4. Dynamic Aspects of Neuronal Cooperativity: Fast Stimulus-Locked Modulations of Effective Connectivity

Ad M.H.J. Aertsen[1] and George L. Gerstein[2]

[1]Max-Planck-Institut für Biologische Kybernetik, Spemannstrasse 38, D-7400 Tübingen
[2]Dept. of Physiology, University of Pennsylvania, Philadelphia PA 19104-6085, USA

Cooperativity among neurons is usually measured by cross-correlation of their simultaneously recorded firing activities under some appropriate stimulus condition. The ordinary 'raw' correlogram presents a time averaged count of (near) coincident spikes, and in general contains contributions from different origins: stimulus-induced modulations of single neuron firing rates and various types of (possibly nonstationary) interaction between the neurons. Appropriate normalization procedures are therefore essential in order to compare results from different experiments; only after such normalization may we interpret residual correlation as indicative of (possibly stimulus dependent) 'effective connectivity' among the neurons.

Recently we reexamined the possibilities for analyzing and interpreting the time course of correlation in two neuron spike trains; in particular we developed procedures to quantify and properly normalize the Joint Peri-Stimulus Time Histogram, a stimulus-locked temporal decomposition of the ordinary cross-correlogram. This normalization allows one to separate purely stimulus-induced correlation from intrinsic interneuronal correlation; moreover it enables one to study the dynamic characteristics of the interneuronal cooperativity.

Another extension of the usual cross-correlation analysis aims at providing a global description of the correlations among larger groups of neurons. The 'gravity representation' gives us a tool that examines the entire set of observed neurons as a single entity, and indicates with high sensitivity which neurons are showing signs of temporal correlations.

These new procedures are described and illustrated by results of their application to physiological multi-unit spike trains from the cat visual cortex (area 17). We present evidence for fast stimulus-locked modulations of 'effective connectivity', pointing at dynamic cooperativity as an emergent property of neuronal assembly organization in the brain.

4.1 Introduction

Since the time of Sherrington, neurobiologists have accepted the doctrine that neurons do not act in isolation, but rather that they form into assemblies for

Springer Series in Synergetics
J. Krüger (Ed.)
Neuronal Cooperativity

various computational tasks. A number of somewhat conflicting definitions of 'neuronal assembly' have been proposed, each implying different functions and properties. Thus, we distinguish (a) Sherringtonian '*neuron pools*', defined by a shared target, (b) *cortical columns*, defined by shared stimulus preference, (c) *Hebbian assemblies*, defined as a net held together by synapses strengthened via the Hebb rule, (d) *correlational assemblies*, defined through correlated time structure in the spike trains of their member neurons; a more exhaustive discussion with many references to the original literature is given by Gerstein et al. (1989).

Two special cases of correlational assemblies can be noted: (1) *pattern generating assemblies*, defined by output of repeating timing sequences in their spike trains (Abeles and Gerstein 1988), and (2) *oscillatory assemblies*, defined by shared coherent oscillatory firing among their member neurons, possibly in response to a particular stimulus (Eckhorn et al. 1988; Gray and Singer 1989; Gray et al. 1989; see also Gerstein 1970).

Correlated spike time structure is readily interpreted as 'effective connectivity' among the observed neurons. Thus, correlational assemblies can be a subset of Hebbian assemblies. On the other hand, neither neuron pools nor cortical columns imply anything about either time structure or effective connectivity. For the rest of this chapter we will be concerned with the physiologically observed properties of correlational assemblies.

In the recent past, it has become possible, in a relatively small number of laboratories, to study directly various phenomena within neuronal assemblies. This has required simultaneous but separable observation of spike trains from as many neurons as possible (currently up to about 30) in a working brain (for reviews see Krüger 1983; Gerstein et al. 1983). The new technology involved in such experiments includes novel types of electrodes and spike shape sorting devices (Schmidt 1984a, b). Above all, it has become necessary to invent new computational tools to analyze and interpret the enormous flow of information coming from such multi-neuron experiments.

The salient and somewhat surprising result of direct assembly observation is that the effective connectivity is *context dependent* and *dynamic* on several different time scales. This suggests that the usual concept of neurons with static interconnections of fixed or only slowly changing efficacy (during learning, for example) is no longer adequate. Instead, we should distinguish between *structural* (or anatomical) connectivity on the one hand, and *functional* (or effective) connectivity on the other. The former can be described as (quasi) stationary, whereas the latter may be highly dynamic, with time constants of modulation in the range of tens to hundreds of milliseconds. It appears that dynamic cooperativity is an emergent property of neuronal assembly organization in the brain, which could not be inferred from single neuron observations. In this chapter we will document these dynamic and context-dependent assembly properties using examples drawn from cat visual cortex recordings. Similar results have been observed from a number of other preparations (Gerstein 1987; Gerstein and Aertsen, in preparation).

4.2 Computational Methods

4.2.1 Cross-Correlation Analysis

The basic tool used to examine timing relations between two spike trains is the cross-correlogram (Perkel et al. 1967). Observed peaks (or troughs) in this measurement, when contrasted with the appropriate control or predictor computations, allow us to define an 'effective connectivity' between the neurons and to parse this into (a) shared modulation by the stimulus, (b) shared input from unobserved, other neurons, and (c) direct connections between the observed neurons (Moore et al. 1970). The 'effective connectivity' is thus equivalent to the simplest possible circuit diagram that would replicate the timing relations between the observed neurons; additional interneurons can obviously not be ruled out. We emphasize that the effective connectivity is a partial subset of whatever actual anatomical connectivity exists, since it represents only the interactions and connections that were sufficiently active to be detectable at the time of observation, and those in a simplified and stylized manner.

The ordinary cross-correlogram is an average measurement over the entire length of the available data. It is appropriate to inquire whether the underlying correlation process is indeed stationary in time (as is required for a time average) or instead has some as yet unknown but potentially significant time structure. There are a number of possible approaches to this problem; here we will use only the Joint Peri Stimulus Time Histogram (JPSTH) – a temporal decomposition of the ordinary cross-correlogram – which will allow us to study the dynamical aspects of the correlation structure that are time-locked to the stimulus events.

