Nonlinear Systems

Nonlinear Dynamics and Neuronal Networks

Proceedings of the
63rd W. E. Heraeus Seminar
Friedrichsdorf 1990

Edited by H. G. Schuster

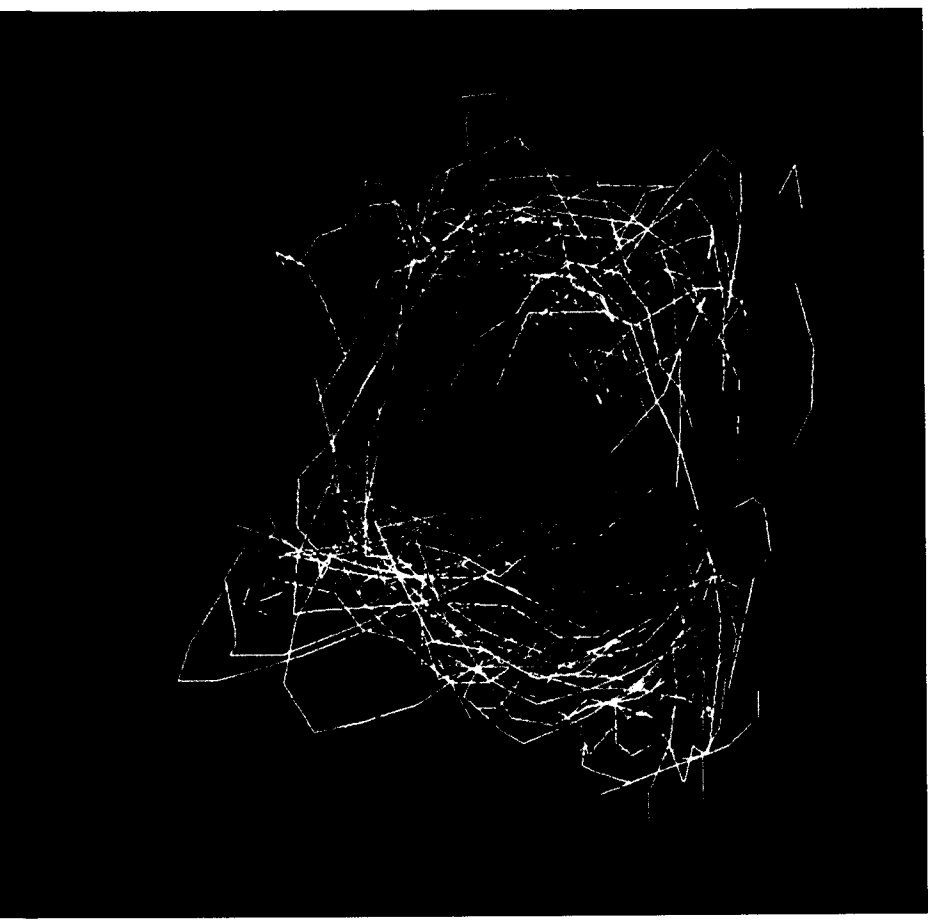

VCH

Dynamics of Activity and Connectivity in Physiological Neuronal Networks

by *Ad Aertsen* and *Hubert Preißl*

Max-Planck-Institut für biologische Kybernetik, D-7400 Tübingen, FRG

Abstract

Multiple neuron spike trains, recorded from different brain areas, were investigated using a newly developed computational tool for dynamical correlation analysis. Findings from this investigation for the first time provided experimental evidence of fast stimulus-locked modulations of neuronal interaction, indicating that the functional connectivity among neurons may be highly dynamic and context-sensitive.

Possible mechanisms underlying such dynamic network organization have been investigated by simulation and analytic approaches. In particular, we discuss the influence of the dynamics of activity in a model network on the functional connectivity among neurons embedded in this network.

In order to arrive at a quantitative description of the dynamics of neuronal activity, we investigated the applicability of methods from nonlinear dynamical system theory. We show that the 'correlation dimension' is not an appropriate tool for the characterization of spike train measurements.

1. Introduction

Within neuroscience there exists a widespread consensus that neurons do not act in isolation, but organize in the form of more or less coherently active groups or populations for the purpose of performing various computational tasks. This cooperative behavior was stressed already by Sherrington (1941), and was made more explicit by Hebb (1949), who introduced the notion of *cell assembly* for such collective groups of neurons. A number of somewhat deviating definitions of 'neuronal assembly' have been developed in the meantime, each implying somewhat different functions and properties; for a more exhaustive discussion with many references to the original literature we refer to Gerstein et al. (1989).

An essential property of the Hebbian cell assembly, induced by its 'ignition'-like activation, is the coherent nature of the group's activity, expressed in the correlated time structure of the firing patterns of its member neurons. This has prompted experimenters to try and observe the functioning of cell assemblies in the working brain by measuring the correlation among spike trains of simultaneously recorded neurons. In the recent past, it has become possible in a relatively small number of laboratories to observe such simultaneous, separable spike trains

from up to about 30 neurons (for reviews see Krüger, 1983, and Gerstein et al., 1983). These experiments involved new developments in (multi-) micro-electrodes and spike sorting devices (Schmidt, 1984a, b). Above all, they required new computational tools to analyze and interpret the enormous flow of information coming from such multi-neuron experiments.

In the present paper some of these new procedures will be described and illustrated by results of their application to physiological multi-neuron spike trains. These results for the first time provided experimental evidence for fast stimulus-locked modulations of neuronal interaction, pointing at dynamic cooperativity as an emergent property of neuronal assembly organization in the brain. Results from a theoretical investigation, aimed at unraveling the possible mechanisms underlying such dynamic network organization will be described.

2. Dynamical Properties of Functional Connectivity in the Brain

Temporal coherence among the activity of different neurons is commonly measured by cross-correlating their spike trains, recorded simultaneously under some appropriate stimulus conditions (Perkel et al., 1967). The ordinary cross-correlogram presents a time-averaged count of near-coincident spikes; peaks (or troughs) in the correlogram, when contrasted with the appropriate control calculations, are then interpreted as the signature of 'functional connectivity' between the neurons (Moore et al., 1970). In general, the cross-correlogram is composed of contributions from quite different origins, such as stimulus-induced modulations of firing rates and various types of neuronal interactions. Appropriate normalization is therefore essential when comparing results across different experiments or across different neuron pairs. Only after such normalization may we interpret residual correlation as indicative of (possibly stimulus-dependent) *functional* or *effective connectivity* among the neurons.

2.1 Joint Peri Stimulus Time Histogram (Joint-PSTH)

Recently the possibilities for normalizing and interpreting the time course of correlation among spike trains of different neurons were reexamined in detail (Aertsen et al., 1989). In particular, procedures were developed to quantify and properly normalize the Joint Peri-Stimulus-Time Scatter Diagram (Gerstein and Perkel, 1969, 1972), a stimulus-locked temporal decomposition of the ordinary cross-correlogram. This normalization allows to separate purely stimulus-induced correlation from intrinsic interneuronal correlation; moreover it enables to study the dynamic properties of the interneuronal cooperativity. The classical quantitative measures of effective connectivity (*efficacy* and *contribution*) could be generalized to incorporate stimulus-time-locked variations.

When applied to pulse trains from simple, simulated neuronal networks, these procedures allowed recovery of the underlying circuitry, including the case of stimulus-locked

modulations of effective connectivity in the presence of strong masking by direct stimulus modulations of individual firing rates. Significance of findings could be quantitatively assessed by a newly developed significance test (*surprise*; Palm et al., 1988).

