In: Physics of Cognitive Processes
Caianiello E.R. (ed.), Singapore: World Scientific Press, 1987

COHERENT ACTIVITY IN NEURONAL POPULATIONS: ANALYSIS AND INTERPRETATION

Ad Aertsen, Tobias Bonhoeffer

Max-Planck-Institut für Biologische Kybernetik

Spemannstrasse 38, D-7400 Tübingen, FRG

and

Jürgen Krüger

Neurologische Universitätsklinik

Hansastrasse 9, D-7800 Freiburg, FRG

ABSTRACT

To study neural interaction it is necessary to simultaneously record spike trains from a population of neurons under different experimental conditions. Evaluation of the data is normally done by crosscorrelating the spike trains of all possible pairs of neurons and inspecting the results for possible signs of coherence. With recording from relatively large populations (10-30) of neurons becoming experimentally feasible, this procedure starts to be prohibitively tedious. Moreover it does not really address the issue of interest: cooperativity in larger groups of neurons, possibly involved in so-called assemblies.

In this paper we discuss a recently developed technique for multi-unit analysis, explicitly designed to overcome these difficulties: 'gravitational clustering'. The basic idea is the following: each of the recorded N neurons is associated with a charged particle in a fictitious (N-)space. Due to the time varying charges, determined by the respective neurons' spike trains, the particles mutually exert forces which cause them to move: coherently firing neurons result in particles aggregating into clusters.

We also describe a slightly modified version of this approach: the 'dynamic correlation matrix'.

Both techniques will be illustrated with the results of their application to multi-electrode recordings from the cat's visual cortex.

1. INTRODUCTION

Much of our understanding of the functioning of the nervous system is based on electrophysiological recordings of the activity of individual neurons, the single unit spike trains. There is, however, ample evidence that the nervous system is more than simply a collection of independent elements. Especially at more central levels, it is thought to involve functional assemblies of neurons, acting together in a coherent manner, mediated by anatomical, presumably 'plastic' connections [20,12,6,31]. In order to gain understanding of processes within and between neuronal assemblies it is necessary to observe simultaneously and separably the activity of many neurons during appropriate experimental manipulation of the whole organism.

Within this context it is understandable that we are witnessing a growing interest in the study of neuronal interaction. Mainly due to recent developments in multi-neuron recording technology (multi-micro-electrodes, spike separation techniques), the simultaneous recording of spike trains from a population (10-30) of neurons has become a feasible enterprise (reviews are given in [15,27]).
The search for 'neuronal assemblies' raises an interesting sampling problem: What is the size of such an assembly and how large a fraction of an assembly does one have to observe in order to draw conclusions about the assembly as a whole? Since this issue remains largely unsettled, it is generally felt advisable to try to record from large groups of neurons. A relatively new technique which seems to be well suited to achieve this goal is the method of 'optical recording', which uses voltage sensitive fluorescent dyes [18,19]. Provided that this technique, when applied as a multi-unit recording device, reaches a spatial resolution which is in the single neuron range (the present resolution is - depending on the preparation - at best some 30 microns), it should be possible to record simultaneously and separably from some hundred neurons at a time. This means that a substantial part of the network is available for observation, as compared to relatively low fractions attainable with multi-electrodes. The advantage of this higher 'sampling' ratio for the analysis of network properties can hardly be overestimated.
Another potential advantage of the optical recording technique is that the neurons one is recording from, are literally visible. This in principle opens the way for a long overdue attempt to combine multi-neuron recording with a more thorough analysis of anatomical structure than currently possible with multi-electrodes 'poking in the dark'.

Facing the massive flow of data resulting from evermore sophisticated multi-neuron experiments, it must be granted that, unfortunately, theoretical developments regarding analysis and interpretation of these data have not been able to keep pace. The principal tool for analysis in use continues to be the cross correlation of pairs of spike trains (two neurons at a time), a method which was developed already in the sixties [34]. Although proven to be quite adequate for few-unit (2-3) recordings, this approach is being strained beyond extent in a situation where recording from some ten neurons simultaneously has become a relative routine in a number of laboratories. In this context the present paper discusses recently developed techniques of multi-unit analysis: *'gravitational clustering'* [17,14,1], and the related technique of *'dynamic correlation matrix analysis'*. These techniques will be described heuristically and illustrated with the results of their application to multi-electrode recordings from the cat's visual cortex [25].

2. CONVENTIONAL ANALYSIS OF MULTI-UNIT RECORDINGS

In the following we will look somewhat more closely at a typical multi-unit recording, in this case from Area 17 in the cat. In these experiments use was made of a linear array of twelve glass covered Pt-Ir electrodes, arranged in organpipe-like manner with a distance between electrode tips of 160 microns. The array was introduced into the left hemisphere of the visual cortex at an angle of 45° such that the electrode tips were situated approximately vertically underneath each other (for details see [25]). The actual recording sites could be reconstructed after the experiment on the basis of lesions made by some of the electrodes (generally 3 to 4). In the present example the array of recording sites indeed was shown to be approximately perpendicular to the cortical layers; individual sites were situated in the following layers: electrode number 1 in the white matter underneath layer VI, 2 in layer VI, 3 in VI, 4 in VI, 5 in V, 6 in IV, 7 in IV, 8 on the border of IV and III, 9 in II/III and, finally, 10 in I. The remaining two electrodes in this particular case did not yield a useful recording and are left out of consideration. Five out of ten electrodes (nrs. 2, 4, 5, 6 and 8) each gave a reliable single unit recording, the remaining five (nrs. 1, 3, 7, 9 and 10) were judged to be probably single unit, with possibly some additional spikes from a second unit.

The stimulus in this experiment consisted of a light bar (length 3°, width 14') moving at constant velocity in a direction perpendicular to its orientation (distance travelled 3°, duration 1.8 sec). At the end of the movement the bar remained stationary for 0.4 sec, after which it moved in the opposite direction (same distance and duration). Finally the bar was rotated over 22.5° (duration 1 sec) after which another, identical cycle of movement back and forth, perpendicular to the bar orientation started. This scheme was repeated periodically, so that after 8 cycles of 5 seconds the original direction of movement was reached again. The complete stimulus sequence, lasting 40 seconds, was presented repeatedly, in general some 20 to 30 times. In the particular example discussed below, the stimulus was presented to the right (contra-lateral) eye.

2.1 Representation of Multi-Unit Spike Trains

Spike trains recorded during 10 consecutive stimulus sequences are shown as dot displays in Fig. 1, with different neurons represented by different colours. Time runs horizontally, covering the complete stimulus sequence of 40 sec, vertical lines in the dot displays signify the different movement cycles of 5 sec each. Dot displays are arranged in the way the corresponding recording sites were located in the cortex: the pink dots at the top represent spikes from electrode 10 which occupied the most superficial position in the cortex (in layer I), the yellow ones at the bottom correspond to spikes from electrode 1 which was down in the white matter below layer VI. In the lowest part of Fig. 1 we have schematically indicated the stimulus program with arrows showing the direction of bar movement. From this Figure it becomes quite clear that not all directions are equally successful in eliciting a response; for instance the stimulus where the bar is moving to the upper left corner under an angle of 45° appears to be clearly favoured in the recordings 5 to 8. The more or less unanimous preference for direction of bar motion is consistent with the earlier mentioned vertical alignment of recording sites.

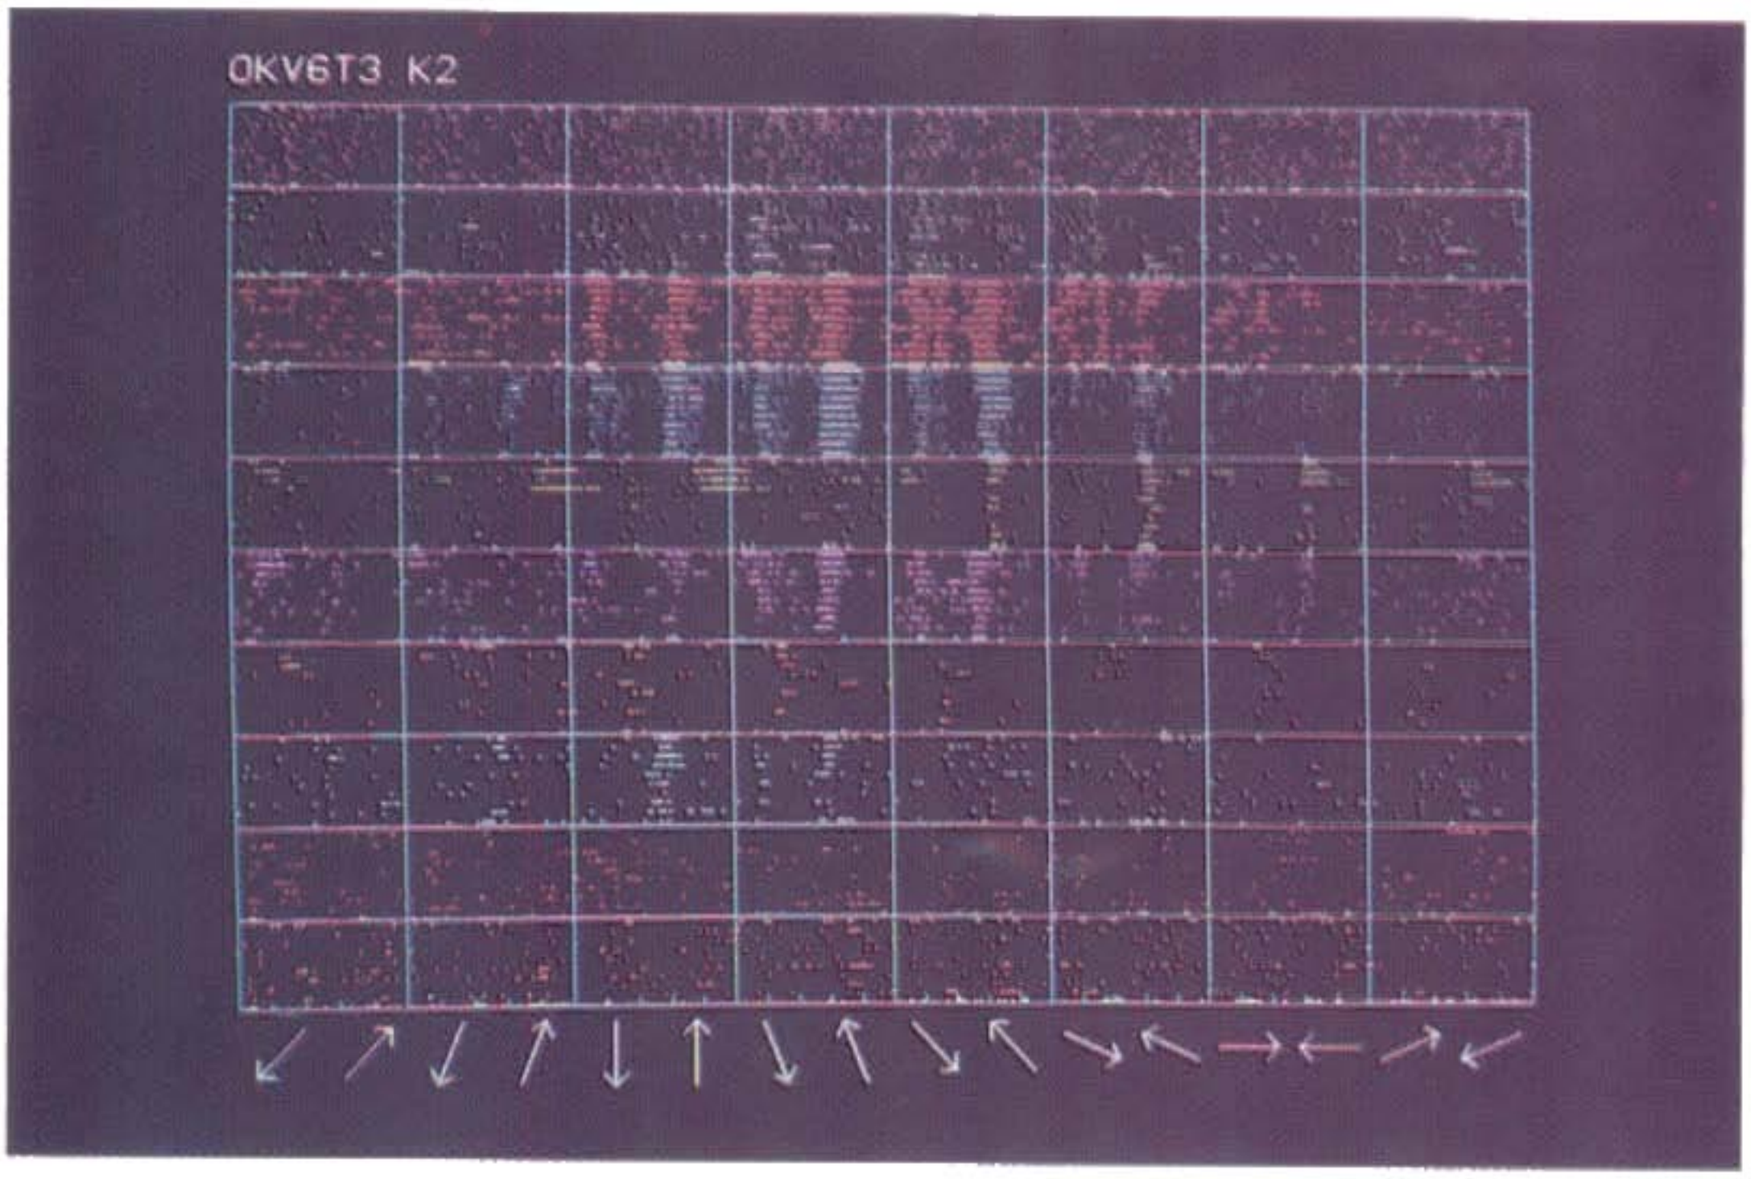

Figure 1. Dot displays of spike trains from a twelve electrode recording in the cat's visual cortex. Two electrodes were omitted because they did not yield a useful recording. Different colours represent different electrodes. The stimulus program (bar moving in different directions) is indicated schematically at the bottom. Further explanation is given in the text.

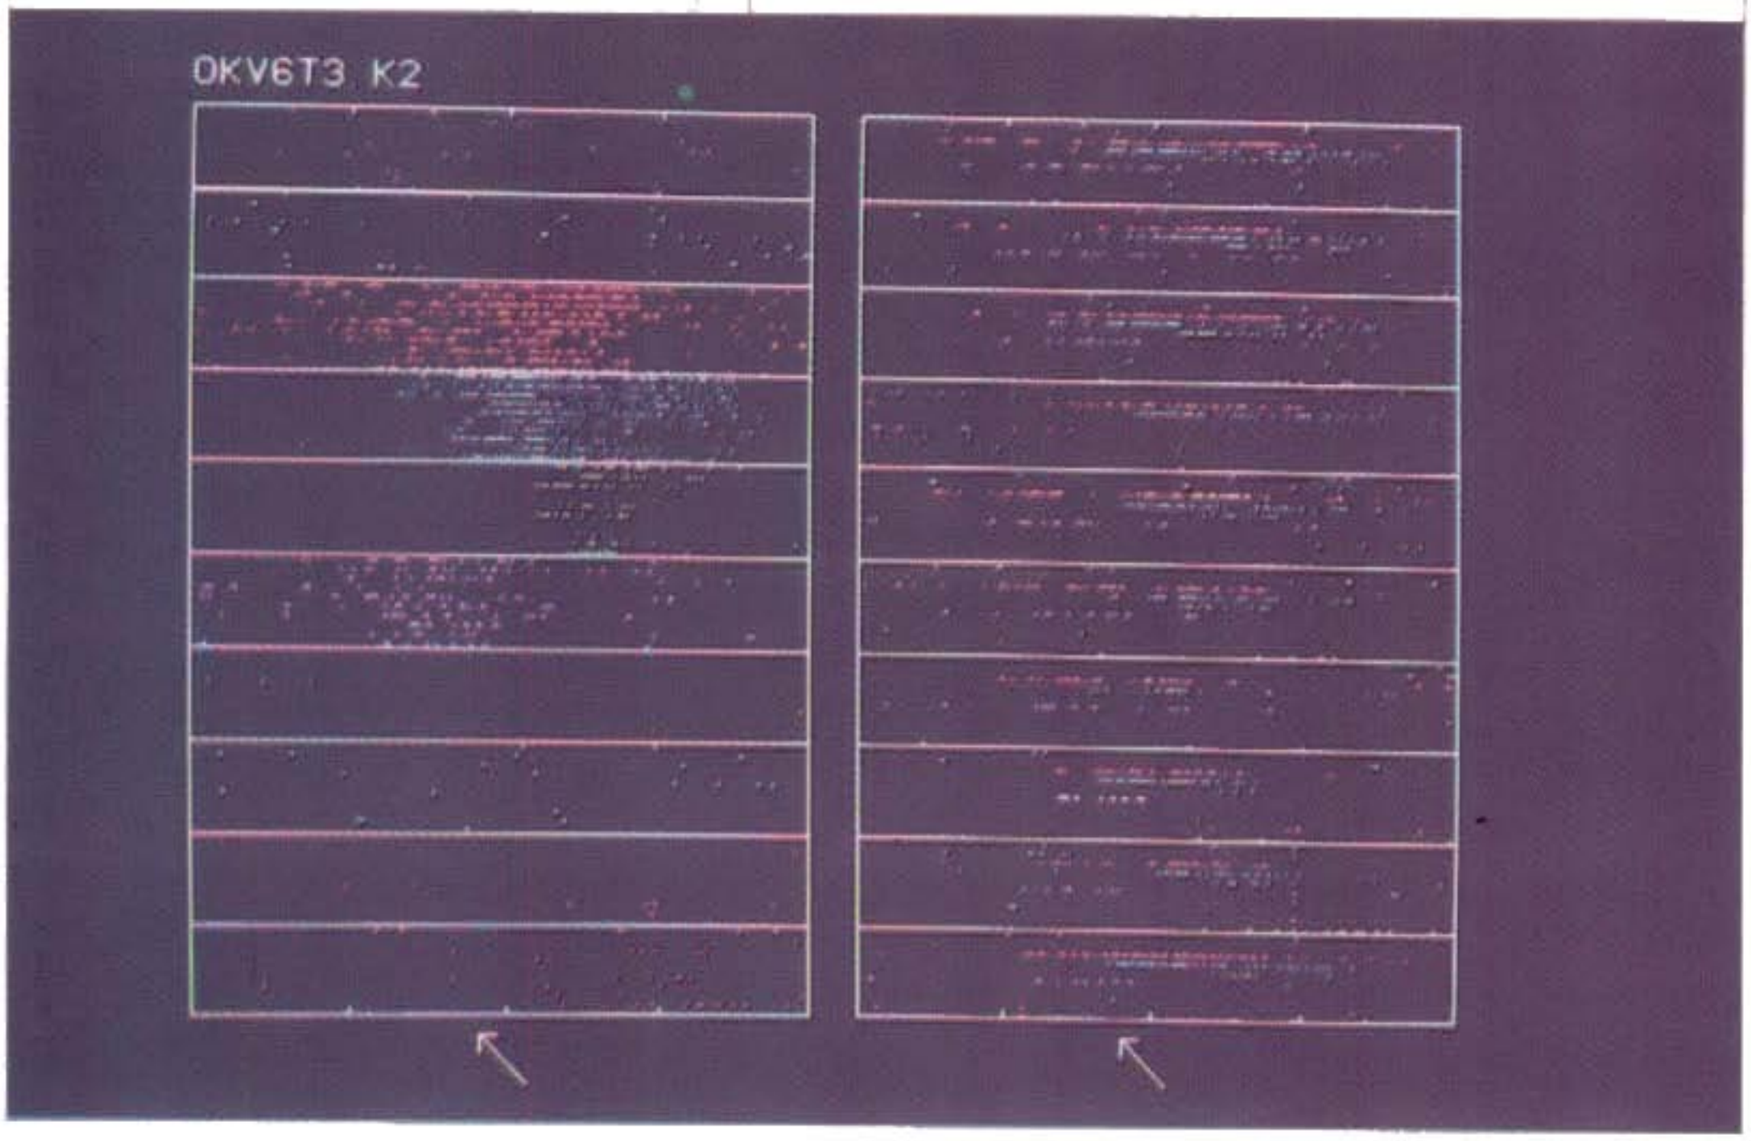

Figure 2. Dot display showing a selected part of the data from Fig. 1: one direction of bar movement indicated at the bottom of the Figure. At the left the data are displayed as in Fig. 1 (sweeps per unit), whereas at the right the adjacent traces in each box represent synchronous activity of the different neurons (units per sweep), different boxes corresponding to different stimulus presentations. The colour code is the same as in Fig. 1.

