



PERGAMON

Neural Networks 13 (2000) 71–89

Neural
Networks

www.elsevier.com/locate/neunet

Contributed article

Pattern segmentation in a binary/analog world: unsupervised learning versus memory storing

C. Lourenço^{a,b,*}, A. Babloyantz^a, M. Hougardy^a

^aCenter for Nonlinear Phenomena and Complex Systems, Université Libre de Bruxelles, CP 231, Campus Plaine, Boulevard du Triomphe, B-1050 Brussels, Belgium

^bDepartamento de Informática, Faculdade de Ciências, Universidade de Lisboa, Campo Grande, 1700 Lisbon, Portugal

Received 10 August 1998; accepted 24 August 1999

Abstract

We discuss the problem of segmentation in pattern recognition. We adopt the model and the general approach in the landmark paper by Wang, Buhmann and von der Malsburg (*Neural Computation*, (1990), 2, 94–106), and expand their model in a number of ways. We review their solution to the segmentation problem in associative memory, which consists in feature binding being expressed by synchrony relations between oscillators or populations of neurons. We extend the model by introducing a law of synaptic change, which allows the network to learn by structuring itself in response to stimuli with relevant features. We discuss the problem of interference between pattern completion and the learning of new memories. We also propose a form of multiplexing of input information taking advantage of the time-structure of the neurons' response. It is based on the assessment of analog as well as of binary properties of the stimuli and provides for an enhancement of the network's processing capacity. The relevance of the results for biological systems is pointed out. © 2000 Elsevier Science Ltd. All rights reserved.

Keywords: Segmentation; Pattern recall; Associative memory; Learning; Olfaction; Oscillatory networks; Neuronal dynamics; Self-organization

1. Introduction

Associative memory is by now a well-established model of memory encoding and retrieval. According to this paradigm, data of long-term memory are encoded in the synaptic weights of the connections between neurons. Short-term memory and memory recall are viewed as patterns of activity of the neural network. This activity can be stationary (Hopfield, 1982; 1984) or oscillatory (Baird, 1986; Doya & Yoshizawa, 1989; Freeman, Yao & Burke, 1988; Li & Hopfield, 1989).

One of the drawbacks of associative memory is its reduced capability of generalization (von der Malsburg, 1983). Patterns in nature are usually complex and highly structured, composed of multiple sub-patterns that may have occurred already, and that are likely to occur again in different arrangements. Classical associative memory does not take this fact into account. Instead, it treats a complex pattern as a whole, retrieving all or nothing of it.

Since a complex pattern is not likely to recur in exactly the same form, the ability to learn from experience is thus compromised. This applies to different sensory modalities, such as vision, audition and olfaction (von der Malsburg & Schneider, 1986; von der Malsburg & Buhmann, 1992; Wang, Buhmann & von der Malsburg, 1990). We shall focus on olfaction. The main reason is that, within the context of memory storing and retrieval, Hamming distance seems to be the natural topology of this sensory modality. Other modalities demand for additional knowledge of the structure of perceptual space, which are unnecessary to olfaction (Wang et al., 1990). Odor molecules can appear in different combinations. From psychophysics (Laing & Frances, 1989; Laing, Panhuber, Willcox & Pittman, 1984), we know that any new mixture of odors is treated as a unique percept; but when the separate components are known in advance, their discrimination becomes possible. Although the clearest examples are at the low level of sensory processing, processes of segmentation should occur also at higher cognitive levels (Wang et al., 1990). For efficiency, the structure of the external world should somehow be reflected in the interior organization of the brain. Once the need for a segmentation formalism is accepted, one must look at how it occurs in practice. One

*Corresponding author. Departamento de Informática, Faculdade de Ciências, Universidade de Lisboa, Campo Grande, 1700 Lisbon, Portugal. Fax: +351-21-7500084.

E-mail address: csl@di.fc.ul.pt (C. Lourenço).

may start by investigating, already at the level of perceptual space, the kind of information over which segmentation can be performed. Complex modalities, such as vision, would require an increased number of laws of perceptual grouping into segments (von der Malsburg & Buhmann, 1992; Wang et al., 1990). In the simpler case of olfaction, it suffices to identify sub-patterns previously stored in memory.

At the system level, one should investigate the appropriate data structure to encode the memory patterns and their segmentation. This concerns the brain as well as models that implement memory. The solution to this problem comes in the form of tags that are placed upon the activity of neurons (Wang et al., 1990). Since segmentation is already present at the level of input, these tags serve as a means to establish relations between particular groups of neurons and features of the input. If a complex pattern is presented to the system (eventually corrupted), then all neurons corresponding to that pattern become active. Their ensemble activity denotes a successful retrieval. Yet, because it is a composite pattern, some kind of differentiation must exist at the level of neuronal dynamics, within the whole set of neurons that have to be active in response to that pattern. One refers to this grouping and tagging as feature binding.

As a mechanism of binding, several authors have proposed that the encoding of a common feature is expressed by positive temporal correlations between neuronal signals (Damasio, 1989; Eckhorn et al., 1988; Engel, König, Gray & Singer, 1990; Engel, König & Singer, 1991a; Engel, Kreiter, König & Singer, 1991b; Gray & Singer, 1989; Gray, Engel, König & Singer, 1989; Gray, König, Engel & Singer, 1990; Schillen & König, 1994). This requires nontrivial dynamics of the neurons. Neurons that respond to the same feature will be positively correlated; neurons not responding to the same feature will be zero- or anti-correlated. A very important property is that the temporal correlations can be spontaneously created within the network, and therefore need not be stimulus-locked. In this manner, the previous learning of segments makes the segmentation of information possible during the phase of retrieval, even when the input data do not have the perceptual structure necessary for that segmentation.

Learning corresponds to changes in the network structure. In contrast to the retrieval phase, segmentation properties already present in the input are needed during the learning phase, in order to adequately change the network structure. In the formulation that we follow, it is the distribution of synaptic weights that is allowed to vary. Correlations and anti-correlations arise within the network through mutual excitation and inhibition, and this process is conditioned by the synaptic values. Hence, a natural prescription to encode segmentation can be summarized as: all neurons belonging to the same segment are coupled with excitatory links; every two neurons not belonging to the same segment have inhibitory links between them. Weights are created in a Hebbian fashion.

At the retrieval stage, there is the possibility of pattern completion. If a stationary input pattern is close to a stored pattern, then, after a short transient, oscillations are observed in the network. All neurons that should be active in response to that pattern oscillate during at least a fraction of the total time. The fact that they oscillate in phase or out of phase with each other depends on whether or not they belong to the same segment. During each fraction of the total time, neurons of a particular segment oscillate in phase, whereas all the other neurons remain silent. If a segment is preponderant in a composite pattern, then the corresponding neurons may oscillate during a fraction of time that is greater than the fractions of the weaker segments. In the modality of olfaction, this is analogous to the case where a strong odor dominates in a mixture (in this case, the strength of an odor is measured by the number of receptor cells that it excites, not by its concentration in the inhaled air). Overall, confusion between segments is avoided via segmentation in the time domain, as described.

Let us clarify the type of associative memory we are dealing with. We consider the case where the clamped external stimulus acts as information input to the network. Whereas in the well-known dynamical version of associative memory by Hopfield (1982) we have a mapping from the space of initial conditions of the network to the space of attractor solutions (actually the same space), here the mapping is from the space of inputs to the space of attractor solutions. Let us suppose that a specific uncorrupted input pattern leads to a certain pattern in neuronal activity. This neuronal pattern thus corresponds to a memory retrieval or input pattern classification in an ideal case. In order to guarantee the main features of associative memory, moderate distortions of an ideal input pattern or the elimination of few neurons should not change the attractor of neuronal activity by much. A distance, Hamming or other, can also be defined in input space. If the input remains within a certain distance from an ideal input pattern, then the same neuronal pattern should always be retrieved. In this arrangement, the initial condition of the network plays a less important role. Still, if some reasonable conditions are obeyed by the input acting on neurons, we can have the equivalent of a content-addressable memory. The simplest such example is the one where each input cell acts on its own neuron, stimulating it with a certain level of excitation, or otherwise with zero-level excitation. If some bits of information are corrupted in the input pattern, then the network dynamics should suffice to correct the fault and restore the desired output pattern completely.

The model discussed in this paper goes along the lines of the latter version, including the use of the simple one-to-one input stimulation scheme. Both time and neuron state are continuous variables, and neurons may present oscillatory states as well as quiescent ones. Nonetheless, thresholding is often used to provide a binarization of input or of network state.

For alternative dynamical modeling approaches to the

problem of odor recognition, see also the recent work reported in (Hendin, Horn & Tsodyks, 1998; Hoshino, Kashimori & Kambara, 1998). The issue of segmentation is more directly addressed by Hendin et al. (1998). However, the segmentation mechanism in the latter reference has a random nature, relying explicitly on noise and rapid multiple sniffing, as opposed to the deterministic, somewhat more reliable transitions between segments in the approach that we follow.

