Pattern Segmentation in Associative Memory

DeLiang Wang
Joachim Buhmann
Computer Science Department,
University of Southern California, University Park,
Los Angeles, CA 90089-0782 USA

Christoph von der Malsburg
Program in Neural, Informational, and Behavioral Sciences,
University of Southern California, University Park,
Los Angeles, CA 90089-0782 USA

The goal of this paper is to show how to modify associative memory such that it can discriminate several stored patterns in a composite input and represent them simultaneously. Segmention of patterns takes place in the temporal domain, components of one pattern becoming temporally correlated with each other and anticorrelated with the components of all other patterns. Correlations are created naturally by the usual associative connections. In our simulations, temporal patterns take the form of oscillatory bursts of activity. Model oscillators consist of pairs of local cell populations connected appropriately. Transition of activity from one pattern to another is induced by delayed self-inhibition or simply by noise.

1 Introduction ⸻⸻⸻⸻⸻⸻

Associative memory (Steinbuch 1961; Willshaw *et al.* 1969; Hopfield 1982) is an attractive model for long-term as well as short-term memory. Especially the Hopfield formulation (Hopfield 1982) provides for both levels a clear definition of data structure and mechanism of organization. The data structure of long-term memory has the form of synaptic weights for the connections between neurons, and memory traces are laid down with the help of Hebbian plasticity. On the short-term memory level the data structure has the form of stationary patterns of neural activity, and these patterns are organized and stabilized by the exchange of excitation and inhibition. Since in this formulation short-term memory states are dynamic attractor states, one speaks of attractor neural networks. Neurons are interpreted as elementary symbols, and attractor states acquire their symbolic meaning as an unstructured sum of individual symbolic contributions of active neurons. The great virtue of associative memory

is its ability to restore incomplete or corrupted input patterns, that is, its ability to generalize over Hamming distance (the number of bits missing or added). Let us just mention here, since it becomes relevant later, that associative memory can be formulated such that attractors correspond to oscillatory activity vectors instead of stationary ones (Li and Hopfield 1989; Baird 1986; Freeman *et al.* 1988).

Associative memory, taken as a model for functions of the brain, is severely limited in its applicability by a particular weakness — its low power of generalization. This is a direct consequence of the fact that associative memory treats memory traces essentially as monolithic entities. An obvious and indispensable tool for generalization in any system must be the decomposition of complex patterns into functional components and their later use in new combinations. A visual scene is almost always composed of a number of subpatterns, corresponding to coherent objects that are very likely to reappear in different combinations in other scenes (or the same scene under a different perspective and thus in different spatial relations to each other). Associative memory is not equipped for this type of generalization, as has been pointed out before (von der Malsburg 1981, 1983, 1987). It treats any complex pattern as a synthetic whole, glues all pairs of features together, and recovers either the whole pattern or nothing of it. Two different arrangements of the same components cannot be recognized as related and have to be stored separately. There is no generalization from one scene to another, even if they are composed of the same objects. Since complex scenes never recur, a nervous system based on the associative memory mechanism alone possesses little ability to learn from experience.

This situation is not specific to vision. Our auditory world is typified by complex sound fields that are composed of sound streams corresponding to independent sources. Take as an example the cocktail party phenomenon where we are exposed to several voices of people who talk at the same time. It would be useless to try to store and retrieve the combinations of sounds heard simultaneously from different speakers. Instead, it is necessary to separate the sound streams from each other and store and access them separately. Similar situations characterize other modalities and especially all higher levels of cognitive processing. The basis for the type of generalization discussed here is the specific and all-pervasive property of our world of being causally segmented into strongly cohesive chunks of structure that are associated with each other into more loose and varying combinations.

There are two attitudes which an advocate of associative memory could take in response to this evident weakness. One is to see it as a component in a more complex system. The system has other mechanisms and subsystems to analyze and create complex scenes composed of rigid subpatterns that can individually be stored and retrieved in associative memory. The other attitude tries to build on the strengths of associative memory as a candidate cognitive architecture and tries to modify the

model such as to incorporate the ability to segment complex input patterns into subobjects and to compose synthetic scenes from stored objects. We subscribe to this second attitude in this paper.

