

Rate coherence and event coherence in the visual cortex: a neuronal model of object recognition

Hartmut Neven¹ and Ad Aertsen^{1,2}

¹ Max-Planck-Institut für Biologische Kybernetik, Spemannstrasse 38, W-7400 Tübingen, Federal Republic of Germany

² Institut für Neuroinformatik, P.O. Box 102184, Ruhr-Universität-Bochum, Universitätsstrasse 150, W-4630 Bochum, Federal Republic of Germany

Received January 14, 1992/Accepted in revised form March 25, 1992

Abstract. We propose a function-oriented model of the visual cortex. The model addresses an essential task of the visual system: to detect and represent objects. These are defined as sets, which reappear in the input with invariant inner relations. A network, incorporating an idealized description of anatomical and physiological data, is presented with a movie showing various moving objects. In the course of time, as a result of Hebbian plasticity, a connection scheme develops which embodies in its forward and lateral connections the information necessary to perform the operations involved in object recognition. We demonstrate that coherent neural activity can exploit this information. Two types of coherence have to be distinguished in this respect. Rate coherence performs invariance operations and association, while event coherence accomplishes segmentation tasks. The model reproduces and explains experimental findings made both in physiological recordings from the visual cortex and in psychophysical studies.

1 Introduction

In spite of an overwhelming body of experimental data, our understanding of how the mammalian visual system performs object recognition is still poor. Even the exact definition of the problem still poses difficulties. Yet, there are numerous theoretical and experimental results indicating how parts of the problem could be solved. To mention just a few:

– The classical work of Hubel and Wiesel has shown that an elementary level of description, enabling the cortex to handle visual information, can be given in terms of feature sensitive cells.

– Recent work, focussing more strongly on the time structure of the neural signals, emphasized the

phenomenon of synchronisation in the cortex, which, by its nature, is able to solve the segmentation problem.

– Theoretical work on associative memory, starting from Hebb's ideas, has revealed how neurons might cooperate in order to store information and to perform pattern completion.

Other issues, such as the invariance problem, still wait for a biologically plausible explanation. Also the integration of these various aspects into a single model is still an open question. The present study undertakes to develop a functional model of the visual cortex, which integrates the basic anatomy and known neural mechanisms, and provides biologically plausible suggestions for their role in object recognition.

2 Specification for the task of object recognition

The basis for our work is an idealisation of the visual situation (for pictorial guidance see Fig. 1). The world consists of *elements* and *space-time* relations between them. Some of the spatial relations are variable over time, while others are not. A set of points whose internal spatial relations do not change over time is called a *visual object* (we restrict our discussion to rigid bodies). The objects are projected onto the retina. Since, by definition, an object is entirely defined by its inner relations, the absolute position of the projection on the retina should be of no importance when determining the nature of the object. An equivalence class, the elements of which are equal up to a transformation, is called an '*orbit*', a term adopted from algebra (Fig. 1). One should be aware that at the level of projections on the retina, we do not have strictly separated orbits, e.g. when looking at a pyramid from below, it may degenerate into a projection of a cube. These cases being extremely rare, however, we ignore them as statistical pathologies, and still speak of orbits of object projections. Now the task of object recognition by the visual system can be specified as follows. The visual system has to represent the orbits of the various object projections and, upon presentation of a picture, to

Correspondence to: Ad Aertsen, Institut für Neuroinformatik, P.O. Box 102184, Ruhr-Universität-Bochum, Universitätsstrasse 150, W-4630 Bochum, Federal Republic of Germany




Fig. 1a, b. Objects, projections and representations. **a** The objects cube and pyramid are sets with internally stable spatial relations. The pyramids at different positions form an orbit under congruence transformations in euclidean space. **b** Projection on the retina. Notice that now orbits are fuzzy sets, only separated statistically. **c** Representations of the projection orbits

activate the representations of those orbits whose elements created the present set on the retina. (This specification was motivated by philosophical, mathematical and psychological considerations: Democritus, about 400 B.C.; Klein 1872; Hebb 1949).

To give a brief overview, we propose that the visual system implements a visual program that obeys the above specification by performing the following steps. First, it generates descriptors for sets repeatedly appearing in the input with stable inner relations. Secondly, it treats those descriptors which encode features belonging to the same orbit as essentially the same: *invariance operation*. Thirdly, it binds together the orbits of feature detectors whose features appear together more often than one would predict under the assumption of independence: *scene segmentation*. Finally, it uses these groups as clues for *association*.

3 Forward and lateral connections in the primary visual cortex

In search of a biological implementation of the functional specification, we first studied the following question. Consider a network of neurons that mimics the anatomy of the primary visual cortex, with synapses displaying Hebbian plasticity. Given that this network sees a 'natural' input consisting of a movie showing contours of moving objects. What sort of connection scheme will develop?

On its way to the primary visual cortex, the retinal input is processed by the center-surround system (for a review see e.g. Kandel and Schwartz 1985). In view of the fact that the cells in this early processing stage show high firing rates in regions where the local variation of light intensity is high, it has been suggested that this system performs the operation of contour extraction (Marr and Ullman 1981). This operation yields a luminance-invariant and data-reduced version of the retinal image, which serves as input to the primary visual cortex. To model this cortical input, we proceeded as indicated in Fig. 2a. First, we set up a database of fifty two-dimensional line patterns, such as circles, squares and question marks. A program selects randomly up to three patterns from this database, places them at random initial positions in a plane with a hexagonal grid




Fig. 2a-c. Model of the primary visual cortex. **a** Scene from the input movie. The hexagons α , β , γ denote receptive fields. **b** Layout of area 17. Neurons at four cortex positions are shown, every pyramid denotes a neuron. Observe the multiplicity ($N = 6$) of receptive fields: the neurons at one position have equal feedforward receptive fields, their profiles may develop differently, though. Note that this multiplicity bears no relation to cortical layering. The cuboid box denotes the competition volume which encompasses the full depth of the cortical sheet. **c** The firing of the neurons is described by an interval of 10 consecutive '1'-signals

structure, and shifts each of them independently along a randomly generated path in this input plane. A scene of the resulting 'movie' is shown in Fig. 2a. Per unit of time, each pattern can move over 0, 1 or 2 grid positions. After a randomly chosen time interval of up to forty time steps, the next set of patterns is selected and the procedure is repeated. In the generation of the

pattern trajectories, we confined ourselves to translations in order to keep the model as simple as possible. Moreover, it is possible to rephrase scale and rotation invariance (at least in the absence of other transformations) in terms of translation invariance by using an appropriate logarithmic map; such a map is presumably implemented in the projection from the retina to the primary visual cortex (Schwartz 1977).