4.2.2 Joint Peri Stimulus Time Histogram

The Joint Peri Stimulus Time Scatter Diagram is constructed (Gerstein and Perkel 1969, 1972) as a two-dimensional display where the firing of one neuron relative to each repeating stimulus occurrence is measured along the x-axis while that of the other neuron is measured along the y-axis. Points are entered on this plane at positions corresponding to all the (delayed) coincidences of the firings of each of the two neurons for each successive stimulus repetition. The resulting scatter diagram may show increased point densities parallel to the axes, representing stimulus related firing modulations of one or both neurons. There may also be increased point densities along and near the diagonal of the plane, representing correlated firing of the two neurons. With appropriate binning of the scatter diagram one obtains a matrix histogram which measures the point density as a function of the location in the scatter diagram: the *Joint Peri Stimulus Time Histogram* (JPSTH; Aertsen et al. 1989). Various summations and marginal distributions recover the two ordinary single neuron *PST histograms* (along the x- and y-axes), the *PST coincidence histogram* (along the diagonal), and finally the *cross-correlogram* (perpendicular to the diagonal). The PST coincidence his-

togram is of particular interest here, and represents the stimulus time-locked average of near-coincident firing of the two neurons in the same sense that the ordinary PST histogram represents the stimulus time-locked average of firing by each individual neuron.

These are all 'raw' measurements, i.e. they do not take into account the effects of possible stimulus modulation of the two individual neuron firing rates. Changes in these rates through the repeating stimulus cycle would necessarily lead to corresponding changes in the near-coincidence rates. For the present purpose this is an uninteresting 'background' effect that we would like to eliminate from the correlation in order to study those components which indicate temporal changes in the underlying neuronal interaction. We have devised the appropriate normalizing calculations as well as a confidence test (see Aertsen et al. 1989; Palm et al. 1988 for details). The *normalization* involves subtracting from each bin in the JPSTH plane the product of the two corresponding bins in the individual PST histograms, and dividing the result by the product of the standard deviations of these same PST bins. Subsequent summations over appropriate diagonal and para-diagonal bins in the plane give us the corrected, 'neuronal interactions only', PST coincidence histogram and cross-correlogram. Thus we can examine fast, stimulus-related changes in the interactions between two observed neurons.

We stress that neither the normalized PST coincidence histogram nor the normalized cross-correlogram could have been obtained by any correction procedure applied to the ordinary cross-correlogram. The proper normalization fundamentally requires the temporal decomposition of coincident firings as done in the JPSTH, its subsequent dynamic correction (subtraction, followed by scaling) and reintegration across or along the diagonal. The nature of these operations means that their order cannot be interchanged. Consequently, the usual procedure for correcting the ordinary cross-correlogram (i.e. subtracting the so-called 'shift predictor') would, in general, lead to different results.

4.2.3 Gravitational Clustering

Unfortunately, once we simultaneously observe even some ten neurons, there is such a combinatorial increase of pairs for ordinary or joint cross-correlation analysis (in general, N neurons give rise to $N(N-1)/2$ different pairs) that the experimenter (not the computer) is overwhelmed. The *gravity representation* (Gerstein et al. 1985a, b; Gerstein and Aertsen 1985; Aertsen et al. 1986, 1987) images each of N observed neurons as a massless particle in an N-space; particles are initially placed at the corners of a hypercube so as to be equidistant. Each particle is given a charge (like electric or gravitational charge) which is a low-pass filtered version of the spike train from the neuron it represents. At each action potential the charge is incremented, and subsequently decays with an appropriate time constant. Suppose that there is a time-varying force along the line joining any two particles that is proportional to their instantaneous charge product. Now assume a viscous medium, so that the velocity of each particle is

proportional to the vector sum of all forces upon it. All particles start equidistant from each other, but with these rules, those particles that experience a temporal correlation of high charge (because of correlated firing of the neurons the particles represent) will tend to aggregate in the N-space. Similarly, particles representing neurons which fire with less correlation than expected from their firing rates (i.e. from some form of inhibitory interaction) will tend to move apart. (For further details the reader is referred to the original literature.)

Using methods similar to those for the JPSTH (Aertsen et al. 1987), it is possible to make adjustments for direct stimulation-induced effects on the individual neuron firing rates, so that we may isolate the neural coordination effects (i.e. the effective connectivity) by observing the clustering process. Thus the gravity representation transforms temporal firing correlation among all the simultaneously observed neurons into a spatio-temporal clustering of the corresponding particles. These clusters are easily detected by standard methods, and identify any and all neuronal assemblies defined through correlated firing.

4.3 Results

4.3.1 General Results

With the above tools, we have examined multi-neuron recordings from a number of laboratories. The observations are thus drawn from a number of different preparations and sensory or association systems, and were recorded by somewhat differing techniques. (A more extensive survey of results will be published elsewhere; Gerstein and Aertsen, in preparation.) Notwithstanding the diverse origin of these data, they share a number of interesting properties regarding correlational assemblies.

1) Such assemblies are easily detected. Even with the relatively small numbers of neurons that it is currently practical to observe simultaneously, we frequently find several distinct assemblies.
2) The organization of an assembly may be context-dependent and highly dynamic. Different stimuli can create different correlational structure – interpreted as organization – among the observed neurons. The effective connectivity for some neurons (as measured by the gravity representation) can change both quantitatively and qualitatively. The membership of an assembly may change; some individual neurons may change their 'allegiance'. These measurements suggest that there is a particular organization in an assembly as averaged over many presentations of the same stimulus. But the organization may be different in a similar average over presentations of a different stimulus. Since stimuli are generally presented in an interleaved manner on a time scale of several seconds, these observations suggest that the effective connectivity (or corresponding organization) of an assembly is dynamic on this time scale. Such dynamic reorganization can also sometimes be demonstrated as a

function of recent stimulus history, so that the average organization during presentation of a particular stimulus may depend on the sequence of preceding stimuli. Note that the time scale for all these changes is far faster than that usually associated with development or learning.