There are three issues that we have to address. The first concerns the type of information on the basis of which pattern segmentation can be performed; second, the data structure of associative memory and attractor neural networks has to be modified by the introduction of variables that express syntactical[1] binding; and third, mechanisms have to be found to organize these variables into useful patterns.

There are various potential sources of information relevant to segmentation. In highly structured sensory spaces, especially vision and audition, there are general laws of perceptual grouping, based on "common fate" (same pattern of movement, same temporal history), continuity of perceptual quality (texture, depth, harmonic structure), spatial contiguity, and the like. These laws of grouping have been particularly developed in the Gestalt tradition. On the other end of a spectrum, segmentation of complex patterns can be performed by just finding subpatterns that have previously been stored in memory. Our paper here will be based on this memory-dominated type of segmentation.

Regarding an appropriate data structure to encode syntactical binding, the old proposal of introducing more neurons (e.g., a grandmother-cell to express the binding of all features that make up a complex pattern) is not a solution (von der Malsburg 1987) and produces many problems of its own. It certainly is useful to have cells that encode high-level objects, but the existence of these cells just creates more binding problems, and their development is difficult and time-consuming. We work here on the assumption (von der Malsburg 1981, 1983, 1987; von der Malsburg and Schneider 1986; Gray *et al.* 1989; Eckhorn *et al.* 1988; Damasio 1989; Strong and Whitehead 1989; Schneider 1986) that syntactical binding is expressed by temporal correlations between neural signals. The scheme requires temporally structured neural signals. A set of neurons is syntactically linked by correlating their signals in time. Two neurons whose signals are not correlated or are even anticorrelated express thereby the fact that they are not syntactically bound. There are first experimental observations to support this idea (Gray *et al.* 1989; Eckhorn *et al.* 1988). It may be worth noting that in general the temporal correlations relevant here are spontaneously created within the network and correspondingly are not stimulus-locked.

As to the issue how to organize the correlations necessary to express syntactical relationships, the natural mechanism for creating correlations and anticorrelations in attractor neural networks is the exchange of excitation and inhibition. A pair of neurons that is likely to be part of one segment is coupled with an excitatory link. Two neurons that do

---

[1]We use the word *syntactical structure* in its original sense of *arranging together*, that is, grouping or binding together, and do not intend to refer to any specific grammatical or logical rule system.

not belong to the same segment inhibit each other. The neural dynamics will produce activity patterns that minimize contradictions between conflicting constraints. This capability of sensory segmentation has been demonstrated by a network that expresses general grouping information (von der Malsburg and Schneider 1986; Schneider 1986).

The system we are proposing here is based on associative memory, and performs segmentation exclusively with the help of the memory-dominated mechanism. Our version of associative memory is formulated in a way to support attractor limit cycles (Li and Hopfield 1989; Baird 1986; Freeman *et al.* 1988): If a stationary pattern is presented in the input that resembles one of the stored patterns, then the network settles after some transients into an oscillatory mode. Those neurons that have to be active in the pattern oscillate in phase with each other, whereas all other neurons are silent. In this mode of operation the network has all the traditional capabilities of associative memory, especially pattern completion. When a composite input is presented that consists of the superposition of a few patterns the network settles into an oscillatory mode such that time is divided into periods in which just a single stored state is active. Each period corresponds to one of the patterns contained in the input. Thus, the activity of the network expresses the separate recognition of the individual components of the input and represents those patterns in a way avoiding confusion. This latter capability was not present in previous formulations of associative memory. The necessary couplings between neurons to induce correlations and anticorrelations are precisely those created by Hebbian plasticity.

Several types of temporal structure are conceivable as basis for this mode of syntactical binding. At one end of a spectrum there are regular oscillations, in which case states would be distinguished by different phase or frequency. At the other end of the spectrum there are chaotic activity patterns (Buhmann 1989). The type of activity we have chosen to simulate here is intermediate between those extremes, being composed of intermittent bursts of oscillations (see Fig. 2), a common phenomenon in the nervous system at all levels.