Topology of the model. Our primary visual cortex model (Fig. 2b) reflects basic anatomical and physiological facts. The afferents are excitatory and topologically ordered. The cells with modifiable synapses are excitatory and densely connected laterally, while inhibition acts locally and is assumed to be unmodifiable. The cortical sheet is composed as a hexagonal grid of $L \times L$ positions, in each of which there are N neurons with identical receptive fields, indicated by the grey hexagons in Fig. 2a. In our simulations, a receptive field consists of 37 positions, arranged in a hexagon to mimic a circle. For the other parameters we chose $L = 20$ and $N = 6$. (Note that N is solely intended to model the multiplicity of receptive field profiles at each location; in particular, it has nothing to do with cortical layering.)

Dynamics of the activity. Upon presentation of a scene, a volume of $l \times l \times N$ neurons (in our case $l = 5$; see the dark cuboid in Fig. 2b) is stepped through the cortical sheet over all L^2 positions. Each such presentation with the associated L^2 update positions counts as one time step. At each of these positions, the potentials π_i of the neurons in the cuboid are computed according to

$$\pi_i = \sum_j s_{ij} \chi_{ij}, \quad (1)$$

in which $\chi_{ij} \in \{0, 1\}$ denotes the input at position j in the receptive field of neuron i . The synaptic weights s_{ij} are positive and normalized according to $\sum_j s_{ij} = 1$. The neurons with the strongest potential is selected to fire. The firing of a neuron is described by a binary process: during the following ten time steps, the neuron emits a '1'-signal, modeling a short epoch of increased firing rate. In a number of simulations, this duration of increased firing was varied to include shorter intervals, down to only one time step.

Modification of connections. Connection strengths of the firing neurons are updated according to a normalized Hebb rule (Hebb 1949; Oja 1982):

$$s_{ij}(t+1) = \frac{s_{ij}(t) + \gamma \pi_i(t) \chi_{ij}(t)}{\sum_j s_{ij}(t) + \gamma \pi_i(t) \chi_{ij}(t)}, \quad (2)$$

in which the gain γ is an order of magnitude smaller than the average synaptic weight. Thus, synapses are modified by a competitive learning algorithm, incorporating the effects of local inhibition (for a description of competitive learning see Rumelhart et al. 1986).

3.1 Development of forward connections

After presentation of a sequence of 30,000 scenes, the initially uniform forward connections had developed

into inhomogeneous 'receptive field profiles'. Figure 3 shows these profiles for the first of the N 'layers' of model neurons. Observe that the receptive field profiles are mostly bar-shaped. This basically reflects the fact that line drawings locally look like bars. The distribution of synaptic weights is very stable. Already after the presentation of 10,000 scenes it had essentially converged to the asymptotic result in Fig. 3. Generally at that early stage, a receptive field profile already showed a bar in the same orientation, the only difference being that the contrast between strong and weak synaptic weights was less pronounced. About 25% of the receptive field profiles remained uniform throughout the simulation. Finally, the number of profiles that were neither bar-shaped nor remained uniform was less than 1%. Inspection of the receptive field profiles of the remaining $N - 1$ neuron layers not shown here reveals that at each position several different orientations are represented. For equal cell potentials (as is initially the case with uniform forward connections), our implementation of the competition algorithm introduced a bias in that it preferentially selected the higher situated of the $l^2 N$ neurons within the competition volume. Simulations with random initial conditions, however, lead to essentially the same results.

A more detailed inspection of the emerging features reveals a number of further interesting aspects. First, it is important to notice that for almost all neurons, the receptive field profile represents (part of) the projection of one object only, even though the network was trained with scenes which generally contained the projections of several objects simultaneously. This is due to the simple statistical effect that the probability for any particular composite input pattern to arise is considerably lower than that of a pattern evoked by the projection of just a single object. For example, the occurrence of the particular composite pattern seen by the neurons with receptive field β (Fig. 2a) requires two conditions to be fulfilled simultaneously. A necessary, though not sufficient, condition is the presence of the two individual patterns, but in addition they have to appear in the appropriate configuration. Since the trajectories of the individual patterns are mutually independent, the probability of such a combined event is the product of the probabilities of the singular events, i.e. that each of the patterns individually occurs at the right position. Thus, a receptive field profile which captures 'inter-features' as in β (Fig. 2a) has a much smaller probability to develop than one which captures 'intra-features' as in α or γ . As a consequence, almost only those receptive field profiles will develop which describe the projection of an object, i.e. a pattern with invariant inner relations. We observe a similar further differentiation within the class of intra-feature receptive fields. The emerging features here are not specific to any particular object, but rather to broad subsets of the whole pattern ensemble. Again, since a higher order features as in α is more specific to a certain pattern and to the position within that pattern than a simple feature as in γ , it has a lower probability to develop.

Summarizing, the feedforward connections that




Fig. 3. Distribution of the feedforward synaptic strengths in the primary visual cortex model after training. Each hexagon represents the feedforward receptive field profile of one neuron, every little circle within a hexagon denotes a synapse and its grey value the synaptic strength. The grey scale is linear and runs from 0.0 (white) to 0.1 (black)

develop in the network in the course of a competitive learning process embody the most significant part of the spectrum of probabilities for the various possible input pattern configurations that the network was exposed to. As a consequence, the various types of emerging receptive field profiles reflect the correlation structure of the input space, characterized by its various classes of intra- and inter-feature relations. Thus, the first task of our proposed visual program has been accomplished: repeatedly appearing features in the input space are represented by neurons in their feedforward connections.

3.2 Development of lateral connections

In order to analyse the developing correlation structure across the cortical sheet, we computed the covariances of the output activity of the model neurons over the last 10,000 timesteps, i.e. after the feedforward network had settled into its asymptotic configuration and, consequently, the covariances were most clear. Figure 4 shows, in a grey density code, the resulting covariances for a selected reference neuron (at the position indicated by the arrow) to all other ones. The analysis for twenty other reference neurons, randomly selected among the class of bar-shaped receptive field-profiles, gave essentially the same results. This implies that the characteristics of the covariance distribution are neither dependent on the position of the reference neuron, nor on its preferred bar orientation. (In order to avoid a

multiplicity of pictures, we overlaid the covariances with each of the $N = 6$ neurons at a single position into one display; details are given in the caption to Fig. 4.)