In order to study the effects of bar movement direction more closely, the data as displayed in Fig. 1 could be 'cut' into different vertical slices by selecting an appropriate time window, such that in every slice one has only one direction of stimulus movement. Using this procedure of editing the spike trains ('cut' and 'glue') we have artificially produced multi-unit recordings, in which the stimulus is a bar periodically moving in one and the same direction. In Fig. 2 we have displayed the response to the earlier mentioned preferred stimulus direction (45°, top left), by selecting a time window between 22 and 24 sec in the 40 sec sweep. On the left hand side the data are displayed in the same way as in the previous Figure (organized in sweeps per unit), whereas on the right hand side adjacent traces in each box represent the simultaneous activity of different units during one stimulus presentation (units per sweep), the different boxes corresponding to subsequent stimulus presentations. The latter representation, where colour now is essential to identify the different units ('Neurochrome' [10]), in fact is the more natural one, since it reflects most properly the activity as it happens to occur in the brain. Consequently, it is also the most appropriate one to look for possible time patterns in multi-unit activity. In Fig. 2 one observes how different neurons react differently to stimulation, both regarding the magnitude and the time course of the responses. Whereas the left hand part of Fig. 2 clearly shows the average dynamic properties of each single neuron's response to the ensemble of identical stimulus presentations, the right hand part emphasizes the intricate inter-relationships between the individual responses of different neurons during each single stimulus sweep. Note also the considerable variability of multi-unit firing patterns over subsequent presentations of the same stimulus (different boxes in right part of Fig.2). In this context we note that, when looking for temporal structure in multi-unit activity it should - apart from proper visualization - also seriously be considered to present the data to the investigator truly as signals in the time domain e.g. as identifiable, audible signals ('Neurophone' [2]).

2.2 Cross Correlation Analysis of Pair Interaction

The commonly used technique to investigate coherence among simultaneously recorded activity of different neurons is to select from the observed group two neurons at a time, and analyze this pair's firings for possible temporal relations. This analysis is based on the cross correlation function of the two simultaneously recorded spike trains [34]. Departures from background in the correlogram are taken as indicative of a *'functional connection'* between the two neurons, where this 'connection' may take different forms: excitatory or inhibitory connection, shared input, stimulus coupling [30]. We use here the term 'functional connection' as purely descriptive: for some, as yet not specified, reason there is a statistical relationship between the probabilities of firing of the two neurons.

An example of this type of analysis is shown in Fig. 3, which presents correlograms with different time resolutions for the pair of units 7 and 8, calculated for the activity evoked by presentation of the '45°, top left'-stimulus (cf. Fig. 2; for the correlograms we used 16 stimulus sweeps instead of only the 10 shown in the dot display). In the upper correlogram (Fig. 3a), the one with the largest time scale (-4 to 4 sec, 80 ms/bin), one clearly sees the dominating effect of the stimulus: periodical waxing and waning of synchrony of firing, with a period of 2 sec, corresponding to the period of stimulus presentation in our 'edited' spike trains. Even at this resolution, though, one also notices that the synchrony around zero time shift (middle peak: simultaneous activity) is slightly higher than in the adjacent peaks (corresponding to a time shift of one stimulus period). This is shown in more detail in the lower correlo-

grams which have increased time resolution (Fig. 3b and 3c: 10 ms/bin, Fig. 3d and 3e: 1 ms/bin). In each of these pairs of correlograms the top one (3b resp. 3d) corresponds to zero time shift (simultaneous recording), whereas the lower one (3c resp. 3e) shows the synchrony for a time shift of one stimulus period (the so-called *'shift control'* [16]). Comparison of simultaneous correlogram and shift control, especially in the lower pair, clearly shows that, during this particular stimulus, the firings of units 7 and 8 exhibit a degree of synchrony which significantly goes beyond the much higher degree of synchrony induced by stimulation per se. The fact that this 'extra' synchrony has its peak at the origin suggests that at least the main contribution is due to shared neuronal input, whereby the rather irregular and unsymmetric form of the elevation does not exclude other neuronal sources of coherence.

Careful inspection of the shape of correlograms and the evaluation of quantitative characteristics, such as peak width, height, displacement from zero etc., are used to make inferences about possible connectivity in the network. An account of an analysis in these terms was given by Krüger [26]. It should be noted here that under conditions of relatively sparse firing, which is the normal case in cortex, the method of cross correlation of spike trains was shown to strongly favour the detection of excitatory connections (signified by an elevation in the correlogram) above inhibitory ones (a trough in the correlogram) [3].

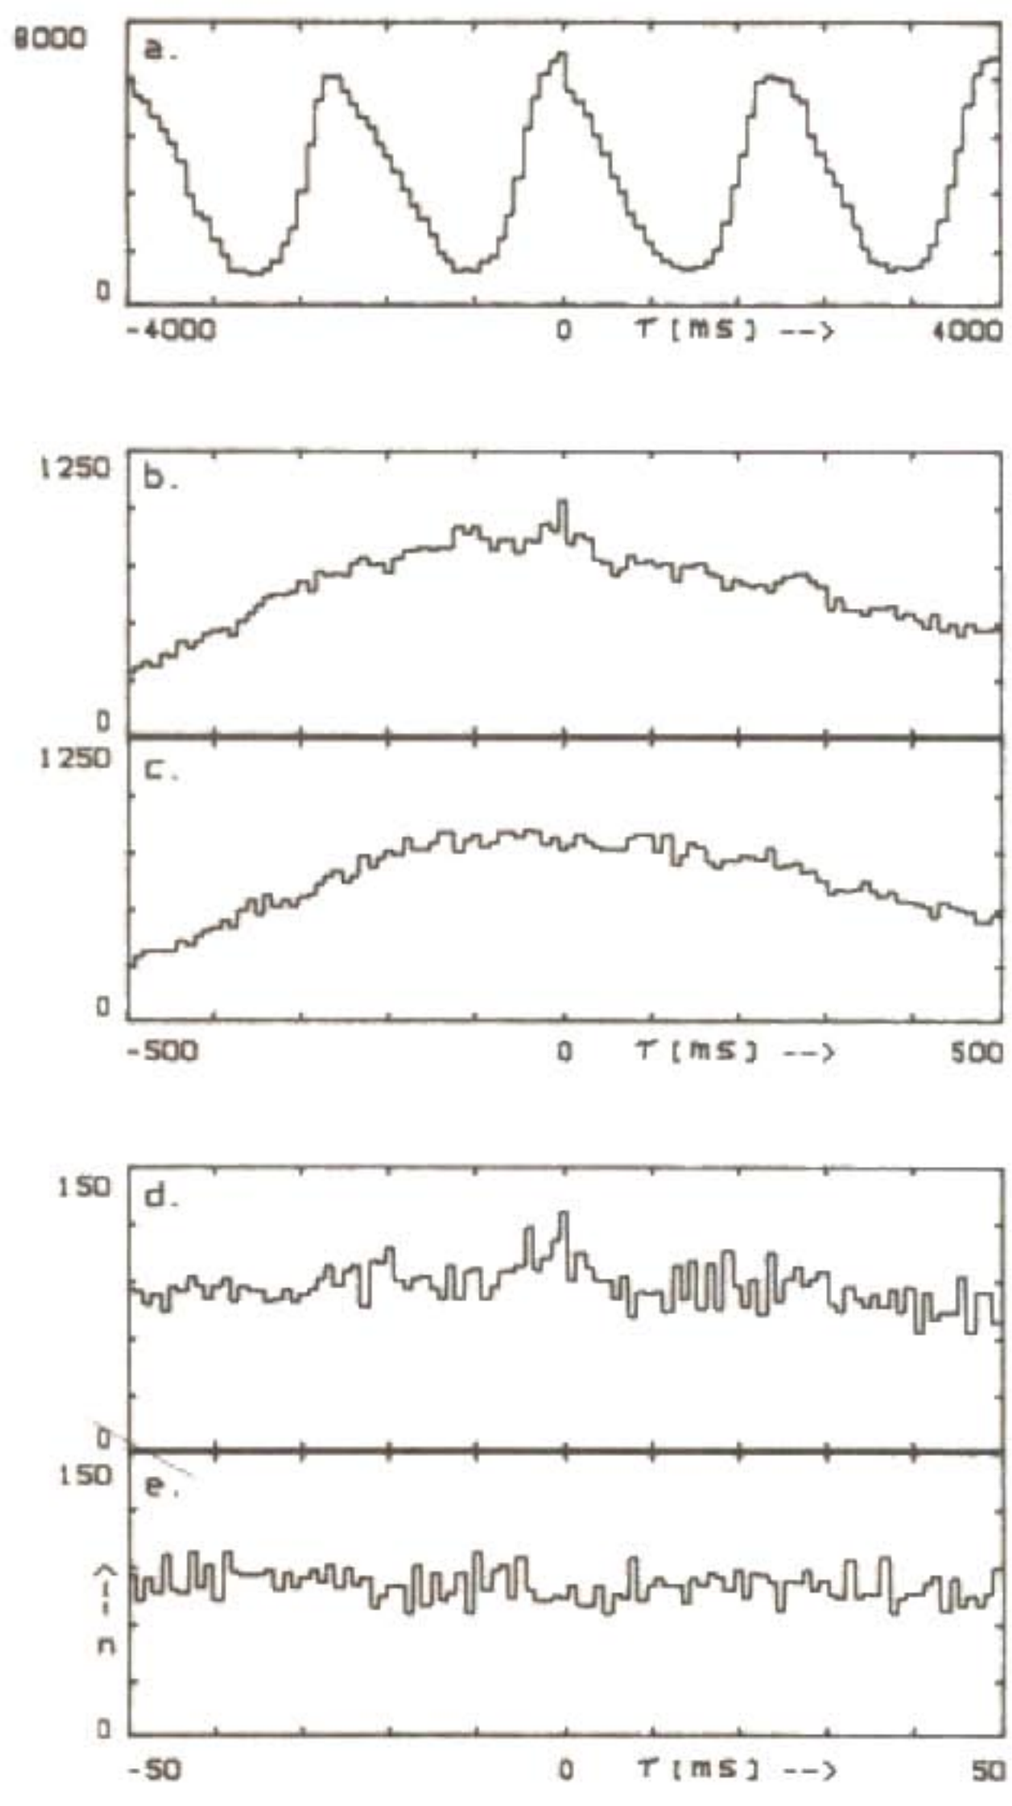

Figure 3. Cross correlograms of the spike trains from neurons 7 and 8 with three different time resolutions (a. 80 ms/bin; b. and c. 10 ms/bin; d. and e. 1 ms/bin). Data are the same as in Fig. 2 (only one direction of bar movement). Two correlograms (b. and d.) display the coherence between neurons 7 and 8 in a simultaneous recording, whereas c. and e. show their synchrony with a time shift over one stimulus-period ('shift control'). Further explanation is given in the text.

The problem about this approach is that the number of correlograms which have to be examined grows nonlinearly with the number of recorded neurons: for N neurons one has N(N-1)/2 different pairs, which in the present case of 10 neurons (cf. Figs. 1 and 2) implies 45 correlograms. Going to 30 neurons (this is about the maximum number of simultaneous recordings that came forward in the last few years, see e.g. [28]) the number of correlograms increases to 435. These numbers have to be multiplied still further, e.g. with the number of different stimuli (in our case, for instance, we have 16 directions of bar movement) and with the desired number of different time resolutions. After examination of all the individual pair correlations, the inference has to be recombined by the experimenter into a single conceptual picture. Quite evidently this procedure becomes prohibitively tedious for growing numbers of recorded neurons (imagine the pairwise analysis of a 100-unit recording from a fluorescent dye experiment!). Furthermore one may question whether the pairwise analysis of correlation addresses the real issue of interest: cooperativity in larger groups of neurons, possibly involved in so-called 'assemblies'.

3. 'GRAVITATIONAL CLUSTERING'

In order to overcome the difficulties discussed above, recently a new analysis technique was proposed: *'gravitational clustering'* [17]. A formal description of this approach has been given elsewhere [1,14,17]; we will restrict ourselves here to a short, rather more heuristic description.

The basic idea is as follows: with each of N neurons we associate a point-particle in a (fictitious) N-space. Initially these particles are located such that mutually they all have the same distance. This is achieved most easily by placing them on the vertices of an N-dimensional hypercube. Now each particle is given a time varying charge, which is determined by the respective neuron's spike train: each spike induces a charge increase which decays in the course of time (we have used an exponential decay with different time constants, e.g. 8 ms). Because we assume that in our fictitious space a kind of Coulomb law is valid (for equations see Appendix) the particles mutually exert forces, which cause them to move. As a consequence of temporal overlap of charge histories some of the particles will cluster, whereas others will on the average preserve their original positions.

The rules of the charge functions can be defined such that the forces lead to aggregation of those particles which correspond to neurons that tend to fire in synchrony (excitation) or, on the other hand, where firing in one is associated with silence in others (inhibition). Different clusters correspond to different assemblies. By this procedure we have translated the amount of temporal coherence among different spike trains into a distance measure in N-dimensional space: the higher the coherence among neurons, the smaller the distance between the corresponding particles and vice versa.

Since the gravitational representation translates synchrony into clustering, it is to be expected that the high degree in synchrony induced by stimulation per se (cf. Fig. 2 and 3) will lead to profound clustering of almost all particles, thereby effectively masking the much smaller amount of coherence of neuronal origin. Since at this point our main interest is in the latter contribution to multi-unit coherence, this purely stimulus-related component has to be taken care of. In a first order approximation this was accomplished by resorting to a differential type of analysis which may perhaps be characterized best as being analogous to subtracting the 'shift control' from

the simultaneous correlograms in Fig. 3. We will return to this methodological issue in Sect. 4; a more extensive discussion on procedures to normalize for directly stimulus-induced nonstationarities c.q. synchrony will be presented elsewhere (Aertsen et al., in preparation). All results in the remainder of this Section were normalized for direct stimulus effects by this differential correction procedure.

3.1 Coherence as Proximity: Analysis of Distance

The original problem of analyzing temporal coherence now has been traded for the problem how to investigate the N-space for clusters of particles. Our first approach was to plot the different inter-pair distances as a function of time: downgoing curves will signify clustering particles while horizontal curves indicate that no net attraction occurred. This was done for the experiment described above. Again we selected the activity during the (45°, top left)-direction of bar movement (cf. Fig. 2) and subjected these spike trains to the gravitational analysis. The resulting pair distances, plotted against time are shown in Fig. 4, where colour was used to identify the pairs. Note that time runs horizontally from 0 to 32 sec: for the gravitation analysis we used 16 (instead of 10, cf. Fig. 2) stimulus presentations of 2 sec each. In this Figure four different bands of curves can be distinguished :

(1) The particles corresponding to neurons 5,7 and 8 cluster most strongly: their pair curves have the strongest negative slopes. Among these three especially the neurons 7 and 8 appear to have a high degree of coherence in their firing patterns (cf. Fig. 3). It should be noted here that in the gravitational algorithm, in order to avoid oscillations around a singularity, the attractive force between any two particles is 'switched off' the moment these particles have approached each other to a certain fraction (here 10%) of their initial distance. The interactions with the other particles, however, continue to be present so that eventually the clustered particles may drift apart again, whereupon their own pair force is switched on again. From the curves in Fig. 4 it can be observed that for these three particles this switching of pair forces actually took place several times.

(2) The three curves which lie closely together and slightly above the former three curves represent the somewhat smaller coherence between neuron 6 and the three previous ones.

(3) Still higher in the Figure one observes a yellowish band of curves which reflects the lower, but still significant degree of coherence between neurons 9 and 10 and the earlier four.

(4) Finally, the remaining particles (1, 2, 3 and 4) more or less preserve their mutual distances which indicates that the coherence among the corresponding neurons is negligible. The slight decrease in distance between the particles in this group on the one hand and those in groups (1), (2) and (3) on the other hand is an artifact: the distance between two particles may slightly decrease without the two themselves having anything to do with each other, simply because one of the two particles belongs to a certain cluster and the other one does not. A simple example in two dimensions may illustrate the point: imagine three particles on the corners of an equilateral triangle. If two of the particles are attracting each other and therefore decrease their mutual distance, this will also affect the distances between these two and the third, non-moving particle: these will also decrease slightly (to a minimum of about 87%). This decrease, however, does not reflect a genuine coherence. Small decreases in pair distances therefore should not be overinterpreted, in fact, they add to the obvious lower limit for effects to be significant, set by the diffusion type influence of the stochasticity of neuronal firing [13].