We will describe the new method of the *Joint Peri Stimulus Time Histrogram (Joint-PSTH)* here briefly by illustrating its application to a piece of data from the Krüger laboratory (Krüger, 1982): spike trains of two neurons in the cat visual cortex (area 17) during presentation of a moving bar stimulus (cf. Fig.1). For a more detailed account and mathematical formulations the reader is referred to the original literature (Aertsen et al., 1989; Palm et al., 1988; Aertsen and Gerstein, 1990).

The Joint Peri Stimulus Time Scatter Diagram is constructed (Gerstein and Perkel, 1969, 1972) as a two-dimensional raster display of all the (delayed) coincidences of the firings of two neurons, with the firing of one neuron relative to each repeating stimulus occurrence measured along the x-axis, and that of the other neuron measured along the y-axis. The resulting scatter diagram, accumulated over the responses to repeated stimulus presentations, may show different kinds of 'spatial' variations of dot density. Variations parallel to the x- and y-axes reflect stimulus-related firing rate modulations of one or both neurons individually. Variations along and near the diagonal of the scatter diagram reflect modulations in co-firing of the two neurons.

The dot density can be quantified by appropriate binning of the scatter diagram. This results in a matrix histogram: the *Joint Peri Stimulus Time Histogram* (Joint-PSTH), an example of which is shown as the grey-coded matrix in Figure 1a. Various summations and marginal distributions of this matrix (cf. the corresponding 1-dimensional histograms in Fig. 1a) recover the two ordinary single neuron *PST histograms* (along the x- and y-axes), the *PST coincidence histogram* (along the diagonal), and finally the usual *cross-correlogram* (perpendicular to the diagonal). In particular, the PST coincidence histogram 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. Inspection of the results in Figure 1a reveals that both neurons show considerable modulation of their firing rates as the bar stimulus moves across the receptive field. In addition, the cross-correlogram shows a peak straddling zero, and riding on a considerable and time varying backround. Finally, the diagonal PST coincidence histogram, which is centered on the peak in the cross-correlogram, shows that there is considerable variation in the magnitude of the contributions to this peak troughout the stimulus cycle.

2.2 Functional Connectivity and Dynamical Correlation

Before we can make the conceptual transition from these observed 'raw' variations in co-firing to modulations in functional connectivity, we now have to take into account the effects of stimulus-induced modulations of the two individual neuron firing rates. The appropriate

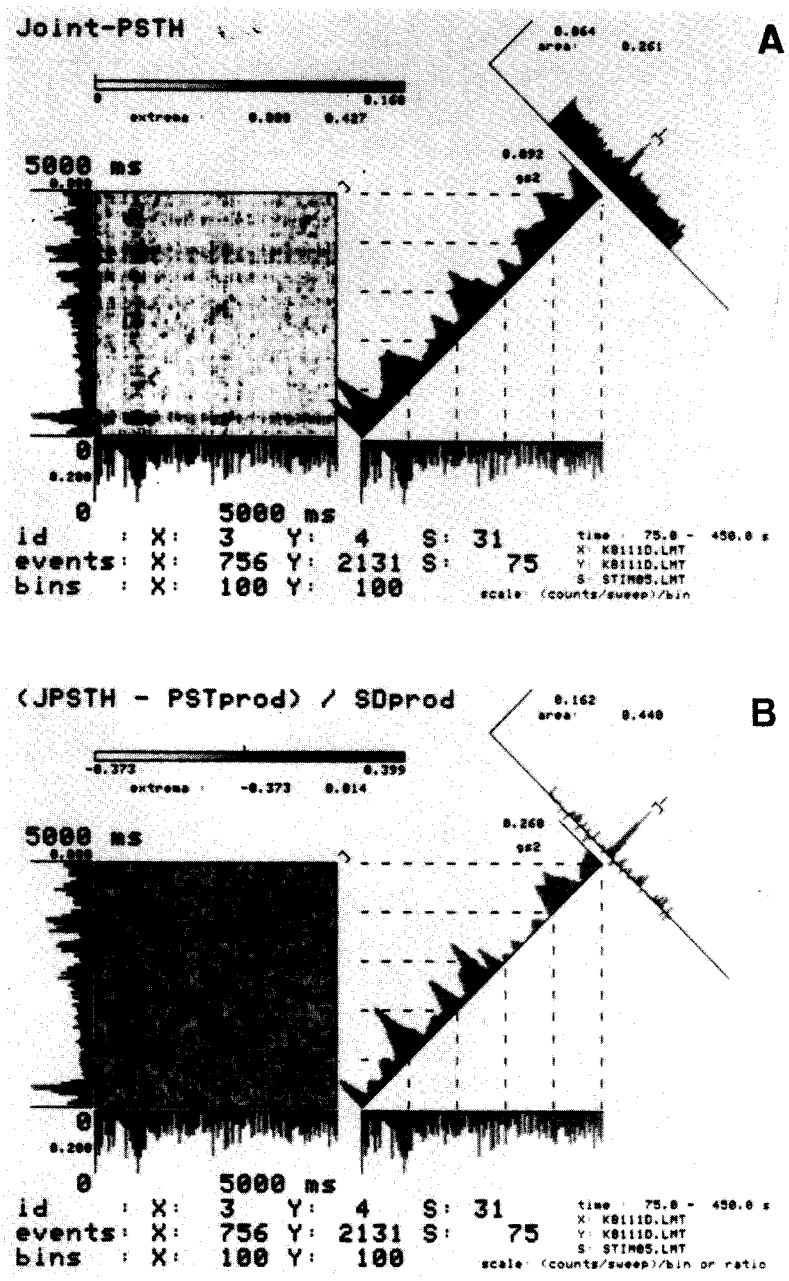

Figure 1. Stimulus-locked dynamic correlation of neuronal firing for two neurons recorded simultaneously from the cat visual cortex (area 17) during repeated presentation of a moving bar stimulus. The upper panel (a) shows the 'raw' Joint-PST histogram and the various associated marginal histograms; the lower panel (b) shows the Joint-PST histogram after normalization for stimulus-induced nonstationarities in the single neuron firing rates. For further explanation see text.

normalization consists of subtracting from each bin in the Joint-PSTH matrix the cross-product of the two corresponding bins in the individual PST histograms, and dividing the result by the cross-product of the standard deviations of these same PST bins (Aertsen et al., 1989). Subsequent summations over appropriate diagonal and para-diagonal bins in the plane give us the normalized PST coincidence histogram and cross-correlogram, describing 'neuronal interactions only'. Only now may we interpret variations in the residual correlation as stimulus-related modulations of the interactions between the two observed neurons.

This normalization procedure was applied to the data in Figure 1a; the results are shown in Figure 1b. One observes how horizontal and vertical features have been removed, to leave only a more or less uniform background around zero level, against which we discern one prominent, diagonal feature. Consequently, the cross-correlogram shows a clear peak around time shift zero, while all background correlation has been removed by the normalization. Interestingly, the diagonal histogram in Fig. 1b demonstrates that near-coincidence firing is strongly modulated as time proceeds through the stimulus cycle, with time constants in the range of fractions of a second. Clearly these two neurons are repeatedly switching, in a stimulus-locked fashion, from a condition of incoherent firing to a condition of coherent firing. Moreover, the time course of this modulation of coherence could not be inferred from the time course of firing rate modulation of either of the two neurons alone.