3) Within the time of presentation of a single stimulus (now in the averaged sense of the JPSTH) there can be rapid modulations of effective connectivity. The time scale here can be in the tens of milliseconds.

4.3.2 Cat Area 17

We will illustrate these assembly properties with examples drawn from experiments with cat area 17 neurons performed in Krügers' laboratory (Krüger 1982). The electrodes were a 12-fold linear and organ-pipe-like array of glass-coated Pt-Ir wires, with a spacing of 160 microns. They were introduced into the cortex so that their tips fell on a line perpendicular to the cortical surface, i.e. within a single cortical column. The top electrode was near the pial surface, while the lower electrode was just into white matter.

The stimulus for these experiments consisted of a light bar (3° by 14′) moving at constant velocity in a direction perpendicular to its orientation. The distance travelled was 3° in 1.8 s, arranged to cover the entire receptive fields of the neurons under study. At the end of the movement the bar remained stationary for 0.4 s, after which it moved in the opposite direction. Finally the bar was rotated by 22.5° (duration 1 s) after which another cycle of movement back and forth, perpendicular to the new orientation was begun. This protocol was continued until, after 8 cycles of 5 s, the original orientation was reached again. The complete stimulus sequence, lasting 40 s was presented repeatedly (in the example shown, $n = 133$). We present results from an experiment in which the stimulus was presented to the ipsilateral eye only.

4.3.3 Gravitational Clustering

We first demonstrate the application of the gravity representation to these data. Figure 4.1 shows particle pair distances during the gravity condensation for the same set of 8 neurons during 25 repetitions each of two different stimuli (bar orientation 67.5°, direction of movement east-south-east vs. west-north-west, as indicated in the panels). All pairs start at the same distance at the beginning of each run, but the distance between particles of some pairs diminishes rapidly, while the distance between others remains (noisily) constant or even increases. The identities of the pairs are coded in color, according to the table shown at the top left of each panel, and are the same for both panels. The difference between the panels indicates that the correlation structure among the observed neurons was different in the two stimulus conditions, i.e. that the condensation proceded differently. The condensation was calculated with compensation for the direct effects of the stimuli on the firing of the individual neurons. Hence, we

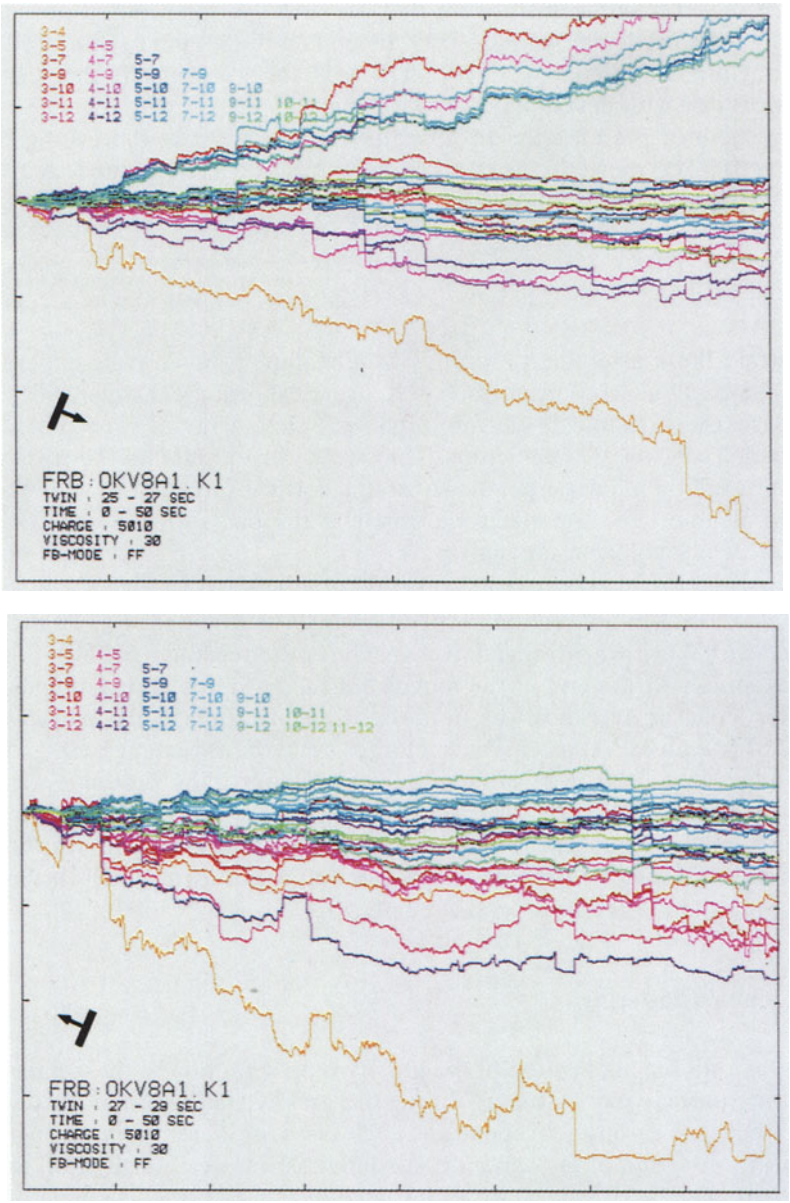

Fig. 4.1. Gravitational clustering: inter-pair distance (y-axis) as a function of time (x-axis) for the 28 pairs of neurons in an 8-neuron recording from the cat visual cortex (area 17) during repeated (n = 25) presentation of a moving bar stimulus. The two panels illustrate the dependence of condensation, and hence of assembly organization, on the stimulus condition: the bar and arrow in each of the two panels indicate the orientation of the light bar (which was the same in both cases) and the direction of motion (which was opposite). The colour coding for the pairs is displayed in the upper left corner of each panel. For further explanation see text