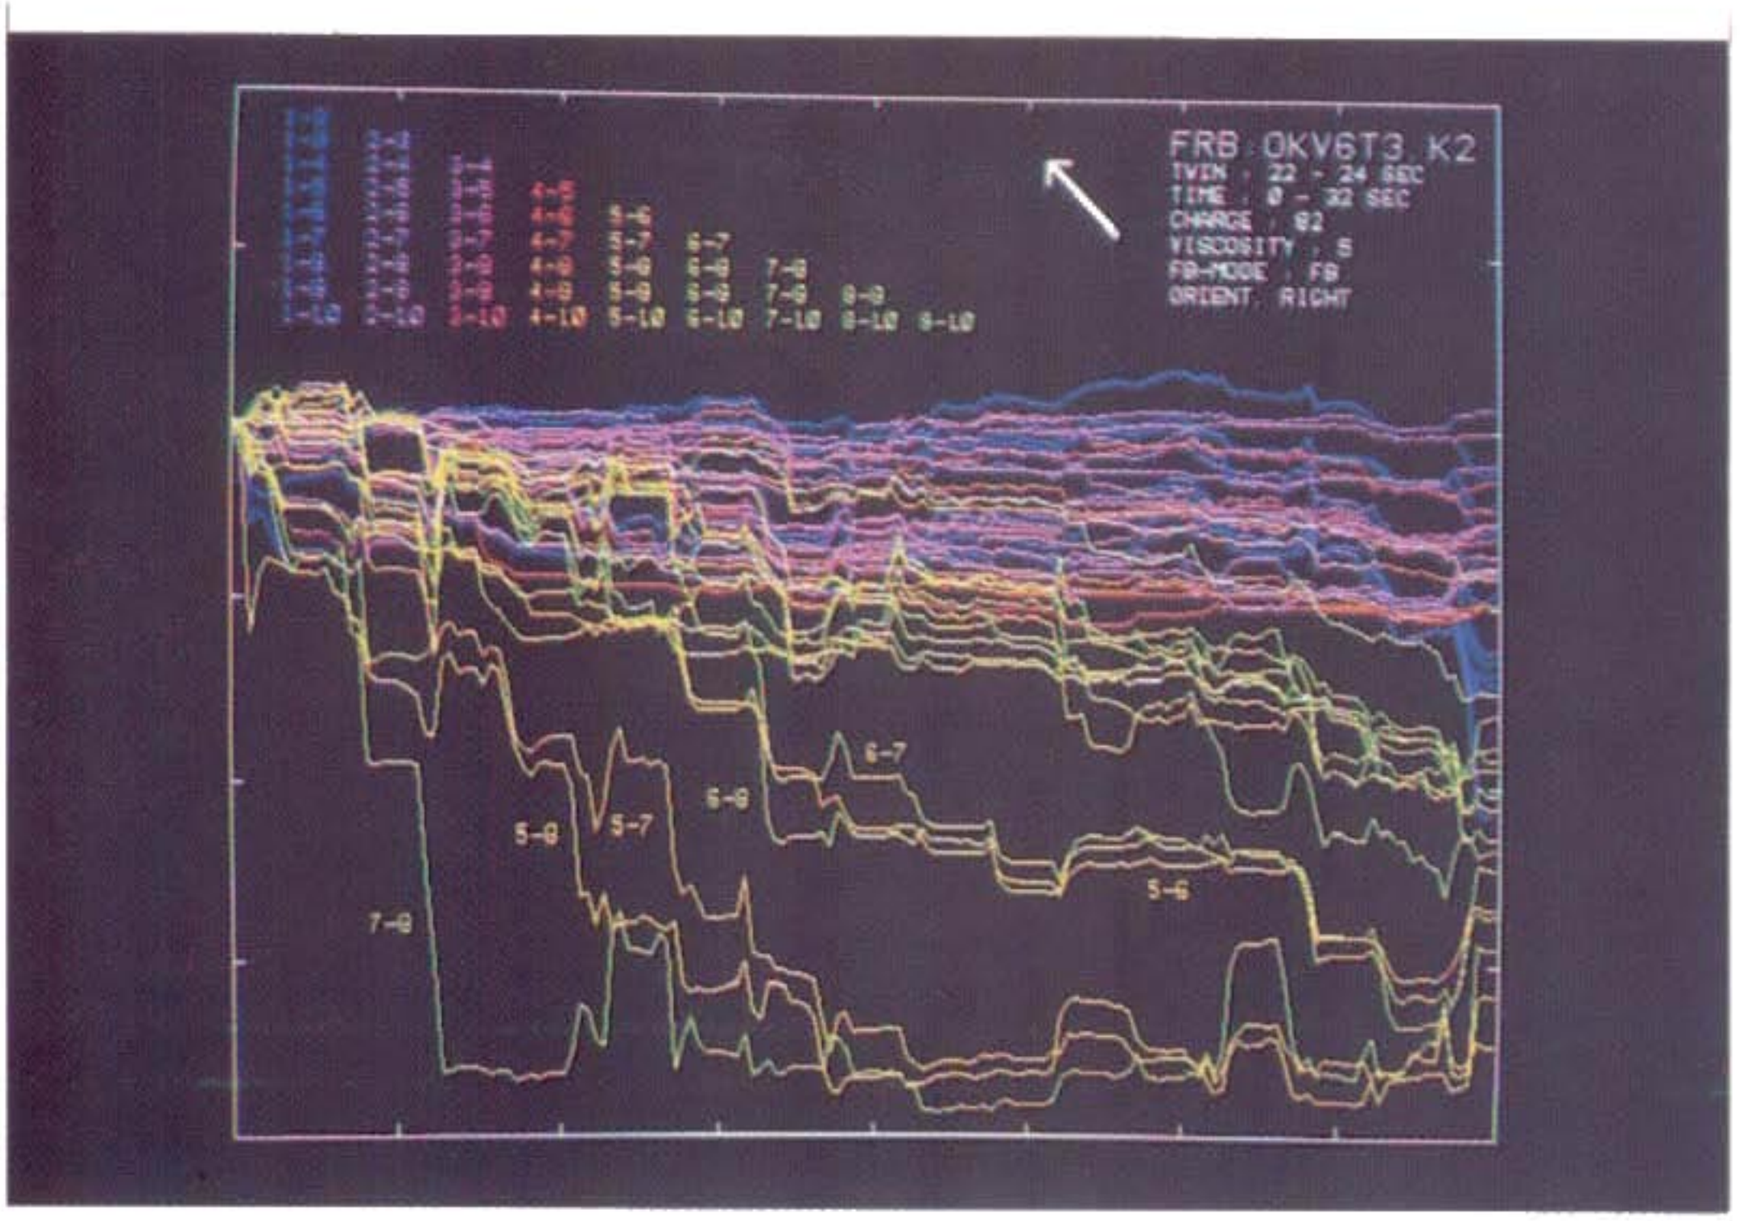

Figure 4. Gravitational clustering: inter-pair distance as a function of time for the 45 pairs of neurons in the 10-electrode recording of Fig. 2. The arrow indicates that the direction of motion of the stimulus was towards the upper left corner of the screen. The colour coding for the pairs is displayed in the upper left corner. For further explanation see text.

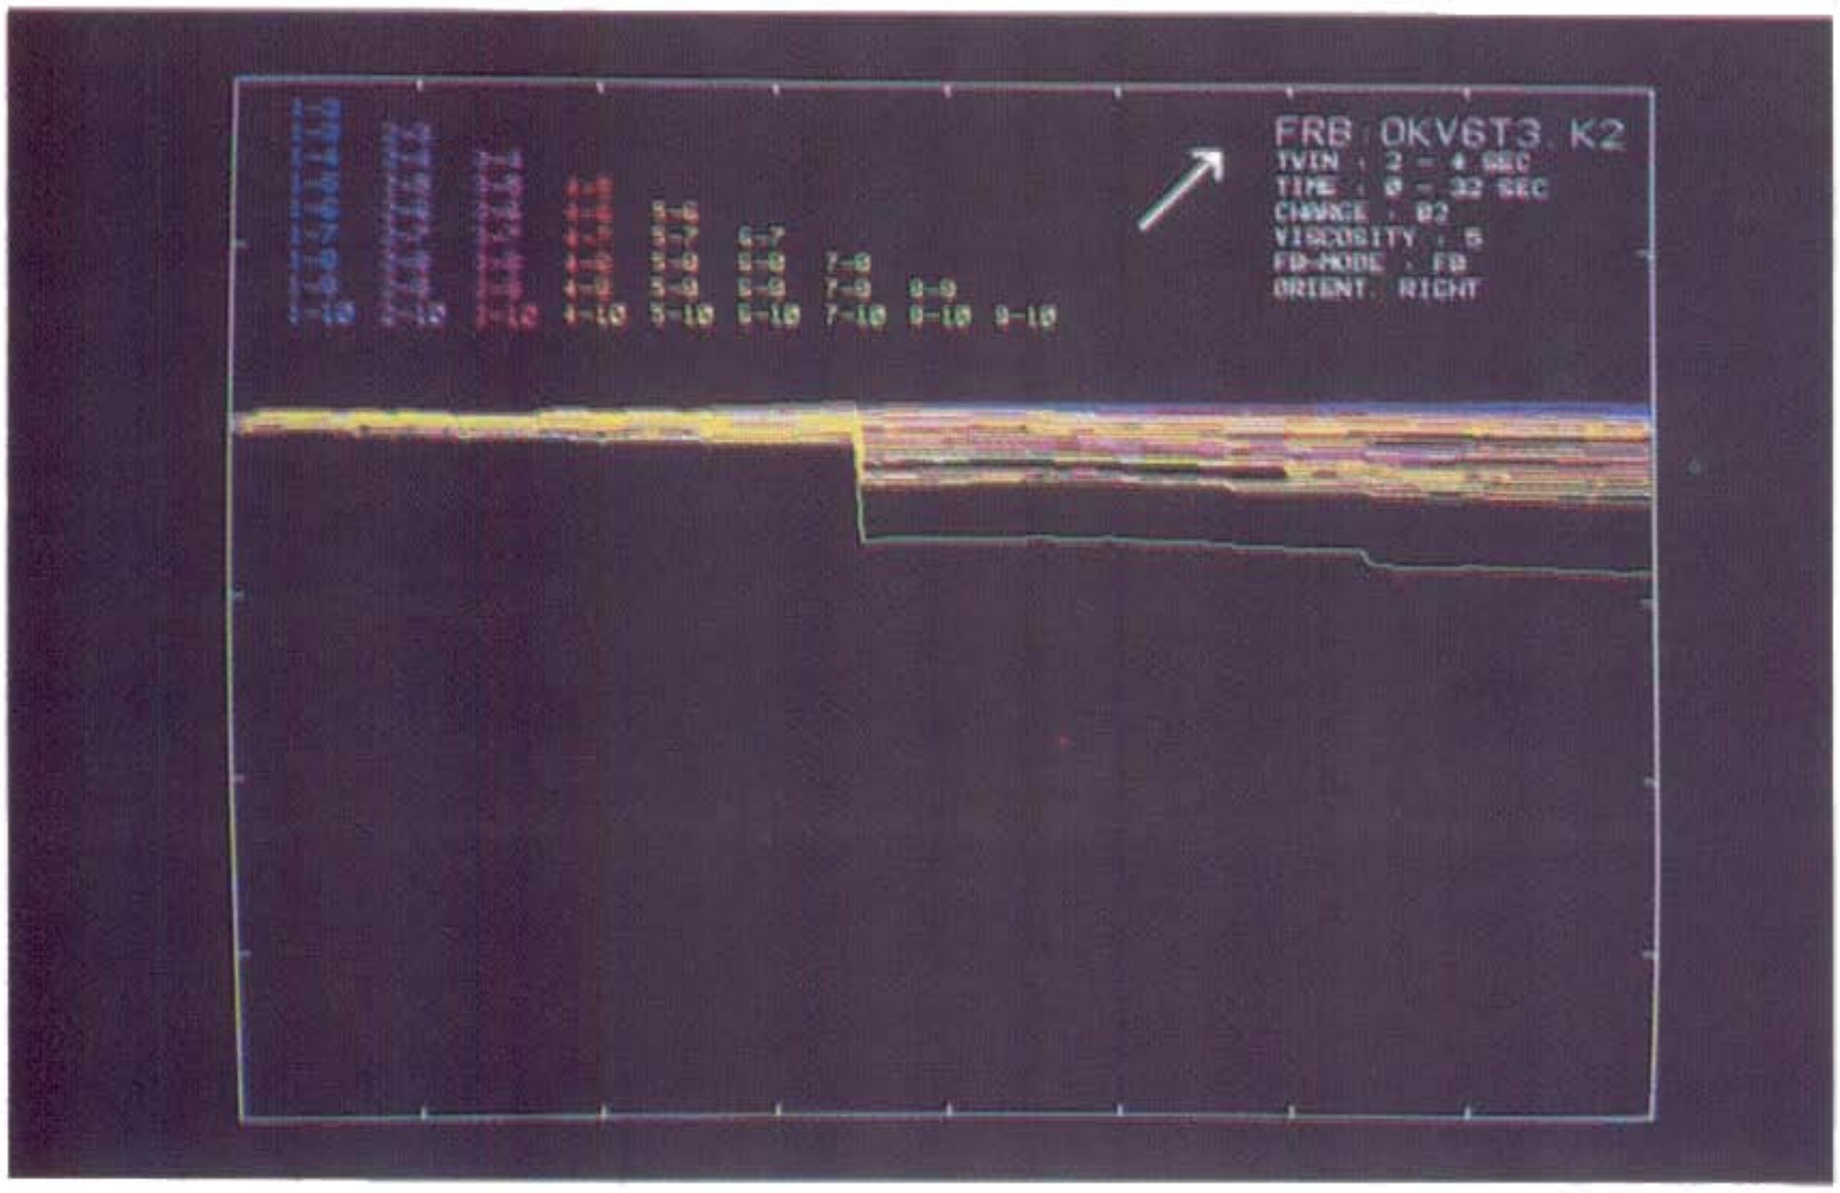

Figure 5. Gravitational clustering: distance plots for the 45 pairs of neurons. The arrow indicates that the direction of bar movement was towards the upper right corner of the screen. Further details as in Fig. 4.

Apart from the global behaviour of the various curves (overall shape), Fig. 4 also entails quite detailed information about the temporal structure of multi-unit coherence. Note, for instance, the temporal variations in slope along each individual curve and the temporal relationship between these variations among different curves. These reflect substantial variations in synchrony among spike trains as time proceeds and are obviously related to the stimulus induced time structure of the multi-unit spike trains, revealed in the 'Neurochrome' (cf. right hand part of Fig.2).

3.2 Stimulus-Dependence of Coherence

Being interested in a possible stimulus dependence of coherence among the multi-unit firings (remember we corrected for directly stimulus-induced coherence), we also analyzed the spike trains corresponding to the other directions of bar movement. Using again the spike train editing procedure, the data displayed in Fig. 1 were 'cut' into 16 vertical slices, each 2 seconds wide, such that each slice corresponds to one particular direction of bar movement; these data were submitted to the gravity procedure. Figure 5 shows the resulting distance curves for one of the other directions: 45° to the top right, which is perpendicular to the previous one. In this case one observes that, apart from some noise related 'diffusion', basically all particles remain their initial positions. Only the curve corresponding to neurons 8 and 9 shows a sudden, small decrease in particle distance after some 14 seconds of recording; this is caused by a coincident burst of firing early in the eighth stimulus sweep, which involved four neurons, with strongest contributions from numbers 8 and 9. On the whole Fig. 5 shows that for this particular direction, perpendicular to a more favoured direction (as far as single unit responses are concerned), at this point there is no evidence for coherence in firing, which goes beyond directly stimulus-induced effects.

In order to get a more complete picture of the influence of direction of movement on the degree of multi-unit coherence, one might plot all the distance curves corresponding to all the directions tested and compare these. Again we are facing a congestion of data, therefore we shall momentarily take a step back and go to a more global measure of coherence. To this end we disregard the dynamics of clustering for a moment and only look at the final result: the inter-pair distances after 16 stimulus presentations. More specifically, we have made a bar graph of the net decrease in pair distance: the higher the functional connectivity between a pair of neurons (i.e. the smaller the final distance between the corresponding particles), the higher we draw the bar for this pair.
In Fig. 6 three of these bar graphs are displayed. The arrow in the upper left corner of each graph represents the direction in which the stimulus moved. The upper graph corresponds to the direction 45°, top left (cf. Fig. 4), the middle graph to the movement in the obverse direction, and the bottom graph corresponds to a direction perpendicular to the other two (cf. Fig. 5). Comparing the upper two graphs it is apparent that the coherence among the earlier mentioned neurons in groups (1), (2) and (3) is comparably strong for movement back and forth along the negative diagonal of the screen, with the exception of neuron 6 which only 'joins in' during movement to the top left (in the middle graph the corresponding large bars for particle 6 are missing). In the orthogonal direction (lower graph) the coherence is generally quite weak, as we had already seen in Fig. 5.