The spatial distribution of covariance exhibits two distinct phenomena. First, the covariances become gradually weaker with increasing distance to the reference neuron; a control run with neurons firing each time only during one time step instead of ten showed the same phenomenon. This dependence on distance reflects the spatial extent of correlation in the input space, which is independent of the duration of the firing interval. Second, the strongest covariances occur between neurons which share the same preferred bar orientation; they become gradually weaker with increasing angular difference. A more detailed inspection shows that we have to distinguish between two cases: collinear and non-collinear bars. High covariance between neurons with collinear preferred bars partly again reflects continuity of the input, and can still be seen, albeit weaker, with a reduced firing interval. For neurons with parallel but non-collinear preferred bars, the strong covariances disappear completely with shortening of the firing interval, indicating that in this case the duration of firing is of crucial importance; we will return to this issue later. Finally, Fig. 4 reveals that a considerable number of neurons develops overlapping receptive field profiles, and that these neurons show high covariances. Thus, redundant coding, i.e. multiple




Fig. 4. Covariances between the firings of the neurons across the cortical sheet in the primary visual cortex model. The receptive field profiles are now characterized only by the strongest synapses. The grey value of the hexagon for each neuron is proportional to the covariance with the reference neuron (the one with the black horizontal bar at the position marked by the arrow). The six hexagons, corresponding to the neurons at one position, are drawn on top of each other, starting with the lightest grey. The grey scale is linear and runs from 0.0 (white) to 0.01 (black)

representation of features, arises quite naturally as an emerging property in our model. Consequently, the presentation of input to the model network will generally give rise to the activation of groups of neurons; these groups form the starting point for any lateral dynamics within the network.

Notice that in our model we did not incorporate the development of lateral connections explicitly. Instead, we studied the covariances of the output activities of the developing model neurons. We assert that these two approaches, of which the latter is obviously highly advantageous for computational reasons, are essentially equivalent for the present purpose. This equivalence is based on the assumption that a Hebbian learning rule of the kind as expressed in (2), when applied to the connectivity among the model neurons within the cortical sheet, would translate the profile of covariances among these neurons' activities, exemplified in Fig. 4, into a corresponding topography of lateral connection strengths.

4 Two types of coherence in neural activity exploit the information contained in the lateral connectivity

Having assigned the first task of our proposed visual program to the feedforward connections, we will now show that the lateral connections contain the information necessary to tackle the remaining tasks, and that coherent neural activity can exploit this information. Physiological evidence leads us to distinguish between two types of coherence, based on the time constants

involved: *rate coherence*, i.e. correlation of firing rates, resulting in a broad peak (in the order of 100 ms or more) in crosscorrelograms of the activities of simultaneously recorded neurons, and *event coherence*, i.e. correlated individual spikes, reflected in narrow (several ms) correlation peaks (e.g. Abeles 1982, 1991; Gerstein et al. 1983; Gerstein 1988; Krüger 1983; Ts'o and Gilbert 1986; Aertsen et al. 1987, 1992; Aertsen and Gerstein 1991; Krüger and Mayer 1991; Vaadia et al. 1991; Vaadia and Aertsen 1992). Furthermore, we assume that only strong connections are able to transmit firing rates, while the effect of weak connections is to synchronize the events in a population of neurons (for a theoretical justification see Abeles 1982, 1991; Boven and Aertsen 1990; Aertsen and Preissl 1991). Rate coherence recruits new neurons into an already active group, event coherence organizes the active group into internally coherent, mutually incoherent subgroups. Or, in the language of set theory: rate coherence performs the operation of union, event coherence that of intersection.

We will demonstrate our ideas by parsing the presumed lateral dynamics into three components, and investigate each one separately in an appropriately formulated model. This avoids the necessity to explore the full lateral dynamics at once, which would require the difficult discussion of a network with more realistic model neurons and connections with delays. The first model uses rate coherence between features belonging to the same 'orbit' to perform invariance operations. In the second model, event coherence between features which appear often together accomplishes segmentation

tasks. The third model implements associative memory on the basis of rate coherence between repeatedly co-occurring features. Finally, we will integrate these three mechanisms into a single functional model.

4.1 Rate coherence between feature detectors belonging to the same 'orbit' performs invariance operations

The set of feature detectors that develops with an input consisting of patterns moving according to a transformation group will generally reflect the symmetry introduced by this particular group. For the translation group this means that a feature detector will have equal probability to appear, independent of position. The distribution of receptive field profiles that develop in our primary visual cortex model (Fig. 3) confirms this. When now a certain pattern feature is moved across the input plane, the associated feature detectors in the network will raise their firing rates one after the other, according to a path that mimics the feature trajectory in the input space (see also Braitenberg 1985, 1986). Moreover, the time intervals of increased firing rate for activated neurons at neighboring positions will overlap to an extent determined by the spike burst duration and the velocity of movement. This overlap leads to the correlation of activity between spatially shifted features which we observed in the covariances (Fig. 4), and which disappeared when the burst duration was decreased. During these overlapping epochs of increased firing rate, the Hebbian learning rule strengthens the intervening connections; a Hebb rule with a time window would produce a similar effect. Thus, once the connections have been strengthened, these neurons will tend to fire simultaneously, thereby destroying the information about where the activating feature actually appeared: the result is translation invariance.

Orbit assemblies. The invariance operation was idealized in the following model. A set of N neurons, encoding different features, is placed at every position in the network (Fig. 5b). The neurons with corresponding features at different positions form an ideal assembly: if one neuron fires, the whole assembly fires (for a discussion of various assembly definitions see Gerstein et al. 1989). Each of the assemblies encompasses the features belonging to one orbit under translation, hence we call them *orbit assemblies*. Upon presentation of a scene (Fig. 5a), the lateral dynamics produce the following result. Proportional to the strength of the input activity, a certain fraction of the orbit assemblies is allowed to fire. In our case this fraction is 50% of the input strength, being simply the number of 'on'-pixels. For each assembly we first determine which member neuron has the largest potential π . Those assemblies will be activated which contain the members with the highest potentials (Fig. 5c). In order to emphasize the essence of the present discussion, i.e. the concept of the orbit assembly, rather than the specific shape of the receptive field profiles, we implemented the receptive fields as random masks. Each of the N neurons at one position was assigned a different receptive field, consisting of




Fig. 5a-c. Model of the invariance operation. **a** Input scene, showing a bottle, an apple and a cake. **b** Idealization of the first primary visual cortex model. Neurons at three positions are symbolized by squares that show the random receptive field profiles. The depth of the cube represents the multiplicity of neurons with identical receptive field location. A neuron receives input only from those pixels in the grey area marked in (a) which correspond to the black regions in its receptive field. The neurons with corresponding receptive field profiles at each of these positions form an orbit assembly. **c** The orbit assemblies that are activated by presentation of the scene in (a). Every active orbit assembly is represented by a single black square.