In the case of a network with pre-determined synaptic weights, one speaks of an information storage algorithm, or the specification of the weights such that the retrieval of an ensemble of patterns be optimally accomplished. This information storage algorithm can be viewed as a prescription. Thus the synaptic weights are imposed upon the network, and are kept fixed. We may also discuss a more complex scenario, one that is likely to occur biologically. One considers two processes that can take place in a neuronal network, at different time-scales: neuronal activity, and a slower evolution of connectivity. On the short time-scale corresponding to neuronal activity, the synaptic strengths can be viewed as constant. On a longer time-scale, the neuronal dynamics may induce a change in the synaptic weights. The slower process corresponds to learning. Synaptic plasticity is a pre-requisite for this type of learning, both in biological systems and model networks, under the Hebbian paradigm. Modernly, Hebbian learning is viewed as a two-way process. That is, synapses can have their efficiency strengthened or weakened as a result of learning history. The former mechanism is the one best studied, and is usually known as Long-Term Potentiation (Bliss & Collingridge, 1993; De Schutter & Bower, 1993; Kanter & Haberly, 1990; Landfield & Deadwyler, 1988; Manabe & Nicoll, 1994; Moser, Trommald & Andersen, 1994). The weakening of synaptic efficiencies has also been reported, both at the slow time-scale of Long-Term Depression (Stanton & Sejnowski, 1989; Stent, 1973) and at the faster time-scale of Depotentiation (Xu, Anwyl & Rowan, 1998). This two-way view of Hebbian learning is essential to the functioning of the model in our paper. It has become standard modeling practice to conceptually divide memory function into separate periods of learning and autonomous retrieval. During learning, external input keeps the neural network in a pattern of firing that encodes some information to be stored. The synaptic weights are changed during this period, as a response to neuronal dynamics. The storage of the pattern is successful if, once the learning period terminated, the desired pattern of activity can be seen in the network, even if the input pattern is incompletely presented or is absent. In this so-called recall phase, synaptic change is usually disallowed.

The model in Wang et al. (1990) features this separation into learning and recall phases. Furthermore, learning actually consists of the storage of computed synaptic weights, which are specified once and for all. In the present paper we implement a form of autonomous synaptic learning of the

Hebbian type, where the network learns from examples. We discuss the case where learning happens simultaneously with pattern recognition, without a rigid separation between learning and recall periods.

2. Physiological aspects

A detailed account of the neurobiology of olfaction is not within our scope. Here, we briefly discuss some biological features of the model. The mammalian olfactory system is adopted as a concrete example.

There is much experimental evidence for the temporal structure of the signals in the form of coherent oscillations. Most of it comes from observations in visual cortex. Additional motivation for the model comes from reported oscillations of electrical activity in the 40 Hz range, in rabbit olfactory cortex (Freeman, 1978) and in monkey somatosensory cortex (Murthy & Fetz, 1991). In insect studies, the major frequency peak during olfactory tasks was found to be close to 20 Hz (Laurent & Davidowitz, 1994; Wehr & Laurent, 1996). Significantly, in the latter references, Laurent and co-workers report on a spatio-temporal structure of neuronal response whereby a given odor elicits oscillations occurring in different time segments in close resemblance to a time-segmentation coding as described in our paper.

The investigation of olfactory memory and learning is facilitated by the fact that the higher processing regions involved are not very far from the very first receptor cells. The pathways are relatively simple. In particular, the input does not have to pass through the thalamus before reaching the cortex, in contrast to other sensory systems.

The initial stage of the olfactory system is composed of the olfactory bulb (OB). It receives input directly from axons of sensory receptors, and transmits in parallel to the anterior olfactory nucleus (AOC) and the prepiriform cortex (PC) (Freeman, 1992). Additional feedback exists between the three parts. Interactions between these parts and transmission from the PC to the limbic and motor systems are in the form of action potentials carried by bundles of axons. Each of the parts has its internal columnar structure. The electrical activity can be divided into local neighborhoods. A local EEG can be defined as the (local) sum of extracellular dendritic currents. For each of the three parts, this EEG features spatially coherent oscillations over the entire domain. A common wave form, the so-called carrier wave (Freeman, 1992), contains from 30 to 90% of the variance in simultaneous recordings from arrays of electrodes. It corresponds to the well-known “40 Hz oscillations”, but it can in practice range from 20 to 90 Hz. On account of this, the oscillations are also called gamma waves (Freeman, 1991). When viewed in EEG tracings, they appear as bursts, i.e. high-amplitude, high-frequency oscillations that stand out from normal EEG, between inhalation and exhalation. The carrier wave is not imposed by the receptors or by other parts of the brain exterior to the olfactory system. It appears




Fig. 1. Illustration of the building blocks of the network. Two oscillators are depicted. Excitatory connections within each oscillator end in triangles; intrinsic inhibitory connections end in circles. Associative connections between oscillators are represented as arrows. Adapted from Wang et al. (1990).

as a cooperative activity that results from self-organization of the neuronal population. In particular, OB bursts clearly show spatial patterns of phase that can only result from such self-organization and not from external forcing (Freeman and Baird, 1987). The carrier wave depends on the existence of interactions between the three parts; but it is also sustained by the internal organization of the parts. These contain mutually connected excitatory and inhibitory neurons. The grouping of neurons into the columns referred to above, results in each column functioning as a local oscillator. Apart from helping sustain the carrier wave, mutual excitation between these oscillators induces a cooperative coupling. Mutual inhibition, on the other hand, provokes a spatial contrast increase in the local amplitudes of the common carrier, thus inducing a modulation in the carrier. It provides for the functional type of separation required by the associative memory and segmentation theory. In the OB, the shape of the carrier wave is not the key factor to signal a particular odor. That shape may even change between successive inhalations of the same odor. Instead, the odor is discriminated by the spatial pattern of modulation of the carrier's amplitude, the same in successive sniffs. Not only is the pattern self-organized, it is also not topographically related to the input (Freeman, 1994).

Olfactory receptors have specialized cilia which respond to different chemicals. Excited cells fire action potentials directly into the OB through axonal connections. The number of activated receptors indicates the intensity of the stimulus. The location of the activated receptors is associated with the nature of the odor. Hence, a spatial pattern of activity is present already at the receptor stage. Yet, this pattern is highly variable, due to unavoidable turbulence in nasal air flow, which causes only a fraction of the receptors sensitive to a particular odor to be excited during a sniff, and the ones selected to vary between successive inhalations in a random manner (Freeman, 1991). Nevertheless, as mentioned above, the same spatial pattern of activity is observed in the OB at all inhalations of the same odor. This implies that some sort of pattern completion must

take place as early as in the OB (and most probably also at later processing stages). A cooperative, self-organized behavior of the entire neuronal population is required to achieve the completion.

The PC and the OB are two candidate parts to which the present modeling could be applied (Wang et al., 1990). One may therefore try to identify the cells that most likely form excitatory and inhibitory groups in each of them. In the OB, these could arguably be the mitral cells for excitation, and the granule cells for inhibition. In PC, pyramidal cells have the role of excitatory neurons, whereas inhibition is provided by inhibitory interneurons.

As referred above, a wealth of experiments confirm the conceptual view of synaptic plasticity, required by associative memory. For experimental evidence of synaptic change specifically in olfactory cortex, (see e.g. Kanter & Haberly, 1990). Apart from the single-cell level, synaptic learning from experience can also be investigated via its consequences at the macroscopic level. Let us consider once more the spatial patterns of activity in the OB. The respective amplitude maps, representing a given odor, are seen to change when the reinforcement associated with that odor also changes (Freeman, 1991). If experience would not influence internal connectivity, then the same spatial pattern would be measured each time the same odor is inhaled, before and after the change in the conditioned association. Further anatomic and physiological arguments, pointing at the associative memory capabilities of the piriform cortex, are collected in the review of Bower, 1994. The patterns of connectivity and the role of the balance between excitation and inhibition, in the formation of sparse and distributed patterns of neuronal activity, are stressed out. Finally, the psychophysical experiments of Laing et al. (Laing & Frances, 1989; Laing et al., 1984) provide good evidence for the segmentation of information in olfactory tasks, against which models could be tested.

