

Models of distributed associative memory networks in the brain*

Friedrich T. Sommer¹ and Thomas Wennekers²

¹Redwood Neuroscience Institute, 1010 El Camino Real, Suite 380
Menlo Park, CA 94025, USA

²Max Planck Institute for Mathematics in the Sciences,
Inselstraße 22-26, D-04103 Leipzig, Germany

Abstract

Although experimental evidence for distributed cell assemblies is growing, theories of cell assemblies are still marginalized in theoretical neuroscience. We argue that this has to do with shortcomings of the currently best understood assembly theories, the ones based on formal associative memory models. These only insufficiently reflect anatomical and physiological properties of nervous tissue and their functionality is too restricted to provide a framework for cognitive modeling. We describe cell assembly models that integrate more neurobiological constraints and review results from simulations of a simple nonlocal associative network formed by a reciprocal topographic projection. Impacts of nonlocal associative projections in the brain are discussed with respect to the functionality they can explain.

Keywords: Cell assemblies; Reciprocal connections; Bidirectional associative memory; Memory systems; Cognitive modeling;

1 Introduction

Building on earlier ideas about cortical function (James, 1890; Lorente de No, 1938), D. O. Hebb (1949) proposed distributed neural representations for

*published in Theory in Biosciences (2003) 122:55–69

psychological concepts such as objects, ideas, situations, etc. Hebb's literally trend-setting contribution was to relate cognitive functions of concept learning and associations between concepts directly to processes in the neural and synaptic dynamics. For concept learning he postulated a simple synaptic modification mechanism that is now referred to as the *Hebb rule*: If the representation of a particular concept is externally stimulated, excitatory cortico-cortical synapses become potentiated that experience coincident pre- and postsynaptic neural activity. The consequence of such a learning procedure is a connectivity graph in which cells representing a learned concept form "cliques" with strong mutual excitation, the so-called cell assemblies. According to Hebb, cell assemblies can be recalled or evoked by what we nowadays call content addressing: Whenever a large enough subset in the assembly is stimulated, activity spreads preferably within the assembly and tends to activate it as a whole. Independent from Hebb, Hayek (1952) developed in his theory of perception quite similar ideas about coincidence based synaptic learning. Hebb and Hayek were among the first to propose a brain theory essentially based on a neural network model. The class of neural network models that implements most directly the ideas of Hebb and Hayek is nowadays referred to as associative memories. Taylor (1956) and Steinbuch (1961) introduced the first formal models of associative memories with the potential to be analyzed mathematically and to be implemented artificially. However, at that time neither the methods for mathematical treatment nor the technology for large-scale artificial implementations were available. A couple of years later it was in fact demonstrated that associative memories can be efficient matched filters for pattern recognition problems (Willshaw et al., 1969; Kohonen, 1972). It was also shown that in formal associative memories with feedback, the memories correspond to fixed-point attractors (Little, 1974; Gardner-Medwin, 1976; Hopfield, 1982) and can be analyzed by methods from statistical physics (Amit et al., 1985). These results on attractor networks dominated the first generation of *computational theories of cell assemblies*, which we will briefly address in section 2.

In the late 1980s, with increasing availability of computer resources, simulation experiments could be used to investigate computational theories that unlike formal neural networks reflected biophysical properties of nerve tissue in a quantitative way. However, cell assemblies were barely the subject of studies in computational neuroscience because there was little experimental evidence for them. In the meantime experimental and data-analysis methods developed and evidence for cell assemblies started to accumulate (Riehle

et al., 1997; Aertsen et al., 1989; Grün et al., 2002). It therefore seems time to recapitulate recent developments in computational theories of cell assemblies, in particular with regard to their ability to reflect neuro-anatomical and -physiological properties, and to provide a framework for cognitive brain research as outlined in a number of conceptual papers (Braitenberg, 1978; Edelman, 1982; Damasio, 1989; Miller, 1991; Pulvermüller, 1992).

2 Early theories of cell assemblies

Progress in the analysis of neural associative memories, mainly by extending methods from statistical mechanics during the 1980s (Amit et al., 1985; Gardner, 1988) led to a first quantitative picture of cell assemblies in the framework of the first-generation computational theories. A negative side effect of this progress was to narrow the scope of the much broader brain theory originally outlined by Hebb to models and mechanisms the analysis methods could handle. The first-generation computational theories are based on formal auto-associative memories Hopfield (1982): These employ units with a sigmoid transfer function and assume a fully and symmetrically connected network. Typically, finite-size effects (effects particular to networks of finite size), realistic neuronal and synaptic dynamics, structured architectures, and biological features of spatio-temporal coding are neglected and did not reach a broader attention until recently. Conceptually, the paper of Hopfield (1982) and the development of first-generation theories it entailed, was of course a huge landmark. However, these theories failed to reflect important properties of real neural networks and requirements of a computational brain theory:

- *Lack of anatomical correspondence:* Neuroanatomy revealed that nerve cells in the brain are neither fully connected, nor just statistically diluted (two cases that can be treated analytically). In fact, regionally varying connectivity densities and long-range connections obeying specific architectural laws seem to be hardwired in the cortex (see below).
- *Lack of physiological correspondence:* (i) There is increasing evidence that neurons employ spatio-temporal codes, e.g., expressed in the exact timing of single spikes in one cell relative to the firing of other cells (Abeles, 1991; Singer and Gray, 1995). However, formal associative memories exclude the description of temporal neuronal codes that go

beyond spike rates. (ii) Electrophysiological and optical recordings further indicate that there is changing activity rather than relaxation to fixed-point attractors even under conditions of stationary stimulation (Arieli et al., 1996). Some evidence for persistent activity, however, has been reported in temporal and frontal areas, cf., e.g., Fuster (1995).

- *Lack of psychological relevance:* The original claim of Hebb’s cell assembly theory was to model psychological processes. To approach this goal the repertoire of memory formation and access in standard associative models seems to be much too limited: (i) Hebb’s cell assemblies represent psychological concepts that are distributed and overlapping to such a degree that they can evoke each other, but associations between different memories cannot be performed in standard attractor networks. (ii) From a psychological point of view, memories are not just monolithic. They can have facets in various modalities and appear to provide a binding function between features (Damasio, 1989; Mesulam, 1998), sometimes in a hierarchical manner. (iii) It seems that a number of learned concepts can coexist in working memory. In standard associative memories only one memory can be activated at a time.

3 Neurobiologically constrained assemblies

3.1 Definition of basic computational units

Using an artificial neural network architecture—like an associative memory model—as a computational model of a neuroanatomical structure requires several assumptions by the modeler. The first assumption concerns the choice as to which neuroanatomical entities should correspond to the basic model units. In the theoretical neurosciences even this choice is anything but univocal. A large body of literature about cell assemblies regards excitatory cells (cortical pyramidal cells) as the basic units in computational models. An alternative choice are local groups of cells that can be considered as functionally homogenous as for instance Mountcastle’s minicolumns (Mountcastle, 1978), cf. the article by Lansner et al. in this special issue.

The concept of minicolumns was confirmed also in electrophysiological studies showing that in many cortical areas small regions are relatively homogeneous with respect to their physiological function. It is questionable, how-

ever, whether this homogeneity holds through all cortical layers (e.g., Miller (1996a)). Therefore, as in most theories of cell assemblies, we will use the first of the described possibilities, i.e., we take excitatory cells as basic units. In the discussion section we draw a connection to the second choice of regarding groups of cells as units.

3.2 The selection of the neural model

Once the basic computational units have been defined, their physiological properties will further constrain the neural model to be used. If the units represent single cells, the formal neurons of the first-generation models of cell assemblies can describe spike rates, but they fail to cover influences arising from the timing of single spikes. However, there is cumulating evidence that these influences are important. The simplest single-neuron model comprising a threshold mechanism for spike generation and after-hyperpolarization, the integrate and fire model (Dayan and Abbott, 2002), can account for some spike timing effects, but still neglects most properties of synaptic and membrane processes with their various nonlinearities and activation time constants. Exploration of the impact of different ionic currents on the synaptic and single-neuron dynamics, and network properties is ongoing, e.g. (Traub et al., 1991; Lisman et al., 1998). In order to account for the most basic effects we employ a two-compartment model for excitatory cells proposed by Pinsky and Rinzel (1994) as a simplification of a more elaborate multi-compartment model studied by Traub et al. (1991). This neuron model as depicted in Fig. 1 consists of a dendritic compartment comprising an after-hyperpolarizing potassium current, a calcium-controlled potassium current, and a slow calcium current. The dendritic compartment is capable of generating slow dendritic spikes and further integrates AMPA- and NMDA-mediated synaptic inputs from other neurons. It is coupled galvanically to a somatic compartment capable of generating fast sodium spikes. This type of neural model retains the basic physiological properties observed in hippocampal/cortical slice experiments (Pinsky and Rinzel, 1994), but is still simple enough to simulate networks of many hundreds of cells. Parameters in our experiments are set as in Pinsky and Rinzel (1994). Differing from the original model we take into account the influence of local inhibitory cells by a sigmoid interneuron that provides negative feedback proportional to the mean excitatory network activity, cf. Fig. 1(left).

3.3 Constraints on network architecture

The choice of basic computational units also prescribes how the neuroanatomical interconnectivity should project on the neural network structure. With single cells as basic units, the true distribution of synaptic contacts in cortex has to be reflected in the connectivity of the network model. For larger modules, e.g., minicolumns, the set of connections between a pair of modules defines an “effective connection” in the network model. Thus, even if the true connectivity between single cells is sparse, there might be enough connections between each pair of modules for direct interaction (Malach, 1994) such that the network of modules might be considered as densely or even fully connected. This fact provides an argument why fully connected associative memories can be biologically plausible even if cell-to-cell connectivity is incomplete (Shaw et al., 1985; Fransén and Lansner, 1998).