Figure 7 summarizes the analysis of coherence for the whole spectrum of directions tested. In this Figure bar graphs were positioned at the tip of the vector which indicates the direction of motion of the stimulus, e.g. the graph for direction

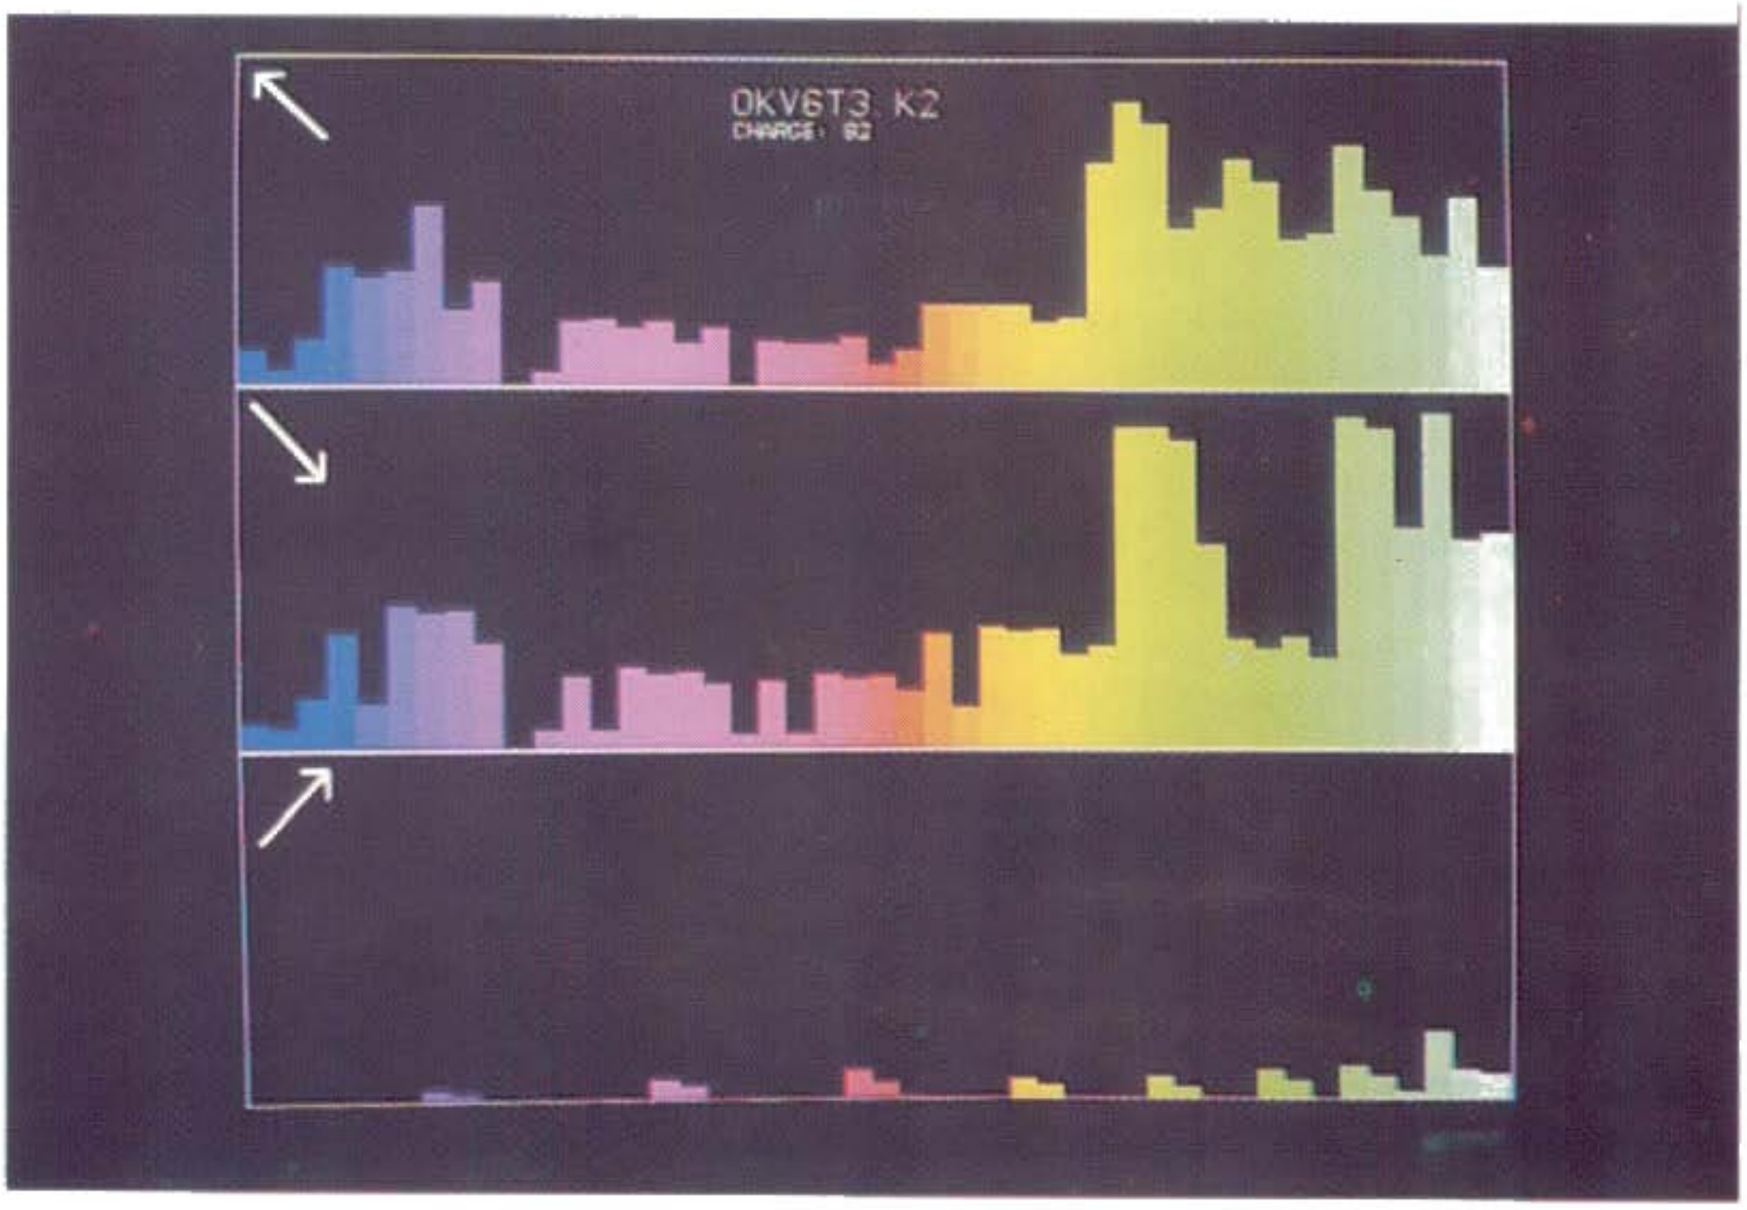

Figure 6. Stimulus dependence of coherence: three bargraphs for different directions of stimulus movement. Colour coding is the same as in Figure 4. The heigth of each bar corresponds to the degree of coherence in the respective pair of neurons. The direction of motion of the stimulus is displayed by the arrows in the upper left corner of each Figure. Further explanation in the text.

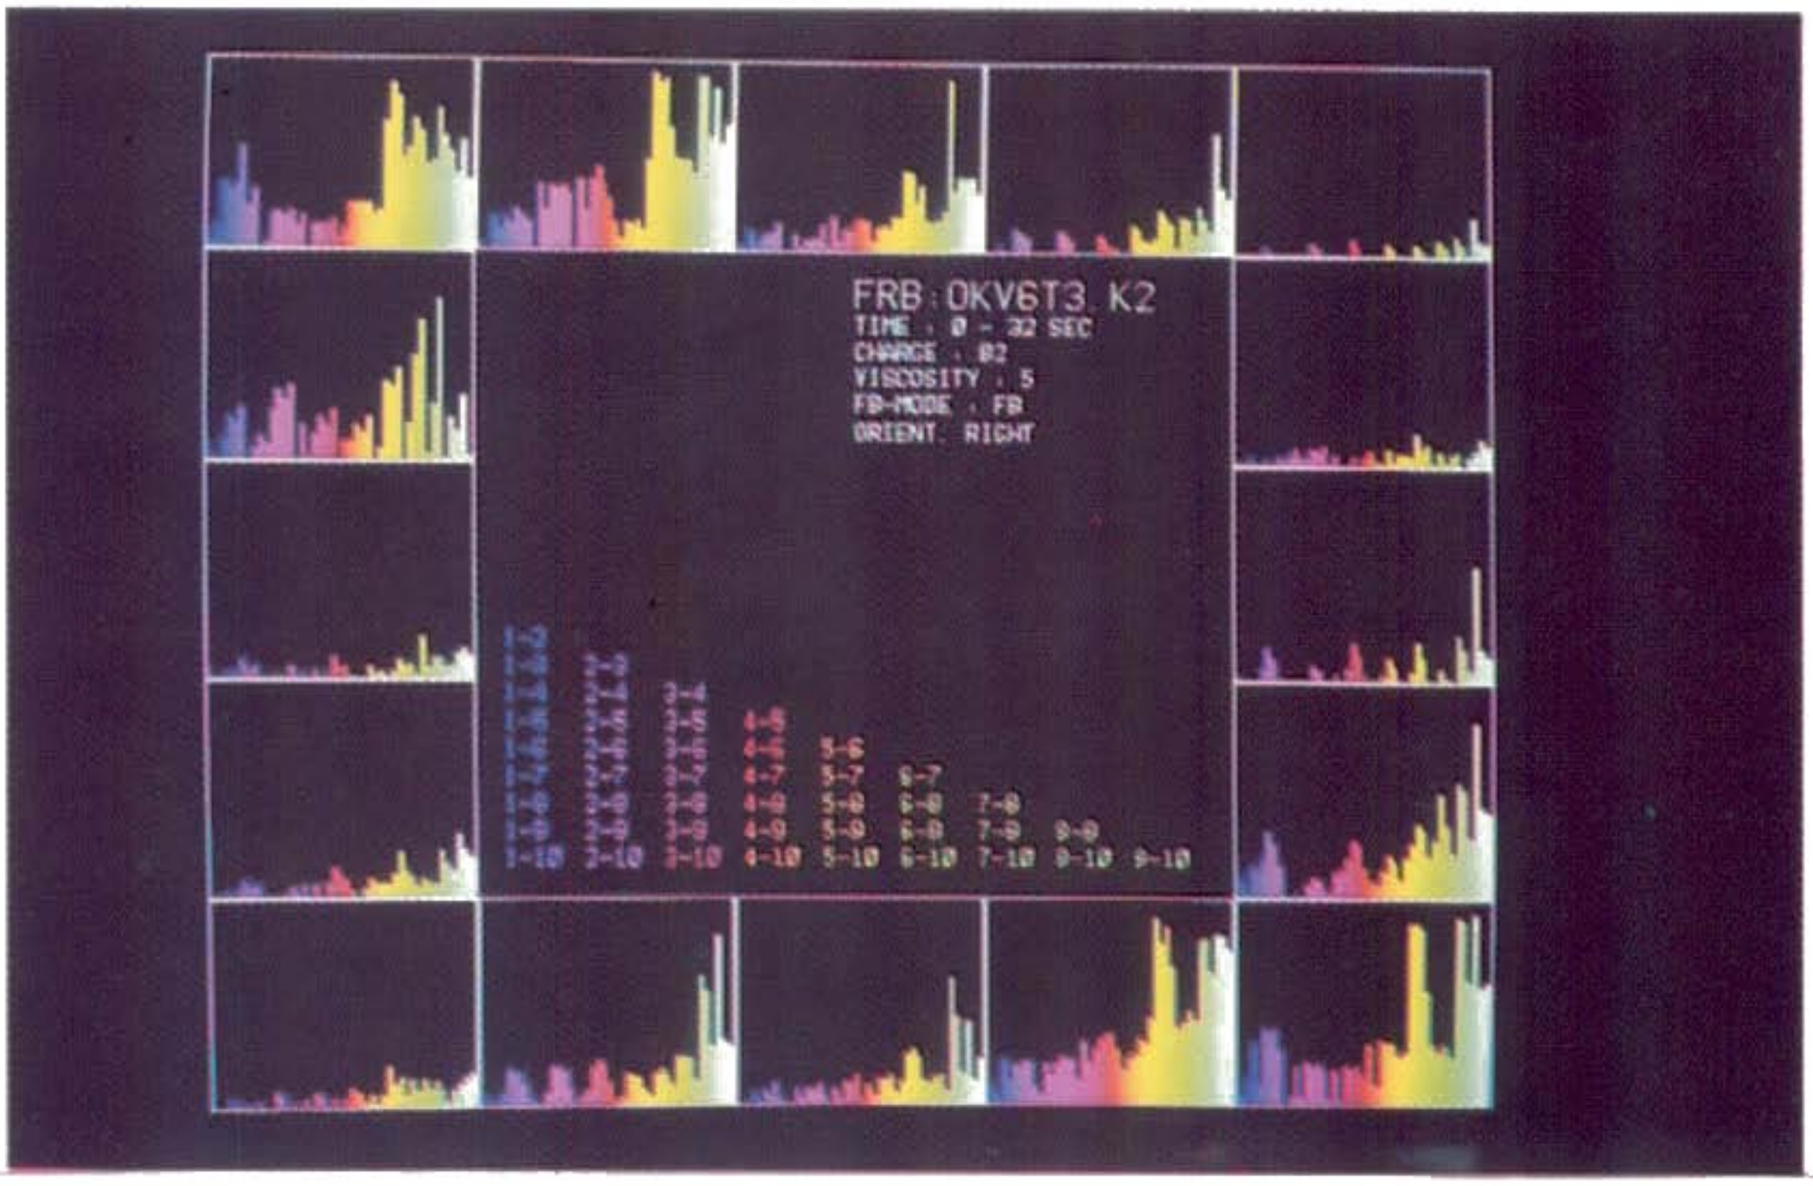

Figure 7. Stimulus dependence of coherence: bargraphs for all 16 directions of stimulus movement. Each graph is positioned on the tip of the vector indicating the direction of motion of the stimulus. Further details as in Fig. 6.

45°, top left (cf. upper graph in Fig. 6) is situated at 45°, top left in Fig. 7. The conclusion we drew from Fig. 6 is confirmed in Fig. 7. The degree of coherence between neurons is shown to be quite strongly stimulus dependent. At this point it should be emphasized once more that this is not simply due to the single neurons showing a stronger response to some directions of movement than to others; the directly stimulus-induced coherence was 'corrected' for by a differential procedure, thus Fig. 7 represents coherence that goes beyond the synchrony in spike trains which is merely due to the temporal correlation of single unit PST-histograms. This suggests that information about (the direction of) a stimulus may very well be represented also in the coherence among firings of different neurons and not only in single unit responses. Although we did not perform a quantitative analysis of direction tuning of single unit responses in our data yet, comparison with Fig. 1 suggests that in this particular experiment the directions with higher coherence globally coincide with the directions of stronger single unit responses. This might be taken to indicate that the differences between the individual bar graphs in Fig. 7 do not so much reflect a genuine stimulus dependence of coherence, but rather suggest that stimulus related rate variations were not adequately normalized for. Indeed it will be argued later (Sect. 4.4) that a true quantitative measure of coherence requires additional dynamic scaling for remaining stimulus effects. Application of this second stimulus 'correction', however, will be shown to confirm and further diversify the stimulus dependence suggested in Fig. 7, rather than reducing it to a mere artifact that can be eliminated by proper scaling.

Finally it should be emphasized that the abstraction of using the bar graphs, i.e. using only the final configuration and not the dynamics of clustering, provides a measure of the summated amount of coherence, integrated over time, and thus necessarily only presents a global picture. Temporal variations of coherence which might signify interesting dynamic reorganizations among different assemblies are integrated out; oscillations may even lead to cancellation and thus complete invisibility of substantial coherence. Furthermore the global picture may be distorted by the earlier mentioned saturation of attraction for very small distances (cf. Fig. 4) which will produce a negative bias for very strong coherence. Evidently for an analysis of temporal variations in multi-unit coherence and their dependence on dynamic properties of the stimulus ensemble one has to consult more detailed and dynamic measures, such as the distance curves (cf. Figs. 4 and 5) and/or other methods of visualization to be described in the next Section.

3.3 Visualization of Clustering: Projection

To envisage a process taking place in N dimensions, so far we have been using a one dimensional method: the analysis of pair distances as a function of time. We have seen that it results in curves which contain a wealth of quite detailed information. However, it may still be difficult for the investigator to imagine, from these curves, the actual particle trajectories in N-space. In order to enhance the spatial impression of the gravitational representation it may be worthwhile to extend our scope to two dimensions: project the N-dimensional particle trajectories on a plane, such that they can easily be visualized e.g. on a video screen.

For the data in our 45°, top left example (cf. Figs. 2 and 4) we have calculated 2-dimensional trajectories from the 10-dimensional ones, using standard geometric projection rules. When plotted on a video monitor this results in an animation movie with particles wandering around on the screen. From this movie six snapshots, taken at equal intervals of 6 sec, are shown in Fig. 8. The projection plane