2 Two Coupled Oscillators

A single oscillator i, the building block of the proposed associative memory, is modeled as a feedback loop between a group of excitatory neurons and a group of inhibitory neurons. The average activity x_i of excitatory group i and the activity y_i of inhibitory group i evolve according to

$$\frac{dx_i}{dt} = -\frac{x_i}{\tau_x} + \mathcal{G}_x \left[T_{xx}\frac{x_i}{\bar{x}} - T_{xy}\mathcal{F}\left(\frac{y_i}{\bar{y}}\right) + S_i + I_i - H_i \right] \qquad (2.1)$$

$$\frac{dy_i}{dt} = -\frac{y_i}{\tau_y} + \mathcal{G}_y \left(-T_{yy}\frac{y_i}{\bar{y}} + T_{yx}\frac{x_i}{\bar{x}} \right) \qquad (2.2)$$

$$H_i = \alpha \int_0^t x_i(\tau) \exp\left[-\beta(t - \tau)\right] d\tau \tag{2.3}$$

where τ_x and τ_y are the time constants of the excitatory and inhibitory components of the oscillator. An appropriate choice of τ_x, τ_y allows us to relate the oscillator time to a physiological time scale. $\mathcal{G}_x$ and $\mathcal{G}_y$ are sigmoid gain functions, which in our simulations have the form

$$\mathcal{G}_r(v) = \frac{1}{1 + \exp[-(v - \theta_r)/\lambda_r]}, \qquad r \in \{x, y\} \tag{2.4}$$

with thresholds θ_x or θ_y and gain parameters $1/\lambda_x$ and $1/\lambda_y$. For the reaction of inhibitory groups on excitatory groups we have introduced the nonlinear function $\mathcal{F}(x) = (1 - \eta)x + \eta x^2$, $(0 \le \eta \le 1)$, where η parameterizes the degree of quadratic nonlinearity. This nonlinearity proved to be useful in making oscillatory behavior a more robust phenomenon in the network, so that in spite of changes in excitatory gain (with varying numbers of groups in a pattern) the qualitative character of the phase portrait of the oscillators is invariant. H_i in equation 2.3 describes delayed self-inhibition of strength α and decay constant β. This is important to generate intermittent bursting. The synaptic strengths of the oscillators' feedback loop are T_{rs}, $r, s \in \{x, y\}$. Equations 2.1 and 2.2 can be interpreted as a mean field approximation to a network of excitatory and inhibitory logical neurons (Buhmann 1989). Notice that x_i, y_i are restricted to $[0, \tau_x]$ and $[0, \tau_y]$, respectively. The parameters $\bar{x}$ and $\bar{y}$ may be used to control the average values of x and y. In addition to the interaction between x_i and y_i, an excitatory unit x_i receives time-dependent external input $I_i(t)$ from a sensory area or from other networks, and internal input $S_i(t)$ from other oscillators.

Let us examine two oscillators of type 2.1–2.3, coupled by associative connections W_{12}, W_{21} as shown schematically in Figure 1. The associative interaction is given by

$$S_1(t) = W_{12}x_2(t); \qquad S_2(t) = W_{21}x_1(t)$$

Two cases can be distinguished by the sign of the associative synapses. If both synapses are excitatory ($W_{12} > 0$, $W_{21} > 0$) the two oscillators try to oscillate in step, interrupted by short periods of silence due to delayed self-inhibition. A simulation of this case is shown in Figure 2a. The degree of synchronization can be quantified by measuring the correlation

$$C(1, 2) = \frac{\langle x_1 x_2 \rangle - \langle x_1 \rangle \langle x_2 \rangle}{\Delta x_1 \Delta x_2}$$

between the two oscillators, Δx_i being the variance of x_i. For the simulation shown in Figure 2a we measured $C(1, 2) = 0.99$, which indicates almost complete phase locking.