The anatomy of connections also influences the choice of the neural network circuit most appropriate as a functional unit. If the density of connections between basic neural units is not homogenous, e.g., patchy, the function of the whole tissue might be better described by the functions of smaller regions of high interconnectivity that operate individually, though not independently, as associative networks. Therefore, the brain structures considered briefly in the next section are local networks of excitatory cells with high connectivity.

3.4 Models of local cell assemblies

Quantitative neuroanatomy revealed that connectivity drops with the distance between pyramidal cells and is roughly 10% within distances of 1 mm (Braitenberg and Schüz, 1991). These data suggest that local circuits of the size of a cortical macrocolumn (about 1 mm² of cortex) are good candidates for associative memories. Therefore, most models of cell assemblies assume spatially localized associative memories as depicted in Fig. 1. Comparable computational models for local cell assemblies have been proposed for various cortical regions including area CA3 of the hippocampus (Marr, 1971; Treves and Rolls, 1992; Hasselmo, 1993; Rolls et al., 1997), dentate gyrus (Levy et al., 1983), and inferior temporal cortex (Fuster, 1995; Amit, 1995).

Our previous studies show that sparse memories in local associative nets can exploit the synaptic memory fast and efficiently in terms of information capacity (Schwenker et al., 1996) and statistical reasoning (Sommer and

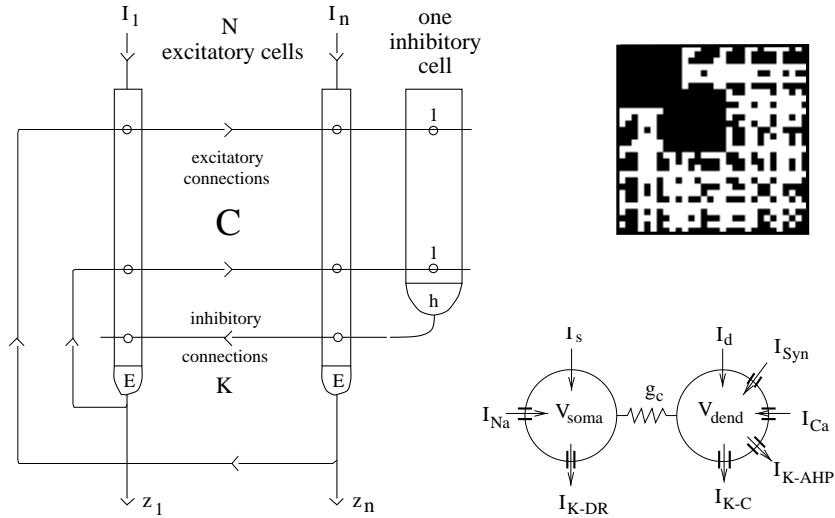


Figure 1: Associative network. The vertical structures $1, \dots, n$ represent excitatory cells. Axon collaterals project back in the same set of cells. The effect of local inhibitory neurons is modeled by a negative feed back loop for the mean network activity. An auto-associative connectivity matrix is depicted on the upper right. The lower right displays schematically a two-compartment Pinsky-Rinzel neuron with corresponding membrane currents.

Dayan, 1998). Based on the numbers given by Braitenberg and Schüz (1991) and Wickens and Miller (1997) we analyzed a local network to estimate the number and size of cell assemblies processable in a macrocolumn: It was found that the local synaptic memory is used most efficiently if local assemblies contain several hundreds of cells each. Then, tens of thousands of assemblies could in principle be stored in a macrocolumn (Sommer, 2000).

We further analyzed temporal properties of recall in a simulation network as the one described in section 3.2 (Sommer and Wennekers, 2000, 2001a). These simulations demonstrate that biologically realistic networks can indeed store patterns at a capacity close to the theoretical limit. The network exhibits gamma oscillations and pattern completion is very fast: Synchronized spike events in an assembly within individual gamma-cycles turn out to provide the elementary local association or retrieval processes contrasting the view of feature-coding by periodic firing (Wennekers and Palm, 1997). However, for a satisfying theory of memory one has to study nonlocal associative memories, too.

4 Reciprocal associative projections

The main emphasis of this paper is to show how associative memory could be expressed in the rich network of long-range projections in the brain which exist in addition to local connectivity patterns. If one important aspect of memory during a cognitive task is to bind groups of neurons distributed throughout brain areas of different functional specialisation (Mesulam, 1998), then associative processes should also be expressed nonlocally, i.e., in topographic long-range projections (Braitenberg's B-system, cf. Palm (1982)).

Most of the known cortico-cortical projections have been shown to be reciprocal (Felleman and Van Essen, 1991). Furthermore, they often start and terminate in patches of the size of cortical columns (Malach, 1994). Recently it has been shown that in reciprocal projections, patches of origins and terminations overlap (Lund et al., 2003). These findings suggest that the most basic circuit extending local associative memories would consist of a pair of cortical patches with bidirectional connections. In the present section we summarize results from a computational study of a nonlocal associative net formed by such a bidirectional topographic projection. Some earlier studies suggested versions of such a model although in more abstract settings (Sommer et al., 1998; Renart et al., 1999; Sommer and Palm, 1999).