3. Model

The system consists in a fully connected network of elements such as the ones depicted in Fig. 1. These elements may have two different states: a resting state, and an oscillatory one. The oscillatory state is also called "bursting". The oscillations are obtained via a feedback loop between a group of excitatory neurons and a group of inhibitory neurons. The average activities of these groups are denoted by x_i and y_i , respectively. As to the equations describing the evolution of the x_i and y_i , we follow the model of Wang et al. (1990):

$$\dot{x}_i = -\frac{x_i}{\tau_x} + \mathcal{G}_x \left[T_{xx} \frac{x_i}{\bar{x}} - T_{xy} \mathcal{F} \left(\frac{y_i}{\bar{y}} \right) + S_i^{\text{exc}} + S_i^{\text{inh}} + I_i - H_i \right] \quad (1)$$

$$\dot{y}_i = -\frac{y_i}{\tau_y} + \mathcal{G}_y\left(T_{yx}\frac{x_i}{\bar{x}} - T_{yy}\frac{y_i}{\bar{y}}\right) \quad \dot{H}_i = \alpha x_i - \beta H_i,$$

where

$$\mathcal{G}_u(\psi) = \frac{1}{1 + \exp\left(-\frac{\psi - \theta_u}{\lambda_u}\right)}, \quad u \in \{x, y\}$$

$$\mathcal{F}(\psi) = (1 - \eta)\psi + \eta\psi^2.$$

The time-constants of the excitatory and inhibitory parts of the oscillators, respectively, τ_x and τ_y , can be chosen on physiological grounds. The average values of x and y are controlled by the parameters \bar{x} and \bar{y} . All input to x and y is filtered through sigmoidal gain functions \mathcal{G} . θ_u and $1/\lambda_u$ are, respectively, thresholds and gain parameters of the sigmoids. The nonlinearity \mathcal{F} renders the oscillatory character of the network somewhat invariant to the number of oscillators in a pattern. The synaptic strengths of the feedback loop of each oscillator are quantified by the $T_{uu'}$, $u, u' \in \{x, y\}$. The summed excitatory input to oscillator i , from all oscillators of the network, is given by

$$S_i^{\text{exc}}(t) = \sum_{j=1}^N \omega_{ij}x_j(t),$$

whereas inhibitory input is given by

$$S_i^{\text{inh}}(t) = \sum_{j=1}^N \omega_{\text{inh}}x_j(t) = \omega_{\text{inh}} \sum_{j=1}^N x_j(t),$$

for a network of N elements. ω_{ij} and ω_{inh} represent excitatory and inhibitory associative synapses, respectively. The latter are taken as constant throughout the network. Symmetric connections $\omega_{ij} = \omega_{ji}$ will be considered. The time-dependent external input, $I_i(t)$, comes directly from a sensory area, or from an intermediate pre-processing region. The array $\{I_i\}_{i=1, \dots, N}$ contains the patterns to be processed by the network.

4. Intermittent bursting—dynamical analysis

The variables H_i have an important role in the segmentation process. They provide delayed self-inhibition, with strength α and decay constant β . This generates intermittent bursting of each group of oscillators belonging to a pattern, and facilitates the transition between simultaneously retrieved patterns. The delayed inhibition can also be expressed by rewriting the H_i equation in integral form:

$$H_i(t) = \alpha \int_0^t x_i(\tau) \exp[-\beta(t - \tau)] d\tau. \quad (2)$$

The following discussion of bursting helps clarify the behavior of the system in response to external input. Let us suppose that a given oscillator i is disconnected from the network, and the S_i and I_i in Eq. (1) are replaced by a constant term Φ providing a fixed amount of excitation.




Fig. 2. Simultaneous time course of x and H for the single oscillator, with a fixed external excitation value. Parameters of Eq. (3) are $\tau_x = \tau_y = 0.4$, $\bar{x} = \bar{y} = 0.2$, $T_{xx} = 1.6$, $T_{xy} = 1.9$, $T_{yx} = 1.3$, $T_{yy} = 1.0$, $\Phi = 0.2$, $\alpha = 0.17$, and $\beta = 0.1$; parameters of \mathcal{G} and \mathcal{F} are $\theta_x = 0.4$, $\theta_y = 0.6$, $\lambda_x = \lambda_y = 0.05$, and $\eta = 0.4$.

Thus we obtain the equations for the single oscillator

$$\dot{x} = -\frac{x}{\tau_x} + \mathcal{G}_x\left[T_{xx}\frac{x}{\bar{x}} - T_{xy}\mathcal{F}\left(\frac{y}{\bar{y}}\right) + \Phi - H\right] \quad (3)$$

$$\dot{y} = -\frac{y}{\tau_y} + \mathcal{G}_y\left(T_{yx}\frac{x}{\bar{x}} - T_{yy}\frac{y}{\bar{y}}\right)$$

$$\dot{H} = \alpha x - \beta H,$$

with Φ representing the total excitation coming from outside the oscillator. The interplay of x and H gives rise to a regular succession of bursting periods, with fast oscillations of x around an average \bar{x} , alternating with resting periods of almost zero activity. This situation is illustrated in Fig. 2. As suggested by the figure, the dynamics can be naturally decomposed into a slow and a fast component. x and y correspond to the fast dynamics. During the bursting, their variation is fast enough so that only their time-averaged values are perceived by the slow process. The variation of H , on the other hand, is slow enough for this variable to be treated as a parameter of the fast dynamics, at least in a first approximation. The theory of dynamical systems now helps to understand the behavior of the fast system (x, y) . In the phase-space spanned by these two variables, more than one asymptotic solution may exist; eventually, two of them can be simultaneous attractors of the (fast) dynamics, and we have bistability. In the present case, a fixed point near $(0, 0)$, and a periodic limit cycle that does not contain the origin, are simultaneously stable when H takes values in an open range (H_{\min}, H_{\max}) . This range can be evaluated from Fig. 2. If we start from the situation where the stationary solution is observed, we see that a steady decrease of H is happening, according to Eq. (3). Inevitably, a minimum value H_{\min} is reached such that a node–saddle bifurcation

(Guckenheimer and Holmes, 1983) of the fixed point takes place. A sharp transition is then observed, and stability is transferred from the fixed point to the high-value oscillations. From there on, the variation of H is inverted and it starts to increase. We note that, immediately after H starts growing, the fixed point near (0,0) becomes stable again. However, the system does not come back into this fixed point before a maximum H_{\max} is reached. At that moment, it is the limit cycle that loses its stability, to the lower fixed point. The fast oscillations are once again a stable solution, right after H starts to decrease. Yet they cannot be observed before the lower fixed point is made unstable once more. This is repeated on a regular basis, producing a phenomenon of hysteresis. H acts as a bifurcation parameter of the fast dynamics, and at the same time its variation is influenced on a slow time-scale by the fast dynamics itself.

In the following, bursting is investigated at the network level.

5. Synchronization of oscillators

As a first step towards the study of collective behavior, Wang et al. (1990) considered the dynamics of a pair of coupled oscillators. This case is represented in Fig. 1. The connectivity is given by

$$(S_1^{\text{exc}} + S_1^{\text{inh}})(t) = (\omega_{12} + \omega_{\text{inh}})x_2(t)$$

$$(S_2^{\text{exc}} + S_2^{\text{inh}})(t) = (\omega_{21} + \omega_{\text{inh}})x_1(t).$$

We remember the choice of symmetric connections $\omega_{ij} = \omega_{ji}$. Two cases can be distinguished, according to the sign of the associative synaptic strength ($\omega_{ij} + \omega_{\text{inh}}$): if the sign is positive, then the oscillators try to oscillate in synchrony, alternating with quiescent periods due to the delayed self-inhibition; if it is negative, mutual inhibition is observed, and time is divided into periods during which only one of the oscillators is active, while the other is at rest. In the latter regime, both oscillators can have significant activity, but never simultaneously.

These associative properties can be generalized to groups of several coupled oscillators. If the net associative connections within a group are excitatory, then that group will have a tendency to oscillate in synchrony. On the other hand, the ensemble of oscillators of that group will inhibit (and be inhibited by) all oscillators from any other group with which they maintain inhibitory connections. In this manner, only one group is active at a given time. Nevertheless, all will eventually be given a chance to oscillate. The active group will necessarily inhibit itself via the delayed feedback, and go into the resting state, thereby no longer inhibiting all the other groups and allowing one of them to burst.

A group of units oscillating in synchrony may behave like a single oscillator, such that Eq. (3) can be used as an approximation. Specifically, we consider a group of M excitatorily coupled oscillators. We may regard the activities x_k

of the remaining $N - M$ oscillators to be zero. Assuming identical excitatory synapses $\omega_{ij} = \omega$ and homogeneous constant input $I_i = I$, Eq. (3) can be used instead of Eq. (1) for the oscillating group. One takes $x_i \equiv x$, $y_i \equiv y$ and $H_i \equiv H$, for all i belonging to the group considered. Furthermore, there is the following correspondence between quantities occurring in each set of equations:

$$\begin{array}{ll} \text{Eq. (3)} & \text{Eq. (1)} \\ \frac{T_{xx}}{\bar{x}} & \leftrightarrow \frac{T_{xx}}{\bar{x}} + (M - 1)(\omega + \omega_{\text{inh}}) \\ \Phi & \leftrightarrow I \end{array}$$

Just as for the single oscillator, it is the variable H that triggers the transition between high frequency collective oscillations and resting periods.

6. Pattern retrieval with inhibited learning

External patterns are presented to the network via the array $\{I_i\}_{i=1,\dots,N}$. We start by considering stationary binary patterns. The I_i have the value 0 or 1, multiplied by some constant amplitude. They provide the required amount of external excitation for oscillations to occur. Successive pattern retrieval is observed when all the units that have to be active for that pattern oscillate in phase. In the presence of an uncorrupted pattern, each oscillator i has its own amount of excitation I_i that helps maintain its active state. In the opposite case, with all $I_i = 0$, no oscillations are seen. In the simulations that follow, the values of the parameters are chosen in a way that renders the network's response as invariant to the size of the patterns as possible.

