

Cell Assemblies, Associative Memory and Temporal Structure in Brain Signals*

Thomas Wennekers and Günther Palm
Department of Neural Information Processing
University of Ulm, D-89069 Ulm, Germany

Abstract: In this work we discuss Hebb's old ideas about cell assemblies in the light of recent results concerning temporal structure and correlations in neural signals. We want to give a conceptual, necessarily only rough picture, how ideas about 'binding by synchronisation', 'synfire chains', 'local and global assemblies', 'short and long term memory' and 'behaviour' might be integrated into a coherent model of brain functioning based on neuronal assemblies.

Keywords: cell assemblies, synchronization, gamma-oscillations, synfire chains, memory, behaviour

1 ASSEMBLIES AND ASSOCIATIVE MEMORIES

1.1 Cell Assemblies

Cell assemblies have been introduced by Donald Hebb with the intention of providing a functional and at the same time structural model for cortical processes and neuronal representations of external events (Hebb, 1949). According to Hebb's ideas, stimuli, objects, things, but also more abstract entities like concepts, contextual relations, ideas, and so on are thought of being represented in the brain by simultaneous activation of large groups of neurons, which are connected by relatively numerous and/or strong mutual excitatory synapses. Single neurons may belong to many different cell assemblies; the determinant of an assembly is the connectivity structure between cells that defines, which cells lend support to each others firing and hence have a higher probability to become coactivated in a reliable manner in response to different versions of the same stimulus: if an external stimulus excites a sufficiently large subset of cells of an assembly, then the whole assembly can 'ignite' or 'fire', because recurrent activity, distributed via the specific mutual connections,

*To appear in: Miller, R. (editor) Time and the Brain. Conceptual Advances in Brain Research, vol. 2, Harwood Academic Publishers.

raise also those cells above threshold, which are not (or only weakly) stimulated externally. This can be viewed as an elementary associative process, where the firing of externally driven cells represents the key information and triggers the firing of cells representing information addressed by, but not yet contained in the key. In a similar way, also some kind of short term memory is supported: if continuously changing subgroups of cells permanently excite other groups, the activation within the assembly may survive for some time after its ignition, even if the external stimulation has already vanished.

The assembly concept furthermore proposes a mechanism for long term memory, that is, the formation of new assemblies in cortical tissue under the influence of electrical activation. Learning is believed to be expressed in activity-dependent changes of synaptic efficacies, and the nowadays widely known ‘Hebbian learning rule’ states that synapses get strengthened, when both connected cells are activated simultaneously within a certain time window. This is motivated by the observation, that events that repeatedly occur together should somehow belong together. Every time they appear in conjunction, they drive certain subgroups of cortical cells; the correlated firing of these sets of neurons (which are not yet assemblies!) should be learned and by that, the respective groups should become associatively connected.

1.2 Associative Memories

The classical associative memory is an abstraction of the assembly concept (Willshaw et al. 1969; Palm, 1982) which models single neurons as simple threshold elements. Figure 1 shows an example of an associative memory comprising 6 neurons with 6 inputs. Circles at the bottom and vertical lines represent threshold-neurons and their dendrites, horizontal lines are axonal input fibres, and small filled circles synapses, which are assumed to be zero initially. Neuronal excitation patterns are represented by binary $\{0, 1\}$ sequences reflecting firing of the cells or not in a given time window. A set of associations, that is, pairs of patterns (x^i, y^i) , $i = 1, 2, \dots$, is stored in the synaptic connectivity matrix. Special cases are the so-called auto-association, $x^i = y^i$ for all pairs i , or hetero-association, $x^i \neq y^i$, $i = 1, 2, \dots$ (Palm, 1991). In figure 1a two associations are stored: first the vector $x^1 = (100101)$ activates the input lines, and $y^1 = (110100)$ the neurons. All synapses, which simultaneously receive pre- and postsynaptic activity are set to 1, indicated by the black filled synapses. Afterwards the second association, $x^2 \rightarrow y^2$ is stored in the same way (grey synapses).

Figure 1b displays an elementary retrieval process. The address pattern $x = (111000)$ activates via the first three input lines some of the previously strengthened synapses (squares). Each threshold-neuron sums over its activated synapses; this determines the ‘potentials’ of the cells, shown as small numbers above the neurons.

A neuron generates an output signal of 1, if its potential is larger than a certain threshold θ . Here, $\theta = 2$ and thus, the output vector is $y = (010011)$, which is obviously exactly the stored vector y^2 . Since the input pattern x is more similar to x^2 than to x^1 , because the number of common ones is 2 in the first, but only 1 in the second case, this is a correct association. In auto-associative networks the output vector can be fed back to the input for further recurrent iteration steps.

1.3 Spiking Neurons and Threshold Control

The classical associative memory presumably covers basic aspects of pattern storage in real brains. Nonetheless, compared with biological neurons and synapses, the inclusion of further physiological details seems desirable, particularly, if we ask for temporal properties of neural systems. Conventional associative memory models employ simple threshold neurons and an iterative, time quantized update scheme. Therefore they can give only a very rough insight into dynamical processes in cell assemblies. For this reason we add some features to the classical model, which mimic the spiking behavior of real neurons as well as properties of spatio-temporal integration on dendrites.

First, we choose a continuous time model in all simulations shown later.[†] Single cells are modeled as ‘spiking neurons’: dendrites and soma of a cell are lumped into a single potential value. Every time this potential reaches a certain threshold θ a pulse-like ‘actionpotential’ is generated. Afterwards a suitable refractory mechanism prevents the cell from firing immediately again. Action potentials, when they arrive at a target cell, evoke ‘postsynaptic potentials’ of realistic form and time-constants, and the responses of different synapses are supposed to add up linearly.

Investigations by Schwenker et al. (1996) of recurrent auto-associative memories reveal the need for a proper continuous and global threshold control in order to avoid that the network is insensitive to any input (θ in figure 1b too large) or that its activity explodes (θ too small). In accordance with common ideas about cortical functioning (e.g. Braitenberg and Schüz, 1991), we assume that pyramidal cells and their excitatory connections essentially carry information processing tasks, like feature extraction, pattern storage, pattern retrieval, etc., and that inhibitory interneurons mainly have regulative functions, one of which probably some kind of activity control similar to the threshold regulation required from theoretical means. Therefore, we start from excitatorily connected associative networks of spiking neurons and embed interneurons into these models, which measure local firing rates and

[†]We only give informal model descriptions in this paper, because the arguments and results in the sequel are to a large degree independent of implementational details. More precise descriptions and equations can be found in Wennekers et al. (1995) and Bibbig et al. (1995). For comparable spiking neuron models see Stein (1967); Gerstner and v.Hemmen (1993).

inhibit the excitatory cells accordingly.

Networks of this architecture typically reveal collective oscillations above some critical level of external excitation supposed excitatory and inhibitory synapses are strong and/or numerous enough. We have shown (Wennekers et al. 1995) that with reasonable assumptions concerning network parameters, firing can be sparse even in states of collective oscillations, that is, single neurons need not fire in every period (see also next section). This work also shows that the more realistic model neurons with a threshold and refractoriness support synchronisation more efficiently than models with purely probabilistic spike encoder (e.g. rate modulated Poisson processes). This is because the threshold crossing process determines precise firing times, whereas those are necessarily imprecise in modulated Poisson models.

2 LOCAL ASSEMBLIES

Experiments on anesthetized cats and alert behaving monkeys have shown that local populations of cells in primary visual areas often respond rhythmically with frequencies in the gamma-range (30-90Hz) (cf. Eckhorn, this volume; experimental review in Singer and Gray, 1995). Most interestingly, those ‘oscillations’, when observed at two distant cortical sites, can reveal a considerable amount of synchronization, which strongly depends on certain non-local stimulus properties and seems to follow simple ‘Gestalt’-rules like proximity, colinearity, etc., even though the recording sites may be located in different cortical areas or hemispheres. These findings have been taken as evidence for the so-called *temporal correlation hypothesis* of sensory integration in the mammalian cortex (von der Malsburg, 1981; Eckhorn et al, 1998; Singer and Gray, 1995), which states, that neurons that fire in response to the stimulation by the same external entity should display correlated - in particular synchronized - firing. This way togetherness between parts of a single object can be signaled even if those are processed over distributed regions of the brain (‘binding by synchronization’). Many attempts have been made to model this ‘binding’ process (c.f. Eckhorn et al. and Borisyuk, this volume; see also Wennekers and Palm (1997) for an overview and a discussion of some principal assumptions of similar models). In the sequel of this section we present our own simulations of oscillatory assemblies in localized patches of cortex and discuss the tentative role of fast gamma rhythms for information processing.

