

On the Relation Between Neural Modelling and Experimental Neuroscience*

Thomas Wennekers and Günther Palm
University of Ulm,
Dep. of Neural Information Processing,
D-89069 Ulm, Germany
phone: +49 731 502 4151
fax: +49 731 502 4156
email: (thomas,palm)@neuro.informatik.uni-ulm.de

Keywords: temporal correlations, gamma-oscillations, global binding, neural networks, brain theory

Abstract

This paper discusses the relation of theory and experiment in neuroscience exemplified by three assumptions often made in models of coherent activation in the cortex: basic feature-coding oscillators, phase-coding and global binding of whole objects. Apparently these assumptions are not very well supported by the experimental evidence. We propose that it is the single synchronized population-burst that matters: spikes of feature-coding cells are temporally clustered in our opinion by recurrent associative processes. In each burst a single stimulus is processed (if there are several). Synchronization is restricted to cortical sites which physically interact. These principles are illustrated by computer simulations.

1 Introduction

The field of neuroscience has rapidly developed in the last 50 years. It is therefore only natural that a theoretical perspective - largely based on mathematics and physics - has begun to evolve, and its relation to the

* Appeared in slightly modified form in: Theory in Biosciences 116:273-289,1997.

ongoing experiments has turned out to be particularly intricate. There are several reasons for this:

- Since many experimental and theoretical developments are very recent many experimenters have not yet appreciated the need of theorizing, and conversely, many theoreticians are not aware of the most recent experimental observations.
- The gap between the level of behavior that has to be explained in terms of the neural machinery and the currently investigated level of biophysics and biochemistry of this machinery, is tremendously large.

In spite of these problems neural modelling has meanwhile become a discipline of its own standing. It can be viewed as a part of computer science, which tries to understand the processing of information in neural networks. A strong technical motivation of this enterprise is to exploit the results for the construction of better (i.e. more flexible, adaptive, fault tolerant, parallel, human) computers by embedding artificial neural networks into the soft- and hardware of modern computers.

In this article, however, our concern is of course the use of artificial neural networks as a modelling tool for theoretical brain research. This approach already has its own history and we can only refer to some important books, review articles and reprint collections for a rough orientation (MacGregor 1987; Anderson and Rosenfeld 1988; Shaw and Palm 1988; Palm 1982, 1990; Abeles 1991).

Instead of trying to discuss all this on some kind of meta-level, we have decided to focus here on one specific issue, which is today heatedly discussed in the neuroscience community. We believe, that the discussion of this issue will reveal a style of argumentation that is typical for the relation of theoretical and experimental paradigms in neuroscience.

2 Gamma oscillations, synchronization and models of binding

Experiments on anesthetized cats and alert behaving monkeys have shown that local populations of cells in peripheral visual areas often respond rhythmically with frequencies in the γ -range (30-90Hz) when they are stimulated by an optimally oriented visual bar or grating (Eckhorn et al. 1988, 1993; Gray and Singer 1989, Gray et al. 1989; Engel et al. 1991; Reviews in Gray

1994, Singer and Gray 1995). These so-called oscillations can be consistently observed in local field potentials (LFP), multiple unit activity (MUA) and often also in single unit spike trains (SUA). Tuning curves derived from peak amplitudes of power spectra of LFP or MUA have similar properties as standard single unit receptive fields at the same cortical site and, most interestingly, ‘oscillations’ at different sites reveal a considerable amount of synchronization strongly depending on certain non-local stimulus properties: for example, crosscorrelations between two recording sites are highest when these are coactivated by a single bar. They are reduced when two independent bars, moving in the same direction, are presented together, and nearly absent when the two bars move in opposite directions, even though the neurons at both sites are stimulated almost equally strongly in all the above cases (Engel et al. 1991, Singer and Gray 1995). Response synchronization of this type has been demonstrated within different visual areas, where it can be detected over a range of several millimeters (Eckhorn et al. 1988). It also occurs between different areas and even between the two hemispheres (Eckhorn et al. 1988, Engel et al. 1991, König et al. 1995). However, a necessary condition for this seems to be that the receptive fields of the observed cells overlap, or have at least some other property in common (e.g. orientation). This has been demonstrated by Eckhorn et al. (1988) and is consistent with other experiments.

On the theoretical side, the above results have been taken as evidence for the so-called *temporal correlation hypothesis* of sensory integration in the mammalian cortex (v.d.Malsburg and Schneider 1986, Eckhorn et al. 1988, Singer and Gray 1995). This states, that neurons which fire in response to the same external object display correlated - in particular synchronized - firing. This could help to solve two problems of sensory processing. First, the *binding (grouping or linking) problem*, e.g. the question, how togetherness between parts of a single object can be signalled even if these are processed over large distributed regions of the brain and, furthermore, if the range of possible feature combinations is assumed to be of combinatorial complexity. (The latter involves that the solution can not be hardwired but must arise from cortex-intrinsic dynamical interactions). Second, the *superposition (or segregation) problem*, which asks, how different objects can be simultaneously represented in the same cortical area by neural activation without becoming mingled. Here it is believed, that neurons corresponding to different objects simply fire in an uncorrelated way. In this way each object can be uniquely labelled as a whole entity. Thus, the two seemingly different problems of grouping and segregation according to the temporal

correlation hypothesis appear as aspects of one and the same mechanism - the correlation and decorrelation of neural signals.