$R^2/2$ pixels in an $R \times R$ square; the resulting set is replicated for all positions in the network (Fig. 5b). In order to avoid the need of a time consuming training with a movie, at the same time to obtain enough specificity without having a multi-layered network, we trained as follows. We presented a sequence of ten

pictures, each one showing an object from a breakfast table; the sequence was presented only once. When an assembly was activated, all its member neurons were updated in identical fashion, copied from the neuron with the highest potential. Learning was applied to enhance the responses to intra-features and to reduce the responses to inter-features. The learning rule (2) stayed the same, only the gain was increased to accommodate for single trial learning. In our simulations, γ was in the order of the average synaptic weight, N was 512, R was 4 and the presented pictures consisted of 32×32 pixels.

After training, we presented scenes containing one or more of the objects from the training set (Fig. 5a). Figure 6 shows some of these pictures on the left and the activated assemblies on the right. The first and second picture show a bottle at two different positions. Observe that in both cases the same orbit assemblies respond. The apple in the third example fires different assemblies. Finally, the last example shows that a scene composed of both objects causes essentially the union of the two sets of assemblies to fire. Notice, though, that we have errors in the sense that one assembly belonging to the union is missing, while another one, not belonging to it, is active (arrows in bottom right panel in Fig. 6). This might be due directly to the fact that part of the bottle was occluded, so that the corresponding features are missing, or that certain detectors respond to a feature made up by the composite pattern. Another, indirect, explanation might be in the competition scheme: the activation of one feature detector might suppress the activation of another one, which would otherwise have made it, or the other way around. Apart from these errors, to which we will return in our discussion on association (Sect. 4.3), the network performs an invariance operation, able to work on unsegmented images. This ability to deal with unsegmented images is a necessary property for any invariance operation to fulfil in order to be of practical use (e.g. Marr 1982).

The need for higher visual cortices. It is a well known fact that the spatial correlation in the visual input decreases with distance. This implies that, in order to have enough correlation for learning, the visual system has to use local feature detectors. However, for local receptive fields, the features which develop in a single layer network will not be very specific for any particular object. Since associative devices with discriminative ability can handle only a small overlap between stored patterns, we assume that more specific and, hence, more global features develop successively in higher visual areas. This is bound to happen, since the higher visual areas have basically the same structure as the primary one, i.e. each area receives a topology-preserving projection from the previous one, and the lateral dynamics in each area spread the activity evoked by the forward input only over a restricted distance (e.g. Tusa et al. 1981; Van Essen 1985). Not every neuron has to be a projection neuron, however. On the contrary, a partial but random selection of neurons providing the projection seems more efficient, because it converges the

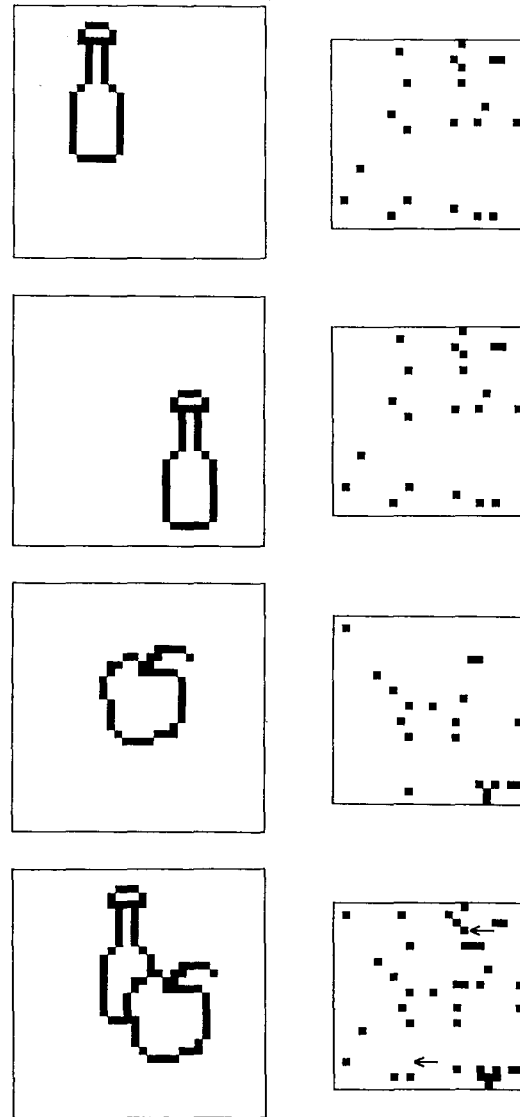


Fig. 6. Performance of the invariance model. Left column: a number of presented scenes, showing different configurations of a bottle and an apple. Right column: the responding orbit assemblies. Observe that the sets of orbit assemblies which are activated by the upper two scenes (bottles, shifted in space) are identical. The arrows in the bottom right panel denote the positions where the actual activity evoked by the combination of two objects differs from the union of the activities evoked by the objects presented in isolation

information carried by the orbit assemblies. Since successive representations of the visual field tend to become smaller and smaller in terms of cortical surface, i.e. they involve decreasing numbers of neurons, the characteristics of subsequent projections have to be such that at each step the multiplicity of feature detectors is reduced. Since in our application we have only a limited number of objects, we can allow ourselves to keep the discussion simple by using only one layer to achieve invariance. This one layer, however, includes lateral interactions: we regard the lateral dynamics as a necessary addition to the conventional, purely feedforward dynamics (for a discussion of problems with invariances in the perceptron

see Minsky and Papert (1988), and in the cognitron (Fukushima 1988) see Menon et al. (1988) and Barnard (1990)). Finally, another important virtue of higher areas is the emergence of associative structures, to which we will return later.