We begin by illustrating the property of segmentation in the time domain. Let us suppose that a composite pattern is presented to the system. It consists of several sub-patterns (hereafter simply called patterns), with no apparent separation between them. These patterns can only be distinguished by the network on the basis of previously acquired knowledge. This knowledge is contained in the distribution of connectivity: if two oscillators correspond to the same pattern, they are mutually connected in an excitatory fashion; if they belong to different patterns, they are inhibitorily connected.

Several variants of the connectivity obey the simple condition in the previous paragraph. The connectivity adopted here is different from the one in (Wang et al., 1990). We set the constant ω_{inh} to a negative value, and let the ω_{ij} control the total value of each associative connection: if ω_{ij} is positive and $|\omega_{ij}| > |\omega_{\text{inh}}|$, then the associative link will be excitatory; otherwise, the link will be inhibitory. The ω_{ij} are fixed by the rule

$$\omega_{ij} = \frac{1}{D_\omega} \left(R' + \frac{S_R}{\sum_l K_{il}} \right) K_{ij}. \quad (4)$$

Table 1
Network parameters in the simulations

Parameters in common	$\omega_{\text{inh}} = -5.0$, $\alpha = 0.17$, $\bar{x} = \bar{y} = 0.2$, $T_{xy} = 1.9$, $T_{yx} = 1.3$, $\eta = 0.4$, $\theta_y = 0.6$, $\lambda_x = \lambda_y = 0.05$
Figs. 3, 7 and 8	$N = 21$, $\tau_x = \tau_y = 0.4$, $\beta = 0.1$, $T_{xx} = T_{yy} = 1.0$, $\theta_x = 0.4$, $R' = 5.0$, $S_R = 1.1$, $D_\omega = 1.0$
Figs. 4–6	$N = 21$, $\tau_x = 0.5$, $\tau_y = 0.6$, $\beta = 0.03$, $T_{xx} = T_{yy} = 1.2$, $\theta_x = 0.25$, $R' = 5.0$, $S_R = 1.1$, $D_\omega = 1.0$
Figs. 9 and 10	$N = 11$, $\tau_x = 0.9$, $\tau_y = 1.0$, $\beta = 0.1$, $T_{xx} = T_{yy} = 1.0$, $\theta_x = 0.4$, $R' = 1.0$, $S_R = 0.3$, $D_\omega = 0.2$, $\theta_K = 3.0$, $\lambda_K = 1.0$, $\varsigma = 0.02$, $\gamma = 1000.0$
Figs. 11 and 12	Same as previous, except $S_R = 0.26$

For each group of oscillators that should respond to the same pattern, excitatory associative links are obtained by setting $K_{ij} = 1$ for each pair of oscillators within that group. For each pair of oscillators belonging to groups that respond to different patterns, K_{ij} is set equal to zero. The term $R' + S_R/(\sum_l K_{il})$ provides a renormalization of synaptic weights. $1/D_\omega$ is a proportionality factor whose meaning will become clear in Section 8. The parameters are chosen in a way to fulfill the condition $\omega_{ij} > 0$ and $|\omega_{ij}| > |\omega_{\text{inh}}|$, or else $\omega_{ij} = 0$ (Table 1). For example, each oscillator of a group of M elements responding to the same pattern, will have two possible values for the associative links: $\omega_{\text{inh}} < 0$ for the connections with oscillators from other groups, and $(\omega_{ij} + \omega_{\text{inh}}) > 0$ for connections within its group, where

$$\omega_{ij} = \frac{1}{D_\omega} \left(R' + \frac{S_R}{M-1} \right).$$

The transition mechanism discussed above, involving

delayed self-inhibition, allows that several different patterns be discriminated by the network in a simultaneous representation. Each group of oscillators that corresponds to a pattern present in the input, is allowed to oscillate on its turn (see Fig. 3). In this manner, temporal segmentation of the information is achieved. Other transition mechanisms could be considered, e.g. stochastic ones. However, delayed feedback makes for more reliable transitions. A more regular succession of states is observed, with each group being given the chance to oscillate without waiting for a long time.

Let us inspect Fig. 3. The dynamics of each oscillator i is conditioned by the signals that it receives from the other members of the network, and also by the input conveyed by its receptor cell I_i . In the experiments of Section 6, the input can only have the values 0 or 0.2. Globally, we observe three different patterns memorized and retrieved, each corresponding to a group of oscillators that are active during a separate period. Segmentation is dependent on




Fig. 3. Discrimination of patterns in a simultaneous representation via a segmentation in the time domain. The network and input dimension is $N = 21$. Each element of the network oscillates in response to external excitation. The individual oscillatory activities are represented by the variables x_i of Eq. (1). The input is separated by the network into different patterns: the network activity is divided into periods of time during which only the oscillators corresponding to one of the patterns are active. Three patterns are memorized. In this experiment, they are simultaneously present in the input, with strengths: 0.2 (0,1,1,0,1,0,1,0,0,1,0,0,1,0,0,0,0,0,1,0), 0.2 (0,0,0,1,0,0,0,0,1,0,1,1,0,0,0,0,1,0,0,0,0), and 0.2 (1,0,0,0,0,1,0,1,0,0,0,0,0,1,1,1,0,1,1,0,1). The thick vertical line at $t = 0$ indicates which receptor cells convey external input: in this case, all of them. Input is kept fixed during the entire simulation. The vertical scale for the activity of oscillators, shown on the right, has a value of 0.3. Parameters of Eqs. (1) and (4) are in Table 1.




Fig. 4. Pattern completion in the case of degraded input. In this figure and the two following ones, external input is 0.2 times a degraded version of pattern (1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,0,0,0). As in Fig. 3, a thick vertical line at $t = 0$ indicates which receptor cells convey external input. The excited cells (cells 4–18) are a subset of the original pattern. Parameters of the network are in Table 1.

intrinsic properties of the network, which relies on its connectivity and internal dynamics to restore information not present in the external input.

The grouping of oscillators is functional, and not topographically conditioned. The index of each oscillator gives no information of its spatial location, which we leave undefined. In view of this, in all the numerical simulations that follow, we shall choose consecutive

indices for the oscillators of the same group, for ease of discussion.

In the experiment of Fig. 3, the memorized patterns have similar sizes. We already observe that, in the case where several groups are excited simultaneously, the larger ones are active more often. When the size difference is larger, this effect is amplified. The smaller patterns may even become totally inhibited by the larger ones, and not oscillate at all, if




Fig. 5. An attempt to restore a pattern with notably degraded input. The memorized patterns are the same as in Figs. 3 and 4, but only cells 10–18 are excited.




Fig. 6. Failure to restore a pattern with extremely degraded input. The memorized patterns are the same as in Figs. 3–5, but only cells 16–18 are excited.

a simultaneous retrieval is attempted. However, the small patterns can be retrieved if excited in isolation, or simultaneously with other small patterns of similar size.

Let us now discuss pattern completion. In the following sequence of experiments, four patterns are memorized in a network of 21 elements. The patterns are (1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,0,0,0), (0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0), (0,0,0,0,0,0,0,0,0,1,0,0,0,0,0,0,0,0),

(0,1,0), and (0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,1). The last three are never excited. The first one is present in the input, but in an incomplete manner. For example, in Fig. 4, the input pattern is 0.2 (0,0,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,0,0). Thus, the first two bits of information are degraded in the external input. Yet, the network manages to restore the desired pattern completely, via the cooperative dynamics that results from the associative connectivity.




Fig. 7. Activity of the network when excited by modulated input. The thick vertical line at $t = 0$ now corresponds only to the receptor cells that are “on”, without expressing their individual intensity. See text for the distribution of input intensity. Parameters of the network are in Table 1.

Figs. 5 and 6 show attempts at retrieval of the same pattern, with more degraded versions of the input. An increasing number of bits is missing in the external excitation. In Fig. 5, the desired pattern is not restored in a perfect manner. However, a kind of “soft” pattern completion is possible. The success or failure of the retrieval depends on the way the network dynamics is possible. The success or failure of the retrieval depends on the way the network dynamics is evaluated. Retrieval may fail if the evaluation of the dynamics is too strict. For instance, if the network dynamics is assessed during a short time-interval, this interval may miss completely the oscillations of x_1 to x_9 . A longer time-scale evaluation may “capture” the oscillations of all units, thus viewing the pattern retrieval as successful. Let us consider a concrete example. Suppose that, in Fig. 5, the response of the network is monitored from $t = 0$ to $t = 100$. In some of the periods of activity of the pattern given by x_1 to x_{18} , the units x_1 to x_9 do not oscillate. A strict criterion for the successful retrieval of the complete pattern x_1 to x_{18} would be the fact that all these oscillators display, or not, the same activity from $t = 0$ to $t = 100$. According to this rule, the retrieval fails in Fig. 5. A less limiting criterion for successful retrieval simply demands that all oscillators display simultaneous nonzero activity during at least a fraction of the time-interval considered. According to the latter criterion, the network achieves the retrieval of the complete pattern.

In the example of Fig. 6, the mutual entrainment of the oscillators is clearly insufficient to complete the pattern. Hence, the retrieval fails.