2.1 Local Dynamics in a Primary Area

We consider a local patch of cortical tissue in a primary cortical area of roughly the size of a cortical column. This patch will contain cells tuned for different stimulus

orientations, directions, velocities, etc. For simplicity we assume that cells can be ordered by their orientation preference and that similarly oriented cells have a higher probability to be connected. We neglect other tuning properties of cells, which might be reflected by the local cortical circuitry and also do not model the laminar structure of real cortices. This means, at first approximation we consider a one-dimensional topographically ordered model network of $N = 128$ excitatory spiking neurons with connectivities restricted to some neighborhood of a given cell (roughly $N/4$, with Gaussian spatial decay of probabilities for synapses). Here, topography is meant to represent orientation tuning, but note that a spatial interpretation is also possible. Embedded into the network of excitatory cells are inhibitory interneurons, which receive input from excitatory cells in a neighborhood (of roughly $N/8$) and inhibit the excitatory cells accordingly. Inhibitory cells have graded responses and represent local pools of interneurons. All synaptic response functions have transmission delays of 1ms, a rise time of t_r and fall time of t_f , where $t_r = 1ms$, $t_f = 3ms$ for excitatory and $t_r = 2ms$, $t_f = 5ms$ for inhibitory synapses. Axonal conduction delays are not included, since we only consider a localized patch of cortex.

An external ‘bar’ or ‘grating’-stimulus will excite cells in a local cortical network differently strongly depending on its orientation (and maybe other properties); this is taken into account as an external input current into the excitatory cells, which is centered at some ‘orientation’ (say neuron $N/2$) and falls off in a Gaussian way. Furthermore each neuron also receives a certain amount of white noise background activity modeling spontaneous spiking of background cells.

Figure 2a displays the activation dynamics of our network for a ‘bar’ stimulus as explained above. Neuron $N/2=64$ receives the strongest input (e.g. is well tuned to the bar), and the input strength decays to smaller and larger neuron numbers (e.g. to less well tuned cells). Shown as ‘LFP’ in the figure is the ensemble average over excitatory postsynaptic potentials averaged over all cells, which roughly corresponds to physiologically derived local field potentials (LFP), and the spike trains of all excitatory cells as a raster plot over time (SUA, single unit activity). The LFP signal clearly shows oscillatory activity, although in a waxing and waning manner similar to physiological recordings. Diamonds above the LFP indicate spiketimes of an arbitrary neuron (here neuron number 49, also indicated by the dashed horizontal line in the raster plot). Apparently the neuron is not rhythmic, but emits spikes only in some periods of the collective rhythm. A closer investigation of all single unit spike trains reveals, that this is the generic case in our model. No neuron fires periodically. In particular we find the same categories of cells as reported by Eckhorn and Obermüller (1993) for experimental data: ‘locked’ cells, which spikes display auto-correlation-histograms (ACHs) with oscillatory side peaks and which are also correlated with the collective rhythm, ‘lock-in’ cells, which are coupled to the LFP but are not rhythmic by themselves (as revealed by the flatness of their ACH),

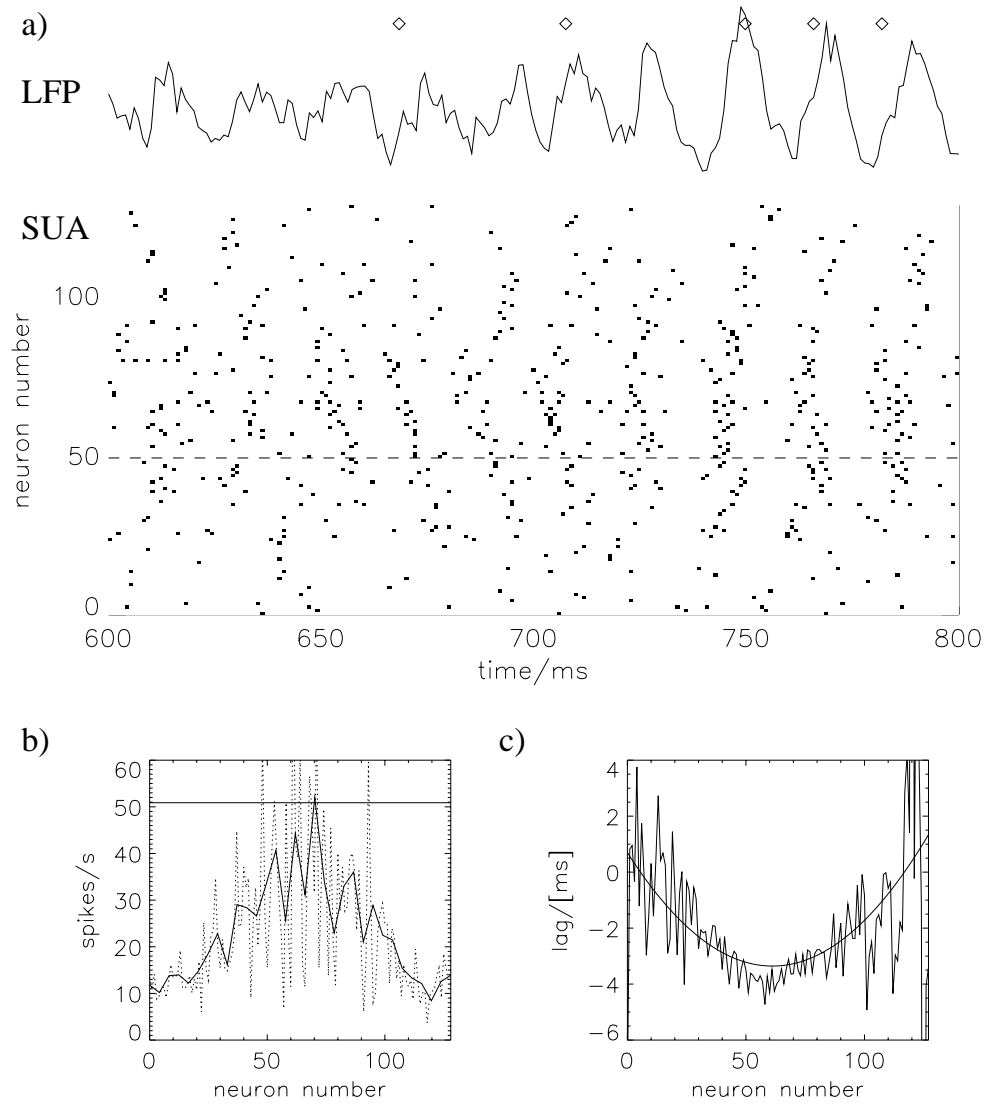


Figure 2: Activation dynamics of a local cell assembly in response to a bar stimulus. a) Local field potentials (LFP; computed as a lowpassed version of the total number of spikes per millisecond) reveal clear oscillations in the gamma-range. The rasterplot of single unit activity of the 128 excitatory neurons (SUA) nonetheless shows that cells are only loosely synchronized. b) Firing rates of all single neurons (dashed curve; average taken over five seconds) reveal 'tuning' properties, but most cells fire at rates well below the frequency of the LFP oscillation (51Hz). Thick line in b): further spatial average over seven nearest neighbours. c) On *average* suboptimally stimulated cells fire later in each LFP-cycle than optimal ones. Here, lag is measured relative to maxima of LFP-oscillation by spike-triggered averaging.

and ‘non-participating’ cells which are neither rhythmic nor coupled to the global oscillation, although they are driven by the stimulus and fire significantly!

As a consequence of this irregular firing, single cells in our model can be tuned, which means they change their firing frequencies, when the stimulus orientation (respectively the center of the Gaussian shaped input) changes. This can be seen in figure 2b, where the firing rate of all neurons is displayed (dashed line; rates are averaged over a total sample length of five seconds) together with a further average over local neighborhoods (thick line, roughly equivalent to multiple unit activity, MUA). The horizontal line at 51/s indicates the ‘oscillation’ frequency derived from the power spectrum peak of LFP. Almost all cells have lower rates in accordance with experimental results by Kreiter and Singer (1996). Furthermore the shape of the rate function in figure 2b reflects the Gaussian input strength, and this in turn the single cell tuning; deviations from the Gaussian shape are mainly due to random lateral connections. Although tuning is one of the most basic properties of cortical neurons, it cannot be observed in many network simulations, which aim to explain ‘binding by synchronization’. Those often operate in parameter-regimes, where single cells (or other kinds of simulated ‘units’) have very similar firing frequencies equal to the oscillation frequency of the collective rhythm, and almost perfectly synchronize, when they are sufficiently strongly connected. As a consequence ‘tuning curves’ are flat in the synchronized regime.