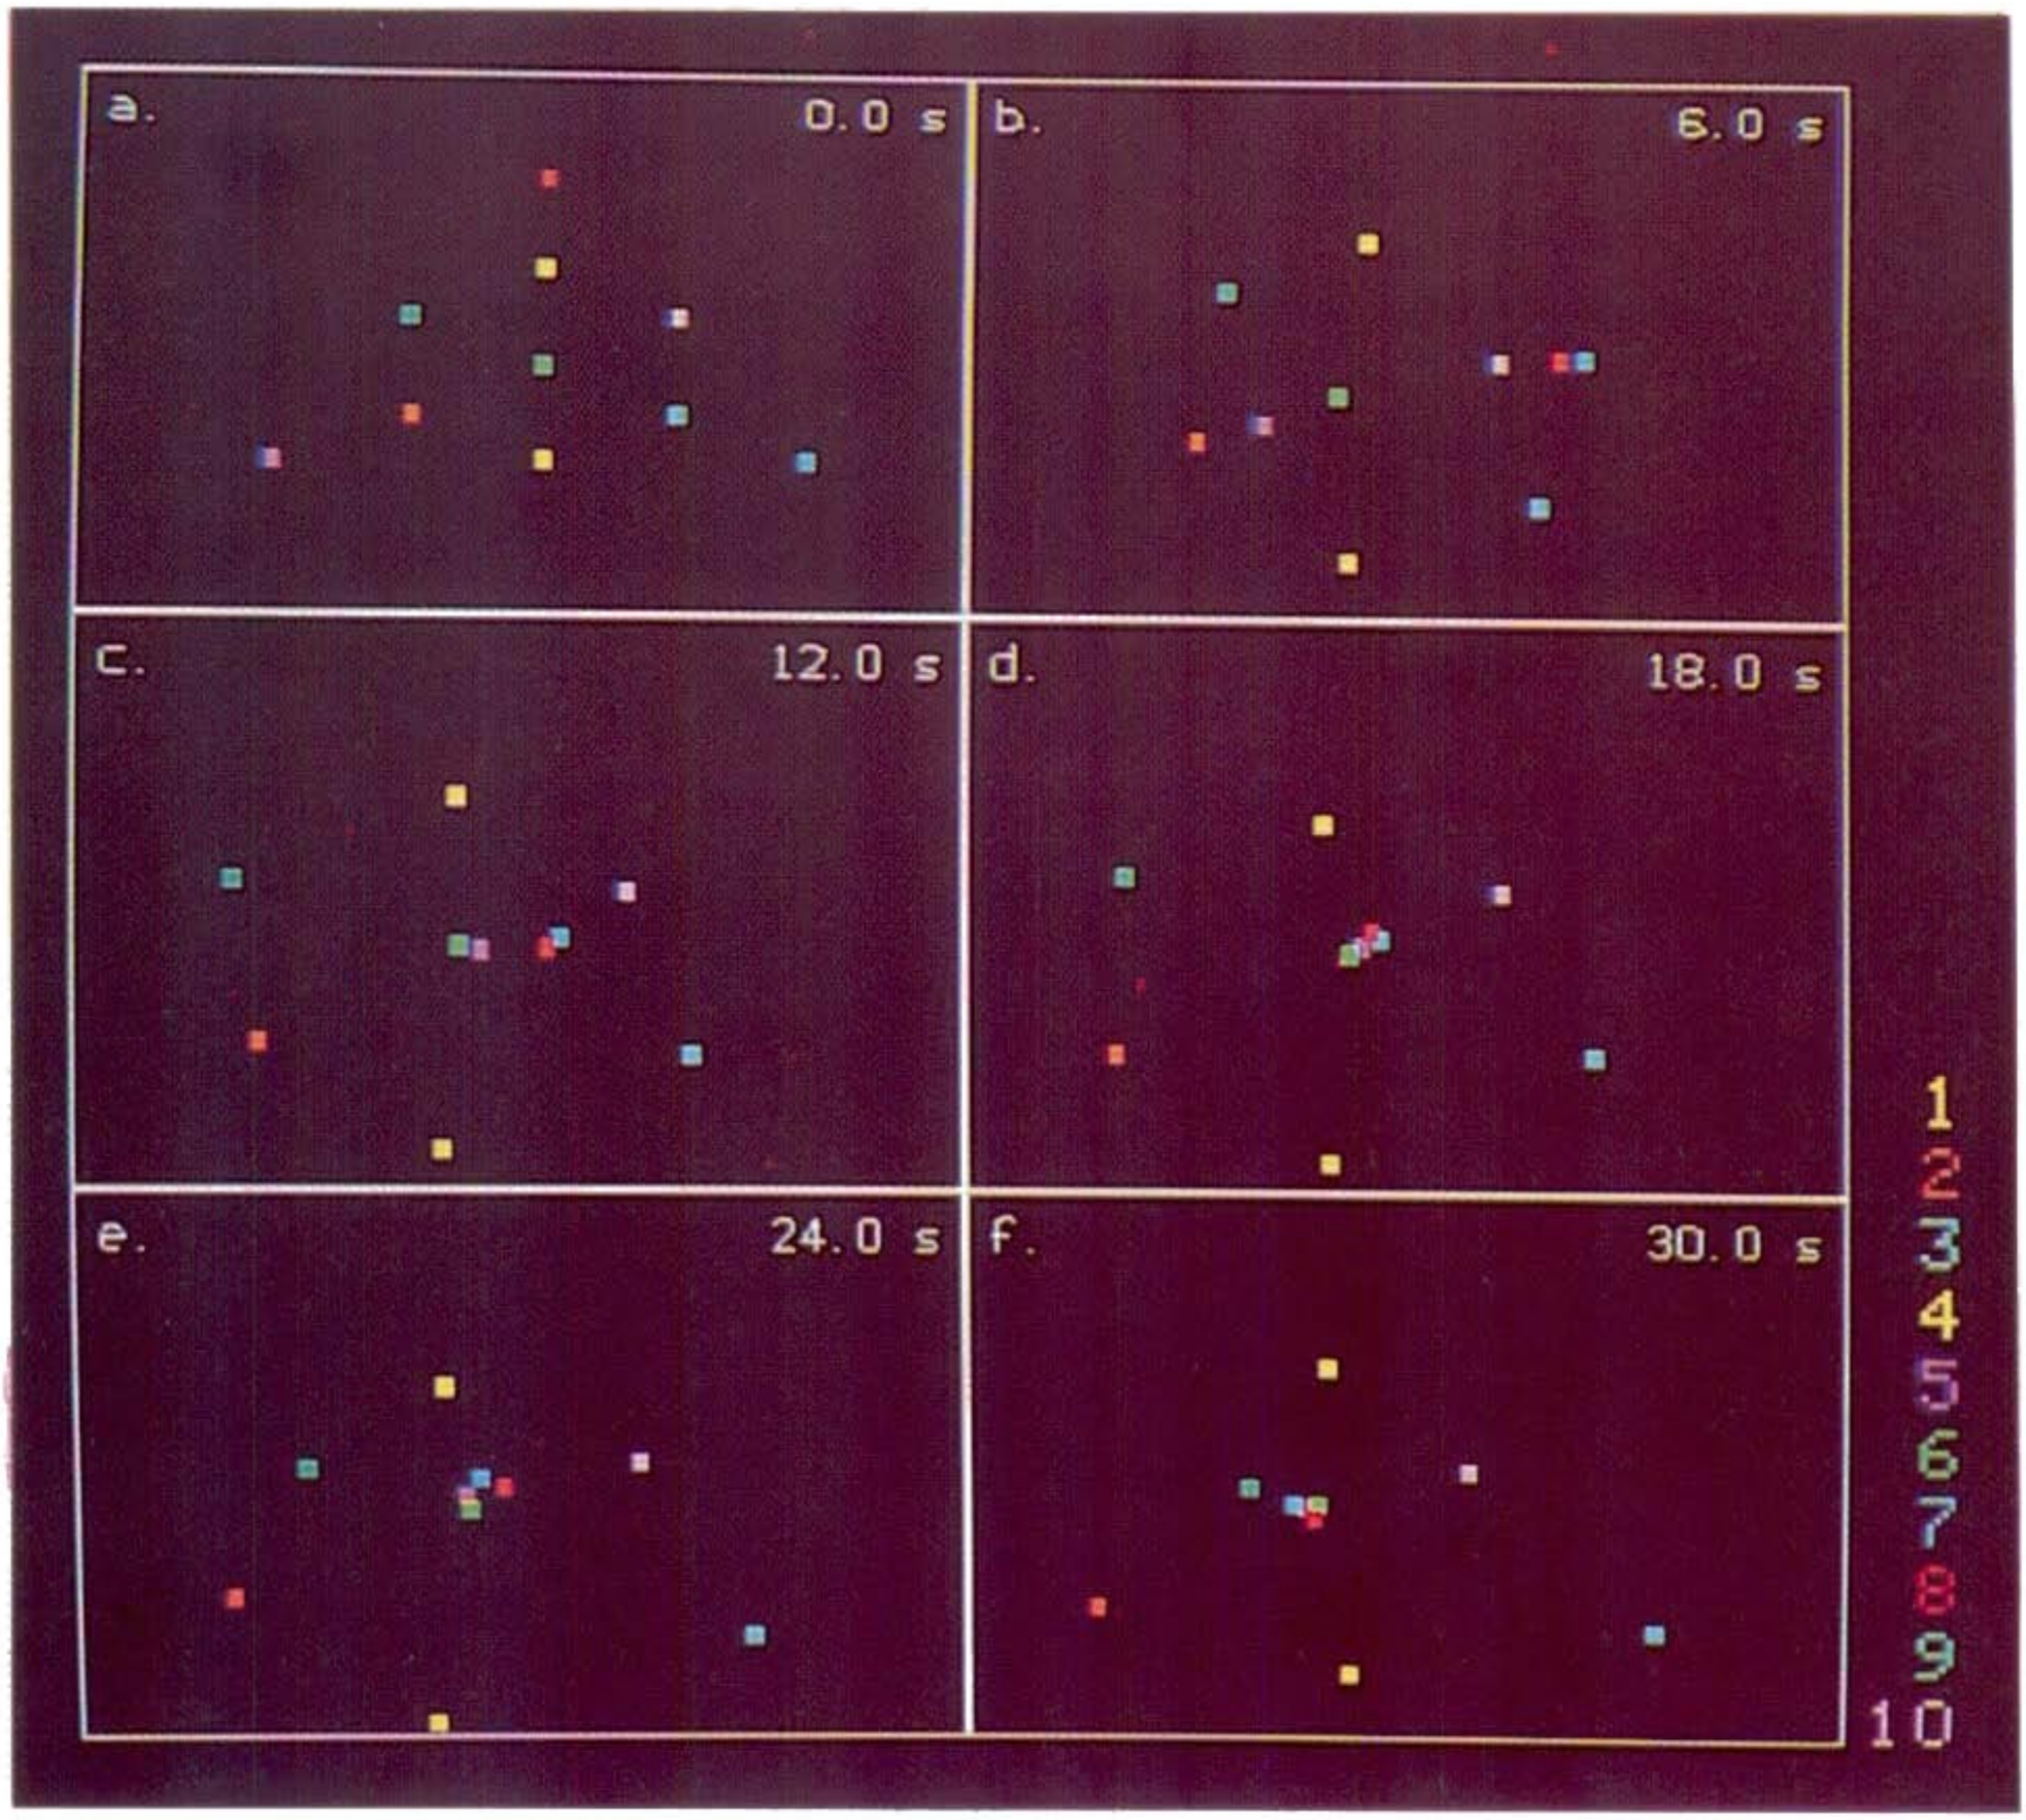

Figure 8. Geometrical projection: six 'snapshots' of the particle trajectories in N-space projected onto a 2-dimensional plane. The snapshots were taken at equal time intervals; time is indicated in the upper right corner of each Figure. Colours code for neurons (cf. Figs. 1 and 2). Further explanation is given in the text.

was chosen such that the initial N-dimensional particle configuration with equal mutual distances is preserved at least to a certain extent (cf. Fig. 8a) (obviously a 2-dimensional projection of 10-space still entails a considerable information loss). Observation of Fig. 8 shows that the information which had to be distilled from Fig. 4 by close inspection of the distance plots, now is much more evident: The particles corresponding to the neurons 5,6,7 and 8 are rather far apart in the beginning but already after (at most) 12 seconds (cf. Fig. 8c) they have formed a cluster at the center of the screen. The finding from Fig. 4 that the neurons 7 and 8 are very strongly connected is also clearly reflected in Fig. 8: already after (at most) 6 sec (cf.Fig 8b) they have formed a tight cluster. As time proceeds the cluster of four continues to collapse, moreover it starts to attract the particles 9 and 10, which apparently are less tightly connected to the former ones. The remaining particles on the average keep their mutual distances and wander around on the screen in a Brownian-motion like manner, which indicates that the corresponding neurons are firing incoherently. It may be noted that, although in the gravitational algorithm we used charge rules such that a distinction is made between the case of a neuron i driving a neuron j and the obverse case of j driving i [14], nevertheless the particles in Fig. 8 are attracting each other in a highly symmetrical way. This indicates that also the coherence among the neurons in our example is quite symmetrical, i.e. although there are clear associations of firing between certain neurons, there is no indication of a specific directionality in these associations. This in turn points to the coherence being the sign of shared neuronal input, rather than being the result of directed connections between members of the group. This observation is in accordance with our finding concerning the neurons 7 and 8 from inspection of their cross correlograms (cf. Fig. 3).

The results of projecting the particle trajectories onto a plane thus give a strong spatial impression of the aggregation process taking place in the N-space. This is especially so when actually witnessing the movements on a screen in the form of an animation movie; the snapshots in Fig. 8 in this respect give only a faint flavour of the dynamic experience. This spatial impression is much harder to be gained through conceptual reconstruction from the distance curves (such as Figs. 4 and 5) and even more cumbersome from analysis of all correlograms involved (such as those in Fig. 3) Especially the hierarchy in clustering, reflecting a hierarchy in coherence and pointing at a hierarchy in assembly organization, becomes quite evident from this type of visualization of the particle trajectories.

As already noted, mapping the N-space onto a 2-space clearly introduces a high information loss. Therefore the results strongly depend on the choice of the projection plane. An unfortunate choice may result in trajectories which are not of much help in analyzing the coherence of the investigated neuronal group. Evidently it is not a desirable situation when the outcome of an analysis technique is so highly dependent on an 'inspired guess' (or, in its absence, on a great deal of trial and error). A promising alternative would be to use the so-called 'structure preserving' nonlinear mapping introduced by Sammon [35] or modifications thereof [5,36]. A related possibility might be to use the self organizing mapping proposed by Kohonen [24]. These methods best preserve the shortest distances in the original N-dimensional distribution; therefore one would expect to see the same clusters as in the direct geometric projections shown here. It remains to be seen, however, how the relationships between clusters will be preserved in this type of approaches. There are a number of unsolved problems involved in this type of data dependent and iterative mappings, such as the dependence of the map on starting configuration, and the

distortion of trajectories. Moreover these methods generally involve heavy computation.
The utility of projection schemes is to provide visual feedback to the investigator in a readily interpretable form involving information about the entire assembly of neurons. An alternative to this type of visualization would be to develop methods that describe the evolution of clusters and their shapes directly in the N-space. An undistorted description in terms of e.g. hyperspheres and hypertubes in N-space might be more accurate and provide more insight than projections. Visualization could then be obtained by (mentally) projecting such gestalts (rather than the individual points) to a two- or three-space [37].

4. THE 'DYNAMIC CORRELATION MATRIX' : A HEBBIAN ALGORITHM

We have also investigated a slightly modified version of the gravitational representation, based on the evaluation of a dynamically modifiable connectivity (correlation) matrix. In this approach also the formal relation which exists between gravitational clustering and cross correlation analysis becomes more explicit. To this end we now shortly return to the conventional approach of analyzing the coherence in a group of neurons by cross correlating pairs of spike trains, two at a time.

4.1 Logical Wiring Diagram and Correlation Matrix

In Sect. 2.2 it was already observed that, when applied to the simultaneously recorded activity of larger groups of neurons, the pairwise analysis of cross correlation inevitably leads to considerable numbers of correlograms. It was soon realized that there is the need for a compact representation of results, in order to avoid a state of conceptual plethora when faced with such numbers of correlograms. One possible way, inspired by graph theory and electronics schematics, is to present 'abbreviated' correlograms of the whole group in the form of a so-called *'logical wiring diagram'* [13], with the width (colour, length) of the connecting wires between any two nodes representing a specific aspect of the correlation of the corresponding neurons (e.g. strength or latency). The 'logical wiring diagram' thus aims at giving a global visualization of the 'functional connectivity' within the observed group of neurons. Note that, unlike in normal electronics diagrams, the connection between any two nodes in a logical wiring diagram may consist of two directional 'wires': the 'connection' between neuron i and neuron j needs not be symmetrical (neither does the cross correlogram).

Another compact way to present the correlations within a larger group of neurons, isomorphic to the previous one, is related to linear algebra: the *'functional connectivity matrix'*, in which the value at a particular entry (i,j) denotes a specific aspect (e.g. strength) of the functional connection from i to j. Again, in general the correlation matrix will not be symmetrical: $C(i,j) \neq C(j,i)$. A visual representation of the 'connectivity matrix' is obtained by displaying the matrix values using a grey- or pseudo-colour coding. A question of considerable interest arises when the connectivity matrix is relatively sparse, i.e. when it contains a considerable number of zero (or very small) elements. In that case one may try to reorganize the matrix into some desirable form by means of an appropriate reordering algorithm [7]. One may hope to arrive in this way at a more or less 'natural' form of the connectivity matrix, which

by its appearance allows direct inference regarding the global properties of the connectivity structure (e.g. a band, a block triangular or a block diagonal form).

4.2 Hebb's Rule: 'Learning' the Correlation Matrix

We have adopted the format of the connectivity matrix in a modified realization of the 'gravitational clustering algorithm'. It was observed [1] that the equations defining the gravitational representation bear a close resemblance to the formalism describing plasticity of synaptic connections between neurons, the so called *'Hebb's rule'* [20]. According to this rule, the principal drive for 'learning', i.e. changing synaptic weight, resides in pair interactions of participating neurons (coherent pre- and post-synaptic firing). The image of particles moving around in N-space, subject to mutual attractive and repellent forces as a consequence of dynamically varying charges, by only minor modifications can be translated into a Hebbian-type learning algorithm (formal equations are given in the Appendix). The functional connections (i.e. correlations) between N neurons are represented in a NxN correlation matrix. Initially the matrix is specified to be uniformly zero, which corresponds to our complete lack of knowledge at the outset of the analysis. With every near-coincident firing of two neurons the value of the corresponding entry in the matrix is increased slightly, completely analogous to the charge interaction between any two particles in the gravity representation: Each spike generates a local influence ('charge') which diminishes in time, this influence is 'sensed' by another neuron, provided it has a spike, the influence of which overlaps in time with that of the first one. The increase in 'synaptic' weight for that particular pair then is proportional to the amount of overlap of the influence (charge) histories. In the course of the analysis the spike trains are processed step by step; at each time step the connective weights in the correlation matrix are updated. Thereby the investigator is able to observe how the matrix gradually 'learns' the functional connectivities underlying the observed spike trains. This learning is based on Hebb's rule which mirrors a supposedly similar development in real neuronal networks.

The updating of matrix entries is defined such that it distinguishes between a spike from neuron i, closely followed by a spike from neuron j (entry i,j), and the opposite case: a spike from neuron j, closely followed by a spike from neuron i (entry j,i) (analogous to the 'two-charge rule' in the gravitational algorithm, see [14]). The resulting connectivity matrix needs not be symmetric, which allows for inference of directionality of the 'connections'. Analogously to gravitational clustering, the time constants in the influence ('charge') histories can be manipulated and thus provide the means to focus on specific temporal configurations in spike patterns. Through the selection of the time course of the elementary (i.e. single spike) charge function, the experimenter effectively makes an operational definition for the 'coherence' being evaluated in that particular analysis run. Throughout the work presented here we have used an exponential waveform, with a time constant of 8 ms.

4.3 Normalization: Stationary Activity

Even in the case of completely incoherent spike trains the algorithm sketched above would result in positive matrix values for all pairs of neurons, simply because in any two spike trains one would a priori expect a number of near-coincidences for purely statistical reasons. For stationary Poisson spike trains this can be corrected for by requiring that each neuron's 'charge' history has a total value of zero, when integrated over time [17]. One possible way is to define the elementary charge

function such that it has zero mean value (bi- or poly-phasic waveform); another way, in fact the one we used here, is to subtract from the various charge histories the corresponding time average, which amounts to a simple DC-shift. For uncorrelated, stationary Poisson spike sequences, this zero-mean-normalization of the charge histories will, in expectation, result in a value of zero for the respective entry in the correlation matrix. A relative abundance of spike-spike and silence-silence combinations will lead to a positive 'connection' ('excitation'), whereas a net surplus of spike-silence combinations will give rise to a negative value ('inhibition').

In order to illustrate the working of this correlation matrix algorithm, we have applied it to artificial spike trains generated by a simulated network of 10 neurons, which was analyzed extensively in the original 'gravity' papers [17,14,1]. The network is shown schematically in Fig. 9a, each of the neurons fires spontaneously according to a Poisson statistic with a spontaneous firing rate of about 10 spikes per sec. The connections in the network all have a strength of 0.35 (on a linear scale from 0 to 1), a latency of 1 ms and a width of 4 ms (for details of the simulator see [3]. Fig. 9b shows the correlation matrices at 4 different moments during the 'learning' analysis: after 1, 2, 4 and 8 seconds. Values in the i-th column of the matrix (i numbered from left to right) represent weight factors of 'outgoing' connections, i.e. originating in neuron i; values in the j-th row (j numbered from bottom to top) represent weight factors of 'incoming' connections, i.e. ending in neuron j. Values along the diagonal (entries i,i) have not been updated and thus represent zero correlation. The strength of the correlation is coded in grey: the larger the value, the darker the grey. The four matrices in Fig. 9b have been drawn with identical scaling in order to enable a direct comparison of matrix values and to emphasize the 'growth' aspect of the analysis.