Rhythmic cortical activity has repeatedly been modelled, numerically and analytically. Besides differences in methods and details, most of these approaches explicitly or implicitly contain one or more (most often all) of the following assumptions:

- (i) **Basic oscillators:** Synchronized cell activity often, but not necessarily, co-occurs with gamma-range oscillations (30-90Hz). Based on this notion, object features are thought to be represented by means of some kind of biophysical oscillators. Thus, basic network units are modelled as abstract kinds of oscillators, so-called phasors or rotators (Neven and Aertsen 1992, Sompolinsky and Tsodyks 1994), pairs of graded excitatory and inhibitory cells (which intrinsically display periodic solutions) (Campbell and Wang 1994, Schillen et al. 1994, v.d.Malsburg and Buhmann 1992), or periodically spiking neurons (Eckhorn et al. 1990, Gerstner et al. 1993, Ritz et al. 1994). Binding then is interpreted as synchronization of appropriate subsets of these oscillators.
- (ii) **Phase coding:** Starting from feature-oscillators as building blocks, an obvious solution of the superposition problem is phase segregation: different objects are separated into different phases within a collective oscillation of a fixed period (Campbell and Wang 1994, Horn and Usher 1991, Neven and Aertsen 1992, Ritz et al. 1994). This principle relies further on the assumption of stationary or at least slowly varying input in comparison with the time-scale of synchronization. Sophisticated mechanisms for rapid synchronization have therefore been devised (e.g. Campbell and Wang 1994).
- (iii) **Globality:** Binding is often modelled as a global phenomenon within and between areas. Conditions are studied, such that every feature of an object is recruited into a single globally synchronized cell-assembly. This is most clearly seen in models which synchronize objects (bottles, cars, helicopters) distributed over extended twodimensional network structures (Campbell and Wang 1994, Sompolinsky and Tsodyks 1994, Neven and Aertsen 1992, Tononi et al. 1992, Ritz et al. 1994, Schillen et al. 1994). Fully connected network schemes may also be envisaged as an implicit assumption of globality (Horn and Usher 1991, Gerstner 1995 and references therein).

Obviously, these assumptions specialize the more general correlation hypothesis. In the next section we briefly investigate their empirical justification. Section 4 outlines an alternative interpretation, which is based on rapid associative processes taking place in roughly ten milliseconds. Those processes, if repeated continuously in response to a stimulus, *can* lead to oscillator-like behaviour. However, we propose, that single objects are essentially processed in single periods and conflicting stimuli in different periods. Furthermore, we argue, that cortical sites in and between areas only synchronize, if they physically interact, that is, if they are synaptically connected. In section 5 simulation results illustrating our hypotheses are presented.

3 Brief discussion of the standard hypotheses

We start with the aspect of global binding of whole objects. The search for conditions guaranteeing precise and/or global synchronization in mathematical models is surely of general biophysical interest and thus it has certainly motivated many of the studies cited above. But beside this mainly technical interest, the requirement of global binding seems to be also an outcome of more general arguments, which briefly proceed as follows: visual objects are typically represented and processed not in one area but distributed over several, if not many areas of the brain. Different objects are represented at once. Psychologically, we can easily distinguish those objects, even in ambiguous situations. How does the brain know which features belong to which objects? How can it code togetherness and distinction? The synchronization hypothesis is brought up as a tentative means to solve exactly this problem. Synchronization can provide an answer, but from this viewpoint globality seems to be necessary; if synchronization is only local, i.e. subsets of features are synchronized, but independent or out of phase with other subsets, then the problem is not solved.

We do not know of direct evidence for global binding. This would in principle include synchronization between arbitrary features of arbitrary objects. Of course, coherent oscillations have been observed between different areas and hemispheres of the cortex (Eckhorn et al. 1988, Engel et al. 1991, König et al. 1995), but as initially mentioned, a prerequisite for this is that the observed cells have at least one receptive field property in common; in most experiments receptive fields of recorded cells overlap. Now, suppose globality holds. Then, since synchrony is a transitive relation it should be observable between *any* two cells responding to a particular object, wha-

tever feature they code. But synchronized cells with completely unrelated receptive fields have not been reported so far, although it should be possible to observe this phenomenon (if it exists) in some of the experimental paradigms. This missing evidence may be partly due to restricted experimental paradigms often dealing with simple straight bars or gratings only. Thus cells excited by an object trivially have the property of ‘orientation’ in common. Nonetheless, stimulating with a moving bar of sufficient width, not even the two edges are synchronized in area 17 of the cat (Eckhorn, private communication). Thus binding is not global, not even within area 17.

Investigating experimental recordings of LFP or MUA, one immediately observes that ‘oscillation’ frequencies and amplitudes strongly fluctuate. From this it has early been noted that the oscillation frequency - which in general is different from single cell firing rates (Eckhorn and Obermüller 1993) - is not likely to code for anything whatsoever (Eckhorn et al. 1989, Gray 1994, Wilson and Bower 1991), although there are of course certain systematic dependencies of the frequency on a few stimulus parameters (Wennekers et al. 1994). Due to the large fluctuations, correlograms often do not have more than a few side peaks indicating a rather short decay time of temporal correlations. Hence, adopting the view of neuronal ‘oscillators’ as a physical substrate, these necessarily must be subject to considerable noise, with the main effect of perturbing the regularity of the oscillation. Thus, the genuine property of oscillators, namely periodicity, seems not to be functionally very important.