In contrast to this ‘tight-binding’ situation, in our simulations synchronization is only rather loose. Nonetheless, we should mention another property well in accordance with experimental results. König et al. (1995) investigated how precisely sub-optimally driven cells lock into the collective rhythm. Measuring peak-shifts in cross-correlograms between two sites with orientation preference ϕ_1 and ϕ_2 and varying the orientation ϕ of a stimulus, they found, that sub-optimal cells reveal a systematic phase-lag relative to optimally stimulated cells, which depends linearly on the stimulus orientation ϕ and is, in addition, proportional to the difference $\phi_2 - \phi_1$ between the preferred orientations of the recorded cells. This implies that the lag of cells with a particular orientation ψ relative to the best matching cells (or alternatively to some reference oscillation like the LFP in our simulation) must depend quadratically on ψ . This is, what is clearly seen in figure 2c, where phase lags obtained from fitted peaks of spike triggered averages of the LFP signal are displayed together with a quadratic fit of the lags estimated from the same 5 second data set than before. Of course, the strong scatter in figure 2a shows that at least in our simulations this lag is only an average property of cell firing and by no means implies deterministic delays between firing times of different neurons.

What happens in our network is the following: best matching cells on average fire more often. They can ignite waves of activation spreading to less excited neurons, which have not yet fired. Therefore, these cells fire later. This may be used as a

coding principle: the later a cell fires the less is it directly supported by the external stimulus, but instead may reveal properties of the cortical connectivity matrix. The inhibitory interneurons recognize the increase in activity and suppress further firing for some milliseconds. Due to different sources of randomness, these cycles of excitatory amplification and subsequent inhibition appear to be rather imprecise in amplitude and period duration. Nonetheless they reveal elementary properties, which cannot be explained by tight-binding theories. Two points are particularly important: first, there is a pronounced global gamma-oscillation, but single cells show a broad spectrum of typically slower firing rates and couple only loosely into the global rhythm. Second, in response to different stimulus conditions cells show systematic shifts in their firing times relative to other cells or local field potentials. These shifts, however, are only observable in crosscorrelograms computed over long times; within single gamma-periods the firing of cells appears to be unreliable and the relative timing imprecise.

After some reflection this suggests that it is less “phase coupling” of oscillatory cortical activity in the gamma range that matters in our simulations, than the temporal synchronicity and fine spike timing shifts of sets of similarly tuned cells within single periods. Sets of cells are required for reasons of proper signaling since the single units fire unreliably. The oscillation appears as a byproduct of repeated local processes characterized by a fast spread of activity in excitatory subnetworks followed by a subsequent inhibition phase, which is somewhat delayed, because inhibitory potentials are typically slower than excitatory, and the excitatory activity itself is needed to evoke the inhibition. Properties of the collective rhythm - amplitudes, frequencies - are rather imprecise. Therefore, it seems unlikely that they code for particular stimulus features. Similarly, due to the strong fluctuations correlations decay quickly in time - typically exponentially on a time scale of several ten to less than 100 milliseconds - just as found in experiments. This supports the idea that only short epochs of the signals are relevant for information processing and long coherent wavetrains (as, for example, in holography) are not necessary. In this context experiments by König et al. (1995b) are quite interesting: these show, that for electrode distances above roughly two millimeters (within or between areas) crosscorrelograms almost exclusively have oscillatory sidebands, whereas at shorter distances those can be completely missing. It seems that distant sites require (at least) two gamma-periods for effective mutual interactions.

2.2 Pattern Completion and Gamma Oscillations

The previous section considers local assemblies in primary sensory areas. Higher association areas are usually small. Therefore, we may assume that those areas can be reasonably modelled as fully connected associative memories (Palm, 1982). The

above interpretation of gamma-oscillations viewed as fast excitatory amplification processes succeeded by inhibited phases does also apply in this case. Here, the spread of activation proceeds from externally driven neurons (the address pattern) to associated neurons of the stored assembly or memory pattern. This may involve one or more synaptic transmissions steps, which might be identified with iterative feedback steps in the associative memory model. The growing inhibition, however, suppresses the activation after a short time, interrupts the retrieval process and restarts it after a short relaxation phase. This always keeps the network sensitive to changes in the input signal.

Now, it is most notably that Schwenker et al. (1996) have shown, that iterative retrieval in sparsely coded associative memories is extremely fast provided firing thresholds are adapted to the network activity in each step; then at most 3 feedback steps (and most often only one or two) suffice for perfect retrieval. In the current context this 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 about 10ms, which corresponds well with the activated phase of observable gamma-periods. Furthermore the work of Schwenker et al. (1996) shows that pattern completion is most efficient in terms of storage capacity, when the number of ones in the address pattern is about half that of the stored patterns. This 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 a few ms, all relevant information should be present during those short time intervals. These arguments show that the hypothesis that local information processing is essentially restricted to population bursts of cell pools in single gamma periods goes very well together with iterative retrieval in associative memories. Not only provide synchronized spikes a high memory capacity, but frequencies in the gamma range are also almost the fastest possible operation speed for rhythmic retrieval. A more thorough discussion of these and related topics can be found in Wennekers and Palm (1997).

Figure 3 shows a simulation of the general principle applied to associative retrieval in a higher area. Three different patterns, e.g. local assemblies, have been stored in a fully connected coupling matrix. Beside this the network structure is very similar to that in figure 2, but time constants of synaptic potentials are somewhat different; therefore, the observed oscillation has a different frequency. Note that during the retrieval phase in figure 3e single perfect patterns are recovered from the composite and incomplete input pattern (d) in a very short time. Usually only one pattern is retrieved per elementary associative process. Which pattern becomes amplified depends largely on the noise level, hence random segregation takes place,

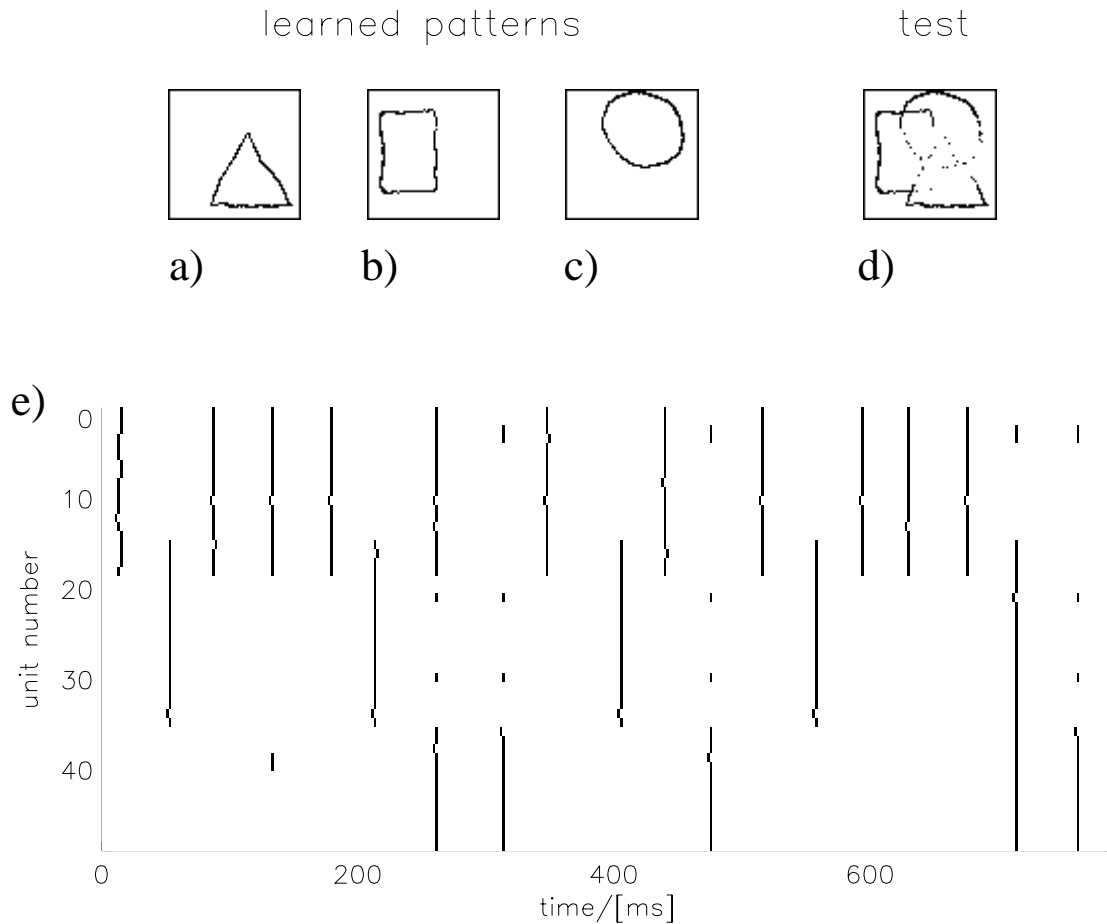


Figure 3: Rhythmic associative retrieval in a fully connected memory comprising 64 spiking neurons. Three input patterns (a-c) - coarse-grained to a size of 8×8 bins - were used as input to the memory neurons. During a preceding learning phase, each pattern was individually presented to the network for 500 steps and learned by means of a Hebbian coincidence rule (data not shown, cf. figure 5). Testpattern (d) was applied during the displayed simulation run. Shown in e) are raster plots of only those neurons, which fire in response to the stimulus; these cells are ordered (as good as possible for three overlapping patterns) by their membership to the different patterns (roughly the upper third of units represents the triangle; middle part: square; lower part plus some cells above: circle). Observe the overall rhythmicity, but a non-periodic activation of single patterns. The mixed patterns in d) are segregated and completed, but not all three phase-coded within single periods. Instead usually only one pattern is processed completely in every period.