4.1 Simulation model

We performed simulation experiments with a bidirectional associative network architecture comprising two cell populations, A and B, each with 50 Pinsky-Rinzel-type neurons, see also Sommer and Wennekers (2001a). There were no (local) connections between excitatory cells within A or within B, but activity in both pools was controlled individually by inhibitory cells as in section 3. Ten pairs of patterns were stored by Hebbian learning as in the Willshaw model. Each pattern had 10 ones at random positions. To test the retrieval properties after learning we stimulated cells in population A for a brief period of 25 ms. After that, the external input was turned off. To assess robustness of the retrieval against spurious input we used subsets of memory patterns for stimulation as well as patterns containing more neurons than the original memory pattern (i.e. additional random bits). This assessment is important because in the cortex additional (spurious) activity is likely, but in most associative memory models the retrieval is considerably more impaired by additional active units than by missing ones. To assess re-

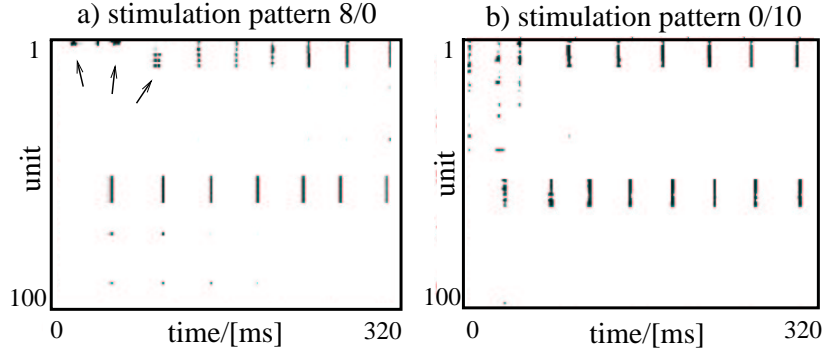


Figure 2: Raster plots of soma potentials during retrieval trials with different stimulation pattern size. Cell 1 to 50 form pool A, cell 51 to 100 pool B. In stimulation pattern x/y , x is the number of cells in the memory pattern that received no stimulation and y is the number of stimulated cells outside the memory pattern. Arrows in the left figure indicate spike bursts.

trieval performance we compared the activity state in the non-addressed cell group (denoted as B) with the respective stored configuration, either during the first synchronized response (one-step mode, see below) or by determining the best (most similar) response during the simulation period (high-fidelity mode). Neuronal and network parameters were identical in both groups with the exception of the external input, which was only supplied to group A.

4.2 Simulation experiments

We conducted experiments using stimulation patterns of different sizes and recorded the network activity over a 500 ms interval after stimulus onset. Typical time courses of soma potentials for different stimulation patterns are displayed in Fig. 2. About 25–40ms after stimulus onset the first synchronized wave of induced activity arrived in cell group B. Subsequently, the activity propagated back and forth through the reciprocal connections, thereby influencing the synchronized activity configurations in the cell groups iteratively. The spike frequencies were in the gamma band (30–90 Hz) just as in the single pool in section 3, cf. also (Sommer and Wennekers, 2000, 2001a). We observed spike synchronization within each cell group but no phase locking between the groups. Neurons that did not receive external input showed regular spiking. In contrast, neurons with external input showed a tendency to burst. This tendency increased as the number of stimulated

cells decreased: In Fig. 2 the small stimulation population in a) produces bursting in the first gamma cycles, while larger stimulation populations as in b) do not. This provides a mechanism of activity balance in the network: In a wide range of stimulation pattern size (between two and twenty cells) recall quality was high without any further (e.g., manual) adjustment of network parameters. Thus, bursting enhances the input fault tolerance of association processes by balancing the network activity for different input pattern size.

4.3 Information capacities and retrieval latencies

To judge the performance of the fastest possible response we determined the retrieval quality in the first wave of activity. We call this *one-step retrieval* since only monosynaptic activity propagation from cell group a to B is involved. To monitor the results provided by the iterative activity flow between the cell groups (the bidirectional retrieval mode) we also detected the quality maximum over the whole recording sequence. Figure 3 shows the results.

Diagrams a) and c) in Fig. 3 display the measured quality values. A quality equal to 1 corresponds to perfectly retrieved patterns. Diagram a) compares the one-step retrieval quality with the performance of the Willshaw model at the same memory load (computed for constant activity threshold setting using the theory from Sommer and Palm (1999)). The fact that the quality curves stay somewhat below 1 indicates that the load in both networks is near maximum, i.e., cross-talk is already setting in, even in associations with perfect stimulus pattern. Due to the additional noise in the biological model its maximum capacity falls somewhat below the values for the Willshaw net. Also the fault-tolerant association capability is impaired, only slightly for lower, but substantially for higher input activities. The question underlying the measurements displayed in diagram b) is whether bidirectional activation cycles in the network can improve retrieval quality. In fact, for very small stimulation sets (2 to 3 cells) and for large input activities (12 to 20 cells) bijective (i.e., iterative) retrieval improved the retrieval quality in between 70 and 80 % of the simulation runs. Figures 3c and 3d compare the qualities and retrieval time of one-step and iterative retrieval in the Pinsky-Rinzel neuron network. The iterative retrieval performs better in all cases. Most importantly, the fault tolerance with respect to higher input activity can be significantly improved by iteration. Of course, the required retrieval time increases from 25–60 ms to 60–260 ms depending on the input condition.