With respect to phase-coding similar problems arise. Since there is no fixed period length, ‘phases’ in the strict mathematical sense of parameterizations of periodic functions cannot be defined. Furthermore spikes in a single period scatter over several milliseconds. Comparing this with period-durations of 15-25ms it is obvious that phase-coding can only be very coarse and timing relations must decay over a few rhythms. Few distinguishable phases would probably be enough in primary areas, where local ambiguities are relatively simple, but would hardly suffice to separate entities on a higher level in a more complex situation. In any case, phases should be observable in experimental crosscorrelation studies. Indeed, they are not! As already mentioned, two bars moving in different directions give rise to nearly flat correlograms not to shifted peaks as they are expected in correlograms of phase-shifted signals (Engel et al. 1991, König et al. 1995, Singer and Gray 1995).

Summarizing these discussions we observe that all these hypotheses are not fully justified by the experimental evidence. If they are taken at face-

value, they are most probably wrong. However, there is also a certain degree of truth in each of these hypotheses so that it is not even unlikely to find statements in experimental papers that could provide a certain basis for assuming them. Thus a theoretical paper that fully relies on the truth of all three hypotheses may effectively claim in the theoreticians' community to be dealing with the biological reality, whereas it would be regarded as quite unrealistic by the experimentalists' community. In the next section we will show, how a relatively small twist in the three hypotheses can render a much more realistic picture.

4 Synchronized spikes and fast associative processes

In our opinion it is less the periodicity of rhythmic cortical activity that matters, but the temporal synchronicity within the single period. Recently we proposed a tentative framework based on processes of rhythmic associative spike synchronization, that seems to be consistent with most experimental findings (Wennekers et al. 1995). It was motivated by theoretical results concerning iterative retrieval in sparsely coded associative memories (Schwenker et al. 1996). Two outcomes of this work are particularly important.

1. Pattern association is extremely fast - at most 3 feedback steps for perfect retrieval - provided firing thresholds are adapted to the network activity in each step (cf. fig. 1).
2. Pattern completion is most efficient, when the number of ones in the address pattern is about half that of the stored patterns (fig.2).

The first point in the current context means the following: Taking a few ms for a single associative feedback step in the cortex, which is determined by synaptic and axonal delays, perfect pattern completion can be performed in less than roughly 10ms, which well fits with the observable γ -periods of 15 to 25 ms. The second point practically means that the active input synapses to any relevant neuron that is about to be 'addressed', should not be less than half the synapses that could be activated by the complete pattern. This implies further, that spikes *must* be synchronous if efficiency is required. Since real neurons have an integration time of (again) a few ms, all relevant information should be present during those short time intervals.

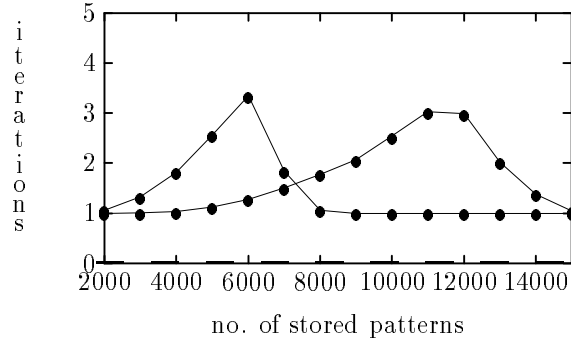


Figure 1: The mean iteration time of the incremental (left curve) and binary Hebbian learning rule with matrix size $n=1900$, $k=13$ ones per stored pattern and $l=6$ ones in address patterns. Note that usually less than only 3 iterations are needed for pattern completion.

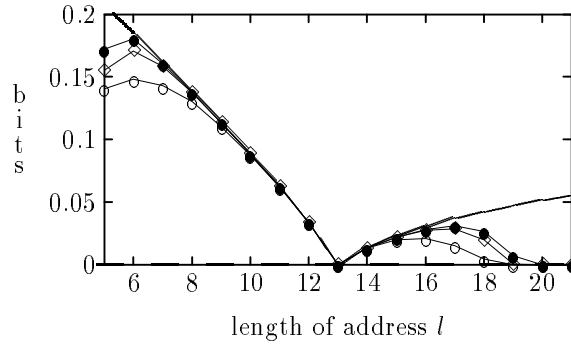


Figure 2: The capacity in bits/synapse of one-step (○), two-step (◇), and iterative retrieval (●) for address patterns of different length; the solid line shows the completion capacity for error free retrieval. The capacity is highest for $l \approx k/2 \approx 6$, below which retrieval becomes progressively faulty.

In experiments one often observes a concentration of spikes in 4 or 5ms. Scattering the same spikes randomly over 15-25ms (the period length) wastes 70% or more of the available information, since less than $5ms/20ms = 1/4$ of all spikes can then be envisaged as addressing the cortical storage matrix at one time (compare also Abeles 1991). Therefore input patterns should be synchronous in the millisecond range.

From these observations and common ideas about anatomy and function of the cortex (e.g. Braitenberg and Schüz 1991), we derive the following mechanisms for rhythmic cortical processing: 1) *Fast associative processes* are carried by sparsely firing pyramidal cells and their excitatory connections. 2) The *threshold control* necessary for efficiency is supplied by interneurons which mainly estimate the average local spike-activity and inhibit the excitatory cells accordingly. (See also Palm 1982 for related ideas on ‘threshold control’).