Comparison of Figs. 9a and 9b shows that already after 2 seconds of 'recording' (i.e. after a mere 20 spikes of each neuron) the connectivity structure starts to show up clearly in the two bottom rows, with additional weak correlations scattered over the entire matrix. These background correlations either represent partial coherence (indirect correlation, e.g. entry 7,8) or are due to purely statistical effects. As time progresses the values in the two bottom rows, corresponding to 'real' connections in Fig. 9a, keep increasing steadily. The other locations in the matrix show this growing behaviour to a clearly smaller extent and rather more erratic. Effectively this implies that, as time proceeds the algorithm, through 'learning', gathers evidence for 'true' correlations, and improves its 'signal to noise ratio' with respect to statistical background correlations. This increase of reliability with time is emphasized in Fig. 9c, where the same matrices from Fig. 9b are shown, but now each matrix was scaled individually, according to its own range of values. Again we notice the dominating 'connections' in the two bottom rows. In this case, however, these values no longer grow (i.e. become darker) in time, a direct consequence of the dynamic scaling applied. In contrast with the relatively stable bottom row correlations, the background correlations appear to be diminishing in time: the noisy background in the matrix gradually cools down to a comfortably quiet landscape, against which the 'true' correlations are gradually gaining significance. It should be noted that in this representation already after one second of recording (only some 10 spikes per neuron, cf. Fig. 9c, matrix at top left) the structure of the network starts to be visible. This makes the dynamic correlation matrix, like the gravitational clustering, a very sensitive tool to analyze coherence in multi-unit spike trains.

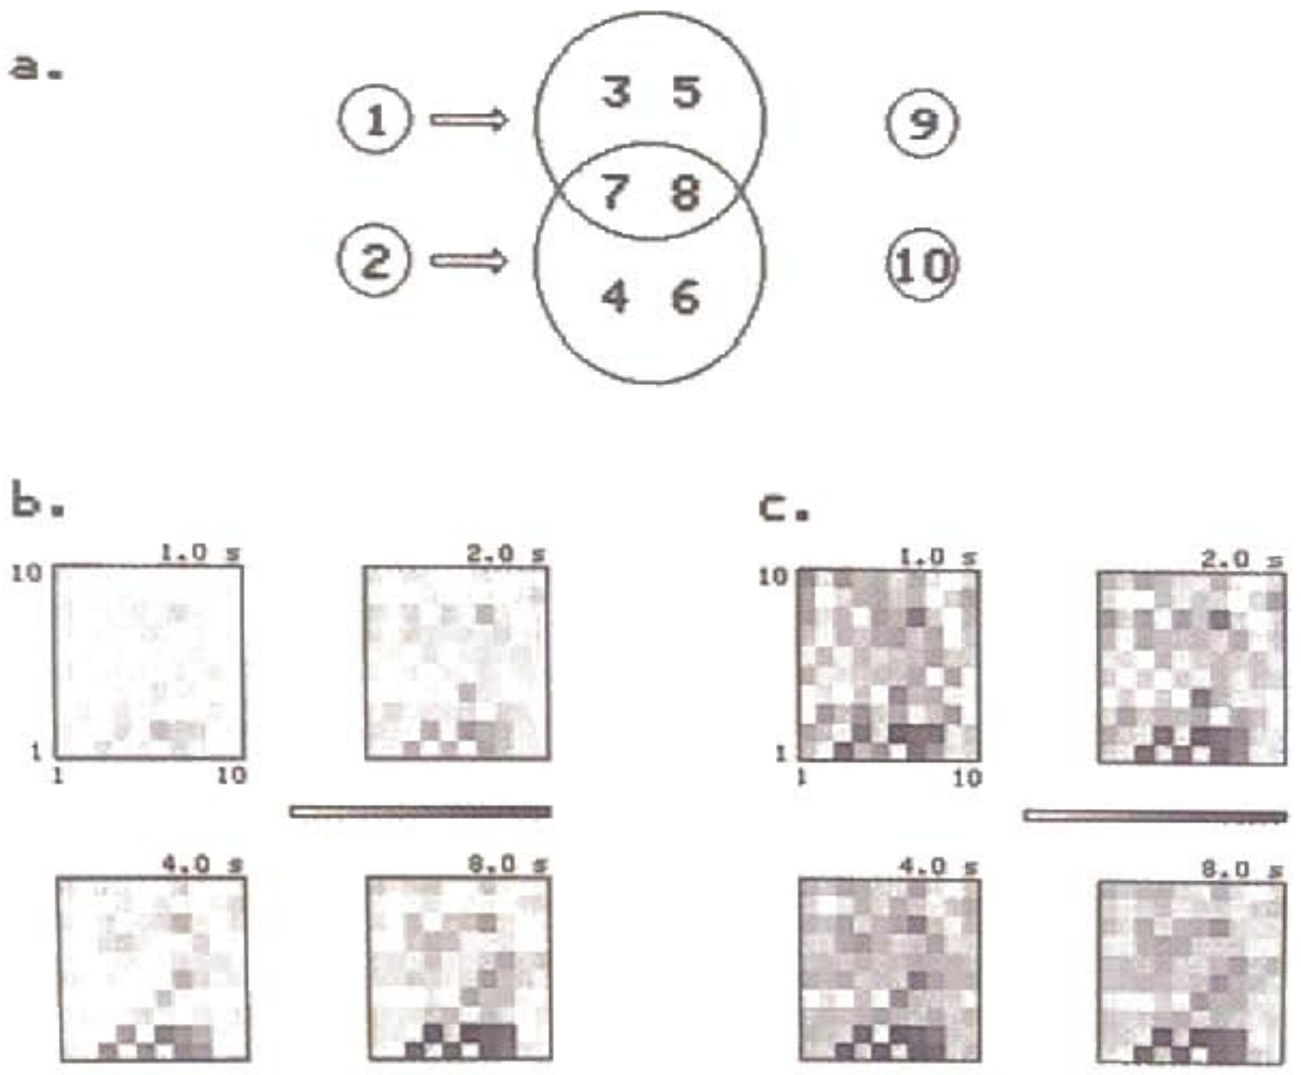

Figure 9. Time evolution of a correlation matrix for a simulated network of neurons. In a. the 'wiring diagram' of the network is given; b. and c. show the evolution of the correlation matrix at different moments in time (indicated at top right). In b. all four matrices are shown with one scaling based on the overall range of values, whereas in c. each matrix is scaled individually according to its own range of values. For further explanation see text.

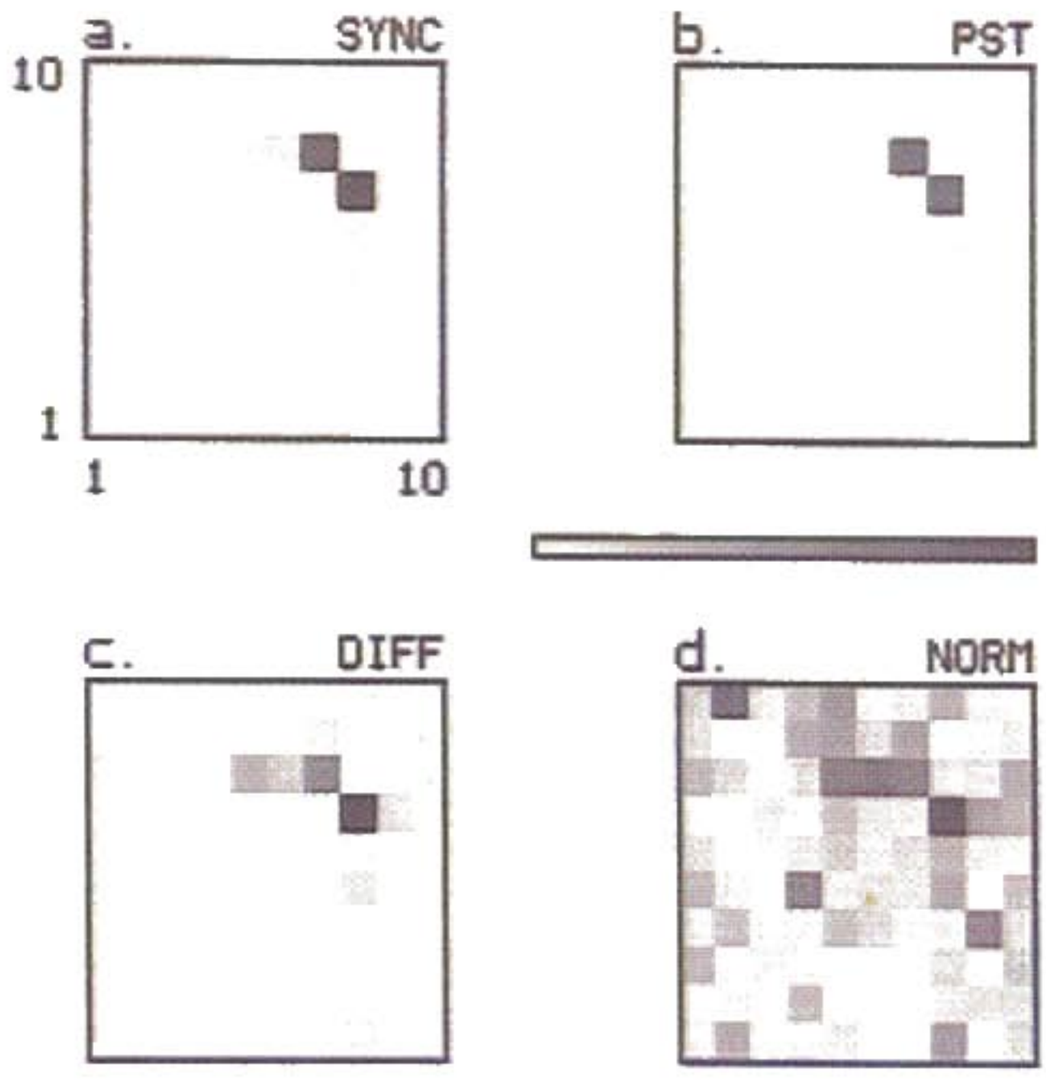

Figure 10. Four versions of 'dynamic correlation matrix' analysis, applied to the spike trains from Fig. 2: a. synchrony matrix; b. PST-matrix; c. differential correlation matrix (first 'correction' for stimulus effects); d. normalized correlation matrix (second 'correction' for stimulus effects). Further explanation is given in the text.

4.4 Normalization: Stimulus-Induced Nonstationary Activity

It is clear from the dot displays in Figs. 1 and 2 that multi-unit activity recorded from the visual cortex under dynamic stimulus conditions does not qualify as a stationary Poisson process. The firing rates of the different neurons show clear variations in time, which appear to be connected to the stimulus: more or less reproducible behaviour over subsequent stimulus presentations. Moreover the time course of these variations appears to be linked to the nature of the stimulus: movement of the bar in different directions leads to different variations of firing rate. This, of course, is all very well known, and in fact is used as the common experimental paradigm to study single neuron response properties. At the same time, however, this stimulus-induced nonstationarity of firing probability presents a nontrivial problem in the analysis of coherence among multi-unit spike trains. First there is a formal problem: the mere fact that spike trains under stimulation become nonstationary causes severe theoretical problems in a field where the mathematical formalism of stationary Poisson processes remains the only mathematical tool to describe spike trains that was developed to any useful extent. Secondly, and more importantly, there is a substantial problem: stimulus-induced nonstationarities provide an additional, strong source of synchrony among spike trains from different neurons (see also our remarks in connection to Fig. 3). In contrast to the simulated example in Fig. 9, neurons may not only fire in close coincidence because they are connected or share a common neural source of input, they may also show coherence in their firing patterns because they are affected in synchrony by some (not necessarily the same) properties of the stimulus (the shared preference for the direction of bar movement among a number of neurons in Figs. 1 and 2 provides a clear example). At this point one may proceed along different, not mutually exclusive ways.

4.4.1 Synchrony Matrix

The first way is simply to treat the spike trains for what they are: time varying signals, transferred from neuron to neuron in a massively interconnected network. Information is thought to be carried both by the time variation of single neuron firing rates as well as by the time varying degree of coherence among spike trains from different neurons. In this view, coherence is considered as a 'code' which signals sensory configurations in the 'outer world' and which at the same time reflects the instantaneous activity state and functional connectivity structure of the neuronal network. This implies that it is worthwhile to look for synchrony of firing, irrespective of what possibly brought it about (stimulation and/or connectivity). This also more or less reflects the way the nervous system is confronted with the problem: how to handle incoming information with time varying characteristics, having no independent knowledge of the external world, or, to put it more strongly: how to 'construct' an external world with its stationary and nonstationary properties, without actually knowing what these are, and only the time varying spike trains to go on.

Based on these considerations, we have analyzed the spike trains from our 45°, top left example (cf. Fig. 2) for synchrony per se, quite irrespective of what caused it (we did use the stationary, zero charge correction, though). The resulting *synchrony matrix* is shown in Fig. 10a, values were coded using a linear grey scale. From this Fig. it is clear that the high synchrony between neurons 7 and 8 (cf. Fig. 3) completely dominates the other pairs. Apart from the synchrony between these two, one also observes a fair amount of co-firing for the pairs (6,8), (5,8), and, although hardly discernible, (5,7) and (6,7). Note that the synchrony of firing for these pairs is

quite symmetrical (cf. Fig. 2). A slightly 'negative' synchrony, i.e. the co-occurrence of firing in one spike train and silence in another one, can be observed (although hardly discernible) for the pair (8,9). This can be understood from the suppression of firing in neuron 9 and the simultaneous activation of neuron 8 (cf. Fig. 2). From the comparison of Figs. 2 and 10a it is also clear that the visibility of synchrony using this kind of measure is strongly influenced by the numbers of spikes involved: the higher the firing rate of a neuron, the larger (in general) the matrix entries involving that neuron. It should therefore be considered seriously whether some kind of scaling procedure, e.g. division by the product of the numbers of spikes of the neurons involved, to arrive at something like a 'synchrony index' might not be the more appropriate way to proceed [9].

4.4.2 PST-Matrix

The second approach to the question of stimulus-induced nonstationarities also involves scaling, but rather for more substantial reasons. This approach interprets coincident firing in different neurons as the sign of 'functional connectivity' and aims at recovering the 'logical wiring diagram' from coherence analysis of the multi-unit spike trains. From this point of view the synchrony brought about by stimulation is considered a nuisance, which masks the effect of primary interest (see also our remarks in connection to the correlograms in Fig. 3); consequently one needs a 'correction procedure' which 'unmasks' the connectivity-related coherence. To this end we developed a procedure, which will be described here heuristically; formal expressions are given in the Appendix, a more fundamental discussion on this methodological issue will be presented elsewhere (Aertsen et al., in preparation).

As a first step we want to measure the synchrony induced by purely stimulus-related modulations of neuronal firing rate. What is asked for, in fact, is an explanatory model of the relation between sensory stimulation and firing probability of the neurons under observation. This in general not being available, we have to suffice with a statistical model, based on the data at hand: different realizations of a multi-variate point process with time varying rates. The best estimate for the stimulus-related nonstationarity that can be obtained from these data is simply the average: for each unit the spike trains recorded during consecutive stimulus presentations are mapped onto a single time interval, covering one stimulus sweep. By this procedure we obtain a new multi-unit spike train, each single train being roughly M times as dense as the corresponding original spike train, the latter of course being M times as long, with M the number of stimulus presentations in the experiment. This new train we call the *multi-unit PST-train*: the time varying density of each single unit component is estimated by the well known peri-stimulus time histogram (PSTH). The multi-unit PST-train presents our best statistical estimate of the relation between stimulation and firing probability. By subjecting this multi-unit PST-train to our correlation matrix analysis we can estimate the purely stimulus-induced component of coherence. The resulting *PST-matrix*, being scaled for the number of stimulus sweeps, for our particular example is shown in Fig. 10b, using the same scaling as in the synchrony matrix (Fig. 10a).

4.4.3 Differential Correlation Matrix

In order to assess the neuronal component of multi-unit coherence one has to compare the synchrony matrix (Fig. 10a) with the PST-matrix (Fig. 10b): any coherence in the multi-unit spike trains which goes beyond purely stimulus-induced

coherence is taken to reflect functional connectivity in the network. Note that this is a one-way statement: from the absence of any difference between the two matrices one cannot infer the absence of functional connectivity; the effects of the latter might simply be buried under (generally much larger) stimulus effects. Not surprisingly (cf. Fig. 3), in the present case the synchrony matrix (Fig. 10a), looks conspicuously like the PST-matrix (Fig. 10b), which once more emphasizes the dominating role played by stimulation. In order to quantify the dissimilarity between synchrony matrix and PST-matrix we chose to calculate the algebraic difference of the two: the *'differential correlation matrix'* equals the synchrony matrix *minus* the PST-matrix (analogous to the subtraction of the 'shift control' in Fig. 3). The result is shown in Fig. 10c, grey scale being adjusted to the range of values in the differential matrix. For the present example this matrix represents the first order estimate of the connectivity contribution to multi-unit coherence. The conclusions from this Figure are completely in register with those from gravitational analysis (cf. Fig. 4), where in fact the same differential procedure was used to get rid of purely stimulus-induced coherence.

4.4.4 Normalized Correlation Matrix

It may be illustrative to note that the actual calculation of these differential measures took a somewhat different path, although numerically the outcome is completely equivalent. In the stationary case (cf. Sect. 4.3) the gravitational representation and the dynamic correlation matrix were calculated on the basis of the respective charge histories, each one corrected for overall rate by undergoing a DC-shift to obtain zero total charge. In the nonstationary case this DC-term, representing the statistical estimate of the time invariant firing rate, was replaced by its time varying analogon: the (properly scaled) charge history of the (periodically continued) PST-train. In other words: the differential correlation matrix measures the coherence between time varying signals which have been *shifted* dynamically to obtain a constant expectation value equal to zero. This, of course, is quite a common procedure in the correlation analysis of ensembles of time varying analog signals (so-called 'cross covariance function' [4,33]. As a matter of fact, this analogy can be pushed one step further: in order to arrive at a *normalized* cross covariance function, i.e. with values between -1 and +1, in conventional signal analysis an additional *scaling* with auto-covariance is applied, equivalent to a transformation to ensembles with zero mean and unit variance. Also in the present case such normalization appears mandatory. First of all, our measure for coherence so far is only a relative measure, not an absolute one. Moreover, in the differential measure we still have rate effects playing an important role for the obvious reason that when the expected firing rate is high, then also the magnitude of the deviation of actual firing rate from expected rate is expected to be high. This implies that, although stimulus induced modulations of firing rates were subtracted, the differential correlation matrix still presents a distorted picture of neuronal coherence, with strongly firing neurons dominating through the sheer amount of their output, and hence, the expected large variability over subsequent stimulus sweeps (compare Fig. 10c with Fig. 2). The same statement holds for the 'differential' gravitational results (cf. Figs. 4-7)

In order to overcome these difficulties, and inspired by the analogy with conventional signal analysis, we have applied additional dynamic scaling to arrive at the *'normalized correlation matrix'*, with matrix values at any moment by definition restricted to the range between -1 and +1. The precise form of the scaling procedure is given in the Appendix, it amounts to a dynamical scaling of entries in the differential correlation matrix with the appropriate time dependent auto covariance of

the (shifted) charge histories (spike trains) involved. The resulting normalized matrix for the present example is shown in Fig. 10d, values range between -0.03 (white) and +0.14 (black). From this Fig. one observes that the normalized matrix indeed differs significantly from the differential matrix (Fig. 10c): the picture now is rather more diverse: the coherence for the original group remains present, but is not dominant anymore, in addition one observes connections entering the scene which so far escaped from being noticed, especially involving the lower electrode numbers. Coherence on the whole assumes only moderate values, with a maximum of 0.14.