4.2 Event coherence between neurons encoding repeatedly co-occurring features performs image segmentation

By definition, the criterion the visual system should use to group features belonging to the same object is to determine which features co-occur more often than predicted on the basis of independence. The lateral connections between neurons with different preferred features contain this valuable information. In our model we perform the grouping by organizing those elements which display coincident firing into temporally coherent assemblies (von der Malsburg and Schneider 1986). For a more elaborate discussion of the *time labeling* idea see von der Malsburg (1981, 1986); we give here one additional theoretical argument. We can think of the cortex as a database, representing the elements in the world and their relations. The primary relations in the world are space-time relations. There are sets whose spatial relations are most important (e.g. rigid objects), in others the temporal relations are dominant (e.g. a rhythm), while in yet others we have a mixture of the two (e.g. the stations of a journey). In order to keep overall consistency, a relation-preserving mapping seems an appropriate strategy: features which occur at neighboring moments in time should be mapped onto neurons which fire at neighboring moments in time.

The generation of synchronized events. Synchronized events have been observed in the cortex in many different guises. Inspired by recent physiological observations (Eckhorn et al. 1988, Gray and Singer 1989; Gray et al. 1989; but see already Hubel and Wiesel 1965 and Gerstein 1970), we used coupled oscillators as a numerically and analytically convenient vehicle to produce such temporally coherent events. Other phenomenological oscillator models, primarily designed to reproduce the findings on oscillations in visual cortex, are given by Sompolinsky et al. (1990, 1991) and Schuster (1991); microscopic models were developed by e.g. Sporns et al. (1989); König and Schillen (1991); Schillen and König (1991), and Eckhorn et al. (1990, 1991). We emphasize, however, that many systems displaying synchrony could do the same. Alternative possibilities that have been proposed are the 'synfire chain' (Abeles 1982, 1991), 'correlation assemblies' (Johannesma et al. 1986; Aertsen et al. 1987; Gerstein et al. 1989) and 'pattern generating assemblies' (Abeles and Gerstein 1988). Finally, coupled chaotic elements (Kaneko 1990; Bauer and Martiensen 1991) would provide yet another means.

We assume that invariance was already achieved and start with a set of patterns $\{\xi_n^p\}_{n \in \{1, \dots, N\}}^p$ with $\xi_n^p \in \{0, 1\}$. P indicates the number of patterns, N

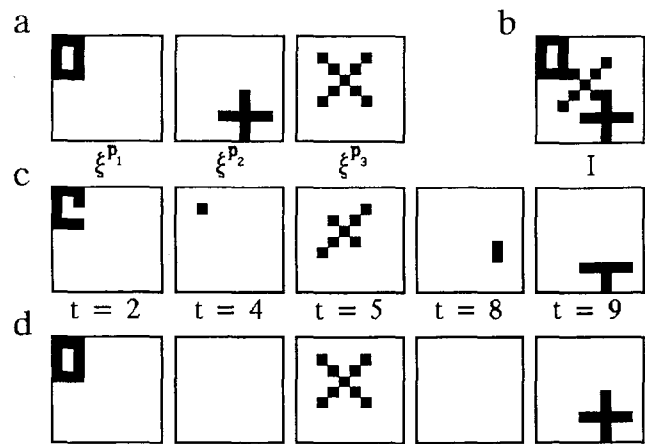


Fig. 7a-d. Performance of the segmentation model. **a** Three separate patterns forming the composite input **b** of the oscillator layer. **c** Some instances of the output of the oscillator layer. Observe that through the interplay of synchronisation and desynchronisation, the composite scene is segmented in time, however, at the expense of 'damage' to each of the segments. **d** The corresponding output of the associative memory. Notice the completion of the retrieved patterns, which match the constitutive scene elements, shown in (a)

the number of pattern elements. Note that the patterns here refer to patterns in some feature space, not primarily in ordinary pixel space. The probability that $\xi_n^p = 1$ is typically smaller than 0.1 in order to ensure small overlap (*sparse coding*). Out of a subset S of patterns (Fig. 7a) we form the composite input pattern

$$I = \bigvee_{p \in S} \xi^p \quad (\text{Fig. 7b}),$$

where \bigvee denotes the 'or'-operator. To every input element is assigned a neural element, modeled by an oscillator described by a phase variable φ which obeys the differential equation $d\varphi/dt = \omega$. Strictly speaking this is a *phasor*, i.e. an oscillator with constant amplitude. For simplicity we assume that all oscillators have the same intrinsic frequency ω . The oscillators are fully connected by diffusive coupling terms, i.e. only depending on the phase difference $\varphi_j - \varphi_i$ (Kuramoto 1984). We consider only the dynamics of neural elements whose input is one ($I_i = 1$). This leads to the following set of coupled differential equations:

$$d\varphi_i/dt = \omega \sum_{\{j | I_j = 1\}} J_{ij} \sin(\varphi_j - \varphi_i) \quad \{i | I_i = 1\}. \quad (3)$$

To be consistent with our first model, we treat the sample patterns like a time series and define the coupling constants J_{ij} by the covariances:

$$J_{ij} = \text{Cov}(\xi_i, \xi_j) = 1/P \sum_p \xi_i^p \xi_j^p - 1/P \sum_p \xi_i^p \cdot 1/P \sum_p \xi_j^p. \quad (4)$$

A positive coupling strength, arising between elements which are often co-active in the set of input patterns, synchronizes the corresponding oscillators. In contrast, a negative coupling, arising between elements which are

negatively correlated in the input patterns, desynchronizes them. Each oscillator is made to produce an event whenever it passes through a small angle segment according to

$$O_i^{\text{osc}}(t) = 1 \text{ if } \varphi(t) < \varepsilon \quad (5)$$

$$0 \text{ otherwise.}$$

Some instances of the time series that results when the oscillator layer is presented with the composite input pattern in Fig. 7b are shown in Fig. 7c. We observe that different subgroups of oscillators are co-active at different times. In particular, oscillators encoding elements belonging to one of the elementary patterns tend to emit synchronized signals. However, these elementary patterns do not appear completely: parts of the patterns are missing.