The second case, mutual inhibition between the oscillators ($W_{12} < 0, W_{21} < 0$), is shown in Figure 2b. The two oscillators now avoid each other, which is reflected by $C(1, 2) = -0.57$.
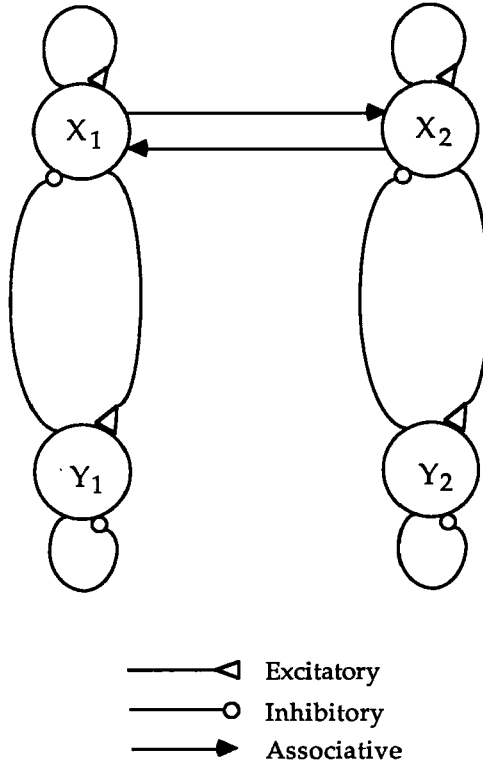
Figure 1: Diagram of two mutually connected oscillators.

Alternatively, both oscillators could be continuously active but oscillate out of phase, with 180° phase shift. That mode has been simulated successfully for the case of two oscillators and might be applied to segmentation of an object from its background; for more than two oscillators with mutual inhibition phase avoidance behavior turns out to be difficult to achieve.

3 Segmentation in Associative Memory

After this demonstration of principle we will now test the associative capabilities of a network of N oscillators connected by Hebbian rules. We store p sparsely coded, random N bit words $\xi^\nu = \{\xi_i^\nu\}_{i=1}^N$ with pattern index $\nu = 1, \ldots, p$. The probability that a bit equals 1 is a, that is, $P(\xi_i^\nu) =$

$a\delta(\xi_i^\nu - 1) + (1 - a)\delta(\xi_i^\nu)$ with typically $a < 0.2$. The synapses are chosen according to the Hebbian rule

$$W_{ik} = \frac{1}{aN} \sum_\nu (\xi_i^\nu - a)(\xi_k^\nu - a) \tag{3.1}$$

With connectivity 3.1, oscillator i receives input $S_i(t) = \sum_{k \neq i} W_{ik} x_k(t)$ from other oscillators.

In the following simulation, 50 oscillators and 8 patterns were stored in the memory. For simplicity we have chosen patterns of equal size
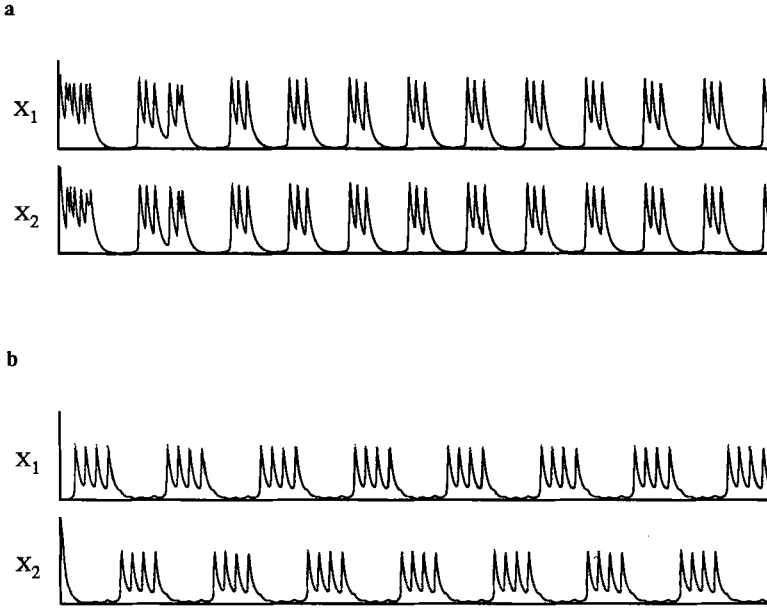