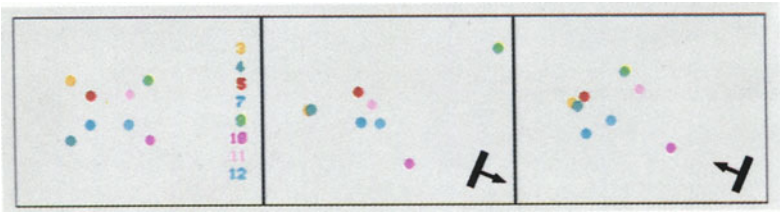

Fig. 4.2. Geometrical projection: 'snapshots' of the particle trajectories in 8-space, projected onto a 2-dimensional plane for visual inspection. The data are the same as shown for the two different stimulus conditions in Fig. 4.1. The leftmost panel shows the initial projected positions of all particles, which were the same in both cases. The colour coding for the particles, indicated in the leftmost panel, is the same in all three panels. The final projected clusterings for the two stimulus conditions are shown in the right-hand two panels. As in Fig. 4.1, the bar and arrow in each of the two panels indicate the orientation of the light bar (which was the same in both cases) and the direction of motion (which was opposite). Note that the final condensations are quite different in the two cases (and consistent with the distance curves in Fig. 4.1). For further explanation see text

may conclude that the effective connectivity among our observed neurons is different with the two stimuli.

Additional aspects of the condensation process are presented in Fig. 4.2 for the same data from two stimulus conditions. Here we see an appropriate projection from the 8-space (in which the particles actually move) to a plane which we can examine by visual means. The projection process necessarily loses information; in order to evaluate apparent clusterings in the projection, we must also examine the pair distances in Fig. 4.1. We show these projections simply as a convenient didactic tool; the clustering process should more properly be examined in the full 8-space in which the particles move, thus avoiding ambiguities and the need to consult the pair distance graphs. Initial projected positions of all particles are shown in the left hand panel of Fig. 4.2. The final projected clusterings shown for the two stimulus conditions in Fig. 4.2 (right-hand two panels) are clearly quite different (and are consistent with the distance information in Fig. 4.1). As noted before, the gravity condensation process has been carried out with compensation for the direct effects of the stimuli on the firing of the individual single neurons. Thus we again demonstrate that the effective connectivity is dependent on the stimulus conditions; the organization and membership of the neuronal assemblies that we have detected is significantly different for the two stimulus conditions shown here.

4.3.4 Joint Peri Stimulus Time Histogram

We now turn to modulations of effective connectivity on a faster time scale. Here we will examine the firing of two neurons from the same data set under one particular stimulus condition making use of the Joint Peri Stimulus Time Histogram. We will evaluate the stimulus time-locked average modulations of the correlations between the firing of these two neurons.

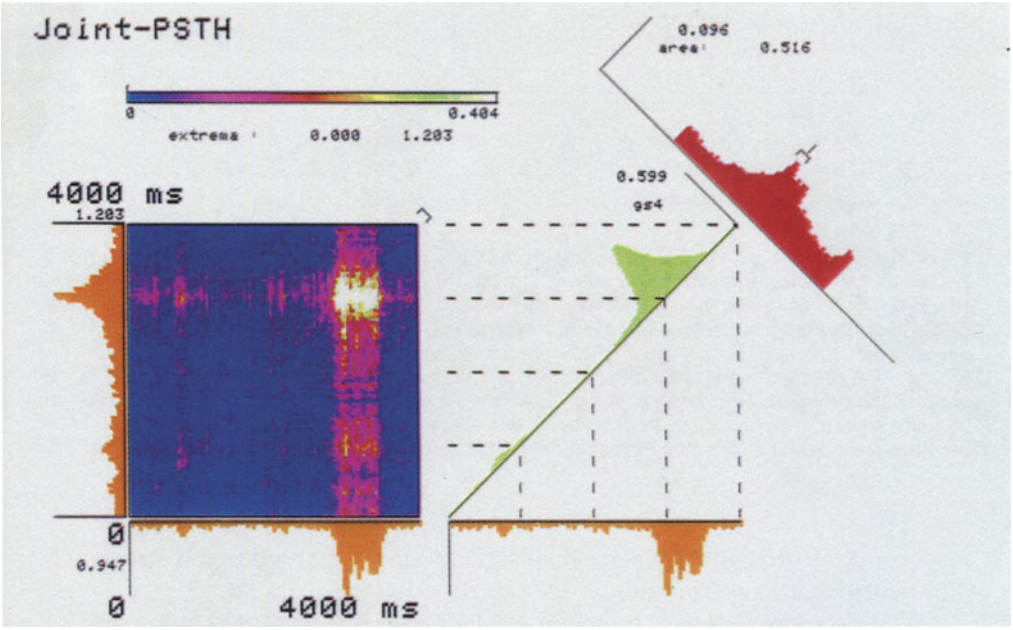
Joint-PSTH
0.096
area
0.516
0
0.404
extrema
0.000
1.203
0.599
gs4
4000 ms
1.203
0
0.947
0
4000 ms