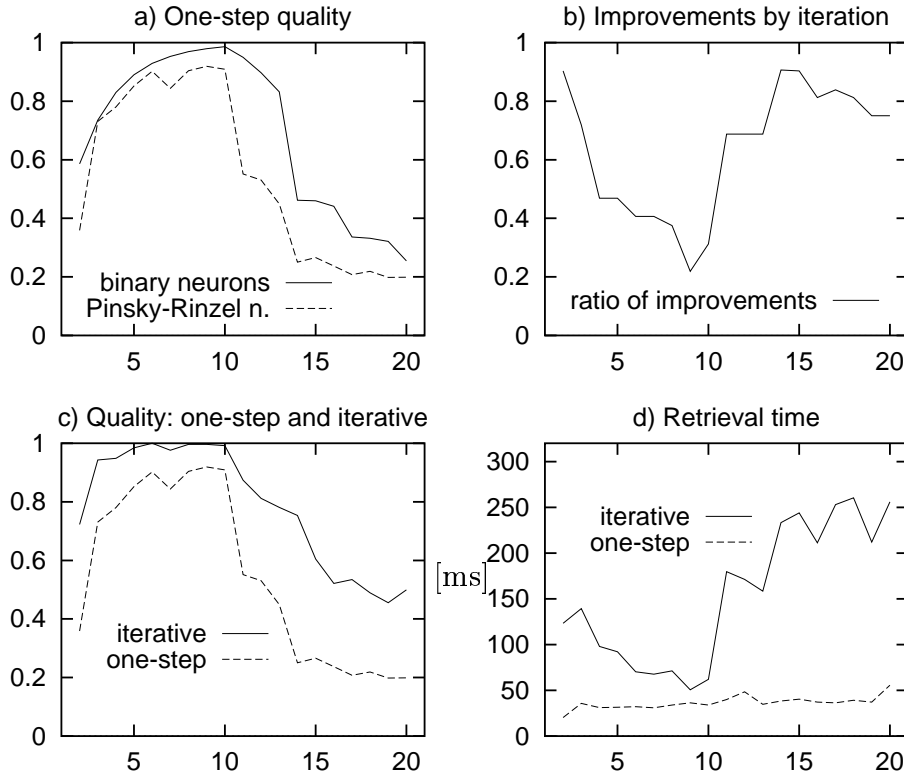


Figure 3: The x-axes label the number of cells that receive addressing input. At an input activity of ten the stimulation pattern was identical to one of the stored patterns. At lower or higher activities it deviated by miss or add errors, respectively. The displayed quantities are averaged for every input condition over all memory patterns and 32 runs with different sets of memory patterns.

5 Discussion

5.1 Associative memory in a reciprocal projection

In the previous section we described how associative memory could be expressed in possibly far-reaching reciprocal projections. We studied simulations of a bidirectional associative network storing pattern pairs in the connections between two pools of conductance-based neurons. We investigated how brief afferent stimulation (25ms) of cells in pool A could be used to complete a stored pattern pair. The simulation demonstrates that the memory load can be increased up to a point where associative memory is efficient in

terms of information capacity. This is in agreement with statistical mechanics analyses by Renart et al. (1999) of a biologically constrained bidirectional memory employing rate coding neurons. Our model with conductance-based neurons makes some specific predictions about the possible recall modes and the latencies associated with them. At higher memory loads two different modes of retrieval are discernable: One can either use the spatial pattern of the first activity wave in population B that occurred after 25–40 ms, or employ the pattern settling after 60–260 ms of bidirectional iteration (up to ten gamma cycles). We performed retrieval tests in a wide range of situations, from cases where only very few cells were stimulated up to the case where a lot of neurons received current injections, a large fraction of them not belonging to the recalled pattern. These extreme cases are biologically relevant since they might correspond to ambiguous situations, where not exclusively one of the stored associations is addressed. Both retrieval procedures performed differently, in particular in the ambiguous situations: It was interesting to find the slower retrieval mode clearly to outperform the first response retrieval. The quality of first response retrieval could be improved in 70 to 80 % of the performed retrieval trials. Because of this result we characterize the two retrieval modes as the fast mode and the high-fidelity mode. We would expect expression of the high-fidelity mode particularly for ambiguous stimulation. It should be accompanied by oscillatory activity and would require new input to be attenuated in order to listen to the reverberating “echoes”.