7. Higher-order coding of information

In order to expand the possibilities of short-term memory in this model, we investigated how the network deals with modulations of the input. These modulations are expressed as a varying magnitude of the intensity multiplying the 0 and 1 bits of external excitation. The intensity variation can be viewed as an “extra dimension” of the input. Below we will argue that, apart from recognizing binary patterns, the network can encode information contained in the modulation of the input. Thus, the network possesses analog processing capabilities.

We consider again the example where four binary patterns are memorized in the connectivity of a network with 21 elements: (1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,0,0,0), (0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,1,0,0), (0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0), and (0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,1). Let us consider a modulation of the first pattern, with amplitudes (0.1,0.1,0.1,0.1,0.1,0.1,0.1,0.15,0.15,0.15,0.15,0.15,0.15,0.2,0.2,0.2,0.2,0.2,0.2,0,0,0). The network’s response is plotted in Fig. 7. As before, the binary pattern is successfully retrieved. In addition, information related to the detailed pattern intensity is encoded in the duration of the bursts corresponding to each receptor cell. As can be seen in

Fig. 7, the higher the intensity, the longer the associated burst. While keeping the input values fixed, we measured the burst duration T_{burst} as a function of the input intensity I . The duration of a burst is defined as the time-interval during which x is greater than a threshold $\varsigma = 0.02$. T_{burst} obeys to a statistical distribution, which is characterized by an average value \bar{T}_{burst} and a standard deviation σ_T . In the conditions of Fig. 7, we obtained the relations

I	$\bar{T}_{\text{burst}} \pm \sigma_T$
0.1	4.6 ± 0.6
0.15	6.5 ± 0.4
0.2	8.7 ± 0.4

We see that the dispersion of T_{burst} around its mean value is small. Hence, the value of T_{burst} for an oscillator x_i provides an almost direct measure of the input intensity I_i .

For a more vivid example of higher-order coding, we split the original 18-bit pattern into three smaller patterns of 6 bits each (via the forcing of inhibitory links between each pair of oscillators belonging to different patterns). Each of the patterns is excited by modulated input. The modulation takes the form (0.065,0.065,0.14,0.14,0.27,0.27) for each of the six-unit patterns. The groups of oscillators respond as shown in Fig. 8. Once again, the oscillators that receive the most intense external excitation are the ones that display the longest bursts. The bottom of the figure emphasizes the distribution of burst duration within the same pattern. With the network connectivity and distribution of input intensity as in Fig. 8, we obtained the following values for burst duration:

I	$\bar{T}_{\text{burst}} \pm \sigma_T$
0.065	3.9 ± 0.5
0.14	6.1 ± 0.6
0.27	9.6 ± 0.4

This experiment combines the segmentation into binary patterns with the assessment of the distribution of input intensity according to burst durations. This provides a kind of multiplexing of the information, thus increasing the network capacity.

The retrieval of a binary pattern is invariant with respect to modulations of the input, as long as the latter remains above some threshold of excitability. In turn, the distribution of burst duration can be viewed as an analog pattern conveyed by the network dynamics. A decoding mechanism sensitive to burst duration may “read” the analog pattern, corresponding to the distribution of input intensity (arguably, knowledge about the intensity of input conveyed




Fig. 8. Top: An experiment similar to the one of Fig. 7, but with the 18-bit pattern segmented into three smaller ones. See text for details of the modulation. Bottom: Enlargement of a portion of the figure. Notice the considerable variation in burst duration due to excitation of varying amplitude.

by each receptor may be relevant for some olfactory tasks). According to the processing task, binary or analog information can be extracted from the network dynamics. Thus, the short-term memory capabilities of the network are enhanced.

The analog pattern can also be “binarized”. In the example of Fig. 8 (top), the analog decoder could group the oscillators according to common burst duration. This would allow to define new “binary” patterns, given by $(1,1,0,0,0,0,1,1,0,0,0,0,1,1,0,0,0,0,0,0,0,0)$, $(0,0,1,1,0,0,0,0,1,1,0,0,0,0,1,1,0,0,0,0,0,0)$, and $(0,0,0,0,1,1,0,0,0,0,1,1,0,0,0,0,1,1,0,0,0,0)$. This can be interpreted as a new type of segmentation. However, it has the transitory nature imposed by possible fluctuations of the

input. It should not be confused with the more permanent segmentation, the one that is encoded in the network connectivity, and that refers to the separation of oscillators into groups of simultaneously active units. The latter would of course reveal the patterns $(1,1,1,1,1,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0)$, and so on. As we saw, the two types of segmentation may coexist.

8. Dynamical learning

The distribution of synaptic weights imposed upon




Fig. 9. Pattern segmentation. A 5-unit pattern separates into two smaller ones. The changes in external stimulus, that force this process, are coded in the vertical lines. As in the figures of Section 6, a thick line at some position indicates that the corresponding oscillator receives external input in the period of time on the right of that line (at least until the next vertical line appears). On the contrary, a thin line indicates absence of external input. The input bits are multiplied by the constant factor 0.2. The vertical scale for the activity of oscillators is now 0.5. Parameters of the network are in Table 1.

the network, as discussed above, may seem somehow artificial. In addition, it does not incorporate possible contributions from autonomous, unsupervised, learning. In the present section, we introduce a form of Hebbian learning for the network. The synaptic strengths are allowed to change according to the dynamics of the network. In the spirit of Hebb's postulate (Hebb, 1949),

the variation of synaptic strengths depends essentially on the covariance of the activities of oscillators at both ends of the connections. We introduce the following law of synaptic change:

$$\dot{\omega}_{ij} = (R^l + R_i)K_{ij} - D_{\omega}\omega_{ij} \quad (5)$$




Fig. 10. Preservation of long-term memory. Pattern presentations, after long periods of no stimulation, bring about successful recalls. Parameters are as in Fig. 9.

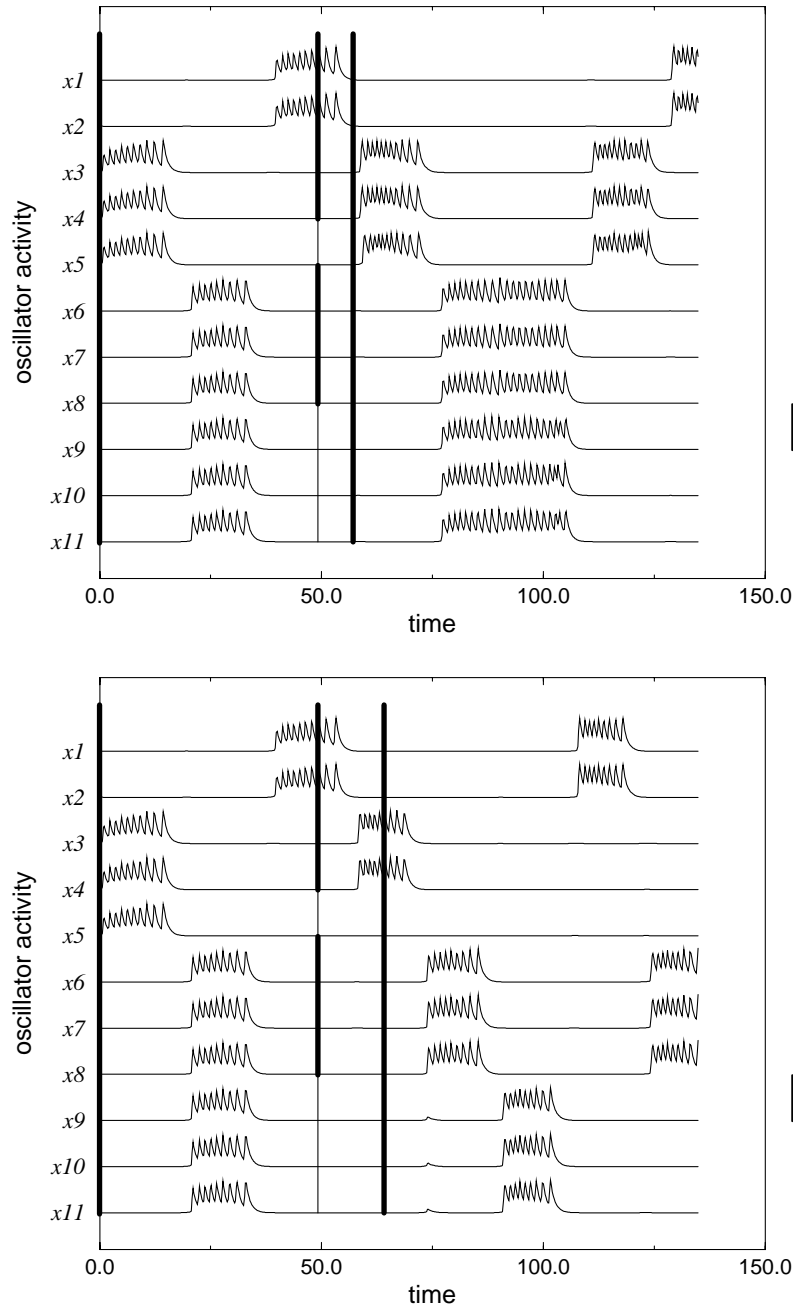


Fig. 11. Top: Dependence of segmentation on the time-history of the input, in a network of 11 elements. Here, segmentation does not occur. Input changes at $t = 49$ and $t = 57$. Bottom: Response of the network, with a different time-history of the input. Segmentation is verified. Input changes at $t = 49$ and $t = 64$. “On” and “off” receptor cells are coded with thick and thin lines, respectively. The input bits are multiplied by the constant factor 0.2.