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 technical means. (A more extensive survey of results will be published elsewhere; Gerstein and Aertsen, in preparation.) Notwithstanding the diverse origin of these data, we have observed similar phenomena in a surprisingly large number of cases (e.g. Aertsen and Gerstein, 1990; Gochin et al., 1990; Vaadia et al., 1990). We conclude that correlation of firing among neurons may change dramatically and fast, not necessarily in parallel with firing rate modulations of any of the individual neurons involved, and sometimes even without any appreciable sign in the probability of firing of any of the individual neurons.

In addition, we observed (e.g. Aertsen et al., 1987; Aertsen and Gerstein, 1990) that different stimuli can create different correlational structure - interpreted as organization - among the observed neurons. The effective connectivity can change both quantitatively and qualitatively with stimulus context. Thus, the membership of an assembly may change; some individual neurons may change their 'allegiance'. These measurements suggest that the organization in an assembly, measured by averaging over many presentations of a particular stimulus, may be quite different in a similar average over presentations of a different stimulus.

Summarizing, the salient and somewhat surprising result of our correlation analysis of data from a number of different multiple-neuron experiments aimed at direct assembly observation, is that the effective connectivity among neurons in the physiological network is

context-dependent and *dynamic* on several different time scales, with time constants as low as tens to hundreds of milliseconds.

3. Dynamical Properties of Functional Connectivity in Neuronal Network Models

Two types of explanation for this physiologically observed phenomenon of context-dependent, dynamic assembly organization can be envisaged: a *local* one and a *global* one. On the one hand the synapses might actually modulate their efficacy on such a short timescale (as was originally proposed by von der Malsburg (1981, 1986; see also Bienenstock and Doursat, 1990). On the other hand the modulation of 'effective connectivity' might be an emergent property of dynamic cooperativity in the brain, i.e. a reflection of the dynamics of activity in the entire network, without any rapid changes at the synaptic level taking place at all (Erb et al., 1986). Obviously, these two explanations are not mutually exclusive. In order to investigate the latter possible explanation, i.e. the influence of network activity on the efficacy of intrinsically constant connections, we have studied the behaviour of artificial neuronal networks under different activity conditions, both in simulation and with analytic methods.

3.1 Context-Dependence of Functional Connectivity in an Interconnected Network Model

In order to investigate the question of context-dependence in artificial neuronal networks, we simulated a network of 100 spiking neurons with fixed synaptic connections (Erb et al., 1989). In this model the neurons are connected by excitatory synapses, inspired by neuroanatomical findings that as much as 80% of the cortical synapses are of this type (Braitenberg, 1978). The model also comprises an inhibitory mechanism for stabilization. The synapses are modeled as lowpass filters with delayed response, causing the incoming spike activity at a synaptic junction to be transformed into EPSP's. For each of the neurons the contributions from the various spike trains arriving at the synaptic sites distributed over the dendritic tree are linearly integrated to yield the membrane potential at the cell body. The next stage consists of a stochastic spike generating mechanism, with the probability of firing described by a sigmoid function of the membrane potential; in addition it is modulated by a refractory mechanism. The inhibitory mechanism consists of two parallel, linear branches: a fast one, with a time constant in the range of that for the excitatory synapses, and a slow one with a much larger time constant. The latter branch serves to regulate the global activity in the network towards a preset value (threshold control; Braitenberg, 1978; Palm, 1982). In order to mimic cortical physiology, activity in the network was regulated to a typical level of sparse firing (see Erb et al. (1990) for some interesting problems concerning stability and oscillations in such a network). From this model network 16 neurons were selected for closer analysis, much in the same way a physiological multi-neuron recording would pick up only a fraction of the entire population of neurons. Figure 2a shows the matrix of *structural connections* among the 16 model neurons: the grey value of a particular matrix entry with index (i,j) represents the strength of the synaptic connection from neuron i to neuron j, with darker grey corresponding to stronger connectivity.

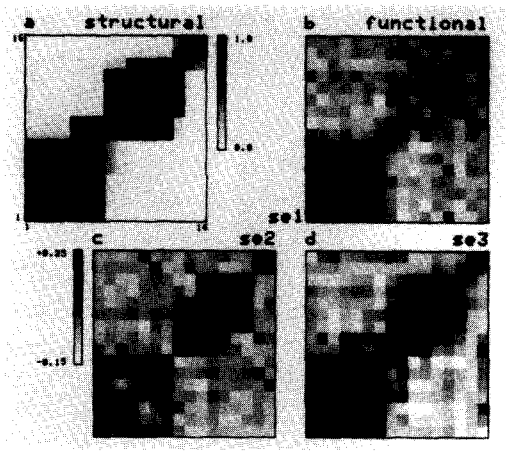

Figure 2. Simulated model network of 100 neurons. (a) Matrix of synaptic connections among 16 selected neurons. (b-d) Functional connectivity among these 16 neurons, measured during presentation of three different stimulus ensembles. Further explanation in text.

Three different stimulus ensembles were presented to the net; in each 'experiment' we recorded the spike trains from the 16 model neurons and used the normalized correlation to measure the 'effective connectivity' among them. Rather than displaying the Joint-PSTH's for each possible pair of neurons, we used a more global measure of functional connectivity, which is analogous to collapsing the diagonal trace in the normalized Joint-PSTH matrix for any particular neuron pair into a single number, the time-averaged effective connectivity (i.e. averaged over the duration of a stimulus cycle; see Aertsen et al., 1987 for more details). These scalar measures for interneuronal connectivity for all possible pairs among the 16 model neurons were then arranged into the form of a matrix, the *functional connectivity matrix*, comparable in format to the structural connectivity matrix in Fig. 2a. In the context of the present Section, where we study the dependence on stimulus context rather than the detailed dynamics, this functional connectivity matrix provides an adequate overall measure of coupling among larger groups of neurons.

Functional connectivity matrices were determined for each of the three sets of 16-neuron spike trains recorded during presentation of the three stimulus ensembles; results are shown in Figures 2b-d. A comparison with Figure 2a clearly shows how, depending on the stimulus-context, different parts of the underlying anatomical network structure become functional, whereas other parts, in spite of their constant and strong structural connectivity remain functionally 'invisible'. Stimulus ensemble 1 (se1) mainly shows connections among the first group of neurons (1-7; Fig. 2b), se2 predominantly shows interactions in the second group (8-14), with weaker signs of the first and third (15,16) groups (Fig. 2c), and, finally, for se3 groups one and two are dominant, with emphasis on group one (Fig. 2d).

We conclude from these simulations that even a net with fixed synaptic connectivity can show strong context-dependence of effective connectivity. The functional reorganization of the network demonstrated here does not have to invoke additional local mechanisms (such as the rapid synaptic modulation referred to earlier), but instead is presumably caused by the

stimulus dependence of the flow of activity through the entire network. This issue will be investigated in more detail in the following Section.

3.2 Effective Connectivity as a Function of Network Activity: Simulation Results

In an attempt to unravel the mechanisms underlying the dynamic changes of functional connectivity, we performed additional, more specifically designed simulations and analytical calculations (Boven and Aertsen, 1990). A scheme of the network we analysed is shown in the inset in Figure 3a. A spontaneously active neuron (1) drives a second neuron (2), which gets additional input from a pool of N independent, spontaneously firing neurons. All synaptic connections α and β have fixed and moderately weak strength, i.e. spikes arriving at any particular junction give rise to EPSP's with constant, sub-threshold magnitude. We studied the behavior of the efficacy of the connection between neurons (1) and (2) as a function of the activity arriving at (2) from the pool, for constant connectivity α and different values of pool coupling β.