This basic circuitry naturally leads to a rhythmic network activation in the gamma frequency range. Every period consists of a short population burst (association phase) followed by a period of inhibition (relaxation phase). The occurrence of oscillations in comparable network structures has repeatedly been demonstrated in other modelling approaches (e.g. Eckhorn et al. 1990, Wilson and Bower 1991, Gerstner et al. 1993). Important here is that in our interpretation subsequent periods can already process different patters. In principle a new input can be sampled from the sensory stream every 20ms.¹ This idea bears similarity to that of ‘correlation-amplification’ proposed by Koch and Schuster (1992), which was intended to give a simple explanation of the biologically detailed simulations of Bush and Douglas (1991). However, both papers did not directly address computational aspects.

Associative memories in the strict and simplified theoretical sense are certainly not a good model for primary visual areas, since those are not likely to store many thousands of distinct patters, but perhaps act more as filter or feature-detector hierarchies with the well-known broad tuning of single cells to stimulus properties. Nonetheless, the main idea of reverberant excitatory processes succeeded by inhibited phases does also apply in these areas. Short and long-range excitatory interactions here could recruit best matching cells into a short population burst, perhaps favourizing cell-constellations that

¹Recently Dong and Atick (1995) measured temporal power-spectra of natural time-varying images; these indicate that rates of some 10Hz would indeed be best for sampling natural scenes.

resemble ‘Gestalt’-principles. The onset of inhibition evoked by excitatory activity in turn could suppress cells less well matching or partially conflicting with the stimulus features. Synchronization, of course, would also increase signal to noise ratios, thereby improving feature-detection.

We propose that the elementary computational processes at the base of gamma-range oscillations are temporally confined to single short association cycles corresponding with distinct gamma-periods. The rhythmicity is then understood as an expression of repeated elementary computations. This should be contrasted with the idea of ‘phase-coding’, which assumes that spike-trains of feature coding cells (or other local signals like MUA or LFP) are intrinsically periodic and become mutually aligned at well defined time-differences, which depend on whether cells represent the same or different objects. This notion of ‘phases’ is meaningless in our interpretation. Nonetheless *timing relations* between spikes on the scale of a few milliseconds in individual population bursts can well carry information and contribute to cortical computations. Such spatio-temporal patterns need not be repeated periodically. Systematic timing relations can arise, if cells are differently strongly excited, which should be the case under almost all stimulus conditions. We would expect that spikes of suboptimally driven cells lag behind optimally excited ones, because suboptimal cells on average spend a longer time below but near firing threshold. Hence, optimally stimulated cells, which have a larger probability to fire, can easily and quickly trigger the firing of those suboptimal cells. This phenomenon may provide a natural explanation for firing delays of suboptimal cells as observed in the visual cortex by König et al. (1995a). An interesting point in the context of dynamically controlled associative memories is that the interpretation of both cell-classes, early and late firing cells, is conceptually different: the early cells can be envisaged as constituting the address-pattern, whereas the later firing cells represent the retrieved additional information. (Of course a strict distinction of two cell classes is artificial and the assumption of a smooth transition between both extremes more plausible, but the argument essentially remains the same in the second case). The early cells are mainly input driven and because they receive the strongest external input, they signal ‘safe’ information about the stimulus. Contrary, the later cells may not even be able to fire in a significant manner supplied with relatively weak external input alone. Their firing times strongly depend on further supporting input from the early cells via recurrent collateral connections. Thus, they provide information about the connectivity structure of the network, or, in a more technical sense, about the associative coupling matrix

and the stored memory patterns. These ideas argue for possible mechanisms of the generation of temporal short-time relations between spikes and their tentative functional use. With respect to the detection of spatiotemporal patterns, we should finally mention, that spike-patterns in the millisecond range can be easily detected by single cortical cells utilizing propagation delays and spatiotemporal integration properties of dendritic structures. It is much harder to imagine mechanisms that detect and take advantage of repeated patterns, which last for more than several ten milliseconds. Thus, even if phase-coding should occur in the cortex, it is likely that all temporal structure a single neuron can effectively make use of, is still concentrated in a single gamma-period. But since the work of Schwenker et al. (1996) and the simulations of the next section show that synchronisation is extremely fast, the function of ‘oscillations’ then is again unclear.

Now, we consider the question how far ‘binding’ might reach. Taking experimental evidence together, the most natural assumption seems to be, that it is restricted to cortical sites within or between areas, which are directly and sufficiently strongly connected via collaterals or interareal fibres. Those sites often do have similar receptive field properties, which appear to be necessary for synchronization (Bullier et al. 1993, Salin and Bullier 1995 and references therein). Furthermore physiologically measured correlation lengths fit well with anatomically observed connectivity ranges (Eckhorn et al. 1988, Bullier et al. 1993, Salin and Bullier 1995).

Between different cortical areas or hemispheres transmission delays can become significantly large. We believe that associative processes in this case can switch to a more iterative nature with volleys of spikes propagating back and forth between areas, and influencing their respective targets with a lag of one period. Comparable interarea-interactions have, for example, been investigated in model studies by Bibbig et al. (1995). They also seem to be supported by experiments of König et al. (1995 b).

Network simulations (cf. next section and Bibbig et al. 1995, Wennekers et al. 1995) indicate that the duration of excitation-inhibition cycles can be strongly fluctuating. Caused by many different sources of randomness, this has to be expected also in the cortex. Therefore, correlations between distant cells with non-related receptive fields and probably few monosynaptic connections can be easily destroyed. Thus the strict transitivity of ‘synchronization’ can be broken such that it may not be observable in experiments.

We conclude that binding is most likely a local phenomenon, if we define locality by physical connectivity and not simply by proximity. Since higher sensory and association areas of the cortex are usually smaller than primary

sensory areas, this does not necessarily exclude that whole objects can be bound together in those areas; what we reject is the hypothesis that they are coherent over the whole primary visual cortex or even the brain.

