" leads to the same result.

Therefore, the significant values in the surprise matrix should be approximately proportional to the square of the normalized JPSTH. This relationship amounts to a "sharpening" of peaks (or troughs) in the surprise matrix, as compared with the corresponding peaks (or troughs) in the normalized JPSTH matrix, as indeed can be observed in our Figs. 4, E and G, F and H, and 5, Cand E.

APPENDIX 2

Quantitative assessment of connectivity: "efficacy" and "contribution"

In this appendix, we will generalize the concepts of *efficacy* and *contribution*, which were first introduced by Levick et al. (23). These concepts are meaningful only in a context in which one neuron drives another. We will assume that neuron a (x-axis of the JPSTH) drives neuron b (y-axis). The signature of this situation in the normalized JPSTH is an above-diagonal feature, with displacement from the principal diagonal corresponding to the connection's latency.

It will be convenient to adopt a notation in terms of probabilities rather than correlation densities. This is trivially realized, because after appropriately binning the spike trains, we obtain a (0, 1)-process, for which expectation values (correlations) equal probabilities

$$E[n_{a}(u)] = p (a = 1 \text{ in } u, u + \Delta) = p_{u}(a)$$

$$E[n_{b}(v)] = p (b = 1 \text{ in } v, v + \Delta) = p_{v}(b)$$

$$E[n_{ab}(u, v)] = p (a = 1 \text{ in } u, u + \Delta \text{ and } b = 1 \text{ in } v, v + \Delta)$$

$$= p_{w}(a, b) \qquad (A2.1)$$

As in the rest of this text, the time indices (u, v) refer to the time since the most recent occurrence of the stimulus marker.

The classical quantities to describe a direct connection are efficacy (effectiveness) and contribution (23). *Efficacy* was defined as the fraction of driver spikes that are related to the driven spikes by virtue of a precise time delay (number of near-coincidences divided by total number of driver events). *Contribution* was defined as the fraction of driven spikes that are time related to the driver spikes (number of near-coincidences divided by total number of driven events). Note that the quantities, as defined by Levick et al. (23), are numbers and not functions of time; the implicit assumption is that they are never modulated by the stimulus.

We now propose a time-dependent generalization of these notions. Because in our model we have assigned neuron a to be the driver and neuron b to be the driven, our first approximation to a time-dependent generalization is

$$e = \frac{p(a, b)}{p(a)} = p(b|a) \tag{A2.2}$$

and

$$c = \frac{p(a, b)}{p(b)} = p(a|b)$$
 (A2.3)

For brevity, we have omitted the time arguments (u, v); it is understood, however, that Eqs. A2.2 and A2.3 are restricted to (u, v)-values corresponding to the above-diagonal feature in the normalized JPSTH; this is where the model applies initially. We will lift this restriction later.

On closer inspection of Eqs. A2.2 and A2.3, it becomes clear that these expressions actually *over*estimate the magnitude of efficacy and contribution. This is most easily realized when we gradually decrease the strength of the connection; in the limit of completely independent neurons a and b, we would have

$$e = p(b|a) = p(b)$$
 (A2.4)

and

$$c = p(a|b) = p(a) \tag{A2.5}$$

which, in the case of spontaneous activity and/or other unobserved drivers, would not yield the values of zero that are intuitively appropriate for this case of "unconnected" neurons. Clearly, in the original definitions (*Eqs. A2.2* and *A2.3*), we are also counting "accidental" coincidental events: those *b*-events that would have occurred anyway near *a*-events in the absence of the (*a*, *b*)-connection. Evidently, we need a "correction" term to compensate for these "accidental" coincidences.

Let us for the moment concentrate on the efficacy; the contribution measure can be treated in completely analogous fashion, and we will give the corresponding results later. A possible estimate of the portion of Eq. A2.2 that represents accidental events could be p(b). The probability of an a-driven event would then be obtained by subtracting p(b) from the conditional probability p(b|a). This would resolve the paradoxical result of Eq. A2.4 for independent neurons; the resulting efficacy in that case would be zero. Actually, this correction is precisely what is accomplished by the subtraction of a PST-based predictor. This has been common usage in the context of usual cross-correlograms (29; for reviews see Refs. 11 and 19): to measure the truly driven time-related events, one takes the area of what is left of a peak in the cross-correlogram after subtraction of the predictor correlogram.