$$\dot{R}_i = S_R - \left(\sum_l K_{il} \right) R_i,$$

where

$$K_{ij} = \mathcal{G}_K[\omega_{ij} + \gamma(x_i - s)(x_j - s)]$$

and

$$\mathcal{G}_K(\psi) = \frac{1}{1 + \exp\left(-\frac{\psi - \theta_K}{\lambda_K}\right)}.$$

The parameter values are chosen such that the factor $(R' + R_i)$ is always positive. The main term associated with Hebbian covariance is in the argument of the sigmoid \mathcal{G}_K . s is a small and positive parameter that defines a threshold for the oscillators to be considered active. Typically, we let $s = \bar{x}/10$. R' can be viewed as the (constant) concentration of some neuromodulator (a more extended discussion of the effects of neuromodulators appears in Section 9). The evolution equation for R_i , on the other hand, expresses the fact that the concentration of a neuromodulator is varying

locally with a constant positive source term S_R and a decay term taking synaptic renormalization into account. D_ω measures the decay rate of synaptic strengths. The above law implements the general learning rules: (a) if two oscillators are simultaneously active, the excitatory links between them are reinforced, or saturate at some higher value; (b) if two oscillators are at rest, their connection is not altered; (c) the excitatory links between an active oscillator and one at rest, are reduced or can practically disappear - in this case, the resulting connection will be inhibitory.

In a first approximation, suppose the K_{ij} in Eq. (5) can only have the values 0 and 1 (these are the limits of the sigmoid \mathcal{G}_K). Then, the stationary solution of Eq. (5) is given exactly by Eq. (4). There can be situations where the external stimulus does not change, or does not appear in novel combinations, for a long period. In this case, there is no induction of learning via synaptic change, and the behavior of the network is analogous to the one seen in Section 6.

Let us discuss how novelty can be conveyed by external input. We assume that some input pattern has been presented to the network previously, but did not feature any segmentation structure. If, at a given occasion, that pattern is presented to the network in segmented form, then the network will incorporate in its links the knowledge about that segmentation. In practice, a segmentation of an input pattern consists in the stimulation of the network by only a part of the bits of the originally larger pattern. Thus, while being a part of the larger pattern, the smaller segment can also occur isolated in the input. Hence, it could be viewed as a pattern of its own. In later presentations, the remaining part of the original pattern can also occur isolated, although this is not a necessary condition for segmentation.

On the level of synapses, segmentation is achieved via the elimination of excitatory connections between oscillators that should become part of distinct sub-patterns. The difference in the activities of the oscillators at both ends of a connection, which forces Hebbian learning, is itself caused by differential external stimulation.

A simple example of this process is shown in Fig. 9. In a network of 11 oscillators, total stimulation starting from $t = 0$ reveals the existence of a 5- and a 6-bit patterns memorized. From $t = 118$ to $t = 180$, only a fraction of the 5-bit pattern is stimulated. This results in this pattern being segmented into two smaller ones, with 2 and 3 bits respectively, as can be seen from the response to another total stimulation on $t = 180$. Unsupervised learning has therefore occurred.

Fig. 10 is a sequel of Fig. 9. It shows the property of “freezing” of the associative links, over long periods of no external stimulation. The associative connections between groups of oscillators are preserved. This allows, in later recalls, to have perfect pattern retrievals that maintain the learned segmentation structure. In this case, freezing is due to the natural dynamics of Eq. (5).

At this point, one must consider an important effect caused by the introduction of the dynamical change of the links. There is a conflict inherent to the choice that the network must make, between pattern completion and pattern segmentation. If stimulated by an incomplete version of a previously stored pattern, should the network complete the original pattern, or segment it? One way to avoid this conflict is by admitting the existence of a so-called modulation of the learning state. The network activity is divided into different periods, respectively, for pattern learning and for pattern recall. A biological mechanism that makes this possible, is briefly discussed below.

9. Biological view of the modulation of state

Certain substances can have diffusive modulatory effects on network function. They are called neuromodulators. In contrast to neurotransmitters, they are not directly responsible for information transmission. A particular neuromodulator, acetylcholine, might just provide the required modulation of the learning state that we mentioned above (Bower, 1994). This substance is diffusely released throughout the olfactory cortex. In the following we summarize the properties of acetylcholine that are relevant for the function of associative learning and recall (Bower, 1994). This discussion is based on experiments where acetylcholine is placed directly onto the cortical network (Hasselmo and Bower, 1992; Hasselmo, Anderson & Bower, 1992).

Acetylcholine selectively suppress intrinsic associative synapses. By increasing cell excitability, this neuromodulator amplifies the influence of afferent input on cortical cells. Furthermore, increased cell excitability causes a net growth of the gain of synaptic modification, which in living tissue is related to the size of the postsynaptic response during synaptic modification. Finally, by maximizing the efficiency of inhibitory neurons of olfactory cortex, acetylcholine could inhibit spurious influences on synaptic modification by eliminating weak neuronal activities. In our case, these spurious influences correspond to attempts at completing some pattern by an associative mechanism dependent on information previously stored in the connectivity. The net action of acetylcholine would therefore be to increase the importance of the afferent input in restructuring the network, as compared to the network’s intrinsic dynamics.

Although physiology rules out a complete suppression of intrinsic synaptic transmission during learning, it is argued in (Bower, 1994) that acetylcholine may effectively contribute to a large inhibition of normal network recall dynamics during learning periods. Also, a regulatory mechanism is proposed by which the concentration of acetylcholine, and thus the magnitude of its effects, varies smoothly according to the needs of any particular learning sequence. In this manner, the olfactory cortex would not necessarily have to be in one of the two extreme states of learning- and recall-only.

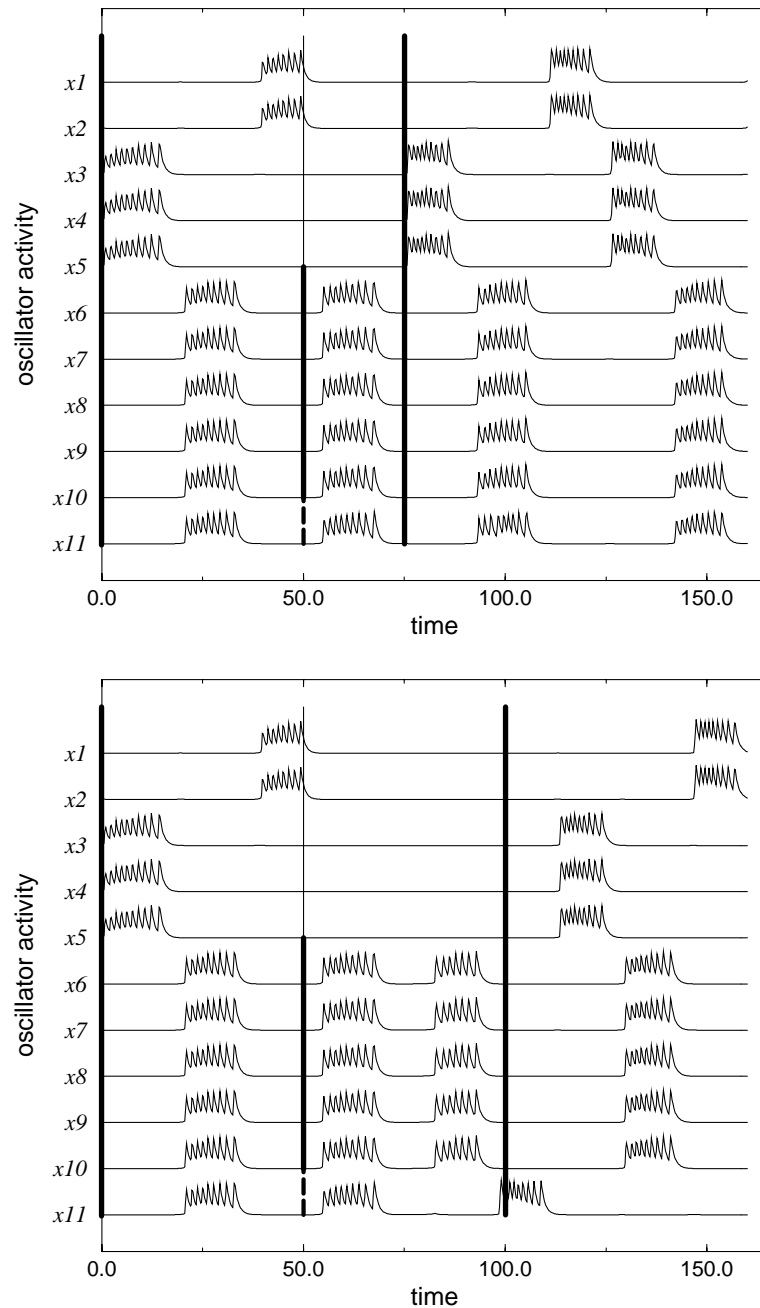


Fig. 12. Top: Restoration of an analog-degraded pattern during a learning period, between $t = 50$ and $t = 75$. Bottom: Segmentation due to inhomogeneity of the input intensity during a longer learning period (from $t = 50$ to $t = 100$). The input bits 0 and 1 are multiplied by the constant factor 0.2, except for I_{11} which has an intensity degraded by 0.1% (dashed line).