5 Network simulations

In this section we briefly describe simulations illustrating our main proposals. Figure 3 shows the implemented model network consisting of 3 layers R, P and C which may be loosely identified as retina, peripheral visual area and more central visual association area.

The input area R is nothing but a two-dimensional binary pixel array of size 64×64 at which static patterns can be applied (cf. figures 4a and 5a-d).

P contains 64×64 topographically ordered excitatory spiking neurons which receive input from R via Gaussian coupling kernels of size 5×5 (depicted as I in figure 3). The excitatory units are also laterally coupled extending over 11×11 neighbourhoods (II). These connections have asymmetric Gaussian profiles introducing orientation selectivity. Orientations are random among cells, neglecting the complex structure of real visual cortexes. Each excitatory cell in P is inhibited by an individual interneuron (grey circle, V) that computes the local average of the excitatory spike-activity over a 7×7 neighbourhood of the particular cell.

The visual association area C consists of 8×8 excitatory neurons which are modelled as a fully connected associative memory. Input to the excitatory cells in C is an average (7×7) over spike activity in P (III). So the input pattern to the association area is a coarse-grained version of the original input at R. C contains only a single inhibitory cell that samples the overall spike-activity and inhibits all excitatory cells with equal strength (VI). The globality of the inhibition reflects the smaller size of ‘higher’ areas.

All cells in P and C implement membrane leakage in form of first order low-pass filters with time-constants of 3ms for excitatory and 5ms for inhibitory cells. Since the latter are intended to represent cell-pools their output is given by a sigmoid function of the membrane potential. Spiking neurons are chosen to model the excitatory cells: when their potential reaches a firing threshold, a spike of 1ms duration is emitted. A refractory mechanism (dynamical threshold) prevents cells from firing again immediately. (Details of the neuron model can be found in Bibbig et al. 1995, Wennekers et al. 1995).

Figure 4 displays the activation dynamics in the topographic area P




Figure 3: Scheme of the network model. R: binary input layer of size 64×64 (retina); P: topographically ordered peripheral area of excitatory and inhibitory cells (64×64 each); C: fully connected associative area (8×8 excitatory cells, one global inhibitor). I&III depict feedforward connections (receptive fields), II&IV recurrent connections between excitatory cells (context fields; local in P, global in C). V&VI indicate inhibitory interneurons (grey circles) and their connections.

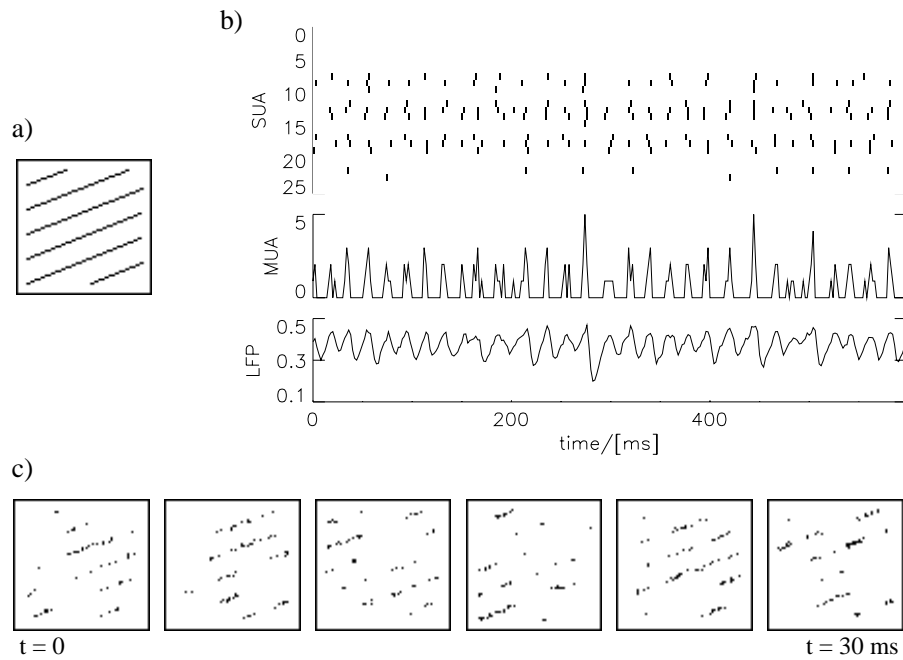


Figure 4: Activation in the peripheral area (P) when the grating stimulus (a) is applied. In (b) 25 single unit spike trains (SUA), their sum (MUA) and the averaged membrane potential (LFP) of a 5×5 square around unit (30,37) are shown. Some units fire regularly, others lock randomly into the population spikes. Thus amplitudes and frequencies are highly irregular. (c) Overall spikepatterns of all excitatory cells at different times (coarsly 1.5 periods). Synchronization is not global, but restricted to connectivity regions. Each plot rebins spikes from 6ms.

when a grating stimulus (a) is applied to R. First observe that the activity is rhythmic (fig.4b), most pronounced in local averages of cell potentials (here called LFP) or spikes (MUA). Single units (SUA) discharge regularly when they are excited sufficiently strongly. If their external drive is below threshold they can only fire, if they receive further input via lateral connections. Thus they lock into the collective bursts irregularly and their spikes lag behind those of optimally driven cells, both in agreement with experimental observations (Eckhorn and Obermüller 1993, König et al. 1995a).