but no phase-segregation of different patterns within single periods. The ‘oscillation’ itself has no direct functional significance except keeping the network sensitive: as soon as a pattern has been retrieved it is suppressed by the somewhat delayed inhibitory response. Afterwards a new retrieval process can take place. This way stationary attractor states are avoided in favour of more flexible recognition processes. Of course, a more orderly (apparently non-random) retrieval of the three patterns could be achieved, if we had added a further adaptation in every single cell with a much lower gain and a larger time constant (say 50ms) as in the simulations of binocular rivalry by Fahle and Palm (1991). In that case cells belonging to a pattern that fired most previously are suppressed most strongly and the chance for the firing of a pattern increases with the time elapsed since it has fired the last time.

3 SYNFIRES CHAINS

So far we dealt with temporal synchrony and gamma oscillations, which often, although not necessarily, co-occur in primary visual areas. Abeles (1991) has described another type of spatio-temporal correlations in frontal areas, which (at first glance) are not characterized by synchronized firing of cells, but consist of precisely timed sets of spikes of one or several cells with well defined relative time-delays; those spike-patterns (repeating triplets, quadruplets, etc.) - also termed synfire activity in the sequel - occur significantly more often than it would be expected by chance, given the hypothesis that spike trains are independent Poisson processes (Abeles, 1991; Abeles et al., 1993b). The occurrence of synfire activity in frontal cortical areas is furthermore clearly correlated with behavioral events (Abeles et al., 1993b). This observation relates the phenomenon to cognitive processes, although the explicit relation is still a matter of discussion.

In order to explain those precisely correlated spike events, which can extend over time-scales of up to hundreds of milliseconds, Abeles (1991) introduced the concept of ‘synfire chains’ (SFCs). The main idea is that they arise from ordered sequences of *synchronously* firing pools of neurons, which iteratively excite well defined other pools, whereby a chain of activation evolves and propagates through the network. This idea can be formalized within the framework of associative memories: to this end it suffices to envisage every single, synchronously firing pool of cells as a memory pattern and store the whole set of linearly ordered (not necessarily non-overlapping) patterns pair by pair just as described in figure 1. The contribution of Aertsen in this volume considers experimental and biophysical properties of synfire chains. We will discuss two theoretical aspects, which focus on the tentative role of synfire activity for cortical information processing. Both view synfire chains as storage elements: the first describes SFCs as a long term store for learning, recognition and replay of

spatio-temporal patterns (cf. also Wickelgren, 1969); the second takes them as a possible physical substrate for short term memory (STM).

3.1 Synfire Chains as Temporal Storage Elements

As mentioned above, synfire chains can be interpreted as an extension of the standard associative memory from static or structural patterns to spatio-temporal ones (cf. Palm, 1982, chapter 11). The examples in this section are a consequent elaboration of this idea. In fact, the regeneration of ordered sequences of patterns has been repeatedly demonstrated in SFC-models earlier (Abeles et al., 1993a; Aertsen et al., 1996; Bienenstock, 1995). Nonetheless in this section we go a step farther and show that SFCs cannot only be used to recover sequences, but also to learn time-patterns and recognize them in a fault-tolerant manner.

We demonstrate the main ideas in form of an example. To this end imagine an excitatorily connected associative network of spiking neurons as considered before. Inhibitory interneurons do no harm as long as the inhibition is not too strong to forbid the stable propagation of synfire activity. Assume that a sequence of P patterns is stored in linear order in the coupling matrix. For sake of simplicity, we further replace each pattern by a single representative cell; hence, the synfire chain network can be thought of as consisting of P cells coupled feedforward in linear order. Some mathematical analysis reveals, that this structure can show stable propagation of activity moving from the first to the last neuron in the chain, provided some global threshold level is chosen appropriately (Wennekers and Palm, 1996). Now, we want to learn, recognize and replay a certain time-pattern $a(t)$, where $a(t)$ may have more than one, for example m , components. To do this, we assume that synaptic connections C_{ij} , $i = 1, 2, \dots, P$, $j = 1, \dots, m$ exist from each component of a to each neuron in the SFC. These synapses are used to store samples of the pattern $a(t)$ at certain times provided by the ordered firing of the SFC neurons. Learning can proceed in a Hebbian way: suppose that by some mechanism the learning pattern $a(t)$ and the activation of the SFC (node 1) start simultaneously. Then, each time t_i , a neuron of the SFC fires, it suffices to store the actual values $a(t_i)$ in the synapses of the respective neuron i . This works in a single trial ('one-shot' learning).

Figure 4c displays an example simulation (with $P = 30$ and $m = 2$). The spike-raster of the SFC-neurons is shown to the right. These spikes represent the postsynaptic part of the Hebbian learning rule. Generation of the (presynaptic) input time-signal $a(t)$ requires further explanation: in principle arbitrary, sufficiently smooth signals can be used as input. Here, those are extracted from simple line-drawings as indicated in figure 4a/b. A spot (rectangle in figure 4a) starts at time zero at some corner of the object and travels along it with constant speed v . The velocity vector $\vec{v}(t) = d\vec{r}/dt = (dx/dt, dy/dt)^T$ along the curve $\vec{r}(t)$ is taken as the two-dimensional

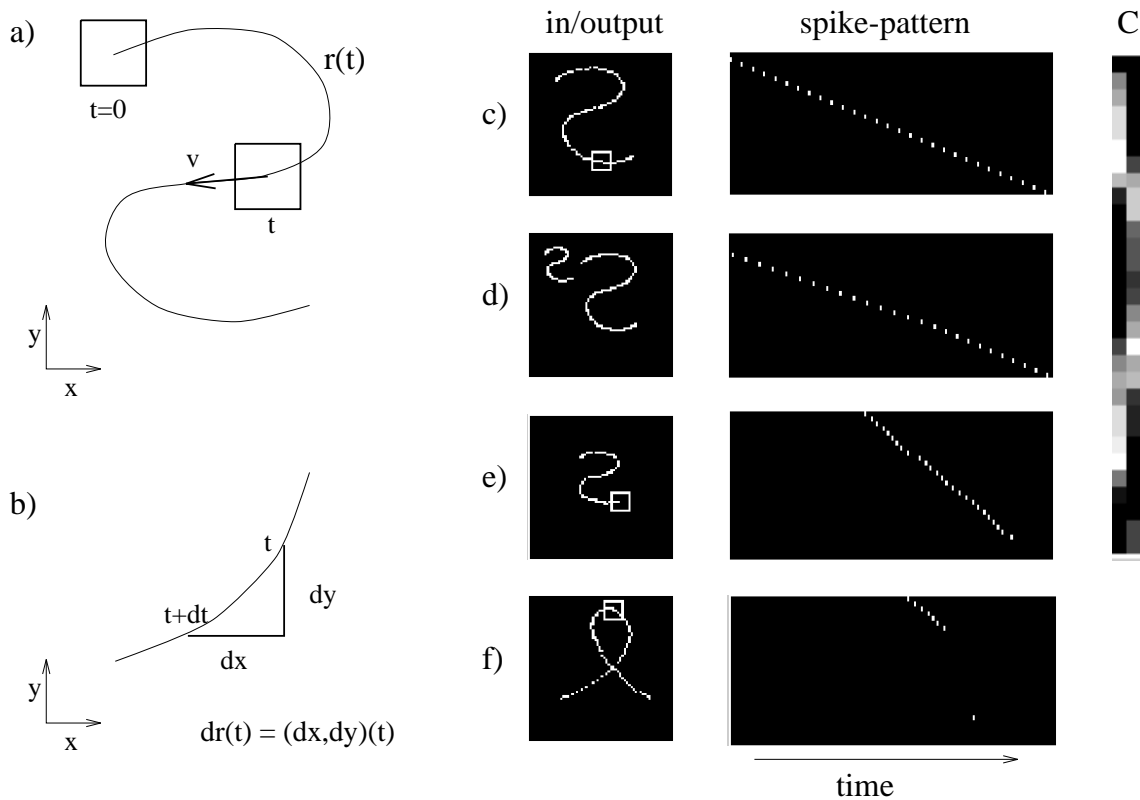


Figure 4: Synfire chains can store, recognize and replay spatio-temporal patterns. a)&b) Time patterns to learn are derived from simple line-drawings: in a) a spot moves along the curve $r(t)$ with constant velocity; the two components of the velocity vector, proportional to $dr(t)$ in b), provide inputs to each neuron in a synfire chain (SFC). c) displays learning of a trajectory: to the left a snapshot of the input space is shown and to the right the spike-raster of the synfire chain neurons. Learning in c) is Hebbian: movement (the spot) and SFC both start at $t=0$. Each time a SFC-neuron fires it stores the actual input values $dr(t)$ in its synapses. The finally learned synaptic input matrix C is shown to the right. d) displays two-fold replay of the trajectory by repeated activation of the SFC with two different gain values. Here, the previously learned synapses are used in reverse direction and control the movement (velocity vector) in input space. e) Recognition of a distorted input trajectory. f) Recognition fails for a different pattern. For explanations see text.