Once again, however, this procedure makes an overestimate, this time of the number of "accidental" coincidences; hence, we are *under*estimating the connection's efficacy. This becomes especially clear when we examine another limiting case: full dependence of b on a, i.e., all b-events are due to a-events, and there are no "accidental" b-events. In this case, taking p(b) to represent the "accidental" coincident events would in fact assign all (driven) b-events to be "accidental" and, hence, would measure the efficacy of the connection to be an erroneous zero rather than the correct positive value. In the general case, p(b) is an overestimate of having an "accidental" coincidence, because besides these "accidental" coincidences, it also incorporates the *a*-driven events.

What we actually need is an estimate for the probability p(b') of a *b*-event as it would be in the hypothetical situation that the (a, b)-connection were not present. Let us sort the total (a, b)-coincidences into two classes: those that are driven, and those that are "accidental." This corresponds to the assumption that the processes responsible for these events are additive. The probability of truly driven coincidences can then be written as p(a, b) - p(a, b'). The second term, by definition, represents independent events and can be written as a product: p(a)p(b'). To calculate the hypothetical p(b') we take the total p(b) and subtract from it the above given probability of truly driven coincidences. This leads to the following recursive relation

$$p(b') = p(b) - (p(a, b) - p(a)p(b'))$$
(A2.6)

which can be explicitly solved for p(b')

1

$$p(b') = \frac{p(b) - p(a, b)}{1 - p(a)} = \frac{p(\bar{a}, b)}{p(\bar{a})} = p(b|\bar{a})$$
(A2.7)

Thus, under the given assumptions, the hypothetical probability p(b') is equal to the measurable probability of having a *b*-event conditional on *not* having an *a*-event (i.e., \bar{a}). For the efficacy we then obtain

$$e = p(b|a) - p(b|\bar{a}) \tag{A2.8}$$

For the earlier discussed case of independent neurons, the conditional probability p(b|a) simply equals the unconditional p(b); therefore, Eq. A2.8 in this case results in an efficacy of zero, as it should. For strongly connected neurons, the second limiting case discussed earlier, $p(b|\bar{a})$ is zero, so that the efficacy is correctly p(b|a). Hence, unlike the two earlier attempts to formally define efficacy, Eq. A2.8 takes proper care of both limiting cases of very weak and very strong connections.

For the more usual intermediate case of partial connectivity, i.e., some *b*-events are directly due to *a*, and some are not, Eq. A2.8 can be further elaborated

$$e = \frac{p(a, b)}{p(a)} - \frac{p(\bar{a}, b)}{p(\bar{a})}$$

= $\frac{p(a, b)}{p(a)} - \frac{p(b) - p(a, b)}{1 - p(a)}$

Collecting terms, we obtain for the efficacy

$$e = \frac{p(a, b) - p(a)p(b)}{p(a)(1 - p(a))}$$
(A2.9)

Similarly for the contribution measure we obtain

$$c = \frac{p(a, b) - p(a)p(b)}{p(b)(1 - p(b))}$$
(A2.10)

Both of these measures can be evaluated directly from the observed spike data by using the JPSTH and the single-neuron PST histograms by using Eq. A2.1

$$e(u, v) = \frac{\langle n_{ab}(u, v) \rangle - \langle n_a(u) \rangle \langle n_b(v) \rangle}{\langle n_a(u) \rangle (1 - \langle n_a(u) \rangle)}$$
(A2.11)

$$c(u, v) = \frac{\langle n_{ab}(u, v) \rangle - \langle n_{a}(u) \rangle \langle n_{b}(v) \rangle}{\langle n_{b}(v) \rangle (1 - \langle n_{b}(v) \rangle)}$$
(A2.12)

These new measures, as all others defined in this paper, have been derived and should be evaluated bin by bin. Originally, in this section, we restricted the (u, v) time values to only cover the range of the off-diagonal feature in the normalized JPSTH. We started this appendix with the model assumption that neuron a drives neuron b with a definite timing relationship, as demonstrated by the single off-diagonal feature. The implicit assumption is that there is no other indication of driving in the normalized JPSTH: efficacy and contribution outside the feature region are understood to be zero. Because Eqs. A2.11 and A2.12 yield exactly this result, there is no reason to restrict the (u, v)-values to any particular range.