For an explanation we turn to the underlying dynamics. These are a superposition of an overall periodic movement and a relative movement. After a transformation $\varphi \rightarrow \psi = \varphi - \omega t$, which subtracts the global periodic movement, we obtain for the relative dynamics

$$d\psi_i/dt = \sum_j J_{ij} \sin(\psi_j - \psi_i). \quad (6)$$

Notice that the relative movement is governed by a Hopfield-type energy function (Hopfield 1982):

$$H = \sum_j J_{ij} \cos(\psi_i - \psi_j) = \sum_j J_{ij} \mathbf{s}_i \cdot \mathbf{s}_j. \quad (7)$$

The only difference is that here the ‘spin’ variables \mathbf{s}_i are unit vectors, rotating on a circle (in the physical literature known as XY-model), instead of scalars confined to the values -1 and $+1$. In this sense, the present model is the most modest step one can make from a fixpoint attractor network towards networks with an attractor at the next level of complexity, i.e. a limit cycle (for fixpoint attractor networks see Amit 1989). When the system relaxes according to the gradient dynamics, a partition is introduced on the set of neural elements: ‘to be synchronized’ is basically an equivalence relation. It should be noted that the internally coherent subsets resulting from this partition are disjoint, i.e. non-overlapping. Obviously, this will create problems when the input patterns are overlapping: since an element can only belong to one equivalence class, a composite scene with patterns overlapping in feature space will be brutally cut, with incomplete retrieval as a consequence (cf. Fig. 7c). Similarly to the problem mentioned in connection to the invariance operation, this error indicates that an additional, fault-tolerant structure is needed to ensure complete performance.

4.3 Rate coherence between neurons encoding repeatedly co-occurring features performs association

As already mentioned, both the invariance operation and the segmentation mechanism make errors. A natural way to cope with these errors is to employ an associative structure. Such structures arise in the higher visual cortical areas in a very natural way. In these

areas, with most of the positional information already gone, high covariances will also occur between features which do not belong to the same orbit, but which appear often together in the same object. Thereby the function connected with the stronger connections gradually shifts from the invariance operation to that of association. It is important to realize that the visual system needs an associative structure anyway, if only to correct for disturbances, such as caused by noise or partly occluded objects.

To model the associative function, we used a matrix associative memory (Willshaw 1969; Palm 1980). The associative matrix was given by

$$\mathbf{A} = \bigvee_p \xi^p \otimes \xi^p, \quad (8)$$

where \otimes denotes the outer product, \bigvee the ‘or’ operator, and ξ^p refer to the same patterns as before. At every point in time we fed the output of the oscillator layer into the associative memory, the output of which is

$$O_i^{\text{ass}} = \Theta \left(\sum_j A_{ij} \cdot O_j^{\text{osc}} - T \right) \quad (9)$$

where Θ is the Heaviside step function and T is a threshold, in our case typically 33% of the average pattern activity. Figure 7d illustrates that incorporation of this associative mechanism into the oscillator network indeed solves the completion task: the incomplete patterns resulting from segmentation (Fig. 7c) are now fully retrieved. Thus, the task of segmentation of composite input patterns is accomplished, including the case that the constituting patterns overlap in feature space.

4.4 Object recognition: invariance, segmentation and association unified into a single model

Having dealt with the different aspects of invariance, segmentation and association separately so far, we will now put them together again into a single, functional model. To this end we presented the ten pictures with objects from the breakfast table, one after the other, to the invariance model and, thus, trained the neurons. To every orbit assembly that develops, we assigned an oscillator in the next stage of the model. Thus, the activity patterns of the orbit assemblies provide the input patterns for the oscillator layer. Finally, the separate assembly patterns also serve as memory states for a hetero-associative look-up scheme, associating each output pattern of the oscillator layer with a name (e.g. apple, champagne, etc.).

The performance of this integrated model upon presentation of composite scenes is illustrated in Fig. 8. Figure 8a shows one such scene: a bottle, partly overlapping with an apple in front of it. The assemblies that are recruited by this scene are shown in Fig. 8b (compare bottom line in Fig. 6). The second line (Fig. 8c) displays some instants of the output of the oscillator layer, which groups the activity along the time axis (compare Fig. 7c). Finally, the bottom line (Fig. 8d)

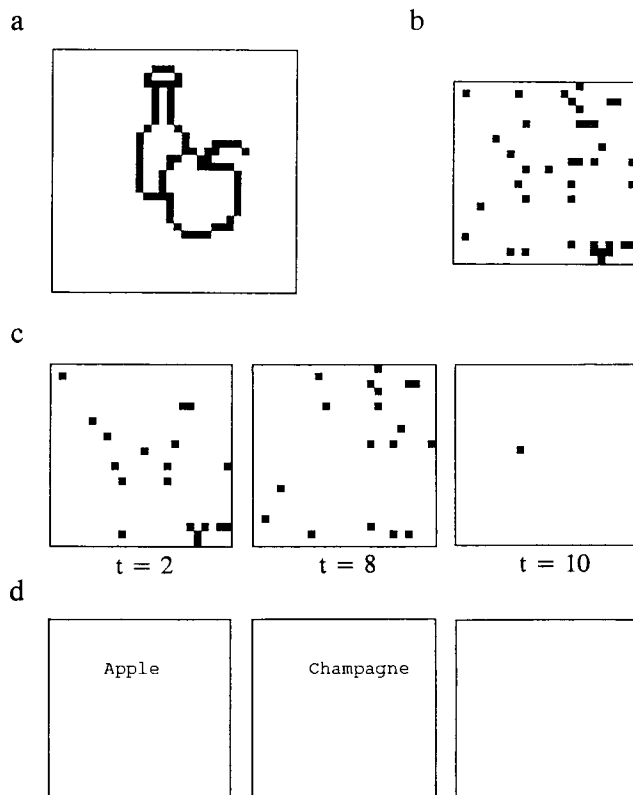


Fig. 8a–d. Object recognition, obtained by integrating the three aspects of invariance, segmentation and association into a single functional model. **a** The presented scene (input to the invariance model), consisting of a bottle and an apple in front of it. **b** The activated orbit assemblies (output of the invariance model and input to the oscillator layer). **c** The organisation of the rate activity into event-coherent assemblies (output of the oscillator layer). **d** Response of the hetero-associative memory. Observe the correct parsing and identification of the objects, even though they partly overlap

shows the responses of the associative memory; these were given when the output of the oscillator layer had an overlap of more than 33% with one of the memory patterns (compare Fig. 7d). Notice that the network correctly parses and identifies the object in the input pattern, even though they partly overlap. Figure 9 shows a similar sequence, but now the bottle is in the foreground and, in addition, both objects have been shifted in space (Fig. 9a). Yet, we observe that also in this case, the procedure leads to the correct result (Fig. 9d). From these and many similar findings, we conclude that our network indeed recognizes objects, i.e. it performs the operations of invariance, segmentation and association, and it does so in an integrated manner.