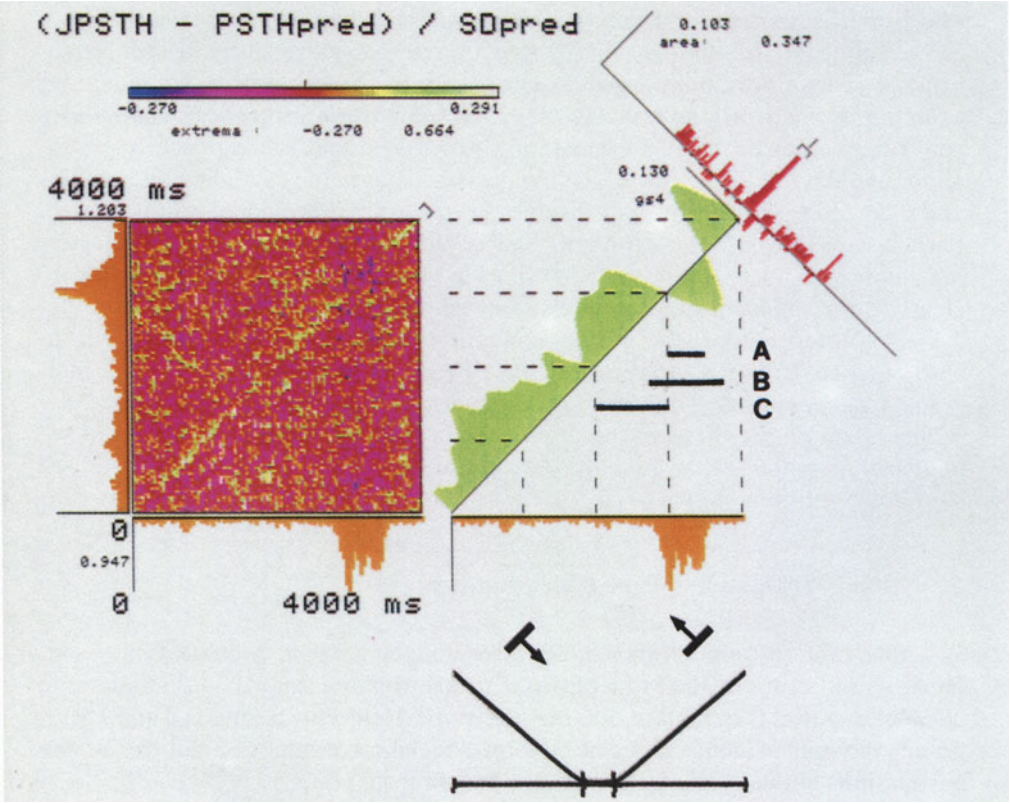
(JPSTH - PSTHpred) / SDpred
0.103
area
0.347
-0.270
0.291
extrema
-0.270
0.664
0.130
gs4
4000 ms
1.203
A
B
C
0
0.947
0
4000 ms

Figure 4.3 shows both the 'raw' (top panel) and the normalized JPSTH (bottom panel) and the associated marginal distributions for the two selected neurons. Within each panel, we show the individual PST histograms at the left of the matrix, and (upside down) twice below. The matrix is the JPSTH, with color used to indicate the density of coincident firing, according to the color bar above the matrix. At the right of each panel are respectively the PST coincidence histogram (along the diagonal rising to the right) and the cross-correlogram (perpendicular to the diagonal and descending to right). We pay particular attention here to the PST coincidence histogram which bins the counts near the diagonal of the joint matrix, and which represents the time-locked average of near-coincident firings of the two neurons in the same sense that the ordinary PST histogram represents the stimulus time-locked average of the individual neuron's firings. The time course of the light bar stimulus movement is indicated by the V-shaped line under the JPSTH matrix; the movement was perpendicular to the bar orientation, and the descending and ascending portions of the V correspond to the two directions of movement (bar orientation 45°, direction of movement south-east vs. north-west, as indicated along the V-shape).

In the 'raw' JPSTH panel at the top we note that the individual PST histograms show a strong time-locked increase in firing for both neurons as the stimulus enters their receptive fields, but only during the second direction of motion. These are direction- (as well as orientation-) selective neurons. The Joint PSTH matrix shows a considerable hill at a location matching the PST peaks, and a corresponding peak is visible in the PST coincidence histogram. In other words, individual, joint, and near-coincidence firings all are increased during a portion of the second direction of movement. In contrast, note that none of these quanti-

Fig. 4.3. Stimulus-locked dynamic correlation of neuronal firing for two neurons recorded simultaneously from the cat visual cortex (area 17) during repeated (n = 133) presentation of a moving bar stimulus. The upper panel shows the 'raw' Joint-PST histogram; the lower panel shows the normalized Joint-PST histogram, calculated by subtracting from the 'raw' JPSTH the cross-product matrix of the individual PST histograms, and by dividing the resulting difference matrix (bin by bin) by the cross-product matrix of the standard deviations of the individual PST histograms. The display format of both panels is the same. The left-hand half of the panel shows the JPSTH matrix and the two ordinary PST histograms along its x- and y-axis (binwidth: 40 ms). Values in the JPSTH matrix are displayed by using color as indicated in the color bar above the matrix. The tic mark above the color bar corresponds to the value zero. All counts were divided by the number of stimulus presentations. The right-hand half of the panel shows the PST coincidence histogram (running along the diagonal from lower left to upper right) and the ordinary cross-correlation histogram (perpendicular to the diagonal, and running from upper left to lower right). The PST coincidence histogram was smoothed using a Gaussian with a sigma of four bins; this particular value (gs4), as well as the location and width of the selected diagonal band are indicated in each panel. 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. Numbers of events: 2118 (neuron 1, x-axis), 3911 (neuron 2, y-axis), 133 (stimulus events). Spike trains were recorded during repeated presentation of a moving bar stimulus; the orientation of the bar and the direction and time course of movement are schematically indicated by the bars with arrows and the V-shaped line underneath the lower panel. In addition, the lower panel shows three different time windows (A, B and C) along the PST coincidence histogram for partial paradiagonal summation, leading to the three corresponding normalized cross-correlograms shown in Fig. 4.4. For further explanation see text

ties is increased as the stimulus passes through the receptive field in the first (and opposite) direction of movement.

The bottom panel of Fig. 4.3 contains the normalized JPSTH, i.e. corrected for the effects of stimulus time-locked modulation of the individual firing rates. We concentrate our attention on the normalized PST coincidence histogram along the diagonal. Note that during the first direction of movement, there is a positive level of residual near-coincidence firing, which is approximately constant, even when the stimulus passes through the receptive field region of the neurons. In contrast, during the second direction of movement, there is a lower than expected amount of near-coincidence firing (the correlation has gone negative) just as the individual firing rates go up in response to the stimulus. Thus, Fig. 4.3 demonstrates that near-coincidence firing can be strongly modulated by the stimulus presentation. These two neurons are repeatedly switching from a condition of ongoing co-firing to a stimulus-related period of strong anti-coincidence: when one fires, the other has a strong tendency to be more silent than expected from the individual firing rates.