Similarly, we restricted this discussion to the case of an off-diagonal feature in the normalized JPSTH. Obviously, the extended efficacy and contribution notions (*Eqs. A2.11* and *A2.12*) could also be applied blindly to *any* normalized JPSTH matrix, including those having a feature that straddles the principal diagonal. Such a feature is the signature of shared input to the two neurons. However, in such a case, blind application of efficacy and contribution will assign meaningless parameters to an inappropriate model. Meaningful interpretation of efficacy and contribution are restricted to the case of off-diagonal features.

With Eqs. A2.11 and A2.12, we have achieved the goal of a generalized time-dependent quantitative descriptor of "effective connectivity": the evaluation across the (u, v)-plane provides a quantitative characterization of the connection, both along the diagonal (time throughout the stimulus period) and at right angles to the diagonal (different delays in the near-coincidence). From these generalized, dynamic measures of effective connectivity, we can return to time-independent measures similar to those originally proposed by Levick et al. (23). This can be done by the appropriate integrations along the diagonal in a broad enough stripe, much the same as was done in going from the normalized JPSTH to the normalized cross-correlogram (cf. Fig. 1, E and G). The result of this procedure reduces each of our dynamic quantities to a single number: its time average over the relevant range of running time u integrated across the time difference v - u.

This work was started independently in several laboratories and brought together during 1986–87 during a Sabbatical stay by G. L. Gerstein at the Max-Planck-Institute for Biological Cybernetics.

We are grateful to Prof. Valentino Braitenberg for this opportunity. T. Bonhoeffer was involved in early phases of this project. We thank V. Staiger for his skillful photography of endless series of display screens and V. Braitenberg, A. Kreiter, and A. Schüz for critical reading of the manuscript.

Portions of the work were supported by System Development Foundation Grant SDF-00013 to G. L. Gerstein) and by Office of Naval Research Grant ONR-N-00014-83-K-387 to M. Habib and G. L. Gerstein.

Address for reprint requests: A. M. H. J. Aertsen, Max-Planck-Institute for Biological Cybernetics, Spemannstrasse 38, D-7400 Tübingen, West Germany.

Received 17 February 1988; accepted in final form 9 January 1989.

REFERENCES

- AERTSEN, A. M. H. J., JOHANNESMA, P. I. M., BRUIJNS, J., AND KOOPMAN, P. W. M. Second order representation of signals. In: Localization and Orientation in Biology and Engineering, edited by D. Varjú and H.-U. Schnitzler. Berlin; Springer-Verlag, 1984, p. 38-41.
- AERTSEN, A. M. H. J. AND GERSTEIN, G. L. Evaluation of neuronal connectivity: sensitivity of crosscorrelation. *Brain Res.* 340: 341-354, 1985.
- 3. AERTSEN, A., GERSTEIN, G., AND JOHANNESMA, P. From neuron to assembly: neuronal organization and stimulus representation. In: *Brain Theory*, edited by G. Palm and A. Aertsen. Berlin: Springer-Verlag, 1986, p. 7-24.
- AERTSEN, A., BONHOEFFER, T., AND KRÜGER, J. Coherent activity in neuronal populations: analysis and interpretation. In: *Physics of Cognitive Processes*, edited by E. R. Caianiello. Singapore: World Scientific Publishing, 1987, p. 1–34.
- BACH, M. AND KRÜGER, J. Correlated neuronal variability in monkey visual cortex revealed by a multi-electrode. *Exp. Brain Res.* 61: 451-456, 1986.