This Figure also points to a problem of the normalized correlation matrix: its sensitivity to noise. Due to the scaling procedure, which in fact amplifies the influence of weakly and/or regularly firing neurons at the expense of strongly and/or irregularly firing neurons, certain matrix entries may be blown up for possibly the wrong reason, e.g. because the respective neuron(s) hardly fired. This is for instance the case in the pairs (2,10) and (9,4), where the matrix indicates relatively high coherence for spike trains containing only very few spikes (cf. Fig. 2). Although this may surely point at a genuine albeit only sporadically detectable coherence, it will be hard to establish statistical significance. On the one hand the normalized correlation matrix is asserted to be our best theoretical approximation to the functional connectivity of the group, on the other hand, however, as a statistical estimate it is clearly suffering from a high variance. Results therefore have to be interpreted with care and dot displays of the actual spikes trains should always be within reach; it is also obvious that the analysis of variance of the estimator and the development of criteria for significance are important topics for further research.

4.5 Results

The results of application of these different correlation matrix techniques to the spectrum of all directions of bar movement tested in the experiment (cf. Fig. 1) are summarized in Figs. 11 to 13. Figure 11 shows the synchrony matrices, Fig. 12 the differential correlation matrices and Fig. 13 the normalized correlation matrices. The layout of these Figures is analogous to Fig. 7: each matrix is positioned at the tip of the vector which indicates the corresponding direction of bar movement. In each Figure the 16 matrices were plotted using one scale, covering the range of matrix values in that particular Figure (ranges are indicated above the grey scale in the Fig.).

In Fig. 11 one observes how larger values in the 'synchrony matrices' are basically restricted to a few directions, associated with movement along or close to the negative diagonal, with the pair (7,8) very much dominating. The PST-matrices (not shown here) give practically the same picture. This once more emphasizes the strong influence of stimulation on multi-unit synchrony, simply through the single unit's dependencies on stimulus properties (in this case direction of bar movement, cf. Fig. 1). Subtraction of direct stimulus influence renders the 'differential correlation matrices' (Fig. 12) clearly more sensitive to smaller values of coherence so that the picture becomes more diverse (more neurons involved); still, basically these results are in register with the synchrony matrices: high values for certain neurons in the case of movement along or close to the negative diagonal, for other directions virtually no signs of coherence (see also the bar graphs in Fig. 7).
Additional scaling leads to the 'normalized correlation matrices' in Fig. 13; these present a much more differentiated picture of coherence and especially of its stimulus-dependence. Whereas the earlier noted preference of certain neurons to cohere for movements along the negative diagonal remains present, albeit less clear,

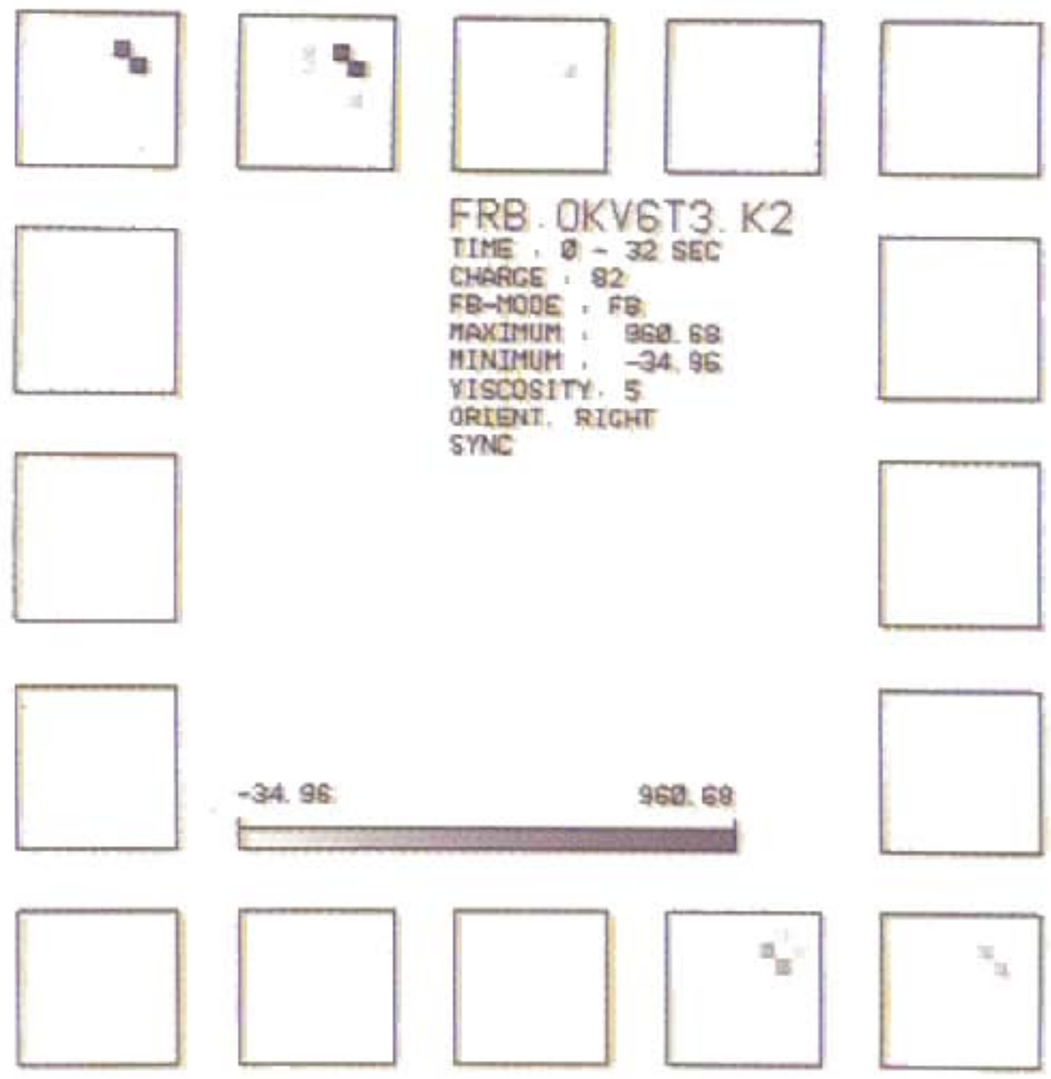

Figure 11. Synchrony matrices for all 16 directions of bar movement (spike trains from Fig. 1). Matrices are positioned at the tip of the vector indicating the direction of motion of the stimulus. All matrices have been plotted according to the same grey scale. Further explanation in the text.

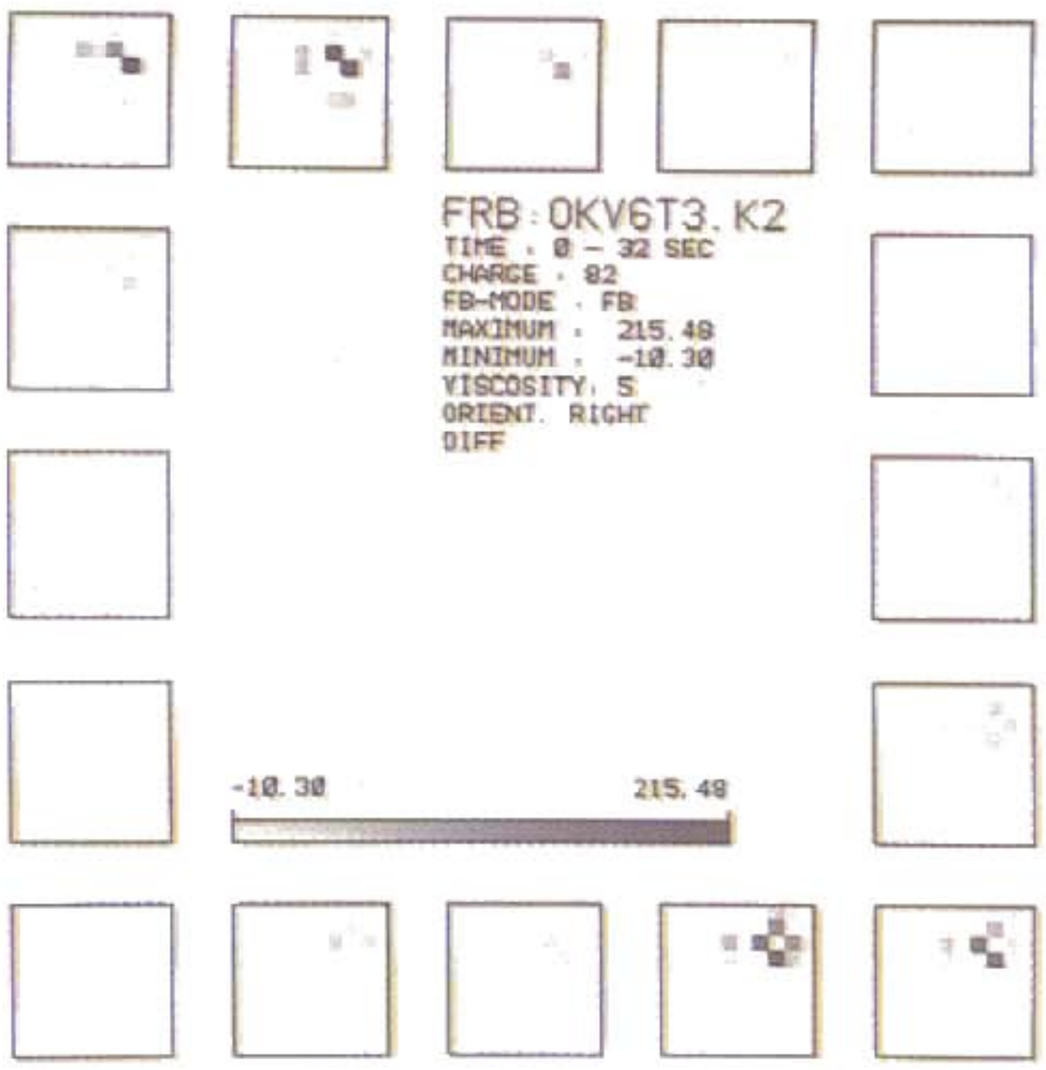

Figure 12. Differential correlation matrices (first stimulus 'correction') for all 16 directions of bar movement (spike trains from Fig. 1). Further details as in Fig. 11.

now also movements in other directions, especially those associated with the positive diagonal, appear to result in increased coherence. Purely horizontal or vertical stimulus movements clearly are less effective in eliciting coherence. In the present example coherence effects are essentially restricted to the electrodes numbered between 5 and 10 (corresponding to recording sites in cortical layers V to I), the actual distribution of coherence varying with direction (e.g. 7 and 8 co-vary strongly along the negative diagonal, 9 becomes involved more strongly with movement along the positive diagonal). Values of the coherence in this experiment attain relatively small values, between -0.05 and +0.20. The asymmetry of this range of values with respect to the value zero may reflect a true asymmetry in the connectivity, although we rather suspect that (a substantial part of) it is due to a preferential sensitivity of our analysis technique for 'excitatory' connections above 'inhibitory' ones (see a discussion on this issue in the context of cross correlograms in [3]. Finally, the matrices in Fig. 13 as compared to those in Fig. 12 also show an increased tendency of not being symmetrical with respect to the diagonal of the matrix.

Figure 13. Normalized correlation matrices (second stimulus 'correction') for all 16 directions of bar movement (spike trains from Fig. 1). Further details as in Fig. 11.

So far our discussion of results from dynamic correlation matrix analysis (Figs. 11-13) has focused on more global properties, such as the degree of (dis)similarity of matrices for different stimuli. Like gravitational clustering, this technique provides a comprehensive and readily digestible overview of the distribution of coherence among larger groups of neurons, and its properties such as dependence on stimulus and time. It should be born in mind, however, that in fact the time development of each individual grey pixel in each of the matrices represents highly condensed information regarding the coherent behaviour of a specific pair of neurons under certain experimental conditions (comparable to the information contained in the

individual distance curves as in Fig. 4). The more detailed analysis of relatively subtle effects like variations of individual matrix entries (c.q. pair distances or projections) is more comparable to conventional analysis of cross correlograms (cf. Fig. 3). An adequate interpretation of results from multi-unit experiments regarding the functional and physiological significance of neuronal interaction requires the continuous interplay of these different levels of analysis.This type of more substantial interpretation of these findings, together with detailed analysis of other multi-unit recordings from similar experiments is subject of current work; preliminary findings seem to be that in general the coherence in these experiments is clearly stimulus-dependent, is confined to relatively low values (roughly between -0.1 and +0.3) and is not so much pointing at small, local circuits but rather at relatively diffuse distributions of coherence, involving rather large fractions of the neurons in the multi-unit recording.
In how far these findings are related to, for instance, the experimental program and the electrode arrangement (linear array, tip distances 160 microns, recording sites positioned vertically with respect to cortical layers) remains subject to further investigation.

5. DISCUSSION

In this paper we have described the recently developed multi-unit analysis technique of 'gravitational clustering' and a technique derived from it: dynamic correlation matrix analysis. These techniques were applied to multi-electrode recordings from the cat's visual cortex, in order to investigate these recordings for possible coherence among the activity patterns of different neurons. This coherence is assumed to reflect the functional connectivity in the neural network recorded from. In the course of this analysis it became clear that the presentation of stimuli provides a strong, if not fully dominating effect on multi-unit coherence, simply through the large, stimulus-induced modulations of single unit firing rates. This evidently poses severe problems to a naive approach of identifying synchrony of firing with neuronal connectivity. Different approaches to this problem were discussed.

The first approach is to apply the measures for coherence simply to the spike trains as they are, thereby measuring the 'synchrony' or 'co-variation' of firing, irrespective of what caused it (stimulation and/or connectivity): the 'synchrony matrix'. In this view multi-unit coherence conveys information to the nervous system about what is going on 'outside' and, at the same time, reflects the instantaneous activity and connectivity status 'inside'. This interpretation of coherence is more or less in line with the proposal that the functional order of a neuronal net may be defined through the order induced by the cross correlations of signals carried by the neurons [22,23]. It is also related to the notions of 'syntactic' and 'semantic' aspects of activity patterns in neuronal populations, as expressed by Johannesma et al. [21].

In the second approach one is mainly interested in the multi-unit coherence as far as it goes beyond purely stimulus driven coherence. The aim is to artificially separate 'outside' from 'inside' by pealing off the strong stimulus-induced coherence. Thereby one hopes to obtain the purely neuronal component of coherence, which may be read as a 'logical wiring diagram', specifying the instantaneous 'functional connectivity' state of the network. In this approach we have argued for a 'stimulus correction' procedure consisting of two stages: subtraction, followed by scaling.

These two approaches are certainly not meant to be mutually exclusive, we rather view them as complementary and in their interrelation providing a conceptual framework for understanding the activity of an entire population of neurons as an entity, rather than as a collection of single elements. A central issue here is to arrive at a quantitative notion of appropriate 'state variables' to describe the neuronal population, analogous to the phase space representation of complex systems, used, for example, in statistical physics. Related to this type of description in terms of micro-variables would be the introduction of the appropriate macro-variables (like e.g. pressure and temperature in thermodynamics) to describe the global properties of the system, disregarding the precise micro-states.
Moreover, the dichotomy between purely stimulus-induced coherence on the one hand and coherence due to 'true' connectivity on the other hand is somewhat artificial to begin with, insofar as, firstly, such a separation requires 'extra-neuronal' knowledge about the experiment, which only the experimenter has, and, secondly, there is assumed to be a strong interaction between the two: connectivity shaping the stimulus-induced coherence (for structural reasons) and stimulus-induced coherence shaping the connectivity (presumably by a Hebb-type mechanism).
Related to this is the paradoxical observation that in order to separate the 'connectivity-coherence' from 'stimulus-coherence', one has to rely on the 'noise' in the system: the more reproducibly the stimulus affects the neurons' firing probabilities, the less room is left for connectivity to make itself known. This is especially clear in the case of highly 'adequate' stimuli, where subsequent presentations evoke strongly driven and virtually identical spike trains: the correlation of multi-unit firing becomes identical to the expected correlation, determined on the basis of PST-trains, and consequently no inference can be drawn regarding possibly underlying connectivity. In other words: if we suppose that a Hebb-type mechanism is operative in shaping the neuronal connectivity, then this implies that the stronger and more reliable it has done so, causing the network to react to external input in a relatively deterministic fashion, the harder it is for an external observer to find out about it. The investigation of functional connectivity by 'pealing off' the 'purely stimulus-induced coherence' through a 'shift control'-type procedure then effectively reduces to a case of 'throwing out the baby with the bath water'.

Throughout this discussion the various approaches were illustrated by results of their application to simulated spike trains and to multi-unit recordings from the cat's visual cortex during presentation of a stimulus ensemble, consisting of a bar moving in different directions. These results show that, like gravitational clustering, the dynamic correlation matrix is truly a very sensitive device to detect possible coherence in multi-neuron activity. This high sensitivity, combined with the multi-dimensionality and the dynamic character inherent to this approach makes these methods a promising tool, which in principle should also be able to detect changes in coherence among spike trains with time constants in the range of one to a few seconds, such as are induced by manipulation of the stimulus and which in our experimental data indeed were shown to be present.

In the present paper we have used the term 'functional connectivity' as purely descriptive: there exists a statistical relationship between the firing of different neurons. This notion of 'functional connectivity' (correlation) is used here as opposed to 'structural' or anatomical connectivity of the neuronal net, involving mono- and poly-synaptic connections. Although, obviously, these two notions are intricately connected, they should not be confused. Functional connectivity may change dramatically on a short time scale, such as induced by manipulation of the sensory

environment, as was demonstrated in the present paper (for other experimental evidence see e.g. [8,13]. It has been proposed that synaptic links in the central nervous system have to be modifiable on the fast time scale of fractions of a second ('synaptic modulation' [29]). Simulation studies, on the other hand, suggest that fast changes in functional connectivity are not necessarily associated with corresponding changes in structural (i.e. synaptic) connectivity [11]. At any rate, the dynamic connectivity structure of neuronal networks has certainly become an element in recent developments in brain theory, and further investigation of the relation between structural and functional connectivity, both theoretically as well as experimentally, appears to be essential.