input function $a(t)$. Hence, every neuron in the SFC has two external inputs, which represent velocity in x respectively y-direction. When, during learning, neuron i fires a time t_i , its synapses C_{i1} and C_{i2} are set to $dx(t_i)/dt$ respectively $dy(t_i)/dt$. The matrix $C = (C_{ij})$ resulting from such a learning process is shown to the right in figure 4; obviously it represents the derivatives of the curve $\vec{r}(t)$ in 4a (white codes for movement in positive and black for movement in negative x or y direction).[‡]

Pattern regeneration is shown in figure 4d. Again the SFC has to operate in the stable regime, but now without external input. Instead, the formerly learned synapses are now interpreted as ‘output’-synapses controlling the movement in the output space. If a neuron in the SFC fires, its synapses determine the instantaneous velocity vector, with which a movement is performed. In figure 4d the SFC is activated twice - only the second spike-raster is shown. By choosing different absolute starting positions and different (arbitrary) gain factors the previously stored object is recovered in two sizes in the input space.

Finally, figures 4e and f display examples for pattern recognition. To this end the thresholds in the SFC-network must be high enough to avoid the stable propagation of excitation without a further external input. An additional temporal input into the actually firing neurons then can lead to the complete recovery of the stored sequence, provided the input pattern matches the synaptic pattern of the actually firing neurons sufficiently closely. At any step only the conjunction of the additional input and that from the previously firing SFC-neuron should lead to the firing of the next cell. If at any position in the SFC, the stored and externally applied patterns do not match, the synfire chain in the recognition network dies out. An example for proper recognition of a distorted version of the stored pattern is shown in figure 4e. Note that 1) the different parameter settings lead to a faster SFC sequence than before (cf. Wennekers and Palm (1996) for speed control of SFCs), 2) the distortions in the test pattern lead to slight fluctuations in the instantaneous speed of synfire propagation and 3) the stored pattern effectively uses only 27 of the total of 30 cells in the SFC (cf. matrix C in figure 4). Hence, under recognition conditions, the last few nodes do not fire in figure 4e. Finally, figure 4f shows failure of recognition for a completely different input pattern.

The above interpretation of synfire activity explicitly takes account of temporal information stored in the network structure. Nonetheless, we should note, that it is not very likely that synfire networks of this kind provide a reasonable substrate for *arbitrary* time-patterns an individual might learn, say for example, complex movement patterns. The main reason for this is that sequences longer than some

[‡]We should mention that the model is not intended as a concrete example for visuo-motor coordination or related tasks, although similarities might exist. Intended is a purely abstract view, just as the standard associative memory at first is an abstract paradigm. Both models may be suited as building blocks for more concrete and complex networks, incorporating static as well as temporal properties of stored entities.

hundred milliseconds need exceedingly large hardware resources (Bienenstock, 1995) and, furthermore, the resulting synfire structure is very inflexible. Motor control (for movements or speech etc.) certainly needs more flexible, probably modular and hierarchically organized structures (cf. also Wickelgren, 1969). However, we believe, that chains with roughly some 100 nodes and perhaps 100 neurons per node may be useful storage devices for elementary ‘spatio-temporal features’ in such architectures: the complete information can still be retrieved in a short time and a time-span of 200 to 500ms would indeed make sense, since this is roughly the duration of syllables or morphemes, which organize speech-production; similar timing intervals have also been proposed to organize other cognitive tasks (see Gibbon and Allan (1984) for a collection of related articles).

3.2 Synfire Chains and Short Term Memory

Primary function of the classical associative memory and its variants is that of a content-addressable storage device. Since information is layed down in synapses, this function may be identified as long term storage (LTM). A second function - possible in recurrent architectures - relates to the short term storage of information: because memory patterns present attractors of the network dynamics (at least in the classical model, see below) it is possible to keep them activated, even if the initial activating stimulus or address pattern already vanished. Different patterns or assemblies, stored permanently in the coupling matrix, may become selectively excited by appropriate stimuli. Persistent firing of the related assemblies afterwards represents the information that the particular stimulus previously occurred in the current behavioral context, either as an external event or internal ‘idea’. This aspect of short term storage in associative memories has recently been reviewed by Amit (1995).

In light of the interpretation of cortical gamma-oscillations as rhythmic fast associative processes followed by a period of inhibition outlined earlier, we face some problems. One of the main points of this interpretation is, that the retrieval state is destroyed almost immediately after it is reached; it lasts just a few milliseconds, long enough to get signaled to target structures. Thereby we *avoid* the usually rather stable, hence hardly perturbable, attractor states in standard models and in turn enable the chance for very fast responses to environmental changes. This model obviously cannot serve directly as an STM device, because it is not clear, how the activation can be transferred from one synchronous population burst to the next in the absence of sufficient external stimulation, without the assumption of exceedingly long excitatory post-synaptic potentials or other facilitating mechanisms on the time-scale of some ten ms (cf. Bibbig and Wennekers, 1996).

STM in associative memories has been demonstrated repeatedly in networks of

time-discrete noisy or noise-free threshold neurons (Little, 1974; Palm, 1982; Amit, 1995). Also in networks of more realistic spiking neurons persisting attractor states can be stable (Abbott and van Vreeswijk, 1993; Gerstner and van Hemmen, 1993). In these models STM is possible because the retrieval states utilize *asynchronous* firing of cells, which can be reached by appropriate choices of network parameters.

However, long-time correlations of the synfire type in persisting activation states are not possible in such networks. It is known that in networks of noisy threshold neurons spikes of different cells in different time-steps become uncorrelated when the system size gets large. This is because the relative contribution of any single neuron to the potential of a target neuron becomes vanishingly small in comparison with the total input (cf. Wennekers and Pasemann, 1996, and references therein). On the contrary, in networks of spiking neurons with refractoriness (as far as demonstrated to date) cells in persistent attractors fire more or less periodically (depending on the noise level). This can lead to long-time correlations, but of a different kind in comparison to synfire activity. In fact, only relatively few cortical cells, especially in prefrontal areas, fire periodically.

It may well be that different areas of the cortex operate in different parameter regimes, such that association areas subservise STM functions by low rate firing, weak synchronisation and synfire activity, whereas sensory areas show gamma-activity with the possibility of quickly changing patterns. But it is also not unreasonable to assume that all three phenomena - cortical gamma, synfire activity and STM - co-exist in one and the same local network. This has not yet been shown explicitly in a single experimental set-up, but gamma is known to be a rather prominent rhythm in many cortical structures (Steriade et al., 1996; Gray, 1994) and synfire activity has also been demonstrated in visual cortex and auditory thalamus (Lestienne and Strehler, 1987; Villa and Abeles, 1990). To date the least evidence has been found for physiological signs of STM in reverberant loops of activity (see perhaps Miyashita and Chang, 1988; Fuster and Jervey, 1980; Fuster, 1989), but if this concept of short term storage is biologically relevant at all, it somehow should be expected to exist at least in wide parts of the association cortexes. None of the above cited models - including our own - can explain gamma oscillations, persistent activity and synfire patterns at the same time without modifications; hence we may ask whether or not there exists a unified view.

To this end synfire chain models seem to be a natural starting point. It is clear, that they can serve as a temporal short term store. Bienenstock (1995) argues, that the control of activity in such networks may even be easier than in more homogeneously connected networks. Whether or not synfire chains are also consistent with gamma-oscillations has been investigated by Abeles et al. (1993a). Their simulation studies show, that oscillatory behavior can indeed be found in two mutually connected networks of reverberating synfire chain models. Preliminary own investigations

of a single reverberating SFC-model with global inhibitory activity control (similar to the networks already described above) show oscillations only near stability boundaries of stable propagation of synfire waves. Here, the mechanism is essentially an instability of the inhibitory control loop while suppressing randomly firing cells that are uncorrelated with the synfire activity.