Another interesting observation concerns the functional role of cell bursting. It was observed in directly stimulated cells when the stimulation set was small. Bursting in that regime actually turned out to constructively support memory retrieval: If the number of stimulated cells is low, they first produce bursts instead of regular spikes. Thereby, the relatively few cells can provide enough excitation in the network to start the recall process.

All in all, our simulation results suggest that efficient associative memory should be possible in reciprocal projections. The memory function we tested just relied on heteroassociative storage in the synaptic structure connecting the reciprocal pathways. However, the bidirectional associative network without local connections we have studied should not be considered in isolation. It is likely that local and nonlocal expressions of associative memory are coexisting and interacting. In fact, we also run simulations with added local connectivity (data not shown) that resulted in an improvement of the retrieval in agreement with the analytic study of Renart et al. (1999). The interaction between local and nonlocal associative memory systems could also

lead to ways of hierarchical processing in the cortex as we will explain next.

5.2 Nonlocal models of cortical memory

The importance of an associative coordination of diverse and multimodal information into concepts (binding) was central in the original theories of Hebb and Hayek. But the idea goes even farther back in history: Already Descartes considered this as pivotal to human cognition when he proposed the convergence of multisensory information in the pineal gland “where the immaterial mind would observe the representation of experience provided by the material brain”. Neuroanatomy and electrophysiology identified a number of regions of multimodal convergence, such as posterior parietal, lateral prefrontal and temporal cortices, and archicortex. Typically, multi-modality is found in areas that are several cortical stages apart from direct sensory inputs, but never in primary sensory areas. In technical terms one would speak of late or perhaps “semantic” sensor fusion as opposed to early sensory feature integration. Mesulam (1998) gave a reason why a role of memory is particularly important in multimodal areas: A permanent convergence of different sensory streams would just be threatened to run into severe crosstalk or capacity problems. Mesulam therefore proposed that the role of crossmodal areas is to create directories for binding distributed modality-specific fragments into coherent memories, see also Bibbig et al. (1995). The function of these areas he describes as content-addressable or associative memories.

A cell assembly theory suited for describing multimodal processing, has to convey how nonlocal cell assemblies are formed based on anatomical structures of the brain. Eichenbaum (1993) pointed out a way how nonlocal cell assemblies could emerge from local representations. Wickelgren (1992) was the first to propose an extended cell assembly theory that, at least qualitatively, reflected the neuroanatomy of the cortex. He proposed what he called neural webs as neural representations of concepts. Neural webs are cell assemblies shaped like spider webs. According to the specific properties of the concept, neurons are recruited in different areas of cortical and subcortical structures that are connected by long-range projections. As the simulation results described in the present work suggest, a single reciprocal connection could be used to form the simplest type or element of a neural web. Participating local neuron sets could further form local associatively formed assemblies, and they could be recruited into different global assemblies. Miller (1996b) put forward a hypotheses about the laminar location

of cortical pyramidal cells. He proposed that cortical laminae II, III and IV contain “library cells”, i.e., cells belonging to cell assemblies. Furthermore he argued that thalamo-cortical projection neurons could have a pivotal role in the save activation of one or more cortical assemblies (Miller, 1996a).

Building on the results presented in this paper we can add more to the picture given so far. Since local assembly completion can be very fast (Sommer and Wennekers, 2001b), a model similar to the one presented here, but with the single neurons replaced by local assemblies, would still exhibit a similar performance as in the simulations shown above. Using the same argument as for the increased connectivity between minicolumns, such a replacement in the model extends the type of cortical circuitry where it could apply: Cortico-cortical pathways, even with sparse connection densities on a cell-to-cell level, could act as associative networks. The formation of cell assemblies in networks of structured connectivity could reflect the learning of hierarchical neuronal representations: Initially, assemblies would form locally representing unimodal objects. Local assemblies might then become basic neural units for a larger scale associative system integrating information of various subregions (and modalities). By such a learning process, sparser and sparser connectivity could be used for the storage of associations. Recall can operate in different subregions in parallel and may lead to associations at a high level representing complex integrated entities.

5.3 Conclusions

We argued that the concepts of cell assemblies and associative memory have been underestimated in theoretical neuroscience over the last ten years because, as we believe, the first-generation computational models were oriented on analysis techniques rather than cortical implementations and brain theory. As a result, associative memory and cell assemblies played only a marginal role in theoretical neuroscience, mainly as a model for local cortical computation. However, as the early promoters of cell assemblies (Hebb, 1949; Hayek, 1952; Braitenberg, 1978; Edelman, 1982; Palm, 1982; Damasio, 1989; Miller, 1991; Wickelgren, 1992; Eichenbaum, 1993; Bienenstock, 1994) pointed out, it is the potential of cell assemblies for multimodal binding that in especially makes them an interesting concept (Mesulam, 1998). To address this issue we have described how a computational model can be built from neuroscientific data and abstract neural networks. We followed this path to derive a computational model of associative memory in a reciprocal

long-range projection. Simulation studies show that associative memory in this model is efficient and supports mainly two different modes for recall, a fast “one-shot” mode and a high-fidelity mode relying on (a few) cycles of bidirectional feedback. Finally, we have outlined how associative memory in long-range projections could be the basis of a more general theory of distributed cell assemblies and, ultimately, a part of a computational theory of memory and cognition.