The formal relations between the techniques of 'gravitational clustering' and the analysis of the 'dynamic correlation matrix' are strong and relatively direct (see also Appendix); the same goes for the kind of information yielded by their application to experimental data, as illustrated by our examples. There are also differences, the main one related to the geometrical nature of the gravitational representation: the influences of different particles (neurons) on another particle (neuron) are combined by vector addition into a single driving force, whereas the correlation matrix keeps pair interactions separated into their respective, modifiable matrix entries. Both approaches have their own virtues, an extensive discussion of which is beyond the scope of the present paper (an example is the quite straightforward relation which can be derived between the dynamic correlation matrix and the 'normal' cross correlation function, see Appendix). Both approaches essentially are highly parallelized calculations of correlations between spike trains. Although so far all calculations were performed on conventional, i.e. serial computer configurations, this parallel nature of both approaches makes them natural candidates to be implemented on a truly parallel machine. This would enable a much faster evaluation of coherence in multi-unit spike trains, with the possibility of having (preliminary) results on-line during the actual experiment; the advantage for the experimenter of having an appropriate feedback cannot be overstressed. In our lab we are currently investigating the possibility of implementing this type of multi-unit analysis techniques on a parallel, hardware realization of an associative network [32]. This would have the additional advantage of closing the conceptual circle: for the analysis of experimental data from real, neuronal networks one uses a machine which in itself represents a model of such a neuronal network.

6. ACKNOWLEDGEMENTS

Calculations were performed using a Fortran software package developed by one of the authors (AA), which was integrated with the software package MATFUN, developed at the Dept. of Medical Physics and Biophysics of the University of Nijmegen, The Netherlands, and implemented on the VAX 750 at the MPI-Tübingen. The latter package was also used for graphics display of the results.
The authors wish to thank Peter Johannesma and Günther Palm for stimulating discussions, Manfred Caeser for his work on the evaluation of histological data, Volker Staiger for skillful assistance in generating the Figures and Shirley Würth for helping out in the final processing of the manuscript.

7. APPENDIX

7.1 Gravitational Clustering

With each one of a group of N neurons we associate a particle in N-space, with initial position $\vec{r}_i(0)$ (i=1,...,N) such that all particles are positioned on the vertices of an N-dimensional hypercube; therefore all initial mutual distances are equal to some constant d_0. After the system starts to evolve, the position of the i-th particle is at any time t given by the N-dimensional vector $\vec{r}_i(t)$.
Let $z_i(t)$ be the spike train from neuron i. The charge function associated with particle i, used to generate inter-particle forces is given by a filtered version of the spike train:

$$Q_i(t) = \int ds\ q(s)\ z_i(t-s) \qquad (A1)$$

In the 'two charge model' [14], designed to allow inference regarding the *direction* of interaction, with each particle we associate two charges: an *effector* charge which generates the propulsive field:

$$Q_{e,i}(t) = \int ds\ q_e(s)\ z_i(t-s) \qquad (A1a)$$

and an *acceptor* charge, used to calculate the force on that particular particle:

$$Q_{a,i}(t) = \int ds\ q_a(s)\ z_i(t-s) \qquad (A1b)$$

In the present paper we have used for the elementary effector charge $q_e(t)$ a decaying exponential starting at and following the spike, and for the elementary acceptor charge $q_a(t)$ a rising exponential that terminates at the spike.
The propulsive field $\vec{E}_{ij}$ at position $\vec{r}_i$ generated by the particle j at position $\vec{r}_j$ is given by

$$\vec{E}_{ij} = Q_{e,j}\ \hat{r}_{ij} \qquad (A2)$$

with the unit vector $\hat{r}_{ij}$ given by

$$\hat{r}_{ij} = \frac{\vec{r}_{ij}}{r_{ij}} = \frac{\vec{r}_j - \vec{r}_i}{\| \vec{r}_j - \vec{r}_i \|} \qquad (A3)$$

From (A2) one notes that the field was chosen to be independent of distance. The force exerted on particle i by particle j is then given by

$$\vec{F}_{ij} = Q_{a,i}\ \vec{E}_{ij} = Q_{a,i}\ Q_{e,j}\ \hat{r}_{ij} \qquad (A4)$$

The total force on i is given by the vector sum of all the pair forces:

$$\vec{F}_i = \sum_{j \neq i} \vec{F}_{ij} = \sum_{j \neq i} Q_{a,i}\ Q_{e,j}\ \hat{r}_{ij} \qquad (A5)$$

As a consequence of this force the particles will move. The dynamic equation is defined by

$$\mu \dot{\vec{r}}_i = \vec{F}_i \qquad \text{(A6)}$$

with μ denoting the 'viscosity' of the medium in N-space. Note that the acceleration term is lacking in (A6): we assume high viscosity and therefore velocity is proportional to force. The displacement, finally, is calculated by numeric integration (Euler) using a time step δ :

$$\vec{r}_i(t+\delta) = \vec{r}_i(t) + \frac{\delta}{\mu} \vec{F}_i(t) \qquad \text{(A7)}$$

7.2 Dynamic Correlation Matrix

The formalism of the dynamic correlation matrix is a simple modification of the equations describing 'gravitational clustering'. With N neurons we associate a NxN-correlation matrix C, with the entry C_{ij} denoting the correlation ('functional connection') *from* neuron j *to* neuron i, i.e. related to the probability of having an i-event *after* a j-event. Note that this implies that C needs not be symmetrical: direction of connectivity is preserved. Initially the matrix C is set to be uniformly zero, signifying our complete lack of knowledge at the outset of the experiment. In the course of time the matrix values are gradually updated, according to a Hebb-type rule for synaptic modification, which is directly borrowed from the gravitational representation: the pair force F_{ij} (Eq. A4) is interpreted as the drive for modification of the 'connection', very much the same as the dynamic equation for the particle motion (Eq. A6):

$$\mu \dot{C}_{ij} = Q_{a,i} Q_{e,j} \qquad \text{(A8)}$$

In our correlation matrix algorithm we have implemented a somewhat more general formulation which also incorporates a leakage term:

$$\mu \dot{C}_{ij} = Q_{a,i} Q_{e,j} - \alpha C_{ij} \qquad \text{(A9)}$$

This extra term, which signifies 'forgetting' of earlier 'experience', is quite common in theoretical work on synaptic plasticity. In the present paper, though, we have always used Eq. (A8) (i.e. $\alpha = 0$), to enable a better comparison with the gravitational algorithm.

The main difference between the gravitational representation and the dynamic correlation matrix is the vector addition (Eq. A5), which is missing in the latter approach. In the gravitational representation the influence of different particles (neurons) on another particle (neuron) is combined into a single force, the correlation matrix keeps all the pairs separated in the respective, modifiable matrix entries. Both approaches have their own advantages, one advantage of the correlation matrix is that certain relations can be derived in a straightforward manner, which, in the gravitational approach, is more cumbersome because of the dynamic interaction of force vectors.

7.3 Relation between 'dynamic correlation matrix' and normal 'cross correlation'

As an example of such a relation it may be illustrative to show the explicit relation between the dynamic correlation matrix and the 'normal' cross correlation function of spike trains. The net 'change' of a connection C_{ij} over some period of time, say from t_1 to t_2, is given by

$$\mu(C_{ij}(t_2) - C_{ij}(t_1)) = \int_{t_1}^{t_2} dt\ Q_{a,i}(t)\ Q_{e,j}(t) \qquad (A10)$$

Substitution of the expressions for the charge functions (Eqs. A1a and A1b) leads to

$$\mu(C_{ij}(t_2) - C_{ij}(t_1)) = \int ds_1 q_a(s_1) \int ds_2 q_e(s_2) \int_{t_1}^{t_2} dt\ z_i(t-s_1) z_j(t-s_2) \qquad (A11)$$

On the right hand side one easily recognizes a term which also appears in the time dependent cross correlation function:

$$R_{ij}(\tau_1, \tau_2) = E\,(z_i(\tau_1) z_j(\tau_2)) \qquad (A12)$$

A considerable simplification is obtained when the correlation is *not* time dependent, i.e. for stationary processes z_i and z_j:

$$R_{ij}(\tau_1, \tau_2) = R_{ij}(\tau_1 - \tau_2) \qquad (A13)$$

This implies that, for *stationary* spike trains, Eq. (A11) can be rewritten as

$$\mu(C_{ij}(t_2) - C_{ij}(t_1) = \int ds_1 q_a(s_1) \int ds_2 q_e(s_2)\ R_{ij}(s_1 - s_2; t_1, t_2) \qquad (A14)$$

where

$$R_{ij}\,(s_1 - s_2; t_1, t_2) = \int_{t_1}^{t_2} dt\ z_i(t-s_1)\ z_j(t-s_2) \qquad (A15)$$

is the estimate of the cross correlation between i and j, obtained by *time averaging* over the interval between t_1 and t_2.
Equation (A14) in fact states that, for stationary spike trains, the net change of the 'connection' C_{ij} over the internal (t_1, t_2) is proportional to the (t_1, t_2)-estimate of the cross correlation function between neurons i and j, 'smoothed' by $q_e(t)$ and 'weighted' by $q_a(t)$. From this relation it becomes clear how, by appropriate choice of q_e and q_a one can obtain directional information from the cross correlogram.

The zero total charge normalization, applied in the case of stationary spike trains (Sect. 4.3), amounts to a DC-shift in the charge functions, equal to the (time averaged) estimate of firing rates $\bar{z}_i$ resp. $\bar{z}_j$:

$$\mu(C_{ij}(t_2) - C_{ij}(t_1)) = \int ds_1 q_a(s_1) \int ds_2 q_e(s_2) \int_{t_1}^{t_2} dt\ (z_i(t-s_1)-\bar{z}_i)(z_j(t-s_2)-\bar{z}_j) \qquad (A16)$$

This DC-shift effectively implies a translation from cross correlation to cross covariance: the departure of the cross correlation from its background level due to the non-zero expectations $\bar{z}_i$ and $\bar{z}_j$.

The *'differential correlation matrix'* applied for *nonstationary* spike trains as a first approximation to 'functional connectivity', with direct stimulus effects subtracted (Sect. 4.4.3), uses precisely the same formalism as given in (A16), with the difference that now the expected firing rates are time dependent, and in the case of periodic stimulation can be estimated through the *PST-trains* (Sect.4.4.2):

$$M\,\bar{z}_i(t) = \sum_{m=1}^{M} z_{i,m}(t) \qquad (A17)$$

with M the number of stimulus presentations, and t ranging from 0 to T, T being the duration of one stimulus sweep. It is clear that in this case (A16) is directly related to the time dependent cross covariance function of z_i and z_j (e.g. [4]). Furthermore, one easily sees that, with $\bar{Q}(t)$ denoting the charge function for the PST-trains, scaled for the number of sweeps M, it holds that

$$\mu(C_{ij}(MT) - C_{ij}(0)) = \int_0^{MT} dt\ (Q_{a,i}(t) - \bar{Q}_{a,i}(t))(Q_{e,j}(t) - \bar{Q}_{e,j}(t)) =$$

$$= \int_0^{MT} dt\ Q_{a,i}(t)\, Q_{e,j}(t) - \int_0^{MT} dt \bar{Q}_{a,i}(t)\bar{Q}_{e,j}(t) \qquad (A18)$$

or: the charge-shift procedure to calculate the differential correlation matrix, is equivalent to subtraction of the *'PST-matrix'* (Sect. 4.4.2) from the *'synchrony matrix'* (Sect. 4.4.1).

The *'normalized correlation matrix'* (Sect. 4.4.4) finally is obtained by additional scaling for auto covariances:

$$C'_{ij}(t_2) - C'_{ij}(t_1) =$$

$$= \frac{\int_{t_1}^{t_2} dt \, \{Q_{a,i}(t) - \bar{Q}_{a,i}(t)\} \, \{Q_{e,j}(t) - \bar{Q}_{e,j}(t)\}}{\left\{ \int_{t_1}^{t_2} dt \, \{Q_{a,i}(t) - \bar{Q}_{a,i}(t)\}^2 \cdot \int_{t_1}^{t_2} dt \, \{Q_{e,j}(t) - \bar{Q}_{e,j}(t)\}^2 \right\}^{\frac{1}{2}}} \qquad \text{(A19)}$$

It is easily shown that, because of the Cauchy-Schwarz inequality, the scaling in (A19) effectively results in matrix values restricted to the range between -1 and +1, i.e. the procedure in (A19) is truly a normalization.

8. REFERENCES

1 Aertsen, A., Gerstein, G. and Johannesma, P., From neuron to assembly: Neuronal organization and stimulus representation. In: G. Palm and A. Aertsen (eds.), Brain Theory, pp. 7-24, Springer, Berlin (1986).

2 Aertsen, A.M.H.J., Erb, M. and Johannesma, P.I.M., The 'Neurophone': Acoustic representation of neural activity patterns (in preparation).

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

4 Bendat, J.S. and Piersol, A.G., Random data: analysis and measurement procedures. Wiley-Interscience, New York, (1971).

5 Biswas, G., Jain, A.K. and Dubes, R.C., Evaluation of projection algorithms. IEEE Trans Pattern Anal Machine Intell, **PAMI-3**, 701-708. (1981).

6 Braitenberg, V., Cell assemblies in the cerebral cortex. In: R. Heim and G. Palm (eds.), Theoretical approaches to complex systems. Lecture Notes in Biomathematics, pp. 171-188, Springer, Berlin (1978).

7 Duff, I.S., A survey of sparse matrix research. Proc IEEE, **65**, 500-535 (1977).

8 Eggermont, J.J., Epping, W.J.M. and Aertsen, A.M.H.J., Stimulus dependent neural correlations in the auditory midbrain of the grassfrog (Rana temporaria L.). Biol Cybern, **47**, 103-117 (1983).

9 Epping, W.J.M., Auditory information processing in the midbrain of the grassfrog. Ph.D. Dissertation. University of Nijmegen, The Netherlands (1985).

10 Epping, W., Boogard, H. van den, Aertsen, A., Eggermont, J. and Johannesma, P., The Neurochrome: An identity preserving representation of activity patterns from neural populations. Biol Cybern, **50**, 235-240 (1984).
11 Erb, M., Palm, G., Aertsen, A. and Bonhoeffer, T., Functional versus structural connectivity in neuronal nets. In: Proceedings 9. Kybernetik-Kongress (DKG), p. 23 (abstract). Göttingen, FRG (1986).
12 Gerstein, G., Functional association of neurons: detection and interpretation. In: F.O. Schmitt (ed.), The Neurosciences: Second Study Program, pp. 648-661, Rockefeller University Press, New York (1970).
13 Gerstein, G., Aertsen, A., Bloom, M., Espinosa, I., Evanczuk, S. and Turner, M., Multi-neuron experiments: Observation of state in neural nets. In: H. Haken (ed.), Complex Systems - Operational Approaches, pp. 58-70, Springer, Berlin (1985).
14 Gerstein, G.L. and Aertsen, A.M.H.J., Representation of cooperative firing activity among simultaneously recorded neurons. J Neurophysiol, **54**, 1513-1527 (1985).
15 Gerstein, G.L., Bloom, M.J., Espinosa, I.E., Evanczuk, S. and Turner, M.R., Design of a laboratory for multi-neuron studies. IEEE Trans Systems, Man and Cybernetics, **SMC-13**, 668-676 (1983).
16 Gerstein, G.L. and Perkel, D.H., Mutual temporal relationships among neuronal spike trains. Biophys J, **12**, 453-473 (1972).
17 Gerstein, G.L., Perkel, D.H. and Dayhoff, J.E., Cooperative firing activity in simultaneously recorded populations of neurons: detection and measurement. J Neurosci, **5**, 881-889 (1985).
18 Grinvald, A., Real-time optical imaging of neuronal activity. TINS, 7, 143-150 (1984).
19 Grinvald, A., Real-time optical mapping of neuronal activity: From single growth cones to the intact mammalian brain. Ann Rev Neurosci, **8**, 263-305 (1985).
20 Hebb, D.O., The organization of behaviour. A neuropsychological theory. Wiley, New York (1949).
21 Johannesma, P., Aertsen, A., Boogard, H. van den, Eggermont, J. and Epping, W. From synchrony to harmony: Ideas on the function of neural assemblies and on the interpretation of neural synchrony. In: G. Palm and A. Aertsen (eds.), Brain theory, pp. 25-47, Springer, Berlin (1986).
22 Koenderink, J.J., Simultaneous order in nervous nets from a functional standpoint. Biol Cybern, **50**, 35-41 (1984).
23 Koenderink, J.J., Geometrical structures determined by the functional order in nervous nets. Biol Cybern, **50**, 43-50 (1984).
24 Kohonen, T., Self-organized formation of topologically correct feature maps. Biol Cybern, **43**, 59-69 (1982).
25 Krüger, J., A 12-fold microelectrode for recording from vertically aligned cortical neurones. J Neurosci Methods, **6**, 347-350 (1982).
26 Krüger, J., Investigation of a small volume of neocortex with multiple microelectrodes: evidence for principles of self-organization. In: H. Haken (ed.), Complex Systems - Operational Approaches, pp. 71-80, Springer, Berlin (1985).
27 Krüger, J., Simultaneous individual recordings from many cerebral neurons: Techniques and results. Rev Physiol Biochem Pharmacol, **98**, 177-233 (1983).
28 Krüger, J. and Bach, M., Independent systems of orientation columns in upper and lower layers of monkey visual cortex. Neurosci Lett, **31**, 225-230 (1982).
29 Malsburg, C. von der, Nervous structures with dynamical links. Ber Bunsenges Phys Chem, **89**, 703-710 (1985).
30 Moore, G.P., Segundo, J.P., Perkel, D.H. and Levitan, H., Statistical signs of synaptic interaction in neurons. Biophys J, **10**, 876-900 (1970).

31 Palm, G., Neural assemblies: An alternative approach to artificial intelligence Springer, Berlin, (1982).
32 Palm, G. and Bonhoeffer, T., Parallel processing for associative and neuronal nets. Biol Cybern, **51**, 201-204 (1984).
33 Papoulis, A., Probability, random variables and stochastic processes McGraw-Hill Kogakusha, Tokyo, (1965).
34 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).
35 Sammon, J.W., Jr., A nonlinear mapping for data structure analysis. IEEE Trans Computers, **C-18**, 401-409 (1969).
36 Terekhina, A.Yu., Methods of multidimensional data scaling and visualization (survey). Avtom Telemekh, 7, 80-94 (1973).
37 Zahn, C.T., Graph-theoretical methods for detecting an describing Gestalt clusters. IEEE Trans Computers, **C-20**, 68-86 (1971).