In the example shown here, the switching from a condition of co-firing to one of anti-co-firing seems to be associated, at least roughly, with the transition of low to high firing rate of the two observed neurons (compare the time course of the PST histograms with that of the normalized PST coincidence histogram in Fig. 4.3). Note, however, also the brief, but strong reduction of co-firing (almost down to zero-level) slightly before 2000 ms, i.e. immediately after the stimulus bar stopped moving (as indicated by the short horizontal portion in the middle of the V-shaped line under the JPSTH matrix). Note also that in this case the modulation in co-firing was *not* associated with a corresponding change in either of the two neurons' individual firing rates. We have observed similar phenomena in other data. For example, the JPSTH in Chap. 11 by Vaadia et al. (Fig. 11.13), taken from two neurons in the prefrontal cortex of a behaving monkey, shows a strong, temporary increase from zero to a considerable degree of co-firing upon presentation of a task-related "go" signal, without any sign of firing rate modulation in the two neurons individually. We conclude that correlation of firing among neurons may change dramatically and fast, even without any appreciable sign in the probability of firing of any of the individual neurons. We have also often observed the converse phenomenon, i.e. strong modulation of single neuron firing rates without a significant modulation in the rate of coincidence firing.

The result shown in Fig. 4.3 has important implications for the use of ordinary cross-correlograms (even after subtraction of the so-called 'shift predictor'). The 'raw' cross-correlogram for the two observed neurons is shown at the extreme right of the top panel of Fig. 4.3, while its normalized version is in the corresponding position of the bottom panel. Both the 'raw' and the normalized cross-correlograms show positive peaks straddling the time origin. However, note that these correlograms are obtained by making sums in paradiagonal bins over the entire 'raw' and normalized JPSTH matrices. In the case of the strong stimulus modulation of near-coincident firing shown here, this is clearly a misleading procedure. If we choose three different time windows along the PST

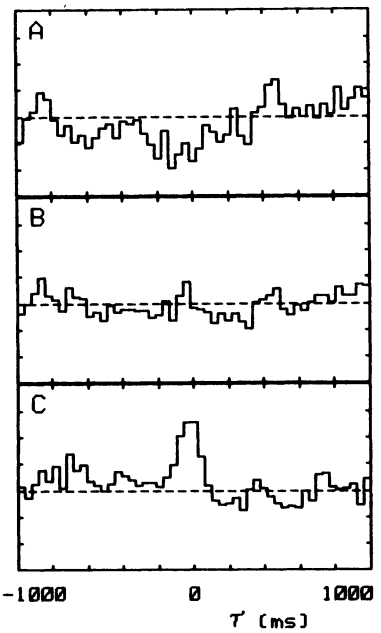

Fig. 4.4. Normalized cross-correlograms for different selections of time window over which to average the data. The three cross-correlograms (A, B and C) show the results for the corresponding time windows along the PST coincidence histogram in the lower panel of Fig. 4.3. Note the dramatic differences in correlogram shape in these three cases, due to the correlation process being dynamically modulated by the stimulus

coincidence histogram for partial paradiagonal summation (as indicated in the lower panel of Fig. 4.3), we obtain the three corresponding cross-correlograms shown in Fig. 4.4. Window A, taken during the time of strongest single neuron firing, coinciding with maximum stimulus-locked anti-correlation, produces a cross-correlogram with a strong valley, which would be interpreted as a strong *mutual inhibition* between the two observed neurons. By contrast, window B produces a flat cross-correlogram, which would be interpreted as *no interaction* between the two neurons. Finally, window C produces a strong peak in the cross-correlogram, which would be interpreted as a *shared excitation* from other, unobserved neurons. The point to be made here is that the ordinary cross-correlogram is an average measure over all the available data. If the correlation process is indeed being modulated by the stimulus (or any other source), such a global average is misleading. Rather, one should select appropriate time windows over which to average the cross-correlation measurement. These time windows, however, must be chosen intelligently, and only after examining the dynamic behavior of the correlation itself, as exhibited in the normalized PST coincidence histogram, rather than by mere inspection of the individual single unit firing rates, as measured by the PST histograms.

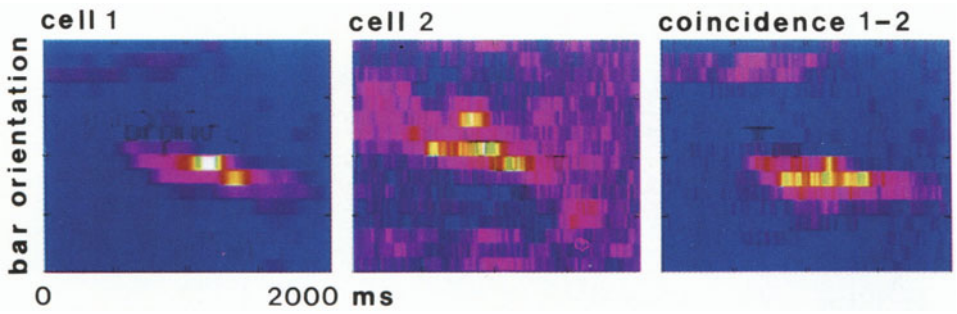

Fig. 4.5. Stimulus coding by correlated firing. The left-hand two panels show the response planes (composed of stacked PST histograms: each horizontal row corresponds to a single PST histogram, obtained for a particular stimulus condition) for the two cat area 17 neurons in Fig. 4.3. The coordinates of the plane are stimulus parameter (16 values for orientation and direction of bar movement) along the y-axis, and time relative to the stimulus along the x-axis. Color is used to indicate response magnitude (according to the same color key as used in the JPSTH in Fig. 4.3). The right-hand panel shows a comparable response plane for the *correlated* firing of these two neurons, made up of stacked PST coincidence histograms (such as in the upper panel of Fig. 4.3). Note that the spatio-temporal tuning domain of the near-coincident firing is more restricted (narrower) than the response domain of either of the two participating neurons alone

4.3.5 Stimulus Coding by Correlated Firing

Finally, we address the issue of stimulus coding by the use of correlated firing. The basic quantity to be examined is the tuning curve; we will compare the tuning of the individual neurons and the tuning of their near-coincidence firing as the stimulus parameter is varied (over bar orientation and direction of movement in this particular experiment).