These considerations on the regulatory action of acetylcholine are not captured numerically in the model that we analyze. Rather, when structural change is allowed in the model network, it has essentially the unsupervised-dynamical character provided by Eq. (5). Nonetheless, a mechanism for imposing a modulation of the learning state, from the exterior, would be easy to implement in the model, as well as biologically funded.

10. Dependence of learning on input timing

The model has additional features that are interesting in the light of the discussion on state modulation. Namely, the processes of pattern segmentation and pattern completion can both be observed even if only “learning” periods are considered.

Let us consider a concrete example. In two different simulations, Fig. 11 (top and bottom) we keep the parameters

and initial conditions of the network equal. Only the time-history of the input varies. In both cases, the network is excited by a segmented pattern during a certain amount of time. This time-interval starts at $t = 49$, for both simulations. However, it lasts longer in the case of Fig. 11 (bottom). In the first experiment, the short sniff conveying inhomogeneous stimulation does not suffice to change the structure of the network. The same three-fold initial pattern persists, as indicated by the response to a complete stimulation of the network after $t = 57$. In the second experiment, segmentation occurs and is responsible for a final ensemble of five patterns, out of the original three. Therefore, a kind of pattern completion takes place on the candidate learning period of Fig. 11 (top), but not on the one of Fig. 11 (bottom).

We suggest that some temporal gating mechanism of the afferent input, incorporating information from the instantaneous distributions of receptor input and of network activity, could enhance or minimize the influence of a particular input pattern on the network structure, according to requirements. Such a mechanism could allow the end of a strict separation between learning and recall phases.

At this point, we notice that the basic conflict—between pattern completion and the learning of segmentation structure—is only partially solved. In Fig. 11 (top), we see that a pattern can be stimulated in incomplete form, during a time-interval where synaptic change is allowed, without necessarily altering the synaptic weights. Yet, this experiment does not represent a situation where a pattern is only partially stimulated and, at the same time, the corresponding oscillators are clearly active (that is, with oscillatory activity well above the threshold ς).

To solve the segmentation/completion conflict, a criterion must be adopted which determines more systematically when structural segmentation is to be performed. This should be achieved for networks operating in unsupervised mode. The inspiration for the required criterion comes once more from biological processing. The basic mechanism is as follows. If a pattern, which is memorized in the distribution of synaptic weights, is excited in incomplete form during a short time-interval, then the network should interpret the incomplete pattern presentation as a degradation of the input due to random fluctuations. Hence the incompleteness of the pattern at the input level would be viewed by the system as incidental. The pattern should be fully restored, with all its oscillators displaying simultaneous activity. On the other hand, if the pattern is repeatedly and consistently presented in partial form, this can no longer be interpreted as incidental. Thus the dynamics of the system should force the synaptic weights to change in order to embody a new segmentation structure. In this manner, learning acquires a statistical character, which is closer to biology. In Section 11, we provide further discussion of this paradigm. In the analysis of corrupted patterns, we will note that analog properties of the system must eventually be considered.

11. Input modulations and change in connectivity

In the discussion of pattern completion, we considered only purely binary input patterns. The bits 0 and 1, in the stimulus, were multiplied by the homogeneous intensity 0.2. We considered the degradation of input patterns in a worst-case scenario. Namely, stimulation was completely absent from one or more receptor cells that should normally be stimulated if the pattern was integral. It was in this context that the conflict between segmentation learning and pattern completion was noted.

The absence of excitation from an individual receptor cell corresponded to a bit 0 attributed to that cell. This is probably an extreme choice. Some amount of excitation may be present, although not enough to elicit an oscillator's response. Moreover, the intensity of external stimulation might take any value between zero and a certain physiological maximum. For example, Fig. 8 might also be viewed as an illustration of pattern restoration when the receptor cells are excited with intensities between 32.5 and 135% of the "ideal" average value 0.2. Analog fluctuations of the input do not preclude a perfect binary response of the network.

The argument can be extended to the case of learning. Let us consider the enlargement in Fig. 8 (bottom). We see that, within the same pattern, there are periods of time during which not all oscillators are simultaneously active (but they do overlap during some time-interval). In the experiment of Fig. 8, the change in connectivity is permanently inhibited. If this was not the case, differential activity might elicit Hebbian synaptic change, as discussed in this paper. Thus one might observe the occurrence of structural segmentation into smaller patterns. Let us now consider the regimes where synaptic change is allowed, as introduced in Section 8. In the latter regimes, the occurrence of segmentation depends on the time-scale of synaptic adaptation. The latter could be adjusted through an adequate choice of parameter values. In the parameter range that we considered for our simulations, the network is very sensitive to analog modulations of the input intensity during learning. In living systems, a certain amount of random noise is unavoidable for the input strength. Therefore, a proper balance should be kept between the amplitude of the noise and the rate of synaptic change, in order to avoid spurious segmentations due to analog degradation. Furthermore, binary pattern degradation—that is, an input cell stimulated with zero intensity—can be viewed as a limit case of analog degradation. Due to the extreme sensitivity of the network equations to modulations of the input intensity, we discuss the completion/segmentation conflict in the context of analog corruption of patterns. Only a very small degradation of the input intensity is allowed. This is a numerical restriction which will desirably be relaxed by choosing different parameter values in future investigation of the model.

The following experiment illustrates our discussion. In the numerical simulations of Fig. 12, three patterns are

initially stored in the network. During a certain time-interval, the largest of the patterns, with 6 units, is stimulated with some amount of analog degradation: oscillator 11 receives 0.1% less input intensity than the other oscillators of the same pattern. We observe that the segmentation of the 6-bit pattern occurs when this learning period is long (Fig. 12, bottom), but not when it is short (Fig. 12, top). Notice that the longer learning period has twice the number of bursts of the shorter one. It can be concluded that the network tolerates a certain amount of analog degradation of the input, during a short time-interval, but it interprets a longer imposition of inhomogeneity of the input as a cue for segmentation.

12. Conclusions

We discussed associative memory by departing from the classical scenario where short-term memory, or pattern recall, is expressed by a stationary distribution of network activity. Instead, the information is encoded in time-dependent activities observed during short periods. Several *dynamical* features are of great importance to memory function, especially to the segmentation of information. Namely, the moment and the order of presentation of patterns to the network, modulations of intensity of the input, and the nature of the network's response to stimuli.

After pointing out a number of abstract as well as physiological aspects of memory, learning and retrieval, we investigated a particular model of memory with a natural application to olfaction. The model was presented in the essence in (Wang et al., 1990), but we introduced modifications along two major directions. Namely, the possibility of changes in the network structure resulting from a *learning* process, and the evaluation of *analog* properties of the input simultaneously with the binary interpretation. Previous studies had only considered purely binary input patterns. Thus the input intensity could only have the values 0 or 1 times some constant magnitude. In contrast, we allowed for intensity modulations of the input. By exploring the fine structure of the time-intervals of bursting, we proposed how the input modulations could elicit a higher-order coding of information. The network could simultaneously code features of the input in an analog and a binary way, thus showing enhanced processing capabilities. In this manner, analog features of the input provide for additional information structure.

We analyzed the dynamical phenomena leading to neuronal synchronization and bursting. The features of completion of corrupted patterns, and segmentation of the information, which are requirements of pattern retrieval, were illustrated with numerical simulations. In Section 8, we studied the operation mode in which the network can extract and memorize the different patterns that make up a composite one. This is achieved in an unsupervised manner. Learning takes place by the network being structured in

response to stimulus history alone. The conflict between pattern learning and pattern recall was considered, as well as the possibility of a modulation of the learning state. The importance of stimulus timing with respect to the internal dynamics of the network was pointed out. We proposed a mechanism allowing the end of a separation between learning and recall phases. This mechanism relies on time-gating of the input, and not on structural blocking of synaptic change. The consequences of analog degradation of the input intensity, during a learning phase, were noted.