Acknowledgement: F.T.S. thanks Pentti Kanerva for detailed comments on the manuscript and all members of the Redwood Neuroscience Institute for stimulating discussions.

References

- Abeles, M. (1991) *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, Cambridge, UK.
- Aertsen, A. M. H. J., Gerstein, G. L., Habib, M. K., Palm, G., and Adrian, E. D. (1989) Dynamics of neuronal firing correlation: Modulation of “effective connectivity”. *J. Neurophysiol.* 61:900–917.
- Amit, D., Gutfreund, H., and Sompolinsky, H. (1985) Storing infinite numbers of patterns in a spin-glass model of neural networks. *Phys. Rev. Lett.* 55:1530–1533.
- Amit, D. J. (1995) The Hebbian paradigm reintegrated: Local reverberations as internal representations. *Behav. Brain Sci.* 18:617 – 657.
- Arieli, A., Sterkin, A., Grinvald, A., and Aertsen, A. (1996) Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses. *Science* 273:1868–1871.
- Bibbig, A., Wenekers, T., and Palm, G. (1995) A neural network model of the cortico-hippocampal interplay and the representation of contexts. *Behav. Brain Res.* 66:169–175.
- Bienenstock, E. (1994) A model of neocortex. *Network* 6:179–224.
- Braitenberg, V. (1978) Cell assemblies in the cerebral cortex. In Heim, R. and Palm, G., editors, *Theoretical approaches to complex systems*, pages 171–188. Springer, Berlin.
- Braitenberg, V. and Schüz, A. (1991) *The anatomy of the cortex. Statistics and Geometry*. Springer, Berlin.
- Damasio, A. R. (1989) The brain binds entities and events by multiregional activation from convergence zones. *Neural Comput.* 1:123–132.
- Dayan, P. and Abbott, L. F. (2002) *Theoretical Neuroscience*. MIT Press, Boston, MA.
- Edelman, G. M. (1982) Selective networks capable of representative transformations, limited generalizations, and associative memory. *Proc. Natl. Acad. Sci.* 79:2091–2095.
- Eichenbaum, H. (1993) Thinking about brain cell assemblies. *Science* 261:993–994.
- Felleman, D. J. and Van Essen, D. C. (1991) Distributed hierarchical processing in the primate cerebral cortex. *Cereb. Cortex* 1:1 – 47.

- Fransén, E. and Lansner, A. (1998) A model of cortical associative memory based on a horizontal network of connected columns. *Network* 9(2):235–264.
- Fuster, J. (1995) *Memory in the cerebral cortex*. MIT Press, Boston, MA.
- Gardner, E. (1988) The space of interactions in neural network models. *J. Phys. A* 21:257–270.
- Gardner-Medwin, A. (1976) The recall of events through the learning of associations between their parts. *Proc. Roy. Soc. Lond. B* 194:375–402.
- Grün, S., Diesmann, M., and Aertsen, A. (2002) Unitary events in multiple single-neuron spiking activity: I. detection and significance. *Neural Comput.* 14:81–119.
- Hasselmo, M. (1993) Acetylcholine and learning in a cortical associative memory. *Neural Comput.* 5:32–44.
- Hayek, F. (1952) *The sensory order*. Chicago University Press, Chicago, IL.
- Hebb, D. O. (1949) *The Organization of Behaviour*. Wiley, New York.
- Hopfield, J. J. (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci.* 79:2554–2558.
- James, W. (1890) *Principles of Psychology*. Holt, New York.
- Kohonen, T. (1972) Correlation matrix memories. *IEEE Trans. Computers* C-21:353–359.
- Levy, W. B., Brassel, S. E., and Moore, S. D. (1983) Partial quantification of the associative synaptic learning rule of the dentate gyrus. *Neurosci.* 81:799–808.
- Lisman, J. E., Fellous, J. M. and Wang, X. J. (1998) A role for NMDA-receptor channels in working memory. *Nat. Neurosci.* 1 (4): 273-275.
- Little, W. A. (1974) The existence of persistent states in the brain. *Math. Biosci.* 19:101–120.
- Lorente de No, R. (1938) The cerebral cortex: Architecture, intracortical connections and motor projections. In Fulton, J. F., editor, *Physiology of the Nervous System*, pages 291–339. Oxford University Press, London.
- Lund, J. S., Angelucci, A., and Bressloff, P. C. (2003) Anatomical substrates for functional columns in macaque monkey primary visual cortex. *Cereb. Cortex* 13:15–24.
- Malach, R. (1994) Cortical columns as devices for maximizing neuronal diversity. *Trends Neurosci.* 3:101–104.
- Marr, D. (1971) Simple memory: A theory for archicortex. *Phil.Trans.Roy. Soc. London B* 262:23 – 81.
- Mesulam, M. M. (1998) From sensation to cognition. *Brain* 121:1013–1052.
- Miller, R. (1991) Designs for a prototype cerebral cortex. In Braitenberg, V., Barlow, H. B., Bullock, T. H., Florey, E., and Grüsser, O.-J., editors, *Cortico-hippocampal interplay and the representation of context in the brain*, pages 11–32. Springer, Berlin.
- Miller, R. (1996a) Cortico-thalamic interplay and the security of operation of neural assemblies and temporal chains in the cerebral cortex. *Biol. Cybern.* 73:258–270.
- Miller, R. (1996b) Neural assemblies and laminar interactions in the cerebral cortex. *Biol. Cybern.* 75:253–261.
- Mountcastle, V. B. (1978) An organizing principle for cerebral function: the unit module and distributed function. In Edelman, G. and Mountcastle, V. B., editors, *The Mindful Brain*, pages 56–100. MIT Press, Cambridge.
- Palm, G. (1982) *Neural Assemblies*. Springer, Berlin.
- Pinsky, P. F. and Rinzel, J. (1994) Intrinsic and network rhythmogenesis in a reduced