5 Comparison with anatomy, physiology and psychophysics

5.1 Summary in a biological language

Even though we modeled in a highly idealized fashion, we want to summarize in a biological language. Feedforward connections, projecting to the pyramidal cells in

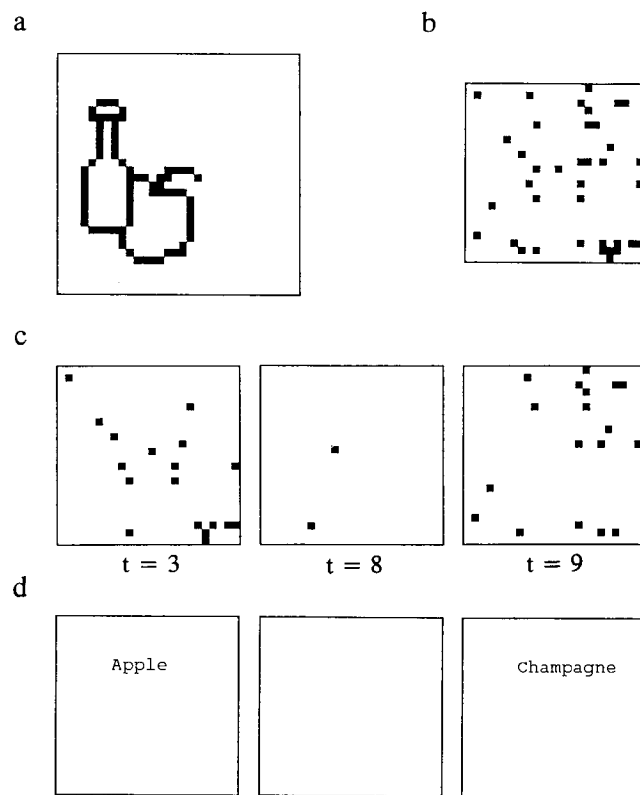


Fig. 9. The same as in Fig. 8, but the objects are shifted in space. Moreover, now the bottle, instead of the apple, appears in the foreground. Observe that also in this case, the objects are correctly identified. Notice that the changes in the scene as compared to Fig. 8, leads to a different timing schedule of segmentation

the primary visual cortex, are moulded by competitive learning. The resulting receptive fields of the pyramidal cells describe constellations in the input with consistently recurring inner relations. The distribution of the receptive field properties reflects the symmetries of the transformation groups at work in the exterior world. In the course of learning, the correspondingly transformed features become strongly connected and form orbit assemblies. Next to these strong connections, also a pattern of weaker lateral connections develops, both of them following the pattern of covariation of neural firing. Local inhibition, exerted by smooth stellate cells, provides for competition and stability.

Upon presentation of an actual scene, those cells whose feedforward receptive fields best match the input pattern raise their firing rates. In addition, via the strong lateral connections, they raise the firing rates of other, not directly addressed members in the orbit assemblies. Hence, positional information is destroyed, while corresponding invariance is gained. Under this envelope of increased firing rates, the weak (and strong) connections organize the neural activity at the level of fine-temporal spike coherence. As a result, those orbit assemblies that are often co-active are linked, leading to a natural segmentation of the scene.

These dynamics repeat themselves in higher areas; these have the same basis structure and receive a topology-preserving input projection from the primary area. At each subsequent step, the set of features becomes more diverse, while the single features become more object-specific and the domain over which they show invariance is increased. In parallel, an error-correcting associative structure emerges, since in these higher areas high covariances of firing also occur between features that do not belong to the same orbit, but often appear together in the same object. Thus, the operation which naturally emerges in such a network is to detect and represent objects, i.e. invariant sets in the visual input. This is accomplished in a distributed fashion, at the level of single neurons as well as at the level of coherently active groups of neurons.

5.2 Experimental findings from anatomy, physiology and psychophysics

Anatomy. The wiring we assume in our model incorporates only a few basic and well documented anatomical rules. The feedforward projections in the visual system are known to be organized topographically (e.g. Tusa et al. 1981; van Essen 1985). No specific demands are made on the lateral anatomical connections between the pyramidal cells; these can simply be distributed isotropically and in a probabilistic fashion (Braitenberg and Schüz 1991). The inhibition is provided by the smooth stellate cells, which are known to act locally (Braitenberg and Schüz 1991; Abeles 1991 and references therein).

Physiology. The primary visual cortex is characterized by neurons with bar-shaped receptive field profiles (Hubel and Wiesel 1977). Just as in our Fig. 3, the centers of these receptive fields are distributed over the entire visual field, with all orientations represented at every location. Unlike the lay-out in our Fig. 3, however, there is an orderly organization of orientation tuning across the cortical surface, albeit that the precise topology of this orientation map is still debated (e.g. Hubel and Wiesel 1977; Braitenberg 1985; Bonhoeffer and Grinvald 1991). We would argue that even this orderly arrangement of orientation tuning will evolve in our model without further assumptions, provided that the lateral dynamics are incorporated into the development scheme. The reason is that neighboring neurons with similar orientation preference will develop strong lateral connections, thereby mutually accelerating their development. Such local cliques, having an advantage in development, will push away the non-ordered sets and, thus, organize the topology (Kohonen 1982; for the question what is learned and what is predetermined in Area 17 see e.g. Wiesel 1982; Fregnac and Imbert 1984; Braitenberg 1985, 1986; von der Malsburg and Singer 1988).