Figure 2: (a) Simulated output pattern of two mutually excitatory oscillators. The parameter values for the two oscillators are the same, $\tau_x = 0.9$, $\tau_y = 1.0$, $T_{xx} = 1.0$, $T_{xy} = 1.9$, $T_{yx} = 1.3$, $T_{yy} = 1.2$, $\eta = 0.4$, $\lambda_x = \lambda_y = 0.05$, $\theta_x = 0.4$, $\theta_y = 0.6$, $I_1 = I_2 = 0.2$, $\alpha = 0.2$, $\beta = 0.14$, $\bar{x} = \bar{y} = 0.2$, $W_{12} = W_{21} = 2.5$. Initial values: $x_1(0) = 0.0$, $x_2(0) = 0.2$, $y_1(0) = y_2(0) = 0.0$. The equations have been integrated with the Euler method, $\Delta t = 0.01$, 14,000 integration steps. (b) Simulated output pattern of two mutually inhibitory oscillators. All parameters are the same as in (a), except that $W_{12} = W_{21} = -0.84$, $\alpha = 0.1$, $\beta = 0.26$.

(8 active units). The first three patterns, which will be presented to the network in the following simulation, have the form

$$\xi^1 = (1,1,1,1,1,1,1,0,0,0,0,0,0,0,0,0,0,0,1,0,\ldots,0)$$
$$\xi^2 = (0,0,0,0,0,0,1,1,1,1,1,1,1,0,0,0,0,0,1,0,\ldots,0)$$
$$\xi^3 = (1,0,0,0,0,0,0,0,0,0,0,0,1,1,1,1,1,1,1,0,\ldots,0) \qquad (3.2)$$

Notice the 25% mutual overlap among these 3 patterns and bits $\xi_{19}^\nu = 1$ for patterns $\nu = 1, 2, 3$.

With this choice of stored patterns we have tested pattern recall and pattern completion after presentation of just one incomplete pattern, the fundamental capability of associative memory. The network restored the information missing from the fragment within one or two cycles. The same behavior had been demonstrated in (Freeman *et al.* 1988).

A more intriguing dynamic behavior is shown by the network if we present all three patterns ξ^1, ξ^2, ξ^3 or parts of them simultaneously. In all simulations external input was time-independent but similar results can be expected for time-dependent input as used in Li and Hopfield (1989). The result of a simulation is shown in Figure 3 where the input is a superposition of patterns ξ^1, ξ^2, ξ^3 with one bit missing in each pattern (see caption of Fig. 3). In this figure only the first 19 oscillators are monitored; the others stay silent due to lack of input and mutual inhibition among oscillators representing different patterns. All three patterns are recognized, completed, and recalled by the network.

In addition to the capabilities of conventional associative memory the network is able to segment patterns in time. The assembly of oscillators representing a single input pattern is oscillating in a phase-locked way for several cycles. This period is followed by a state of very low activity, during which another assembly takes turn to oscillate. In Figure 4 we have plotted the correlations between the first 19 oscillators. The oscillators in one pattern are highly correlated, that is coactive and phase-locked, whereas oscillators representing different patterns are anticorrelated. Oscillators 1, 7, and 13, which belong to two patterns each, stay on for two periods. Oscillator 19, which belongs to all three active patterns, stays on all the time. According to a number of simulation experiments, results are rather stable with respect to variation of parameters.