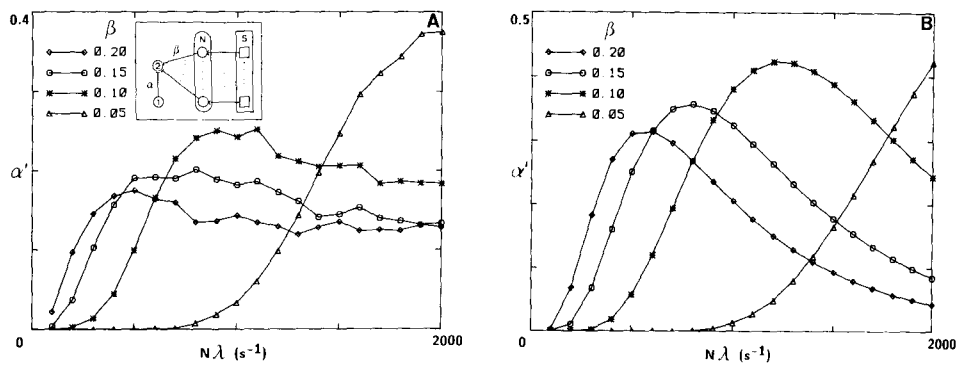

Figure 3. Effective connectivity α' as a function of network activity $N{\cdot}\lambda$, for different values of pool coupling β. (a) simulation results, (b) analytic results.

Each of the neurons is represented as a nonlinear system described by four state variables, the properties of which are determined by specifying conductances, thresholds and time constants (MacGregor, 1987). This representation allows us to monitor both the neuron's membrane potential and its spike firing behaviour. The mean firing rate λ of all spontaneously active neurons was set to about 10 spikes/sec in order to mimic the typically low firing rates measured in cortex. The overall activity from the pool could be varied by changing either the number of neurons N or the firing rate λ of each of the member neurons. The latter manipulation mimics the effect of driving the pool with a stimulus S. The effective connectivity α' between neurons (1) and (2) was measured by crosscorrelating the spike activity from the two: α' equals the number of correlated events N_c (i.e. the net area of the peak in the correlogram) divided by the number of presynaptic events N_1 (Aertsen and Gerstein, 1985). Curves of the behavior of the effective connectivity α' as a function of pool activity $N{\cdot}\lambda$ for various choices of pool coupling β are shown in Figure 3a.

Two observations can be made:

1. The efficacy α' varies strongly with pool activity $N{\cdot}\lambda$, even though the synapse itself is kept at a fixed strength α (= 0.2) throughout all simulations. With increasing pool activity the efficacy of the connection initially increases strongly to reach a maximum, after which it slowly decays again. Dynamic changes of pool activity, reflected in excursions along the horizontal axis of Figure 3a, would thus result in equally fast, but distorted excursions along the vertical axis, signifying corresponding changes in effective connectivity.
2. The dynamic range of the efficacy α' increases with decreasing pool coupling β. Moreover, for smaller β the maximum of the effective connectivity α' is actually higher, but is only reached at larger values of the pool activity. These findings are especially relevant for large, weakly coupled nets, of which the mammalian cortex is a typical example.

3.3 Regulation of Effective Connectivity by Modulation of Network Activity: a Simple Analytic Model

The explanation of the observed phenomena is actually quite simple. Due to the subthreshold nature of the EPSP's, action potentials from neuron (1) by themselves will not reach to make neuron (2) fire. For this they need the coincident arrival of spikes at the other synapses of neuron (2). These are provided by the activity reaching (2) from the remainder of the network. The pool activity thus provides a background level, which, depending on its magnitude, will make activity from the driver neuron (1) more (or less) viable in eliciting activity from its target neuron (2). This effect of the network activity can be described as a control mechanism determining the operating point of the receiver neuron, and, thereby, the efficacy of the otherwise subthreshold connections. This also explains our second observation, i.e. why the dynamic range increases with decreasing pool coupling. Due to the stochastic and pulse-like nature of the activity arriving at the target neuron, the background level of the membrane potential created by the pool activity will be fluctuating, and hence occasionally reaching threshold, causing neuron (2) to fire. Since these 'accidental' firings do not require coincident arrival of spikes from neuron (1), they do not contribute to the efficacy of the (1-2) connection. On the contrary, they are rather counter-productive since they cause spikes from (1) to arrive more often at its target when this is in a non-responsive state. The size of the membrane potential fluctuations, and, hence, the amount of occasional firing of (2) is proportional to (amongst others) the strength of the pool coupling β. This implies that the smaller this coupling, the closer one can regulate the operating point of the target neuron towards a condition of firing upon arrival of activity from the driver neuron, without the danger of too many non-contributing, accidental firings. Rather than a 'noisy' control parameter, one thus obtains a smooth, and thereby more influential regulation of synaptic efficacy.

In order to investigate this potential mechanism more quantitatively, we analysed a simplified model by analytical means. An explicit expression for the effective connectivity can be obtained by assuming (1) that the membrane potential at the site of spike initiation is the

result of linear, spatio-temporal integration of incoming spike trains, and (2) that a spike is generated whenever this membrane potential crosses a fixed threshold level Θ. The contribution to the membrane potential of neuron 2, due to the pool of N independent Poisson processes, each with constant rate λ, connected to neuron 2 with strength β can then be described as 'shot-noise' (i.e. a linearly filtered (time constant τ) Poisson process). For reasonably large N, the resulting distribution of membrane potential values can be approximated by a Gaussian distribution with a mean value $\mu = N \cdot \lambda \cdot \beta$ and variance $\sigma^2 = N \cdot \lambda \cdot \beta^2 / 2\tau$. Following Abeles (1982), the firing rate λ_2 of neuron 2 can then be calculated as:

$$\lambda_2 = 1/\sqrt{2\pi} \int_{(\Theta-\mu)/\sigma}^{\infty} \exp(-y^2/2)\, dy$$

The effective connectivity α' from neuron 1 to neuron 2 equals the increase in firing rate λ_2 upon the arrival of a spike from neuron 1, i.e. upon a shift of the membrane potential over an amount α:

$$\alpha' = 1/\sqrt{2\pi} \int_{(\Theta-\mu-\alpha)/\sigma}^{(\Theta-\mu)/\sigma} \exp(-y^2/2)\, dy$$

The curves corresponding to this analytic expression for the effective connectivity α' are shown in Figure 3b. These curves exhibit essentially the same behavior as the simulation results in Figure 3a: an initial fast increase until a maximum is reached, followed by a slower decay. Also the dynamic range of the effective connectivity shows the same dependence on strength of pool coupling. We conclude that, even with the simplifying assumptions in this analytic model, we have captured the essence of an activity-related mechanism for dynamically linking neuronal groups: it is based on activity variations in the entire network being 'projected down' onto the connections among the recipient neurons.