A second integrative scenario, perhaps with more physiological significance, assumes, that Hebbian cell assemblies are distributed over wide parts of a single or several cortical areas. In such a framework it is possible that rhythmic associative processes may occur locally by the same mechanism and with the same interpretation as discussed above. However, to bridge the inhibitory phases and obtain short term memory it is necessary that local synchronously firing pools of cells excite other cell groups in sufficiently large distance. The locally synchronous activity together with the patchy structure of cortical long-range connections support this hypothesis, inasmuch as the high convergence of activity in target columns under the influence of synaptic plasticity should lead to rather specifically and reliably excited cell pools in those columns similar to synfire nodes (Sommer et al., 1998). Finally a sufficiently distributed ensemble of such mutually connected patches should lead to persistent reverberating synfire-type activity in conjunction with local gamma oscillations (which in this case are not necessarily globally synchronized; see our discussion in Wennekers and Palm, 1997). Furthermore, the results in Sommer et al. (1998) indicate that associative modules connected and operating in a bidirectional manner provide means for associative storage that can be much more efficient and advantageous than single localized auto-associative memories.

4 GLOBAL ASSEMBLIES

At the end of the last section we arrived at the concept of ‘global cell assemblies’. Those widely distributed assemblies may include sub-assemblies in different sensory modalities as well as in higher association areas. Therefore, they may serve as representational schemes for virtually any kind of entity including things, situations, contexts, concepts, etc. However, in some sense the concept of global assemblies only presents a generalization of localized associative memories to networks of such networks, where the ‘super-network’ may itself reveal associative properties. This still holds, if we include spatio-temporal features as exemplified in figure 4, although the possible physical modes of temporal behavior will certainly become very complex in that case. Nonetheless, the main purpose is intrinsically that of representation and association of patterns of excitation, and less that of cognitive reasoning, planning and complex behavior in time. In this section we outline some ideas related to the latter topics.

4.1 Cortico-Hippocampal Interplay

In a series of computational studies we investigated global cell assemblies in some detail (Palm, 1993; Bibbig et al. 1995, 1996; Bibbig and Wennekers, 1996; Wennekers and Palm, 1997; Sommer et al., 1998) In particular we were interested in the role of the hippocampal formation during memory consolidation and retrieval of information in networks of higher neocortical associative areas (see also Klimesch, this volume; Klimesch, 1994; Miller, 1991).

Two sensory pathways, say visual and acoustical, were considered in our model, which both consisted of a primary ‘pre-processing’ area and a higher uni-modal ‘associative’ area. The higher, more central areas, A_1 and A_2 , were bi-directionally connected with a further associative structure, supposedly the hippocampus H (more generally, this could also be a higher cortical association area, maybe a ‘convergence zone’ in the sense of Damasio, 1989). The model structure of the individual sub-networks A_1 , A_2 and H was virtually the same than that considered in earlier sections. External inputs to the sensory streams were static, simple geometrical line-figures for the visual and abstractions of tones or frequency combinations for the acoustical branch. Several inputs could be applied in different combinations. The ‘pre-processing’ areas mainly transformed a specific uni-modal input into a local representation in the respective uni-modal association cortex, A_1 or A_2 . Spatio-temporal properties of these peripheral parts of the model are similar to those of the local assemblies described above (cf. also Wennekers and Palm, 1997 and Bibbig et al., 1995 for details). Therefore, we focus on the combined behavior of the neocortical memories and the hippocampus in the following.

It is known, that activity in higher association areas is often rather sparse. Furthermore, those areas are usually not densely connected with each other. Estimates by Palm (1993) suggest, that the probability for synchronous pre- and postsynaptic activation sufficient for Hebbian strengthening of a synapse is only very small. Without further supportive input, learning of local assemblies in and global assemblies including those areas seems unlikely. Because the hippocampus receives input from virtually all higher neocortical areas and can in turn also influence them, we have hypothesized that this brain structure plays a key role in the learning of global assemblies by providing such support (Palm, 1993). Even if it sends out unspecific activity to the higher areas, probabilities for coincidences increase steeply, because the neuronal threshold process is highly nonlinear. Since, furthermore, the hippocampus is itself an associative structure, which at any moment receives a reduced but global picture of the ongoing neocortical activation, the hippocampal activity will by no means be random, but may organise into reduced representations, say specific local assemblies or chunks, which store the information about the conjunctive occurrence of subevents in different areas of associative cortex (Wickelgren, 1992).

Hebbian strengthening of feedback connections from hippocampal to neocortical local assemblies then may lead to stable representations of global cortical states in cortico-hippocampal loops (cf. Miller, 1991). This may be the basis for more complex hierarchical representations of ideas and concepts that are less clearly related to sensory impressions.

Figure 5 displays a computer simulation of this process. A_1 , A_2 and ‘Central’ reveal rasterplots of spikes in two neo-cortical areas respectively the hippocampus. The above two black lines indicate that certain input stimuli are supplied to each of the sensory pathways during the whole simulation run. From step 0 to 500 these lead to apparently random firing of a fraction (here 50%) of the cells in both areas A_1 , A_2 . The curve ‘Learning’ measures accumulative changes in synaptic efficacies due to synaptic plasticity (Hebbian coincidence rule) in A_1 (solid) and A_2 (dashed). Obviously learning is very slow during the first 450 steps of the simulation. Then, the hippocampal area receives some further pacemaker signal, which in real brains may be supplied by the septal region. This signal in the theta-frequency range serves as a rhythmic threshold control for the hippocampus and - together with input from A_1 , A_2 - leads to spikes of some hippocampal cells, which, after being transferred back to the neocortical areas, almost immediately start to organize the activity in those areas into synchronized population bursts accompanied by a significant increase in learning rates. After several theta periods synaptic efficacies saturate at some maximum level, and the neocortical areas reveal a pronounced gamma-rhythm similar to that in figures 2 and 3. [§]

After learning the hippocampal formation binds distributed sets of local assemblies into global ones, which now can be retrieved from only partial information in a single sensory stream: first, the respective local assembly is completed in its unimodal association area, and then - via the cortico-hippocampal-cortical loops - the hippocampal contextual chunk is reactivated; thereby it addresses and restores the complete widely distributed information in neocortex (Miller, 1991; Wickelgren, 1992; Bibbig et al., 1995) We have also demonstrated, that in case of learning of very many contextual situations, generalization across stimulus properties can take place in the hippocampus, with a qualitative change of its influence from supporting highly specific chunks to a more coarse grained threshold control (Bibbig et al., 1996) Finally we should emphasize that the neocortical areas themselves were not mutually connected in the simulations. It is clear that also those connections would become consolidated during learning, with the consequence that at the end cross-modal retrieval may rely solely on neocortical interactions as in the work of Sommer et al. (1998).

[§]Due to technical constraints (limited computer time) the time-scales of gamma and theta rhythms and probably also the learning rates are not realistic.




Figure 5: Hippocampal support of learning of a global cell assembly. Two associative areas/memories A_1 and A_2 (say visual and acoustical) are bidirectionally connected with the hippocampus ('Central'); this is further controlled in the theta range by the septal pacemaker. Input patterns presented to A_1 and A_2 evoke apparently random firing of cells as long as the pacemaker does not drive hippocampal cells (steps 0 to 400). Therefore, spike coincidences and learning rates are small (solid learning curve: A_1 , dashed: A_2 ; displayed are accumulative changes in overall synaptic efficacies). Activation of the septal pacemaker supports selective firing of some hippocampal cells, which represent the current stimulus combination. These spikes propagate back to A_1 and A_2 and organize the activity in these areas into synchronized population bursts accompanied by a strong increase in learning rates and oscillatory activity (similar to figure 3). This way local assemblies are consolidated in A_1 and A_2 and the global contextual information is stored in the cortico-hippocampal loop.

4.2 Associative Brain Models

Several brain theories based on assembly- or associative memory models as outlined above have been devised (Hebb, 1949; Braitenberg, 1978; Wickelgren, 1981, 1992; Palm, 1982; Damasio, 1989; McGregor, 1993; and more) Interestingly within such general frameworks the question whether persistent assemblies are supported in the brain by stochastic or correlated activity has also been discussed early (e.g. Hebb, 1949; Eccles, 1958).