The random involvement of different sets of cells in different single periods leads to fluctuating burst amplitudes, which in turn give rise to varying period durations. The latter reflect the action of inhibition: the more cells fire in a burst, the more pronounced and therefore long lasting is the inhibitory response. Thus low amplitudes correspond with faster rhythms, again consistent with experiments (Wennekers et al. 1994). We take this as evidence that the observed oscillations and frequencies are to a large extent a network effect and less due to single unit properties.

Figure 4c shows spikepatterns of all 64×64 excitatory cells in P at consecutive times. As can be immediately seen, synchrony is not global. Measuring crosscorrelations of LFP we observe a nearly exponential decay of spatial correlations. The estimated decay constant of 5.46 ($\approx 11/2$ lattice units) implies that correlations do not reach much farther than the hardwired connections. Generalizing this result to cortical situations, correlations should only be observable in the range of long-range collaterals, which indeed is the case (Bullier et al. 1993, König et al. 1995b, Eckhorn et al. 1988).

Wilson and Bower (1991) compared correlations evoked in distant cells in a network similar to our peripheral area P, when either a long bar or two separated but colinear short bars were used as stimuli. Repeating these simulations we found essentially the same results: Correlations are larger in case of the continuous stimulus (long bar) because cells excited by the long bar which were situated in the activationless region between the two short bars can relay information between distant sites. Therefore synchronizing feedback is stronger and correlations are more pronounced. Moreover, when the distance between the two short bars is larger than the coupling width, correlations are virtually zero. In contrast to the simulations by Wilson and Bower, this also happens at distant sites stimulated by a continuous long bar. The same effect is already revealed by Figure 4c for the grating stimulus: here, synchronized firing of cell groups appears to be localized in patches of roughly the size of synaptically connected regions and does

not extend much along the individual bars. This difference to Wilson and Bower's work is probably due to the inhomogenous lateral couplings in our model area P, which give rise to spatially varying average oscillation periods. Under homogeneous conditions local circuits are everywhere the same, have the same intrinsic period duration and are therefore easier to synchronize. The experimental results on the spatial decay of correlations mentioned above indicate that less homogenous conditions are probably more realistic for the cortex.

Figure 5 displays the spike activity in the central area C when the test-pattern (d) is applied to area R. Patterns P1-P3 (a-c) have been stored in a previous learning phase by means of an incremental Hebbian learning rule.² Observe that in fig. 5e each pattern contained in the input is recognized and completed in a very short time. Usually only one pattern is retrieved per elementary associative process. This is not trivial, because the patterns overlap. In fact it is the action of the inhibition that effectively raises the firing thresholds (by lowering membrane potentials) as soon as the amplification of one pattern happens. In this way co-activation of other patterns is prevented just as in the more abstract model of Schwenker et al. (1996). Furthermore the dynamic threshold also avoids latching of the associative area into an attractor state (cf. Koch and Schuster 1992, Horn and Usher 1991, Wennekers et al. 1995). Often attractor states in associative neural network models are quite stable in the sense that strong input is needed to destroy such an attractor and enable a new network response towards another stored pattern. Contrary, a dynamic threshold that rhythmically resets the network activation, keeps the network sensitive to new - and especially weak - input at any moment. Which pattern becomes amplified in an individual retrieval period, mainly depends on fluctuations in the incoming activity (here from area P) and thus is essentially random. Pattern sizes can bias the retrieval; a large pattern is detected more often. Furthermore, by introducing adaptation into the excitatory cells, retrieval can be forced to occur in cycles (cf. Fahle and Palm 1991, Ritz et al. 1994), but the elementary associative processes then happen nonetheless on a scale of roughly 25ms. The pattern-cycle and 'phases' defined by the ordered retrieval take place on a slower scale. This is *not* the usual notion of phase-coding by 40Hz-oscillations.

Finally, we should mention that asynchronous input from layer P to C

²Speedup of Hebbian learning is another yet unmentioned advantage of synchronized activity. This issue has been discussed in detail in Bibbig et al. (1995).