3.4 Regulation of Effective Connectivity by Modulation of Network Activity: Dynamical Correlation

The above described dynamic effect of variations in network activity on the functional connectivity can also be shown more directly by rapidly varying the pool activity as a function of time, and performing a Joint-PSTH analysis of the resulting spike trains from model neurons 1 and 2. The pool neurons were driven by a periodic stimulus S, causing them to fire at a stimulus-locked time-varying rate. Stimulation parameters were selected such, that the spike train of each of the pool neurons could be described by a rate-modulated Poisson process and that the Poisson processes of different neurons remained independent. The dynamical correlation between the spike trains from neurons 1 and 2 was analysed using the Joint-PSTH; results are shown in Figures 4a ('raw' correlation) and 4b (normalized

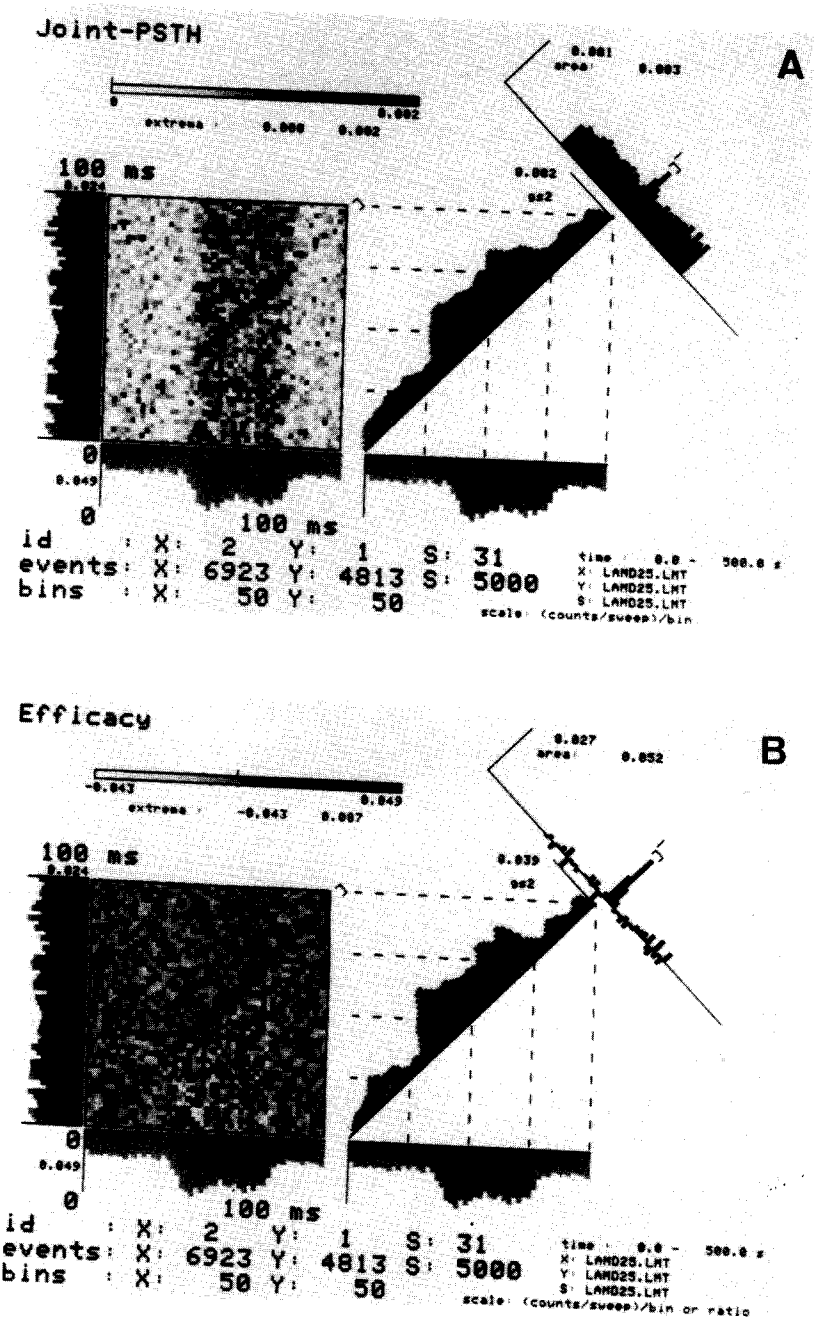

Figure 4. Stimulus-locked dynamic correlation of neuronal firing for neurons (1) and (2) from the model network, shown in the inset in Fig. 3a. As in Figure 1, the upper panel (a) shows the 'raw' Joint-PST histogram and the various associated marginal histograms; the lower panel (b) shows the Joint-PST histogram after normalization for stimulus-induced nonstationarities in the single neuron firing rates.

correlation). Not unexpectedly, the time variation of the pool activity is reflected in a similar modulation of the firing rate of neuron 2 (cf. the PSTH along the x-axis). The 'raw' Joint-PSTH matrix shows vertical as well as diagonal features. Similar to the case in Figure 1, the vertical features can be removed by appropriate normalization of the correlation for variations in firing rate of both neurons 1 and 2. This leads to the *dynamic synaptic efficacy* (cf. Fig. 4b), a measure which is closely related to the normalized Joint-PSTH (Aertsen et al., 1989). As in Figure 1b, only the diagonal feature remains after normalization. Variations in the time course of the 'effective connectivity' reveal that - even though the synaptic coupling between neurons 1 and 2 was kept constant throughout the entire simulation - the modulation of the pool activity results in a strong modulation of the synaptic efficacy of the (1-2) - pathway, with a time constant of only a few milliseconds. The numerical values of the efficacy along its time course closely follow the values predicted by the statically determined curves in Fig. 2a.

The principal cause of the above described dynamic effects resides in the time modulations of the magnitude of the pool activity projecting onto the target neurons. In another, related simulation study it could be shown that very similar and equally rapid effects of pool activity on functional connectivity among target neurons can be obtained by dynamically manipulating the internal correlation structure of the pool activity, while keeping the magnitudes of the pool members' firing rates constant (Bedenbaugh et al., 1988, 1990; Gerstein et al., 1989). Summarizing, we conclude that the demonstrated effect of fast, activity-related changes of synaptic efficacy induced by changes in activity in the entire network indeed provides a possible mechanism underlying the physiologically observed, rapid stimulus-locked modulations of firing correlation in multi-neuron activity.

4. Characterization of the Dynamics of Neuronal Activity

These results once more emphasize the necessity of a quantitative description for the dynamics of neuronal activity. In this context we investigated the applicability of methods from nonlinear dynamical system theory (Preißl et al., 1990). In recent years there has been a growing interest in the description of dynamics in complex systems, both at the theoretical and experimental level. It was shown that one can distinguish between stochastic and deterministic, non-periodic (*chaotic*) processes, the latter characterized by a *strange attractor*; various measures for the characterization of chaotic systems were developed (for a review see Bergé et al., 1986).

4.1 Chaotic Dynamics and Fractal Dimensions

Since the work of Lorenz (1963) it is known that non-linear, dissipative dynamical systems can exhibit chaotic behavior. The attractor (the post-transient state) of such a system may be characterized by a non-integer or *'fractal'* dimension, in which case it is called a *'strange attractor'*. The most popular among the various current definitions for fractal dimensions is

the *correlation dimension*. This can be determined using an algorithm developed by Grassberger and Procaccia (G-P) (1983a, b), based on considerations by Renyi (1962). It involves calculating the correlation integral $C(\varepsilon)$, which determines the number of state vectors which fall within a region with size ε. For small ε, $C(\varepsilon)$ behaves like a power of ε: $C(\varepsilon) \sim \varepsilon^r$. The exponent r is called the 'correlation dimension'; it can be determined as the slope of the *ln* $C(\varepsilon)$ vs. *ln* ε curve in the so-called 'scaling region' (sr), i.e. where this curve is linear.