McGregor (1993) called the two alternatives ‘stochastic’ respectively ‘sequential configuration hypothesis’ and outlined a theory of neocortex according to which entities are represented locally by ‘sequential configurations’ or dynamical modes’ very similar to reverberating synfire chains in local networks. Similar modes in different local modules can support each other by long-range connections (the ‘super-network’, see above), thereby leading to specific collective excitation patterns distributed over almost the whole cerebral cortex. Unfortunately the theory is not yet developed far enough to suggest a functional role for synfire activity or temporal structure in brain signals. Most effects in McGregors paper occur, for example, also in coupled associative modules with exclusively static attractor states or ‘modes’.

Bienenstock (1995) is more explicit with regard to the role of ‘sequential configurations’. He proposed a model of the neocortex, focussing on ‘cognitive compositionality’. This takes synfire chains as basic, quasi atomic, functional elements or excitation modes, which are not necessarily restricted to local modules as in McGregors work. According to Bienenstock elementary or ‘narrow’ chains themselves carry very little information, but can aggregate or synchronize with each other by interactions via mutual, relatively weak, and plastic synapses. Only conjunctive ‘broad chains’, which are composed of flexible sets of synchronized narrow chains carry meaningful information. Those broad chains again are thought of as being distributed over large parts of the cortex. This way an almost infinite variety of complex excitation patterns seems possible, an argument, which has also been put forward by Palm (1982) as an advantage of earlier assembly theories in favour of theories based on grandmother neurons. In some sense also Bienenstocks model is essentially a hierarchical associative memory model; the difference to earlier theories is, that he replaces stationary attractors by SFCs and includes ideas concerning ‘binding by synchronization’, now applied to coupled SFCs and not coupled oscillators.

4.3 Towards Cognitive Operations

Experiments performed by Abeles et al. (1993b) on trained awake monkeys show, that synfire activity in frontal cortical areas is correlated with behavioral events. Those areas are known to contribute to short term and temporal memory, problem

solving and planning in complex behavioral tasks (Fuster, 1989). Therefore, it seems reasonable to ask how cognitive capabilities might be included in assembly theories and which particular role synfire activity may play in this context.

Simple stimulus response schemes may be easily explained: it suffices that an assembly is learned during training that represents a particular stimulus and triggers the corresponding response in case it becomes activated in subsequent tasks. This way, also conjunctions or other logical combinations of different simultaneous stimuli may become associated with a certain response. Those mappings would implement basic ‘rules’ of reactive behavior.

Time-order relations between cues may still be represented statically in form of assemblies “A” and “B” for events, and other assemblies “AB” and “BA” for their occurrence in different temporal order. The latter assemblies may then trigger the desired response. Nonetheless, we face the problem how “AB” is excited from stimuli A and B. Possible solutions require short term storage of the first stimulus and subsequently the excitation of assembly “AB” by both, the external input B as well as the internal representation of A. Clearly, both these aspects - in general terms: short term memory of the temporal stimulus context, and operations on (or influenced by) those internal representations - exceed the simple stimulus-response scheme and are central to any theoretical framework of ‘cognition’. Our assembly theory is well suited for representational purposes, but still lacks elaborated operational components. This should not be a serious problem, in principle, because arbitrary finite automata can be built from very simple neural networks (McCulloch and Pitts, 1943; von Neumann, 1958). Therefore, a general and rough picture for cognitive operations based on assemblies is developed in the following:

First, we assume, that external stimuli or events relevant to a certain task are represented in form of assemblies, which comprise local sub-assemblies in frontal areas; those may have short-term storage properties and reveal synfire type activity as discussed earlier. This set of assemblies or ‘memory chains’, perhaps together with further ones for the storage of internal events (like “AB”, “BA”, see above), can be viewed as a set of ‘(logical) variables’ representing entities relevant to the particular experimental situation. Storage assemblies can be activated or silent; they represent propositions about the experimental situation. For example, activation of assembly “AB” means that event A has occurred and afterwards also B; its inactivation that this is not true.

Second, we assume the existence of control structures able to perform operations on memory chains (e.g. the activation of “AB” if “A” is already active and B occurs). Those structures may implement logical and/or procedural knowledge; they present neuronal ‘programs’. Interestingly, such structures can also be implemented by synfire chain networks. It turns out that a simple modification leads from reverberating synfire chains to models of spiking neurons able to implement arbitrary

finite state automata. All that is required are several possible successor states at any node of the chain, which are activated not only by the firing of predecessors, but also in dependence of specific input-patterns. This way complex graph-like synfire structures can be built, which implement the desired control components.

Third, memory and control networks are assumed to interact with each other and also with further input/output structures like sensory-motor or higher associative areas: firing of specific nodes in the ‘program’-network may trigger elementary motor patterns, which are perhaps themselves represented by synfire chain like activity as discussed in section 3. They also can induce sampling of new input by threshold control in appropriate sensory areas. More locally, firing of program-nodes may excite or inhibit activity in some local memory chains, or gate the transfer of activity from one storage element to another. Contrary, the flow of activity in the program network can be influenced by the actual state of one or more of the memory chains as well. Also external input may induce conditioned transitions in the program-network, such that different behavioral procedures can be performed in different situations in time.

The above assumptions suffice to model arbitrary complex tasks; for example, we have implemented a network of spiking neurons that is capable to perform arithmetics based on generalised synfire chains as described above (Wennekers, 1998). Similarly also other behavioural or cognitive tasks can be realised.

In many behavioural experiments, monkeys need a rather long time to achieve the skills required for good performance. It is plausible to assume that the structures for storage and processing (the ‘cognitive modules’) are acquired during this training phase and are essentially fixed afterwards. This suggests, that these structures presumably are quite specific, that is, built and usable for the particular task only. With respect to local assemblies in frontal areas this would mean that those are not ‘universal variables’, but can only represent specific external events, expressed by their connectivity pattern to a corresponding global assembly. Now it seems possible that repeated training of the same task, but with many different, interchangeable stimuli may lead to some kind of generalisation across stimuli, expressed, for example, as a dissociation of the local assembly from particular features of the stimuli and remaining connections to general stimulus properties. In that case the ‘cognitive module’ loses the ability to refer back to a specific actual global assembly; it becomes more universal by generalisation across stimulus properties. A similar situation occurs also in delayed match to sample experiments, because the sample changes from trial to trial and may even be completely new to the animal. Hence, there might not even exist a fixed global representation for it. Both cases are characterised by the existence of local frontal assemblies which do not correspond to fixed and unique global assemblies. Therefore we need some mechanism that transiently binds the local storage assembly, which is part of the well-trained cognitive

module responsible for the general class of tasks, to the respective representations of the actual sample stimuli. Apparently, the mechanism of memory formation based on cortico-hippocampal interplay (see 4.1) is able to solve this problem. It is interesting in this context that tasks which require the continuous learning of new representations appear to be in some sense ‘more difficult’ than those which solely rely on fixed sets of stimuli (cf. Klimesch, 1994; Mishkin et al., 1984). The latter experimental paradigms may involve only associations along neocortical pathways, which have been strengthened during previous training by repeated presentation of the fixed stimulus set; this should lead to fast and reliable recognition under test conditions. On the other hand, the former paradigm requires a more complex architecture and the acquisition of a new global representation from a single stimulus; hence the recognition may be less reliable and slower.

5. Conclusions

In summary, we have developed a rough picture of cortical function based on cell assemblies incorporating the time-structure in neural signals as well as operational components. Local information processing has been characterised by fast recurrent associative amplification processes serving feature binding and “Gestalt”-principles in sensory areas and pattern recognition or segregation in higher association areas. Local assemblies in a modular architecture can be integrated into global ones by means of numerous uni- and bi-directional synaptic pathways. The formation of such assemblies - that is, the integration of information from different modalities and other internal sources - as well as their consolidation and crossmodal retrieval can well be supported by cortico-hippocampal loops. Synfire chains, i.e. sequential synchronised activity along specific synaptic pathways, have further been shown to provide a basis for short term memory as well as the storage of spatio-temporal features of internal or external events. Finally, by extending the synfire chain concept to “synfire graphs”, the controlled interaction of assemblies distributed over many cortical areas can serve as a basis for operational short term memory organizing in principle arbitrary complex behavioural responses.

Acknowledgement

We are grateful to Andrea Bibbig, Reinhard Eckhorn and Axel Frien for valuable discussions, and to Robert Miller for many comments that led to improvements of the manuscript. This work has been supported by the German Research Foundation DFG (grants Pa 268/8-1 and 268/12-1).