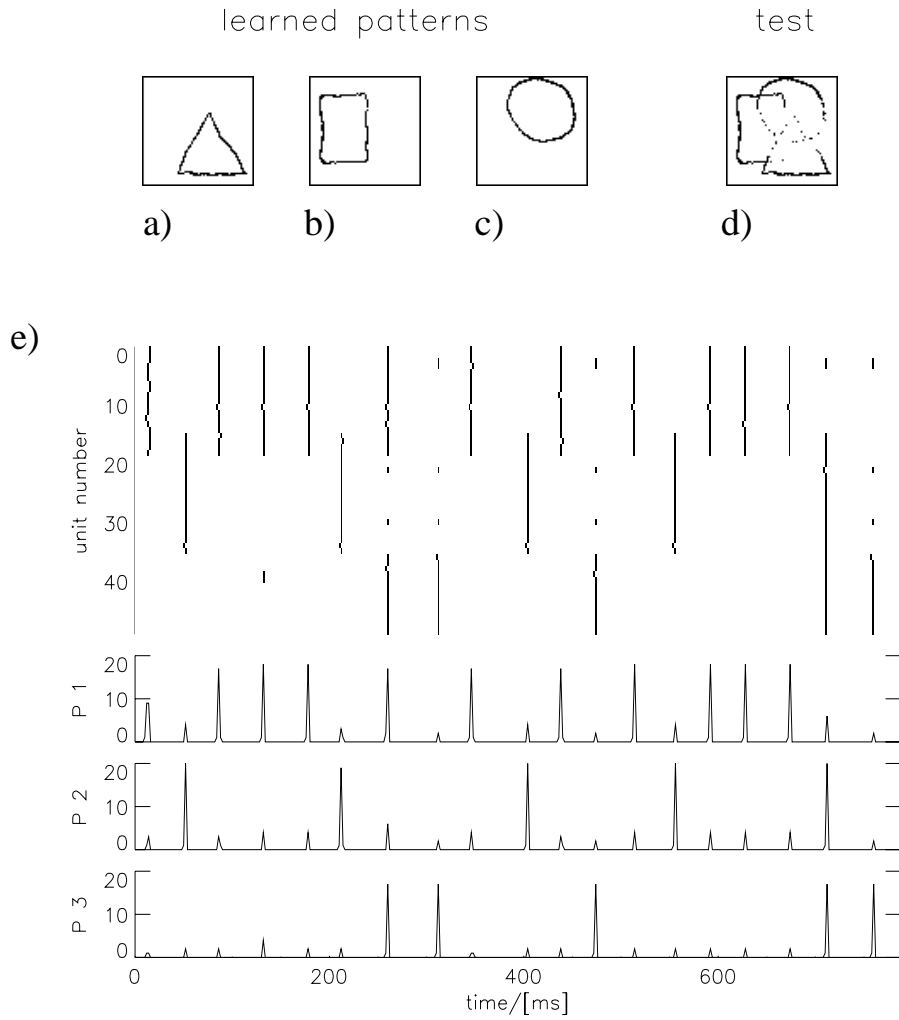


Figure 5: Rhythmic associative retrieval in the central area. Patterns P1-P3 (a-c) were presented to the input area during a preceding learning phase utilizing an incremental Hebbian coincidence rule in the central store C. Testpattern (d) is applied during the displayed simulation run. In (e) spike patterns of all cells in C belonging to central representations of P1-P3 are shown (above) as well as the overlaps of these representations with the current network spike pattern (below). Note that perfect pattern retrieval and random segregation takes place, but no phase-coding.

with the same average firing rates as above hardly evokes any activity in the associative area. When lateral connections and inhibitory interneurons in the peripheral area are switched off, then the spikes in P are asynchronous and not locally clustered in time. Therefore cells in the central area are driven less effectively and respond with very sparse firing rates (data not shown). Under these conditions retrieval of patterns occurs seldom and in a non-rhythmical manner (similar arrhythmical behaviour occurs in the model of Koch and Schuster 1992). Lowering thresholds does not solve this problem, since then even single perfect input patterns address different memory patterns due to overlapping neurons. This problem becomes the more serious the more patterns are stored in area C (cf. Schwenker et al. 1996). Contrary, increasing firing rates in P *can* solve the problem, but the necessary rates correspond to roughly $1/5ms \approx 200Hz$, if we assume a membrane time constant of 5ms as the necessary coincidence interval. This is way above the firing rates of typical cortical neurons and again suggests that at least one important role of synchronized cortical codes is the larger efficiency in exciting target neurons and the increased signal-to-noise ratio (cf. Abeles, 1991) that can be achieved even at (arbitrary) low rates.

6 Conclusions

In this paper we have discussed the relation between experimental evidence and theoretical modelling in neuroscience for the concrete example of high-frequency oscillations in cortical neural activity and their purpose for neural computation and representation. We have criticized the three assumptions of basic oscillators, phase coding and globality of synchronization that are often made in theoretical papers. Finally we have illustrated a very similar but to our knowledge much more realistic picture based on spiking neurons with a refractory mechanism, temporal coincidence and laterally decreasing coherence by some computer simulations.

Acknowledgement

We thank Reinhard Eckhorn, Axel Fries and Michael Erb for many discussions concerning synchronized activity in the cortex and W.A. Phillips for useful comments on an earlier draft of this paper. This work has been supported by the Deutsche Forschungsgemeinschaft (Pa 268/8-1).

References

- Abeles, M. (1991) *Corticonics*. Cambridge University Press. Cambridge.
- Anderson, J.A.; Rosenfeld, E. (1988) *Neurocomputing*. MIT Press. Cambridge.
- Bibbig, A.; Wennekers, T.; Palm, G. (1995) A Neural Network Model of the Cortico-Hippocampal Interplay and the Representation of Contexts. *Behavioural Brain Research* 66: 169-175.
- Braitenberg, V.; Schüz, A. (1991) *Anatomy of the cortex*. Springer. Berlin.
- Bullier, J.; Munk, H.J.; Nowak, L.G. (1993) Corticocortical connections sustain interarea synchronization. *Concepts Neurosci.* 4: 159-174.
- Bush, P.C.; Douglas, R.J. (1991) Synchronization of Bursting Action Potential Discharge in a Model Network of Neocortical Neurons. *Neural Comp.* 3: 19-30.
- Campbell, S.; Wang, D.L. (1994) Synchronization and desynchronization in a network of coupled Wilson-Cowan oscillators. Technical Report OSU-CISRC-8/94-TR43. Dept.of Physics, Ohio State University.
- Dong, D.W.; Atick, J.J. (1995) Statistics of natural time-varying images. *Network* 6: 345-358.
- Eckhorn, R.; Bauer, R.; Jordan, W.; Brosch, M.; Kruse, W.; Munk, M.; Reitboeck, H.J. (1988) Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biol.Cyb.* 60: 121-130.
- Eckhorn, R.; Frien, A.; Bauer, R.; Wölbern, T.; Kehr, H. (1993) High frequency (60-90Hz) oscillations in primary visual cortex of awake monkey. *Neuroreport* 4: 243-246.
- Eckhorn, R.; Obermüller, A. (1993) Single Neurons are Differently Involved in Stimulus-Specific Oscillations in Cat Visual Cortex. *Exp. Brain Res.* 95: 177-182.
- Eckhorn, R.; Reitboeck, H.J.; Arndt, M.; Dicke, P. (1990) Feature Linking via Synchronization among Distributed Assemblies: Simulations of Results from Cat Visual Cortex. *Neural Comp.* 2: 293-307.
- Engel, A.K.; König, P.; Kreiter, A.K.; Singer, W. (1991) Interhemispheric Synchronization of Oscillatory Neuronal Responses in Cat Visual Cortex. *Science* 252: 1177-1178.
- Fahle, M.; Palm, G. (1991) Perceptual rivalry between illusory and real contours. *Biol.Cybern.* 66: 1-8.
- Gerstner, W. (1995) Time Structure of the Activity in Neural Network Models. *Phys.Rev.E* 51: 738-758.
- Gerstner, W., Ritz, R., van Hemmen, J.L. (1993) A biologically motivated and analytically soluble model of collective oscillations in the cortex, I. Theory of weak locking. *Biol.Cybern.* 68: 363-374.
- Gray, C.M. (1994) Synchronous Oscillations in Neuronal Systems: Mechanisms and Functions. *J.Comp. Neurosci.* 1: 11-38.