In most experimental situations the differential equations for the system under investigation and, consequently, the state variables, are unknown. The experimenter usually only has access to consecutive measurements of a single scalar observable y(t). A considerable step forward was made when it was shown that the attractor of the underlying system can be reconstructed from such a time series by a procedure known as *'embedding'* (Takens, 1981; Packard et al., 1980). A vector in an n-dimensional *'pseudo state space'* is defined from the measurement by taking an appropriately spaced 'comb' of values $(y(t), y(t+\tau), ..., y(t+(n-1)\tau))$; different values of t correspond to different vectors in that space. The *'embedding theorem'* of Takens (1981) states that, if y is a smooth function mapping the original attractor to R, and if $n \geq 2m+1$ (m is the dimension of the attractor), the reconstructed attractor in pseudo state space is diffeomorphic to the original attractor: both have the same dimension.

The embedding theorem and the G-P algorithm provide the means of analysing the dynamics of a system under experimental observation. One records a time series, embeds it in an n-dimensional space (usually starting with $n=1$) and determines a 'dimension' d with the G-P algorithm. This procedure is repeated with increasing n. Initially the 'dimension' d will increase with n. However, if the underlying system has a low-dimensional attractor with a correlation dimension r, the 'dimension' d will saturate at this value for increasing n. In the scaling region the *ln* $C(\varepsilon)$ vs. *ln* ε curves for different n will consequently become parallel.

4.2 Fractal Dimensions and Spike Train Sequences

Since it was not obvious how to apply these methods to a system of multiple channel recordings of activity from groups of neurons, we decided to restrict our analysis first to the description of single neuron dynamics. Already at the onset an interesting issue arises. The dynamics of the activity from a single neuron can in principle be determined from two different types of measurement: the membrane potential, which is a *continuous* signal, and the train of action potentials (spikes), described as a *point process* (cf. Figure 5). The first can be measured with intracellular electrodes, the second also by the more common (and less tedious) technique of extracellular recording. Most of single neuron electrophysiology is in fact based on such extracellular spike train recordings. Since it was not clear whether the choice between the continuous signal and the pulse train might influence the outcome of the analysis, we decided to investigate this problem in more detail. To this end we analysed several well-known dynamical systems (Lorenz, Rössler, Henon), each one with a strange attractor and chaotic dynamics, as well as membrane potentials and spike trains derived from

Figure 5. The membrane potential of a single cell with clearly visible action potentials. Two types of data can be obtained from such a measurement: the continuous time course of the membrane potential, and a point process, specified by the time intervals between successive action potentials.

a neural network simulator. The focus in this study was on the issue of continuous versus discrete time series observations.

In order to obtain the two types of data for each of the various systems, we recorded the time course of one of the system variables (continuous signal) as well as the series of time intervals between successive, positive-going level-crossings of a fixed amplitude threshold (pulse train). This procedure of obtaining both continuous and pulse train signals is illustrated in Figure 6. Note that the latter procedure emulates the usual technique of extra-cellularly recording single neuron spike trains. For the pulse trains, the time intervals between the events were regarded as the dynamic system variable. In the following we will focus on the results obtained for the Lorenz system. From the continuous measurement (the time course of the z-component) we obtained the correct correlation dimension of 2.06 (Grassberger and Procaccia, 1983a, b). Rather to our surprise, however, the pulse train measurements gave quite different results, as is illustrated in Figure 7 for 4 different values of the threshold Z. Only for $Z=35$ (Fig. 7a) and $Z=25$ (Fig. 7b) were we able to determine a single correlation dimension: for $Z=35$ it was 1.92, whereas for $Z=25$ it amounted to 1.79. Clearly these two values differ from each other as well as from the correlation dimension of the continuous system. With a threshold at $Z=15$ (Fig. 7c) one obtains two scaling regions (sr1, sr2); this is even more obvious with a threshold at $Z=12$ (Fig. 7d). For $Z=15$ the correlation dimension equals 1.52 in sr1, whereas in sr2 it cannot be determined at all: the 'dimension' d continues to grow with the embedding dimension. For $Z=12$ neither of the two scaling regions allows to determine a correlation dimension.

We made similar observations for the other (x and y) components of the Lorenz-system, as well as for the Rössler- and the Henon-system. In addition we considered other types of pulse train measurements, e.g. the sequence of intervals between maxima in a selected component, and the sequence of intervals between consecutive points of entry into an arbitrarily selected box in pseudo state space. In *none* of these cases was it possible to obtain the correct correlation dimension of the attractor from analysis of pulse train measurements. Finally we analysed both the continuous membrane potential and the simultaneously recorded spike

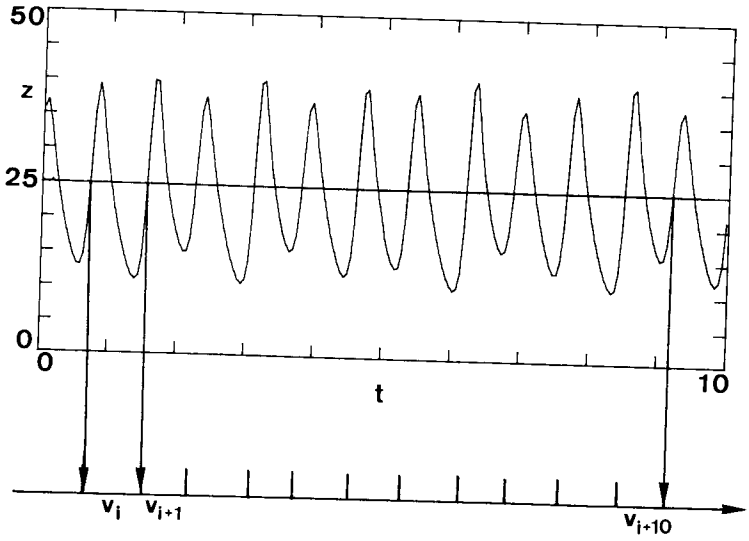

Figure 6. Generation of pulse train observations from the z-component of the Lorenz system: the dynamic variable for the reconstruction procedure is the sequence of time intervals between successive, positive-going level crossings at a fixed threshold value Z=25. Note that in this case the vectors are defined by taking sequences of adjacent intervals.

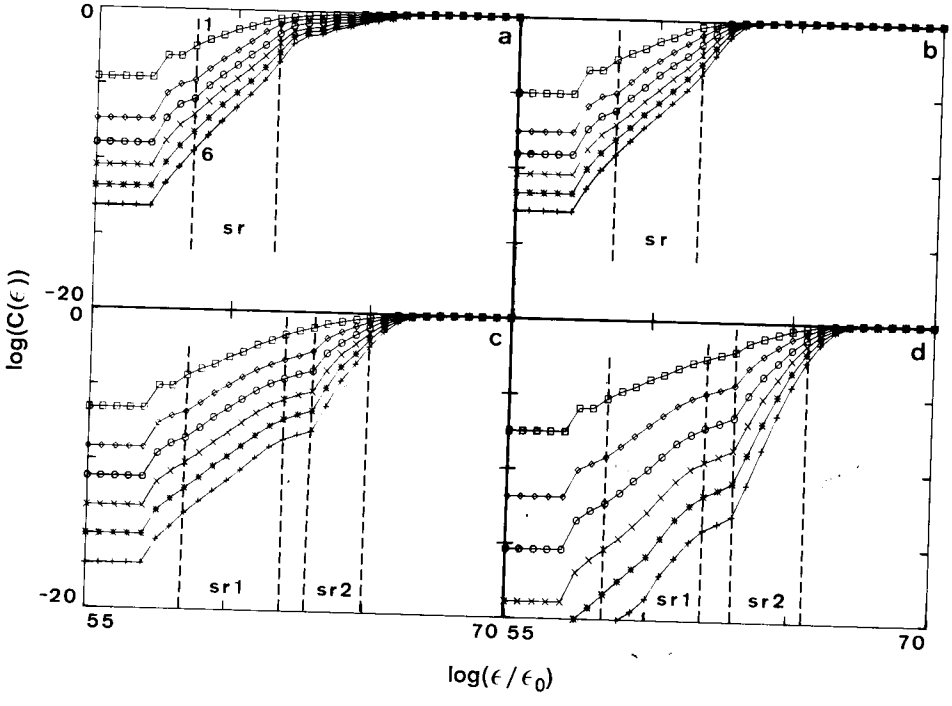

Figure 7. Results of dimension analysis for pulse trains obtained with different thresholds Z on the z-component of the Lorenz system: Z=35 (a), Z=25 (b), Z=15 (c), and Z=12 (d). In each case the reconstruction was performed with 5000 intervals; the embedding dimension ranges from 1 to 6.

train from a neuronal network simulator which employs a pseudo-random number generator (Boven and Aertsen, 1990). Again the two different types of measurements gave different results for the correlation dimension.