REFERENCES

- Abbott, L.F. and van Vreeswijk, C. (1993) Asynchronous states in networks of pulse-coupled oscillators. *Physical Review E*, **48**, 1483-1490.
- Abeles, M. (1991) *Corticonics: Neural circuits of the cerebral cortex*. Cambridge: Cambridge University Press.
- Abeles, M., Vaadia, E., Bergman, H., Prut, Y., Headman, I. and Slovin, H. (1993a) Dynamics of Neuronal Interactions in the Frontal Cortex of Behaving Monkeys. *Concepts in Neuroscience*, **4**, 131-158.
- Abeles, M., Bergman, H. Margalit, E. and Vaadia, E. (1993b) Spatio-temporal firing patterns in frontal cortex of behaving mokeys. *Journal of Neurophysiology*, **70**, 1629-1643.
- Aertsen, A., Diesmann, M. and Gewaltig, M.O. (1996) Propagation of synchronous spiking activity in feedforward neural networks. *Journal of Physiology (Paris)*, **90**, 243-247.
- Amit, D.J. (1995) The Hebbian Paradigm Reintegrated: Local Reverberations as Internal Representations. *Behavioral and Brain Sciences*, **18**, 617-657.
- Bibbig, A., Wennekers, Th. and Palm, G. (1995) A Neural Network Model of the Cortico-Hippocampal Interplay and the Representation of Contexts. *Behavioural and Brain Research*, **66**, 169-175.
- Bibbig, A., Wennekers, Th. and Palm, G. (1996) A Neural Network Model of the Cortico-Hippocampal Interplay: Contexts and Generalization. *Engineering Applications of Artificial Intelligence*, **9**, 145-151.
- Bibbig, A. and Wennekers, Th. (1996) Hippocampal two-stage Learning and Memory Consolidation. In: *Proceedings of the 13th European Meeting on Cybernetics and Systems Research, Vienna, 1996*.
- Bienenstock, E. (1995) A model of the neocortex. *Network*, **6**, 179-224.
- Braitenberg, V. (1978) Cell Assemblies in the Cerebral Cortex. In: *Theoretical Approaches to Complex Systems*, edited by R.Heim and G.Palm, pp.171-188. Berlin: Springer.
- Braitenberg, V. and Schüz, A. (1991) *Anatomy of the cortex*. Berlin: Springer.
- Damasio, A.R. (1989) Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition*, **33**, 25-62.
- Eccles, J.C. (1958) The physiology of imagination. *Scientific American*, **199**, 135-146.

- Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M. and Reitboeck, H.J. (1988) Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biological Cybernetics*, **60**, 121-130.
- Eckhorn, R. and Obermüller, A. (1993) Single Neurons are Differently Involved in Stimulus-Specific Oscillations in Cat Visual Cortex. *Experimental Brain Research*, **95**, 177-182.
- Fahle, M. and Palm, G. (1991) Perceptual rivalry between illusory and real contours. *Biological Cybernetics*, **66**, 1-8.
- Fuster, J.M. (1989) *The Prefrontal Cortex: Anatomy, Physiology, and Neuropsychology of the Frontal Lobe*, 2nd edn. New York: Raven.
- Fuster, J.M. and Jervey J.P. (1980) Inferotemporal neurons distinguish and retain behaviorally relevant features of visual stimuli. *Science*, **212**, 952-955.
- Gerstner, W. and v.Hemmen, J.L. (1993) Associative memory in a network of 'spiking' neurons. *Network*, **3**, 139-164.
- Gibbon, J. and Allan, L. (eds.) (1984) Timing and Time Perception. Special Issue of *Annals of the New York Academy of Sciences*, **423**.
- Gray, C.M. (1994) Synchronous Oscillations in Neuronal Systems: Mechanisms and Functions. *Journal of Computational Neuroscience*, **1**, 11-38.
- Hebb, D.O. (1949) *The organization of behavior. A neuropsychological theory*. New York: Wiley.
- Klimesch, W. (1994) *The structure of long-term memory: A connectivity model of semantic processing*. Hillsdale, NJ: Lawrence Erlbaum.
- König, P., Engel, A.K., Roelfsema, P.R. and Singer, W. (1995) How Precise is Neuronal Synchronization. *Neural Computation*, **7**, 469-485.
- König, P.; Engel, A.K.; Singer, W. (1995b) Relation between oscillatory activity and long-range synchronization in cat visual cortex. *Proceedings of the National Academy of Science*, **92**, 290-294.
- Kreiter, A.K. and Singer, W. (1996) Stimulus-Dependent Synchronization of Neuronal Responses in the Visual Cortex of the Awake Macaque Monkey. *Journal of Neurophysiology*, **16**, 2381-2396.
- Lestienne, R. and Strehler, B.L. (1987) Time structure and stimulus dependence of precisely replicating patterns present in monkey cortical neuronal spike trains. *Brain Research*, **437**, 214-238.
- Little, W.A. (1974) The existence of persistent states in the brain. *Mathematical Bioscience*, **19**, 101-120.

- MacGregor, R.J. (1993) Composite cortical networks of multimodal oscillators. *Biological Cybernetics*, **69**, 243-255.
- Malsburg, C. von der (1981) The correlation theory of brain function. *Internal Report 81-2*. Max-Planck Institut für Biophysikalische Chemie, Göttingen.
- McCulloch, W.S. and Pitts, W. (1943) A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5:115–133.
- Miller, R. (1991) *Cortico-Hippocampal Interplay and the Representation of Contexts in the Brain*. Berlin: Springer.
- Mishkin, M., Malamut, B. and Bachevalier, J. (1984) Memories and habits: Two neural systems. In: *Neurobiology of Learning and Memory*, edited by Lynch, G., Gaugh, J.L. and Weinberger, N.M., pp. 65-77. New York: Guilford Press.
- Miyashita, Y. and Chang, H.S. (1988) Neural correlate of pictorial short term memory, *Nature*, **331**, 68-70.
- Neumann, J.von (1958) *The computer and the brain*. Yale University Press, New Haven, London.
- Palm G. (1982) *Neural Assemblies. An Alternative Approach to Artificial Intelligence*. Berlin: Springer.
- Palm, G. (1991) Memory Capacities of Local Rules for Synaptic Modification. *Concepts in Neuroscience*, **2**, 97-128.
- Palm, G. (1993) Cell Assemblies, Coherence and Cortico-Hippocampal Interplay. *Hippocampus*, **3**, 219-225.
- Schwenker F., Sommer F.T. and Palm G. (1996) Iterative Retrieval of Sparsely Coded Associative Memory Patterns. *Neural Networks*, **9**, 445-455.
- Singer, W. and Gray, C.M. (1995) Visual feature integration and the temporal correlation hypotheses. *Annual Review of Neuroscience*, **18**, 555-586.
- Sommer, F.T.; Wennekers, Th.; Palm, G. (1998) Bidirectional completion of cell assemblies in the cortex. In: J.M.Bower (ed) *Computational Neuroscience: Trends in Research*. Plenum Press, New York.
- Stein, R.B. (1967) Some Models of Neuronal Variability. *Biophysical Journal*, **7**, 37-68.
- Steriade, M., Amzica, F. and Contreras, D. (1996) Synchronization of fast (30-40Hz) spontaneous cortical rhythms during brain activation. *Journal of Neuroscience*, **16**, 392-417.
- Villa, A.E.P. and Abeles, M. (1990) Evidence for spatiotemporal firing patterns within the auditory thalamus of the cat. *Brain Research*, **509**, 325-327.

- Wennekers, Th., Sommer, F.T. and Palm, G. (1995) Iterative Retrieval in Associative Memories by Threshold Control of Different Neural Models. In: *Proceedings of the Workshop on Supercomputers in Brain Research*, edited by H.J. Hermann. World Scientific Publishing Company.
- Wennekers, T. and Palm G. (1996) Controlling the Speed of Synfire Chains. In: *Proceedings of the International Conference on Artificial Neural Networks ICANN96*, edited by von der Malsburg, C. et al., pp. 451-456. Berlin: Springer.
- Wennekers, T. and Palm G. (1997) On the relation between neural modelling and experimental neuroscience. *Theory in Biosciences*, **116**, 273-289.
- Wennekers, T. and Pasemann F. (1996) Synchronous Chaos in High-Dimensional Modular Neural Networks. *International Journal of Bifurcation and Chaos*, **6**, 2055-2067.
- Wennekers, T. (1998) Synchronisation und Assoziation in Neuronalen Netzen. PhD Thesis in Computer Science, University of Ulm, Germany.
- Wickelgren, W.A. (1981) Human Learning and Memory. *Annual Review of Psychology*, **32**, 21-52.
- Wickelgren, W.A. (1969) Context-sensitive coding, associative memory, and serial order in (speech) behavior. *Psychological Review*, **76**, 1-15.
- Wickelgren, W. (1992) Webs, cell assemblies, and chunking in neural nets. *CINS* 3(1), 1-53.
- Willshaw, D.J., Buneman, O.P. and Longuet-Higgins, H.C. (1969) Non-holographic associative memory. *Nature*, **222**, 960.