Switching between one pattern and another can be produced either by noise, or by delayed self-inhibition (the case shown here), or by a modulation of external input. A mixture of all three is likely to be biologically relevant. The case shown here is dominated by delayed self-inhibition and has a small admixture of noise. The noise-dominated case, which we have also simulated, has an irregular succession of states and takes longer to give each input state a chance. Delayed self-inhibition might also be used in a nonoscillatory associative memory to generate switching between several input patterns. Our simulations, however, indicate that limit cycles facilitate transitions and make them more reliable.
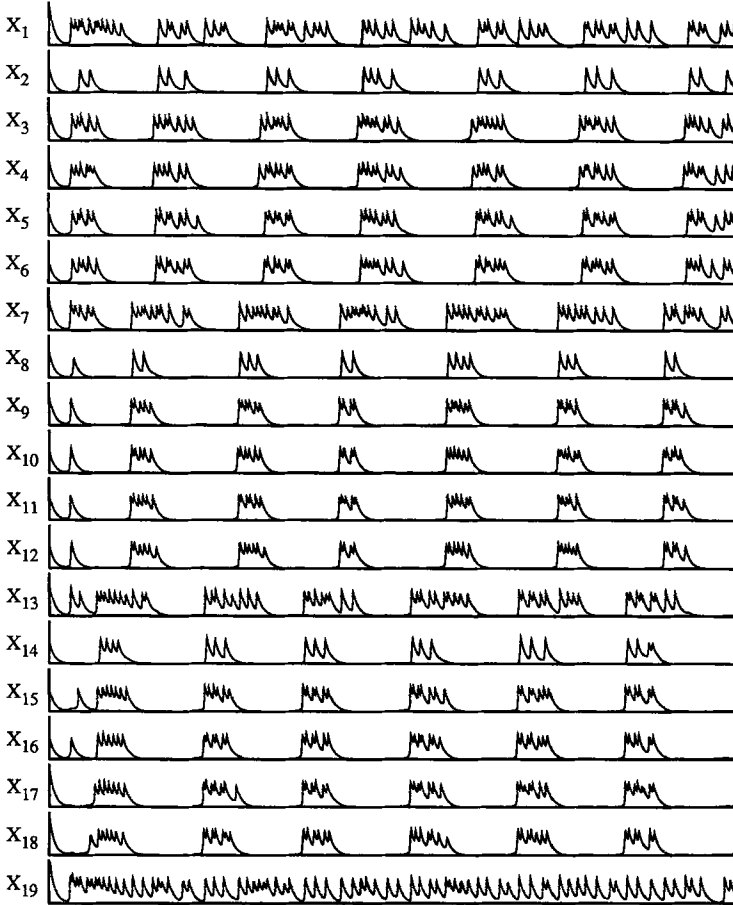
Figure 3: Simulation of an associative memory of 50 oscillators. Eight patterns have been stored in the memory and three of them, ξ^1, ξ^2, ξ^3 (3.2) are presented in this simulation simultaneously with one bit missing in each pattern. Only the output of the first 19 oscillators is shown. The others stay silent due to lack of input. The vertical dashed lines identify three consecutive time intervals with exactly one pattern active in each interval. From the result we see that at any time instant only one pattern is dominant while in a long run, all patterns have an equal chance to be recalled due to switching among the patterns. The parameter values differing from Figure 2 are $T_{yy} = 1.0$, $\alpha = 0.17$, $\beta = 0.1$. We added uncorrelated white noise of amplitude 0.003 to the input to the excitatory groups. Initial value: $x = 0.2(1, \ldots, 1)$, $y = (0, \ldots, 0)$. Input: $I = 0.2 \ (1,0,1,1,1,1,1,0,1,1,1,1,1,0,1,1,1,1,1,0,\ldots,0)$.
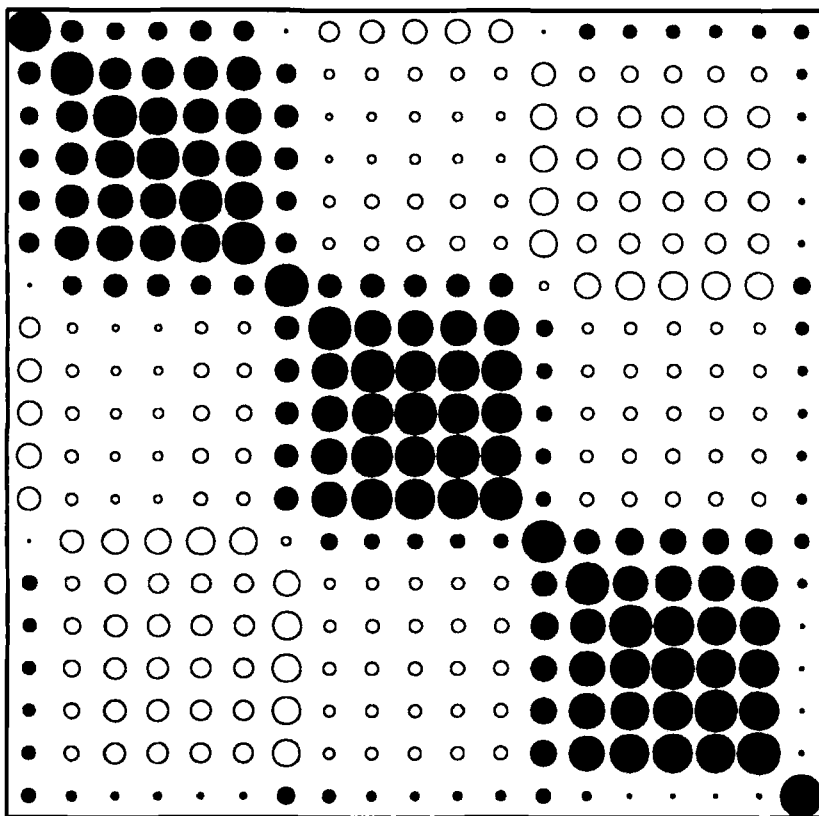
Figure 4: Correlation matrix between the first 19 oscillators (cf. Fig. 3). Filled and open circles stand for positive and negative correlations, respectively. The diameter of each circle is proportional to the absolute value of the correlation.