- Traub model for CA3 neurons. *J. Comput. Neurosci.* 1:39–60.
- Pulvermüller, F. (1992) Constituents of a neurological theory of language. *Concepts in Neuroscience* 3:157–200.
- Renart, A., Parga, N., Rolls, E. T. (1999) Backward projections in the cerebral cortex: Implications for memory storage. *Neural Comput.* 11(6):1349–1388.
- Riehle, A., Grün, S., Diesmann, M., and Aertsen, A. (1997) Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278:1950 – 1953.
- Rolls, E. T., Treves, A., Foster, D., and Perez-Vicente, C. (1997) Simulation studies of the CA3 hippocampal subfield modelled as an attractor neural network. *Neural Networks* 10(9):1559–1569.
- Schwenker, F., Sommer, F. T., and Palm, G. (1996) Iterative retrieval of sparsely coded associative memory patterns. *Neural Networks* 9(3):445–455.
- Shaw, G. L., Silverman, D. J., and Pearson, J. C. (1985) Model of cortical organization embodying a basis for the theory of information processing and memory recall. *Proc. Natl. Acad. Sci.* 82:2364–2368.
- Singer, W. and Gray, C. M. (1995) Visual feature integration and the temporal correlation hypotheses. *Ann. Rev. Neurosci.* 18:555 – 568.
- Sommer, F. T. (2000) On cell assemblies in a cortical column. *Neurocomputing* 32-33:517–522.
- Sommer, F. T. and Dayan, P. (1998) Bayesian retrieval in associative memories with storage errors. *IEEE Trans. Neural Networks* 9(4):705–713.
- Sommer, F. T. and Palm, G. (1999) Improved bidirectional retrieval of sparse patterns stored by Hebbian learning. *Neural Networks* 12:281–297.
- Sommer, F. T. and Wennekers, T. (2000) Modeling studies on the computational function of fast temporal structure in cortical circuit activity. *J. Physiol. Paris* 94(5–6):473–488.
- Sommer, F. T., Wennekers, T. (2001a) Associative memory in a pair of cortical cell groups with reciprocal connections. *Neurocomputing* 38:1575–1580.
- Sommer, F. T. and Wennekers, T. (2001b) Associative memory in networks of spiking neurons. *Neural Networks* 14(6-7):825–834.
- Sommer, F. T., Wennekers, T., and Palm, G. (1998) Bidirectional completion of Cell Assemblies in the cortex. In: *Computational Neuroscience: Trends in Research*. Plenum Press.
- Steinbuch, K. (1961) Die Lernmatrix. *Kybernetik (Biol. Cybern.)* 1:36–45.
- Taylor, W. K. (1956) Electrical simulation of some nervous system functional activities. In *Information Theory*, Vol. 3, pages 314–328. Butterworths.
- Traub, R. D., Wong, R. K. S., Miles, R., and Michelson, H. (1991) A model of a CA3 hippocampal pyramidal neuron incorporating voltage-clamp data on intrinsic conductances. *J. Neurophysiol.* 66(2):635–650.
- Treves, A. and Rolls, E. T. (1992) Computational analysis of the operation of a real neuronal network in the brain: the role of the hippocampus in memory. volume 2, pages 891–898. Elsevier Science Publishers B.V.
- Wennekers, T. and Palm, G. (1997) On the relation between neural modeling and experimental neuroscience. *Theory Biosci.* 273 - 289:116.
- Wickelgren, W. A. (1992) Webs, cell assemblies, and chunking in neural nets. *Concepts*

in *Neuroscience* 3:1–53.

Wickens, J. R. and Miller, R. (1997) A formalism of the neural assembly concept 1. constraints on neural assembly size. *Biol. Cybern.* 77:351–358.

Willshaw, D. J., Buneman, O. P., and Longuet-Higgins, H. C. (1969) Non-holographic associative memory. *Nature* 222:960–962.