Figure 4.5 shows three *response planes*; these contain averaged responses for each of the 16 different stimulus conditions (8 bar orientations, 2 movement directions). The averaged time-locked response to each stimulus condition is represented by a separate row in these planes, with color used to indicate response magnitude (according to the same color key as used in the JPSTH; cf. Fig. 4.3). The coordinates of the plane are therefore stimulus parameter along the y-axis, and time relative to the stimulus along the x-axis. Thus the plane exhibits spatio-temporal tuning domains (with spatio- referring to the stimulus parameter along the y-axis, in this case direction of bar movement) as an obvious generalization of a one-dimensional tuning curve in a compact manner.

The two left-hand response planes in Fig. 4.5 represent the stimulus time-locked activity of each of the two neurons individually; these are stacked ordinary PST histograms (cf. Gerstein et al. 1968). The right-hand panel shows a comparable plane made up of stacked PST coincidence histograms. It is clear from these pictures that the spatio-temporal tuning domain of the near-coincident firing is more restricted (narrower) than the response domain of either of the two participating neurons alone. Thus we demonstrate that near-coincident firing can have a more selective stimulus tuning than the individual

neuron's firing. Since we represent 'raw' near-coincident firing in the right panel, this sharpening of tuning is potentially available as a neuronal code for transmitting information. Although we do not demonstrate it here, the normalized version of the coincidence panel in Fig. 4.5 can be used to show the changes of effective connectivity between the two neurons as a function of both stimulus parameter and time relative to stimulus presentation in a compact manner.

4.4 Discussion

In this paper we have examined some properties of real neuronal assemblies as observed by simultaneous but separable observation of spike trains from several neurons. Using appropriate computational tools, we have been able to examine both the 'raw' correlation structure of the spike trains and the effective connectivity. The latter measure gives a simplified model of the organization among the observed neurons, and defines the assembly structure.

Our principal observation is that the correlation structure, and hence the membership as well as the internal organization of the observed assemblies, is context-dependent and dynamic on several time scales ranging from 10 ms to 10 s.

The gravity representation allowed examination of the average organization of the observed assemblies over many presentations of the same stimulus, and allowed comparison with the average organization during repeated presentation of a different stimulus. Since stimuli are usually presented in an interleaved manner, these organizational changes are occurring on a time scale of seconds.

The Joint Peri Stimulus Time Histogram allowed us to examine the stimulus time-locked variations of correlation between the firings of two neurons. Here, by means of the appropriate normalization procedure, we could demonstrate stimulus time-locked changes on a time scale of fractions of a second. Similar results on other data not shown here bring the time scale for this kind of variation of correlation down to the 10 ms range (see e.g. Gochin et al. 1990). These fast modulations of firing correlation may occur with or without associated modulations in the individual neurons' firing rates. We have pointed out that ordinary cross-correlation is an average over all the data, and that, consequently, when the correlation is varying in a time-locked (or other) way, this is at best a misleading calculation. Both appropriate and inappropriate windows with respect to the stimulus structure can be chosen, so that inspection of the PST coincidence histogram is essential before deciding on time windows for the cross-correlation measurements.

Finally we have examined tuning curve structure as parameters of the stimulus are varied, showing both single neuron responses and near-coincidence firing. In some cases the correlated firing was more sharply tuned over the stimulus parameter variation, and hence, at least in principle, might be a useful code for the processing of information.

It is appropriate to seek the mechanisms that could underlie the observed context-dependent dynamics of assembly organization. One possibility is that there are rapidly changing synapses connecting the observed neurons, and that variations in the synaptic weights [of the kind proposed by von der Malsburg (1985)] account for the dynamics. Known mechanisms like pre-synaptic inhibition or NMDA synapses provide possible substrates for this type of explanation, but leave undefined the source of the neuronal activity which controls the synaptic modulation. Alternatively, we may suppose that the synaptic weights are constant, and that we are seeing different, partly overlapping Hebbian assemblies which are differentially ignited by the stimulus context, and which have a complex internal spatio-temporal structure (see e.g. Johannesma et al. 1986). A final class of explanation also assumes that synaptic weights are constant, but that we are seeing the effects of large numbers of unobserved neurons whose summed pool activity plays on our observed neurons, and hence modulates the apparent influence that the observed neurons exert on each other. Considerations of these alternatives are beyond the scope of this chapter, and are partially examined elsewhere (Gerstein et al. 1989; Boven and Aertsen 1989, 1990; Erb et al. 1986, 1989; Bedenbaugh et al. 1988, 1990).

Acknowledgements. We wish to thank Jurgen Kruger for providing us with the cat area 17 spike recordings, and Claudia Martin-Schubert for her help in making the final pictures. GLG was partially supported through ONR N-00014-87-K-0766 and AFOSR-88-0296, AA through grants from the BMFT and from GIF.