- 6. BENDAT, J. S. AND PIERSOL, A. G. Random Data: Analysis and Measurement Procedures. New York: Wiley, 1971.
- VAN DEN BOOGAARD, H., HESSELMANS, G., AND JOHANNESMA, P. System identification based on point processes and correlation densities. I. The nonrefractory neuron model. *Math. Biosc.* 80: 143–171, 1986.
- EPPING, W. J. M. AND EGGERMONT, J. J. Coherent neural activity in the auditory midbrain of the grassfrog. J. Neurophysiol. 57: 1464-1483, 1987.
- ERB, M., PALM, G., AERTSEN, A., AND BONHOEFFER, T. Functional versus structural connectivity in neuronal nets. In: *Strukturbildung* und Musteranalyse. Proc. Neunter Kybernetik-Kongress (DGK), March 19-21, 1986. Goettingen, FRG (abstract), 1986, p. 23.
- GERSTEIN, G. L. AND PERKEL, D. H. Simultaneously recorded trains of action potentials: analysis and functional interpretation. *Science Wash. DC* 164: 828-830, 1969.
- 11. GERSTEIN, G. L. AND PERKEL, D. H. Mutual temporal relationships among neuronal spike trains. *Biophys. J.* 12: 453-473, 1972.
- GERSTEIN, G., BLOOM, M., ESPINOSA, I., EVANCZUK, S., AND TURNER, M. Design of a laboratory for multi-neuron studies. *IEEE Trans. Systems, Man and Cybernetics* SMC-13: 668-676, 1983.
- 13. GERSTEIN, G. L., PERKEL, D. H., AND DAYHOFF, J. E. Cooperative firing activity in simultaneously recorded populations of neurons: detection and measurement. J. Neurosci. 5: 881-889, 1985.
- GERSTEIN, G. L. AND AERTSEN, A. M. H. J. Representation of cooperative firing activity among simultaneously recorded neurons. J. Neurophysiol. 54: 1513-1528, 1985.
- GERSTEIN, G. Information flow and state in cortical neural networks: interpreting multi-neuron experiments. In: Organization of Neural Networks: Structures and Models, edited by W. von Seelen, U. Leinhos, and G. Shaw. Weinheim, FRG: VCH Verlagsgesellschaft, 1987, p. 53-75.
- 16. GLASER, E. M. AND RUCHKIN, D. S. Principles of Neurobiological Signal Analysis. New York: Academic, 1976.
- GOCHIN, P., KALTENBACH, J., AND GERSTEIN, G. Stimulus linked temporal changes in interactions of neurons in rat dorsal cochlear nucleus. Soc. Neurosci. Abstr. 12: 1269, 1986.
- HABIB, M. K. AND SEN, P. K. Non-stationary stochastic point-process models in neurophysiology with applications to learning. In: *Biostatistics: Statistics in Biomedical, Public Health and Environmental Sciences,* edited by P. K. Sen. Amsterdam: Elsevier/North-Holland, 1985, p. 481-509.
- 19. KRUGER, J. A 12-fold microelectrode for recording from vertically aligned cortical neurones. J. Neurosci. Meth. 6: 347-350, 1982.
- KRUGER, J. Simultaneous individual recordings from many cerebral neurons: techniques and results. *Rev. Physiol. Biochem. Pharmacol.* 98: 177-233, 1983.
- KUZNETSOV, P. I. AND STRATONOVICH, R. L. A note on the mathematical theory of correlated random points. *Izv. Akad. Nauk SSSR Ser. Math.* 20, 167–178 (1956). Also in: *Nonlinear Transformations of Random Processes.* New York: Pergamon, 1965, p. 101–115.
- LÉGENDY, C. R. Three principles of brain function and structure. Int. J. Neurosci. 6: 237-254, 1975.
- LEVICK, W. R., CLELAND, B. G., AND DUBIN, M. W. Lateral geniculate neurons of cat: retinal inputs and physiology. *Invest. Ophthalmol.* 11: 302-311, 1972.
- VON DER MALSBURG, C. Nervous structures with dynamical links. Ber. Bunsen-ges. Phys. Chem. 89: 703-710, 1985.
- MELSSEN, W. J. AND EPPING, W. J. M. Detection and estimation of neural connectivity based on crosscorrelation analysis. *Biol. Cybern.* 57: 403-414, 1987.
- 26. PALM, G. Evidence, information, and surprise. Biol. Cybern. 42: 57-68 1981.
- PALM, G., AERTSEN, A. M. H. J., AND GERSTEIN, G. L. On the significance of correlations among neuronal spike trains. *Biol. Cybern.* 59: 1-11, 1988.
- PAPOULIS, A. Probability, Random Variables, and Stochastic Processes. Tokyo: McGraw-Hill Kogakusha, 1965.
- PERKEL, D. H., GERSTEIN, G. L., AND MOORE, G. P. Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophys. J.* 7: 419-440, 1967.
- 30. PERKEL, D. H. AND BULLOCK, T. H. Neural coding. Neurosci. Res. Program Bull. 6: 1968.
- SCHNEIDER, J., ECKHORN, R., AND REITBÖCK, H. Evaluation of neuronal coupling dynamics. *Biol. Cybern.* 46: 129-134, 1983.