In the simulations in our paper, distinct patterns do not present a superposition of active units. This facilitates the numerical implementation of the dynamical learning, while preserving desired properties of associative memory such as the capabilities of pattern completion and generalization. In the original model (Wang et al., 1990) there could be a superposition of active bits between different patterns, but dynamical learning was lacking. In that case, a superposition number of one active unit was reported. Notwithstanding, a high superposition rate does not seem to be mandatory in biological systems, and may even make for inefficient memory systems, see e.g. Rolls (1989) for a discussion on sparse coding and the advantages of avoiding interference between different dynamical patterns). Still, due to the importance of the possibility of pattern overlap, we consider that it should be given priority in further development of the model. We believe that further investigation should commence by a search of more adequate parameter values, especially concerning numerical stability, so that at least moderate pattern overlap be made possible. We only provide examples of learning featuring the breaking of excitatory associative connections. However, the model is also able to reunite patterns by re-establishing excitatory links between them. A mechanism of reunification is useful e.g. to remedy accidental segmentation caused by strongly "anomalous" input. The law of synaptic change includes the possibility of creating or even "rebuilding" excitatory links between groups of neurons. As it comes, the parameter region that we considered in the examples is more favorable for segmentating than it is for reuniting patterns. A more balanced behavior could be obtained by considering other regions of parameter space. In particular, the importance of external stimuli could be enhanced in situations where internal dynamics tends to preclude the re-binding of patterns. Overall, the re-binding mechanism necessarily has a Hebbian, statistical nature, and depends upon stimuli sequences with adequate properties. Specifically, excitatory links will be allowed to "grow" between patterns that repeatedly and consistently are stimulated together for a long enough period of time. This will be interpreted as the elementary patterns really being part of a same pattern. The fact that a sub-pattern has appeared isolated in the past will be regarded as a spurious phenomenon, "corrected" by more normal stimulus history at a later time.

We believe that the discussion in this paper could help

clarify important aspects of the primary and intermediate stages of sensory processing and pattern recall.

Acknowledgements

Carlos Lourenço acknowledges the support of the Portuguese Government (JNICT - PRAXIS XXI BD/3851/94, PBIC/TIT/2527/95).

References

- Baird, B. (1986). Nonlinear dynamics of pattern formation and pattern recognition in the rabbit olfactory bulb. *Physica D*, 22, 150–175.
- Bliss, T., & Collingridge, G. (1993). A synaptic model of memory: Long-term potentiation in the hippocampus. *Nature*, 361, 31–39.
- Bower, J. (1994). The modulation of learning state in a biological associative memory: an in vitro, in vivo, and in computo study of object recognition in mammalian olfactory cortex. In L. Niklasson & M. Boden (Eds.), *Connectionism in a broad perspective* (pp. 157–163). New York: Kluwer.
- Damasio, A. (1989). The brain binds entities and events by multiregional activation from convergence zones. *Neural Computation*, 1, 123–132.
- De Schutter, E., & Bower, J. (1993). Sensitivity of synaptic plasticity to the Ca^{2+} permeability of NMDA channels: a model of long-term potentiation on hippocampal neurons. *Neural Computation*, 5, 681–694.
- Doya, K., Yoshizawa, S. (1989). Memorizing oscillatory patterns in the analog neural network. *Proceedings of the International Joint Conference on Neural Networks*, Washington, 1989 (pp. 27–32).
- Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., & Reitboeck, H. (1988). Coherent oscillations: a mechanism of feature linking in the visual cortex? *Biological Cybernetics*, 60, 121–130.
- Engel, A., König, P., Gray, C., & Singer, W. (1990). Stimulus-dependent neuronal oscillations in cat visual cortex: intercolumnar interaction as determined by cross-correlation analysis. *European Journal of Neuroscience*, 2, 586–606.
- Engel, A., Kreiter, A., König, P., & Singer, W. (1991). Synchronization of oscillatory neuronal responses between striate and extrastriate visual cortical areas of the cat. *Proceedings of the National Academy of Sciences USA*, 88, 6048–6052.
- Engel, A., König, P., & Singer, W. (1991). Direct physiological evidence for scene segmentation by temporal coding. *Proceedings of the National Academy of Sciences USA*, 88, 9136–9140.
- Freeman, W. (1978). Spatial properties of an EEG event in the olfactory bulb and cortex. *Electroencephalography and Clinical Neurophysiology*, 44, 586–605.
- Freeman, W. (1991). The physiology of perception. *Scientific American*, 264, 78–85.
- Freeman, W. (1992). Tutorial on neurobiology: from single neurons to brain chaos. *International Journal of Bifurcation and Chaos*, 2, 451–482.
- Freeman, W. (1994). Chaos in the CNS: theory and practice. In F. Ventriglia (Ed.), *Neural modeling and neural networks*, (pp. 185–216). New York: Pergamon Press.
- Freeman, W., & Baird, B. (1987). Relation of olfactory EEG to behavior: spatial analysis. *Behavioral Neuroscience*, 101, 393–408.
- Freeman, W., Yao, Y., & Burke, B. (1988). Central pattern generating and recognizing in olfactory bulb: a correlation learning rule. *Neural Networks*, 1, 277–288.
- Gray, C., & Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences USA*, 86, 1698–1702.
- Gray, C., König, P., Engel, A., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit intercolumnar synchronization which reflects global stimulus properties. *Nature*, 338, 334–337.
- Gray, C., Engel, A., König, P., & Singer, W. (1990). Stimulus-dependent neuronal oscillations in cat visual cortex: receptive field properties and feature dependence. *European Journal of Neuroscience*, 2, 607–619.
- Guckenheimer, J., & Holmes, P. (1983). *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*, Springer series in applied mathematical sciences, 42. New York: Springer.
- Hasselmo, M., & Bower, J. (1992). Cholinergic suppression specific to intrinsic not afferent fiber synapses in rat piriform (olfactory) cortex. *Journal of Neurophysiology*, 67, 1222–1229.
- Hasselmo, M., Anderson, B., & Bower, J. (1992). Cholinergic modulation of cortical associative memory function. *Journal of Neurophysiology*, 67, 1230–1246.
- Hebb, D. (1949). *The organization of behavior*, New York: Wiley.
- Hendin, O., Horn, D., & Tsodyks, M. (1998). Associative memory and segmentation in an oscillatory neural model of the olfactory bulb. *Journal of Computational Neuroscience*, 5, 157–169.
- Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences USA*, 79, 2554–2558.
- Hopfield, J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences USA*, 81, 3088–3092.
- Hoshino, O., Kashimori, Y., & Kambara, T. (1998). An olfactory recognition model based on spatio-temporal encoding of odor quality in the olfactory bulb. *Biological Cybernetics*, 79, 109–120.
- Kanter, E., & Haberly, L. (1990). NMDA-dependent induction of long-term potentiation in afferent and association fiber systems of piriform cortex in vitro. *Brain Research*, 525, 175–179.
- Laing, D., & Frances, G. (1989). The capacity of humans to identify odors in mixtures. *Physiology and Behavior*, 46, 809–814.
- Laing, D., Panhuber, H., Willcox, M., & Pittman, E. (1984). Quality and intensity of binary odor mixtures. *Physiology and Behavior*, 33, 309–319.
- Landfield, P., & Deadwyler, S. (1988). *Long-term potentiation: from biophysics to behavior*, New York: Alan Liss.
- Laurent, G., & Davidowitz, H. (1994). Encoding of olfactory information with oscillating neural assemblies. *Science*, 265, 1872–1875.
- Li, Z., & Hopfield, J. (1989). Modeling the olfactory bulb and its neural oscillatory processings. *Biological Cybernetics*, 61, 379–392.
- von der Malsburg, C. (1983). How are nervous structures organized? Synergetics of the brain. In E. Basar & F. Flohr & H. Haken & A. Mandell (Eds.), *Proceedings of the International Symposium on Synergetics* (pp. 238–249). Berlin: Springer.
- von der Malsburg, C., & Schneider, W. (1986). A neural cocktail-party processor. *Biological Cybernetics*, 54, 29–40.
- von der Malsburg, C., & Buhmann, J. (1992). Sensory segmentation with coupled neural oscillators. *Biological Cybernetics*, 67, 233–242.
- Manabe, T., & Nicoll, R. (1994). Long-term potentiation: Evidence against an increase in transmitter release probability in the CA1 region of the hippocampus. *Science*, 265, 1888–1892.
- Moser, M. -B., Trommald, M., & Andersen, P. (1994). An increase in dendritic spine density on hippocampal CA1 pyramidal cells following spatial learning in adult rats suggests the formation of new synapses. *Proceedings of the National Academy of Sciences USA*, 91, 12673–12675.
- Murthy, V. & Fetz, E. (1991). Synchronized 25–35 Hz oscillations in sensorimotor cortex of awake monkeys. *Neuroscience Abstracts*.
- Rolls, E. (1989). The representation and storage of information in neuronal networks in the primate cerebral cortex and hippocampus. In R. Durbin & C. Miall & G. Mitchison (Eds.), *The computing neuron* (pp. 125–159). Workingham, UK: Addison-Wesley.
- Schillen, T., & König, P. (1994). Binding by temporal structure in multiple feature domains of an oscillatory neuronal network. *Biological Cybernetics*, 70, 397–405.
- Stanton, P., & Sejnowski, T. (1989). Associative long-term depression in the hippocampus induced by Hebbian covariance. *Nature*, 339, 215–218.

- Stent, G. (1973). A physiological mechanism for Hebb's postulate of learning. *Proceedings of the National Academy of Sciences USA*, 70, 997–1001.
- Wang, D. L., Buhmann, J., & von der Malsburg, C. (1990). Pattern segmentation in associative memory. *Neural Computation*, 2, 94–106.
- Wehr, M., & Laurent, G. (1996). Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature*, 384, 162–166.
- Xu, L., Anwyl, R., & Rowan, M. (1998). Spatial exploration induces a persistent reversal of long-term potentiation in rat hippocampus. *Nature*, 394, 891–894.