References

Abeles, M., Gerstein, G.L. (1988): Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. J. Neurophysiol. **60**, 909–924

Aertsen, A., Gerstein, G., Johannesma, P. (1986): From neuron to assembly: Neuronal organisation and stimulus representation. In: *Brain Theory*, ed. by Palm, G., Aertsen, A. (Springer, Berlin, Heidelberg) pp. 7–24

Aertsen, A., Bonhoeffer, T, Krüger, J. (1987): Coherent activity in neuronal populations: analysis and interpretation. In: *Physics of Cognitive Processes*, ed. by Caianiello, E.R. (World Scientific, Singapore) pp. 1–34

Aertsen, A.M.H.J., Gerstein, G.L., Habib, M.K., Palm, G. (1989): Dynamics of neuronal firing correlation: modulation of "effective connectivity". J. Neurophysiol. **61**, 900–917

Bedenbaugh, P.H., Gerstein, G.L., Boven, K.-H., Aertsen, A.M.H.J. (1988). The meaning of stimulus dependent changes in cross correlation between neuronal spike trains. Soc. Neurosci. Abstr. **14**, 651

Bedenbaugh, P.H., Gerstein, G.L., Aertsen, A.M.H.J. (1990): Dynamic convergence in neural assemblies. Soc. Neurosci. Abstr. **16**, in press

Boven, K.-H., Aertsen, A. (1989): Functional connectivity in neuronal systems: fast modulation of synaptic efficacy reflects dynamics of network activity. In: Elsner, N., Singer, W. (eds.): Dynamics and plasticity in neuronal systems, 444 (Thieme, Stuttgart)

Boven, K.-H., Aertsen, A. (1990): Dynamics of activity in neuronal network give rise to fast modulations of functional connectivity. In: Eckmiller, R. et al. (eds.): Parallel processing in neural systems and computers (Elsevier, Amsterdam) pp. 53–56

Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., Reitboeck, H.J. (1988): Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analysis in the cat. Biol. Cybern. **60**, 121–130

Erb, M., Palm, G., Aertsen, A., Bonhoeffer, T. (1986): Functional versus structural connectivity in neuronal nets. In: Strukturbildung und Musteranalyse. Proc. 9th Cybernetics Congress (DGK). Göttingen, FRG, p. 23

Erb, M., Aertsen, A., Palm, G. (1989): Functional connectivity in neuronal systems: context-dependence of effective network organization does not require synaptic plasticity. In: Elsner, N., Singer, W. (eds.): Dynamics and plasticity in neuronal systems, 445 (Thieme, Stuttgart)

Gerstein, G.L. (1970): Functional association of neurons: detection and interpretation. In: The Neurosciences: Second Study Program, ed. by Schmitt, F.O. (Rockefeller University Press, New York) pp. 648–661

Gerstein, G. (1987): Information flow and state in cortical neural networks: interpreting multi-neuron experiments. In: Organization of Neural Networks: Structure and Models, ed. by Seelen, W. von, Leinhos, U., Shaw, G. (VCH, Weinheim) pp. 53–75

Gerstein, G.L., Perkel, D.H. (1969): Simultaneously recorded trains of action potentials: Analysis and functional interpretation. Science **164**, 828–830

Gerstein, G.L., Perkel, D.H. (1972): Mutual temporal relationships among neuronal spike trains. Biophys. J. **12**, 453–473

Gerstein, G., Aertsen, A. (1985): Representation of cooperative firing activity among simultaneously recorded neurons. J. Neurophysiol. **54**, 1513–1527

Gerstein, G.L., Butler, R.A., Erulkar, S.D. (1968): Excitation and inhibition in cochlear nucleus. I. Tone-burst stimulation. J. Neurophysiol. **31**, 526–536

Gerstein, G.L., Bloom, M.J., Espinosa, I.E., Evanczuk, S., Turner, M.R. (1983): Design of Laboratory for multi-neuron studies. IEEE-Trans. System, Man and Cybernetics **SMC-13**, 668–676

Gerstein, G., Perkel, D., Dayhoff, J. (1985a): Cooperative firing activity in simultaneously recorded populations of neurons: detection and measurement. J. Neurosci. **5**, 881–889

Gerstein, G., Aertsen, A., Bloom, M., Espinosa, I., Evanczuk, S., Turner, M. (1985b). Multi-neuron experiments: observation of state in neural nets. In: Complex Systems – Operational Approaches, ed. by Haken, H. (Springer, Berlin, Heidelberg) pp. 58–70

Gerstein, G.L., Bedenbaugh, P., Aertsen, A.M.H.J. (1989): Neuronal Assemblies. IEEE Trans. on Biomed. Engineering **36**, 4–14

Gochin, P., Gerstein, G., Kaltenbach, J. (1990): Dynamic temporal properties of effective connections in rat dorsal cochlear nucleus. Brain Res. **510**, 195–202

Gray, C.M., Singer, W. (1989): Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. Proc. Natl. Acad. Sci. USA **86**, 1698–1702

Gray, C.M., König, P., Engel, A.K., Singer, W. (1989): Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. Nature **338**, 334–337

Johannesma, P., Aertsen, A., Van den Boogaard, H., Eggermont, J., Epping, W. (1986): From synchrony to harmony: Ideas on the function of neural assemblies and on the interpretation of neural synchrony. In: *Brain Theory*, ed. by Palm, G., Aertsen, A. (Springer, Berlin, Heidelberg) pp. 25–47

Krüger, J. (1982): A 12-fold microelectrode for recording from vertically aligned cortical neurones. J. Neurosc. Meth. **6**, 347–350

Krüger, J. (1983): Simultaneous individual recordings from many cerebral neurons: techniques and results. Rev. Physiol. Biochem. Pharmacol. **98**, 177–233

Malsburg, C. von der (1985): Nervous structures with dynamical links. Ber. Bunsenges. Phys. Chem. **89**, 703–710

Moore, G.P., Segundo, J.P., Perkel, D.H., Levitan, H. (1970): Statistical signs of synaptic interaction in neurons. Biophys. J. **10**, 876–900

Palm, G., Aertsen, A.M.H.J., Gerstein, G.L. (1988): On the significance of correlations among neuronal spike trains. Biol. Cybern. **59**, 1–11

Perkel, D.H., Gerstein, G.L., Moore, G.P. (1967): Neuronal spike trains and stochastic point processes II. Simultaneous spike trains. Biophys. J. **7**, 419–440

Schmidt, E.M. (1984a): Instruments for sorting neuroelectric data: a review, J. Neurosci. Meth. **12**, 1–24

Schmidt, E.M. (1984b): Computer separation of multi-unit neuroelectric data: a review, J. Neurosci. Meth. **12**, 95–111