We conclude that pulse train measurements (like spike trains from single neurons) do *not* allow the characterization of an attractor, even if the underlying system is a deterministic dynamical system: the attractor reconstructed from the pulse trains is not diffeomorphic to the underlying attractor.

5. Discussion

In this paper we have examined some properties of physiological neuronal assemblies as observed by simultaneous, separable recordings of multiple neuron spike trains. Using newly developed computational tools, we have been able to examine both the 'raw' temporal correlation structure of the spike trains and, derived thereof by appropriate normalization, the effective neuronal connectivity, The latter measure gives a simplified model of the functional coupling among the observed neurons, and defines the assembly organization.

5.1 Functional versus Structural Connectivity

The Joint Peri-Stimulus Time Histogram enabled us to examine the stimulus time-locked variations of correlation between the firings of two neurons. By means of the appropriate normalization procedure for stimulus-induced nonstationarities of the individual neuron firing rates, 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 msec 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 point out that this dynamic cooperativity among groups of neurons is an emergent property of neuronal assembly organization in the brain, and could not have been inferred from any multitude of non-simultaneous single neuron observations. At the same time, this highly dynamic behavior also implies that the commonly used approach of measuring cross-correlograms of multi-neuron spike trains with the goal to draw inference regarding the underlying *anatomical* connectivity should be used with great care. Especially when applaying the ordinary, time-averaged cross correlogram to data in which the correlation is varying in a time-locked (or other) way, this may at best be a misleading calculation (see Aertsen and Gerstein (1990) for an illustrative example of such a case). Both appropriate and inappropriate windows with respect to the stimulus time structure can be chosen, so that inspection of the PST coincidence histogram is essential before deciding on time windows for the cross-correlation measurements.

A more profound consequence of our findings concerns the notion of 'effective connectivity' itself. Our principal, and somewhat surprising observation is that the correlation structure in multiple-neuron spike trains, and hence the membership as well as the internal organization of the observed assemblies is context-dependent and highly dynamic. This suggests that the usual concept of neurons with static interconnections of fixed or only slowly changing (during learning) efficacy is no longer appropriate. Instead one should distinguish between *structural* (or anatomical) connectivity on the one hand and *functional* (or effective) connectivity on the other. Whereas the former can presumably be described as (quasi) stationary, the latter may be highly dynamic, with time constants of modulation as low as tens to hundreds of milliseconds. The 'effective connectivity' thus is not so much to be interpreted as an anatomical wiring diagram (as was the implicit, sometimes even explicit goal of many multi-unit experiments), but rather as a descriptive statement about the internal organization of the network as it is functional more or less instantaneously and under the specific stimulus and behavioral conditions of the physiological experiment. The effective connectivity thus describes the interactions and connections that are sufficiently active to be detectable at the time of observation, and those in a simplified and stylized manner. Lacking currently a more adequate 'language', we propose that the notion of 'effective connectivity' should be understood as the experiment- and time-dependent, simplest possible circuit diagram that would replicate the observed timing relations between the recorded neurons. Yet, although it is phrased in the terminology, usually applied when describing anatomical connectivity, and while this anatomical connectivity certainly provides the material substrate on which the 'effective connectivity' thrives and reorganizes in an ongoing interplay with the varying computational demands, the two should not be confused.

5.2 Activity-related Modulation of Functional Connectivity

In our theoretical investigations we could show that one can indeed envisage mechanisms, which might mediate this transition from static, anatomic connectivity to dynamic, functional connectivity. Rapid changes of synaptic efficacy between neurons in a structurally constant network architecture can be induced by 'dynamic convergence' (Bedenbaugh et al., 1990) of activity from the entire network, in particular by temporal variations of its magnitude and/or internal coherence. The proposed mechanism does not have to invoke intrinsic rapid modulation of synaptic efficacy (von der Malsburg, 1981, 1986); obviously it does not exclude it as a possible additional mechanism either. Further evedince, in particular of experimental nature, is required to clear this issue.

The observed rapid, stimulus-driven modulations of effective connectivity are the signature of an ongoing process of dynamical and activity-related 'linking' and 'unlinking' of neurons in the neural network. This dynamical process may have consequences at different levels of observation. At the single-neuron level, this process might provide a mechanism for the physiologically measured context-dependence and intrinsic dynamics of receptive fields in

central sensory neurons (Aertsen et al., 1981; Eggermont et al., 1981; Dinse et al., 1990). At the multiple-neuron level it might account for dynamical coherence variations in a spatially distributed neural code. The recent observations of stimulus-specific oscillatory events with coherence properties that may extend over wide ranges of the cat visual cortex (Gray et al., 1989a, b; Eckhorn et al., 1988) are in fact discussed as a possible candidate for such a distributed code. At yet another level, the modulation of effective connectivity might serve as a fast 'gating' mechanism for synaptic pathways (consider, for example, the modulation of the (1-2) - pathway in Fig. 3 through variation of the activity carried by a diffuse projection onto distal dendrites of neuron 2). Finally, by dynamic modulation of the input activity to an interconnected network, it might provide a mechanism for the successive ignition of selected neural assemblies within that network ('phase sequences' (Hebb, 1949); 'synfire chains' (Abeles, 1982, 1990)). This aspect bears a close relation to the concept of the 'skeleton filter' (Sejnowski, 1981), as well as to the mechanism of 'threshold control' (Braitenberg, 1978; Palm, 1982). The latter was proposed as a method to modulate the activity level in neuronal networks and, hence, to provide the means to dynamically switch from activation of one cell assembly to the next. In fact, the accomplished regulation of synaptic efficacy in our network models is very reminiscent of the mechanism of threshold control, the main difference being that it is the distance to threshold which is regulated, rather than the threshold itself.

5.3 Dynamical Systems Theory and Spike Train Measurements

From our observations on the behavior of the correlation dimension for different types of measurements we conclude that there exists a fundamental discrepancy between the analysis of a continuous process and of pulse trains generated from that same process. Consequently, the correlation dimension analysis in its present form is not an appropriate tool for the characterization of dynamics in spike sequences. The reason is that a pulse train is not the result of a smooth mapping of the underlying continuous process, such as would be required in the embedding theorem (Takens, 1981). This implies that also other, related measures like generalized dimensions and Lyapunov exponents can not resolve this discrepancy, since they require that the reconstructed attractor is diffeomorphic to the original. In addition, a point process, by its very nature, induces a description in terms of intervals between events. This dynamical variable, however, does not generate a proper Poincaré map of the underlying, continuous system. Consequently, also this approach of making qualitative statements regarding the system dynamics seems to be precluded.