For conceptual reasons, only a limited number of states can be represented in response to a static input. A superposition of too many (more than perhaps 10) input states leads to ambiguity and the system responds with an irregular oscillation pattern. The exact number of entities that can be represented simultaneously depends on details of implementation, but a reasonable estimate seems to be the seven plus or minus two, often cited in the psychophysical literature as the number of objects that can be held in the human attention span.

4 Discussion _____

The point of this paper is the demonstration of a concept that allows us to compute and represent syntactical structure in a version of associative memory. Whereas in the attractor neural network view a valid state of short-term memory is a static activity distribution, we argue for a data structure based on the history of fluctuating neural signals observed over a brief time span (the time span often called "psychological moment") (Pöppel and Logothetis 1986). There is ample evidence for the existence of temporal signal structure in the brain on the relevant time scale (10–50 msec). Collective oscillations are of special relevance for our study here. They have been observed as local field potentials in several cortices (Gray et al. 1989; Eckhorn et al. 1988; Freeman 1978). The way we have modeled temporal signal structure, as bursts of collective oscillations, is just one possibility of many. Among the alternatives are continuous oscillations, which differ in phase or frequency between substates, and stochastic signal structure.

  Is the model biologically relevant? Several reasons speak for its application to sensory segmentation in olfaction. A major difficulty in applying associative memory, whether in our version or the standard one, is its inability to deal with perceptual invariances (e.g., visual position invariance). This is due to the fact that the natural topology of associative memory is the Hamming distance, and not any structurally invariant relationship. In olfaction, Hamming distance seems to be the natural topology, and for this reason associative memory has been applied to this modality before (Li and Hopfield 1989; Baird 1986; Freeman et al. 1988; Haberly and Bower 1989). Furthermore, in the simple model for segmentation we have presented here, this ability relies completely on previous knowledge of possible input patterns. In most sensory modalities general structure of the perceptual space plays an additional important role for segmentation, except in olfaction, as far as we know. Finally, due to a tradition probably started by Walter Freeman, temporal signal structure has been well studied experimentally (Freeman 1978; Haberly and Bower 1989), and has been modeled with the help of nonlinear differential equations (Baird 1986; Freeman et al. 1988; Haberly and Bower 1989). There are also solid psychophysical data on pattern segmentation in olfaction (Laing et al. 1984; Laing and Frances 1989). It is widely recognized that any new mixture of odors is perceived as a unit; but if components of a complex (approximately balanced) odor mixture are known in advance, they can be discriminated, in agreement with the model presented here. When one of the two odors dominates the other in a binary mixture, only the stronger of the two is perceived (Laing et al. 1984), a behavior we also observed in our model.

  How can associative memory, of the conventional kind or ours, be identified in the anatomy (Shepherd 1979; Luskin and Price 1983) of the

olfactory system of mammals? In piriform cortex, pyramidal cells on the one hand and inhibitory interneurons on the other would be natural candidates for forming our excitatory and inhibitory groups of cells. They would be coupled by associative fibers within piriform cortex. Signals in stimulated olfactory cortex are oscillatory in nature (in a frequency range of 40–60 Hz) (Freeman 1978) and therefore lend themselves to this interpretation. On the other hand, also the olfactory bulb has appropriately connected populations of excitatory (mitral cells) and inhibitory (granule cells) neurons, which also undergo oscillations in the same frequency range and possibly in phase with cortical oscillations. The two populations are coupled by the lateral and medial olfactory tract in a diffuse, nontopographically ordered way. Thus a more involved implementation of associative memory in the coupled olfactory bulb–piriform cortex system is also conceivable.