- Gray, M.; Singer, W. (1989) Stimulus specific neural oscillations in orientation columns of cat visual cortex. *PNAS* 86: 1698-1702.
- Gray, C.M.; König, P.; Engel, A.K.; Singer, W. (1989) Oscillatory responses in cat visual cortex exhibit intercolumnar synchronisation which reflects global stimulus properties. *Nature* 338: 334-337.
- Horn, D.; Usher, M. (1991) Parallel Activation of Memories in an Oscillatory Neural Network. *Neural Comp.* 3: 31-43.
- Koch, C.; Schuster, H. (1992) A Simple Network Showing Burst Synchronization without Frequency Locking. *Neural Comp.* 4:211-223.
- König, P.; Engel, A.K.; Roelfsema, P.R.; Singer, W. (1995a) How Precise is Neuronal Synchronization. *Neural Comp.* 7: 469-485.
- König, P.; Engel, A.K.; Singer, W. (1995b) Relation between oscillatory activity and long-range synchronization in cat visual cortex. *PNAS* 92: 290-294.
- von der Malsburg, C.; Buhmann, J. (1992) Sensory segmentation with coupled neuronal oscillators. *Biol. Cybern.* 67: 233-42.
- von der Malsburg, C.; Schneider, W. (1986) A neural cocktail party processor. *Biol.Cybern.* 54: 29-40.
- MacGregor, R.J. (1987) *Neural and Brain Modeling*. Academic Press. New York.
- Neven, H.; Aertsen, A. (1992) Rate coherence and event coherence in the visual cortex: a neuronal model of object recognition. *Biol.Cybern.* 67: 309-322.
- Palm, G. (1982) *Neural Assemblies. An Alternative Approach to Artificial Intelligence*. Springer, Berlin.
- Palm, G. (1990) Cell Assemblies as a Guideline for Brain Research. *Concepts in Neuro.Sci.* 1:133-147.
- Ritz, R.; Gerstner, W.; Fuentes, U.; van Hemmen, J.L. (1994) A biologically motivated and analytically soluble model of collective oscillations in the cortex. II. Application to binding and pattern segmentation. *Biol.Cybern.* 71: 349-358.
- Salin, P.A.; Bullier, J. (1995) Corticocortical Connections and the Visual System: Structure and Function. *Physiol. Rev.* 75: 107-154.
- Schillen, T.B.; König, P. (1994) Binding by temporal structure in multiple feature domains of an oscillatory neural network. *Biol.Cybern.* 70: 397-405.
- Schwenker, F.; Sommer, F.T.; Palm, G. (1996) Iterative Retrieval of Sparsely Coded Associative Memory Patterns. *Neural Networks* 9: 445-55.
- Shaw, G.L.; Palm, G. (1988) *Brain Theory Reprint Volume*. World Scientific, Singapore.
- Singer, W.; Gray, C.M. (1995) Visual feature integration and the temporal correlation hypotheses. *Ann.Rev.Neuroscience* 18: 555-586.
- Sompolinsky, H.; Tsodyks, M. (1994) Segmentation by a Network of Oscillators with Stored Memories. *Neural Comp.* 6: 642-57.
- Tononi, G.; Sporns, O.; Edelman, G.M. (1992) Reentry and the problem of integrating multiple cortical areas: simulation of dynamic integration in the visual system. *Cerebral Cortex* 2: 310-335.

- Wennekers, T.; Erb, M.; Eckhorn, R.; Palm, G. (1994) Dependence of oscillation frequency on stimulus parameters in primary visual cortex of monkeys: simulations and experimental results. *Europ.J.Neurosci.Suppl.* 7, 66.03.
- Wennekers, T.; Sommer, F.T.; Palm, G. (1995) Iterative Retrieval in Associative Memories by Threshold Control of Different Neural Models. In: Herrman HJ, Wolf DE and Pöppel E (eds) *Supercomputing in Brain Research*. World Scientific. Singapore, pp 301-319.
- Wilson, M.A.; Bower, J.M. (1991) A Computer Simulation of Oscillatory Behavior in Primary Visual Cortex. *Neural Comp.* 3:498-509.