In addition to the above described problems, the application of concepts like fractal dimension in the field of neurophysiology faces further, as yet unresolved questions. For example, spatially extended systems with many interacting components, of which the brain is surely an example, may exhibit dominating, long transients (Crutchfield and Kaneko, 1988). This makes it impossible for the system to reach the attractor during a reasonable observation interval. Moreover, such systems may have multiple, coexisting attractors (Mayer-Kress and Kaneko, 1989). This argues that the spatio-temporal dynamics of a spatially extended system

cannot possibly be characterized by means of an attractor reconstructed from the measurement of a single variable, irrespective of whether this is a single-neuron pulse train or a continuous signal such as a membrane potential or some form of spatio-temporal summation like the EEG. Consequently, attractor characteristics such as fractal dimensions and Lyapunov exponents would be of little value here to begin with (see also Chate and Manneville, 1987).

Methods from dynamical system theory may potentially contribute new insights to various fields, including neuroscience. However, one should be aware of the fact that fractal dimension analysis is not an appropriate tool for the characterization of pulse train (or interval) measurements, nor does it seem to be an adequate descriptor for the dynamics of a distributed system like the brain. In this context, currently being developed concepts for measuring 'complexity' and prediction of time series (see e.g. various contributions in Atmanspacher et al., 1990) seem to be more promising.

Acknowledgements

This paper presents a survey of recent results from a number of ongoing collaborations. The contributions by Purvis Bedenbaugh, Karl-Heinz Boven, Michael Erb, George Gerstein, and Günther Palm are gratefully acknowledged.
We thank Volker Staiger for skillful assistance in preparing the Figures and for the intracellular recording in Fig. 5, Stefan Rotter for critical reading of the manuscript, Jürgen Krüger for providing us with the spike train recordings used in Fig. 1, and Ron MacGregor for putting a version of his neural network simulator at our disposal.

References

Abeles, M.: Local cortical circuits. An electrophysiological study. Studies of Brain Function, Vol. 6. Berlin, Heidelberg, New York: Springer, 1982

Abeles, M: Corticonics. Neural network of the cerebral cortex. Cambridge University Press, 1990 (in press)

Aertsen, A.M.H.J., Johannesma, P.I.M.: The spectro-temporal receptive field. A functional characteristic of auditory neurons. Biol. Cybern. 43: 133-143 (1981)

Aertsen, A.M.H.J., Gerstein, G.L.: Evaluation of neuronal connectivity: sensitivity of cross correlation. Brain Res.: 340: 341-354 (1985)

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

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

Aertsen, A.M.H.J., Gerstein, G.L.: Dynamic aspects of neuronal cooperativity: fast stimulus-locked modulations of 'effective connectivity'. In: Neuronal Cooperativity. Krüger, J. (ed.). Berlin, Heidelberg, New York: Springer, 1990 (in press)

Atmanspacher, H., Scheingraber, H., Treumann, R. (eds.): Proc. of International Meeting on 'Information Dynamics', Bad Irsee (West-Germany), 1990 (in press).

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

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

Bergé, P., Pomeau, Y., Vidal, C.: Order within chaos. New York: Wiley, 1986

Bienenstock, E., Doursat, R.: Spatio-temporal coding and the compositionality of cognition. In: Lestienne, R. (ed.): Proc. Workshop on 'Temporal Correlations and Temporal Coding in the Brain', Paris, 1990

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

Braitenberg, V.: Cell assemblies in the cerebral cortex. In: Heim, R., Palm, G. (eds.): Theoretical Approaches to Complex Systems. Lecture Notes in Biomathematics, Vol. 21, pp. 171-188. Berlin, Heidelberg, New York: Springer, 1978

Chate, H., Manneville, P.: Transition to turbulence via spatiotemporal intermittency. Phys. Rev. Lett. **58**: 112-115 (1987)

Crutchfield, J.P., Kaneko, K.: Are attractors relevant to turbulence? Phys. Rev. Lett. **60**: 2715-2718 (1988)

Dinse, H.R.O., Best, J., Krüger, K.: A new dynamic model of receptive-field organization. In: Eckmiller, R. et al. (eds.): Parallel processing in neural systems and computers, pp. 61-64. Elsevier Science Publishers, 1990

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

Eggermont, J.J., Aertsen, A.M.H.J., Hermes, D.J., Johannesma, P.I.M.: Spectro-temporal characterization of auditory neurons: redundant or necessary? Hearing Res. **5**: 109-121 (1981)

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

Erb, M., Aertsen, A., Palm, G.: 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. Stuttgart, New York: Thieme, 1989

Erb, M., Aertsen, A., Palm, G.: Stability in physiology-oriented neuronal network models. In: Elsner, N., Roth, G. (eds.): Brain - Perception - Cognition, 527. Stuttgart, New York: Thieme, 1990

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

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

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

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

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

Grassberger, P., Procaccia, I.: Characterisation of strange attractors. Phys. Rev. Lett. **50**: 346-349 (1983a)

Grassberger, P., Procaccia, I.: Measuring the strangeness of strange attractors. Physica **9D**: 189-208 (1983b)

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

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

Hebb, D.O.: The organization of behavior. A neuropsychological theory. New York: Wiley, 1949

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

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

Lorenz, E.N.: Deterministic nonperiodic flow. J. Atmos. Sci. **20**: 130-141 (1963)

MacGregor, R.J.: Neural and brain modeling. San Diego: Academic Press, 1987

Malsburg, C. von der: The correlation theory of brain function. Internal report 81-2. Max-Planck-Institute for Biophysical Chemistry, Göttingen (FRG), 1981

Malsburg, C. von der: Am I thinking assemblies? In: Palm, G., Aertsen, A. (eds.): Brain Theory, pp. 161-176. Berlin, Heidelberg, New York: Springer 1986

Mayer-Kress, Kaneko, K.: Spatiotemporal chaos and noise. J. Stat. Phys. **54**: 1489-1508 (1989)

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

Packard, N.H., Crutchfield, J.P., Farmer, J.D., Shaw, R.S.: Geometry from a time series. Phys. Rev. Lett. **45**: 712-716 (1980)

Palm, G.: Neural assemblies. An alternative approach to artificial intelligence. Studies in Brain Function, Vol. 7. Berlin, Heidelberg, New York: Springer, 1982

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

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)

Preißl, H., Aertsen, A., Palm, G.: Are fractal dimensions a good measure for neuronal activity? In: Eckmiller, R. et al. (eds.): Parallel processing in neural systems and computers, pp. 83-86. Elsevier Science Publishers, 1990

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

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

Sejnowski, T.J.: Skeleton filters in the brain. In: Hinton, G.E., Anderson, J.A. (eds.): Parallel Models of Associative Memory, pp. 189-212. Hillsdale: Lawrence Erlbaum Assoc. Publishers, 1981

Sherrington, C.: Man on his nature. The Gifford Lectures, Edinburgh 1937-38. Cambridge: University Press, 1941

Takens, F.: Detecting strange attractors in turbulence. In: Rand, D.A. and Young, L.-S. (eds.): Dynamical Systems and Turbulence, pp 368-381. Berlin, Heidelberg, New York: Springer, 1981

Vaadia, E., Ahissar, E., Bergman, H, Lavner, Y.: Correlated activity of neurons: a neural code for higher brain function. In: Neuronal Cooperativity. Krüger, J. (ed.). Berlin, Heidelberg, New York: Springer, 1990 (in press)