Our model makes the following theoretical prediction. If the animal is stimulated with a mixture of a few odors known to the animal, then it should be possible to decompose local field potentials from piriform cortex into several coherent components with zero or negative mutual correlation.

Acknowledgments

References

Baird, B. 1986. Nonlinear dynamics of pattern formation and pattern recognition in the rabbit olfactory bulb. *Physica D* **22**, 150–175.

Buhmann, J. 1989. Oscillations and low firing rates in associative memory neural networks. *Phys. Rev. A* **40**, 4145–4148.

Damasio, A. R. 1989. The brain binds entities and events by multiregional activation from convergence zones. *Neural Comp.* **1**, 123–132.

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? *Biol. Cybernet.* **60**, 121–130.

Freeman, W. J. 1978. Spatial properties of an EEG event in the olfactory bulb and cortex. *Electroencephalogr. Clin. Neurophysiol.* **44**, 586–605.

Freeman, W. J., Yao, Y., and Burke, B. 1988. Central pattern generating and recognizing in olfactory bulb: A correlation learning rule. *Neural Networks* **1**, 277–288.

Gray, C. M., König, P., Engel, A. K., and Singer, W. 1989. Oscillatory responses in cat visual cortex exhibit intercolumnar synchronization which reflects global stimulus properties. *Nature (London)* **338**, 334–337.

Haberly, L. B. and Bower, J. M. 1989. Olfactory cortex: Model circuit for study of associative memory? *Trends Neural Sci.* **12**, 258–264.

Hopfield, J. J. 1982. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U.S.A.* **79**, 2554–2558.

Laing, D. G. and Frances, G. W. 1989. The capacity of humans to identify odors in mixtures. *Physiol. Behav.* **46**, 809–814.

Laing, D. G., Panhuber, H., Willcox, M. E., and Pittman, E. A. 1984. Quality and intensity of binary odor mixtures. *Physiol. Behav.* **33**, 309–319.

Li, Z. and Hopfield, J. J. 1989. Modeling the olfactory bulb and its neural oscillatory processings. *Biol. Cybernet.* **61**, 379–392.

Luskin, M. B. and Price, J. L. 1983. The topographical organization of associational fibres of the olfactory system in the rat, including centrifugal fibres to the olfactory bulb. *J. Comp. Neurol.* **216**, 264–291.

Pöppel, E., and Logothetis, N. 1986. Neuronal oscillations in the human brain. *Naturwissenschaften* **73**, 267–268.

Schneider, W. 1986. Anwendung der Korrelationstheorie der Hirnfunktion auf das akustische Figur-Hintergrund-Problem (Cocktailparty-Effekt). Dissertation, University of Göttingen.

Shepherd, G. M. 1979. *The Synaptic Organization of the Brain*. Oxford University Press, New York.

Steinbuch, K. 1961. Die Lernmatrix. *Kybernetik* **1**, 36–45.

Strong, G. W. and Whitehead, B. A. 1989. A solution to the tag-assignment problem for neural networks. *Behav. Brain Sci.* **12**, 381–433.

von der Malsburg, C. 1981. *The Correlation Theory of Brain Function*. Internal Report 81-2, Abteilung für Neurobiologie, MPI für Biophysikalische Chemie, Göttingen.

von der Malsburg, C. 1983. How are nervous structures organized? In *Synergetics of the Brain. Proceedings of the International Symposium on Synergetics, May 1983*, E. Başar, H. Flohr, H. Haken, and A. J. Mandell, eds. Springer, Berlin, Heidelberg, pp. 238–249.

von der Malsburg, C. 1987. Synaptic plasticity as basis of brain organization. In *The Neural and Molecular Bases of Learning*, Dahlem Konferenzen, J.-P. Changeux, and M. Konishi, eds. Wiley, Chichester, pp. 411–431.

von der Malsburg, C. and Schneider, W. 1986. A neural cocktail-party processor. *Biol. Cybernet.* **54**, 29–40.

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