Clearly, the receptive field properties of cortical neurons will be affected by a selective lateral connection scheme as depicted in Fig. 4. This involves particularly the neurons with the strongest lateral connections, i.e.

those for which the feedforward receptive field profiles stand out most clearly in Fig. 4. If we assume that these connections are, in fact, strong enough so that the corresponding neurons mutually transmit increased firing rates, then the complete receptive field profile of a neuron, i.e. the receptive field as it is actually measured in a physiological experiment, will, in fact, be made up of its own feedforward profile, plus the feedforward profiles of those neurons with which it has strong lateral connections (see also Braitenberg 1985, 1986). This has a number of interesting consequences for the functional properties of these neurons. First, the actually measured receptive field of a neuron will generally reach beyond its own feedforward component, its size being largely determined by the extent of the strong lateral connections. Moreover, the cells will tend to display so-called 'complex' behavior: they will respond to bars of the proper orientation, but irrespective of the precise location within the receptive field. Also the orientation-tuning curve of a recorded neuron will be broadened, since the optimal bar orientations of the neurons with which that neuron is strongly connected show some degree of scattering (Fig. 4). A suboptimal stimulus, both in terms of receptive field location and of bar orientation, will necessarily evoke a delayed response, because its influence is transmitted indirectly, i.e. via intermediate neurons. The lateral connection scheme also explains the finding that a moving stimulus elicits a stronger response than a stationary one. Suppression of inhibition results in a reduced acuity. All these phenomena have indeed been described extensively for the primary visual cortex (e.g. Hubel and Wiesel 1977; Orban 1984). Also the recent observations on context-dependence and changes of receptive field properties following manipulations in the periphery (e.g. Gilbert 1990; Gilbert and Wiesel 1991; Pettet and Gilbert 1991) can be explained along the lines of the present model.

Whereas the discussion in the preceding paragraph was restricted to the strong lateral connections, i.e. those that are strong enough to help shape the receptive field profiles, we now incorporate also the weaker lateral connections. Referring to Fig. 4, this means that now also the lighter grey hexagons are included in the discussion. Our assumption that the weak lateral connections are sufficient to synchronize spike activity implies that Fig. 4 can, in fact, be read as a spatial map of where in the primary visual cortex spike synchronization can be measured. When the reference electrode would be positioned at the site marked with an arrow (Fig. 4), coherent spike activity generally decreases with increasing distance to the second electrode. More specifically, spike synchronization occurs between cells with overlapping receptive fields, or, when they do not overlap, if the cells encode similar stimulus features (iso-orientation neurons). This scheme is in full accordance with the experimental findings (e.g. Ts'o et al. 1986; Gray et al. 1989; Eckhorn et al. 1988; Ts'o 1991; Schwartz and Bolz 1991). It is important to notice that synchronous firing is also observed between cells with non-overlapping receptive fields, which do not show

coherence in the rate behaviour. This provides an additional justification for the distinction between the two types of coherence. As assumption without direct experimental support are the desynchronizing connections between elements with negative covariance (however, see Aertsen and Gerstein (1991), showing a neuron pair in the cat primary visual cortex which switches dynamically from synchronizing correlation to desynchronizing correlation as the firing rates increase).

Another important point is that our model only needs *local* synchrony. The reason is that in the scheme we propose, in which the lateral dynamics are embedded in a hierarchical organization, the activity in one area is only locally evaluated by the neurons in the next higher area. In this context one should be aware that synchrony as judged by a neuron refers to the differences in arrival times of activity reaching the cell body of that neuron. This is not necessarily the same as synchrony judged by the experimenter who measures the differences between times at which spikes are generated at different sites. Thus, when judging in laboratory time the degree of synchrony between activities recorded from different sites, one should realize that the time interval within which spikes must occur in order to be considered synchronous by a potentially receiving neuron must increase with the distance between these generation sites (a rough estimate would be proportional to the ratio of distance between the electrodes and velocity of the neural signals). In distributed systems like the cortex we have a time structure similar to the one in special relativity, since the exchange of information about the state of synchronisation itself takes time (Lamport 1978; Mattern 1989).

Our assertion that in higher areas the function of rate coherence shifts from generating invariances to association is supported by the findings of von der Heydt et al. (1984). Using Kanizsa triangles (1979) as a stimulus, they could show that neurons in Area 18 of the monkey responded to illusory contours. Although the actual contours in the stimulus were lying entirely outside the classical receptive field, the neuron nevertheless responded if the presented picture suggested its participation. Such behavior was not observed in Area 17.

Psychophysics. There are many instances in which the behaviour of our model, as well as the assumptions underlying it, are supported by results from psychophysics and the psychology of perception. Since a thorough discussion of these correspondences is beyond the scope of this paper, we restrict ourselves to a few conceptual issues (but the reader is invited to have a look at Fig. 1a (p. 108) in Treisman (1986) and Figs. 1.14 and 2.2a in Kanizsa (1979) with our model in mind). The concept of local feature extraction is supported by many psychophysicists (e.g. Sagi and Julesz 1985; Koenderink 1984a, Treisman 1986). Regarding our specification of the task of object recognition and its effects on the developmental aspects of our model, it is interesting to note that Spelke (1990) explicitly states that the visual system of an infant seems to define

objects as units that are spatially connected, bounded and spatio-temporally continuous. For example, only at a later stage of development is a child able to perceive an object, divided in two parts by an occlusion, as a whole. This indeed suggests that the way humans divide the exterior world into building blocks is learned. Also, our point that functional order in the cortical network is defined on the basis of covariances of neuronal activity has been stressed by a number of authors (e.g. Koenderink 1984b, c; Johannesma et al. 1986). Finally, our model strongly suggests that 'Gestalt' perception (Wertheimer 1923; Kanizsa 1979) naturally emerges as a result of a statistical evaluation of the visual input along Hebbian principles; examples are the 'Gestalt' criteria of continuity and collinearity.

6 Final remarks

We are well aware that we developed and discussed our model while imposing parsing at different levels of description. We treated the development of feedforward and lateral connections separately. The dynamics of activity and the modification of connectivity was treated separately as well. We parsed the task of object recognition into the separate aspects of invariance, segmentation and association, which were then associated with different aspects of the lateral dynamics: rate coherence and event coherence. These, in turn, were studied in different modules of the model. It remains to be shown that all these aspects can be integrated more fully than we did into a single model with realistic neurons and initial connectivity, which then is organized by visual experience, while it processes input and gradually improves performance. However, we maintain that already in its present guise, the model captures the essential features of object recognition by the visual nervous system.

We attempted to tackle the principal tasks associated with object recognition all at once. This reflects our view that the visual system is to be understood as a tight family of processing steps acting closely together, and that it may be misleading to study part of its function in isolation. The function of the *whole* visual system is optimized, not the function of any one subsystem alone, at the expense of its ability to communicate with others (cf. Marr 1982).

A natural extension of our work is to study object recognition in three dimensions, with realistic camera pictures as visual input. At the horizon of such work appears the description of a network which, by simply observing the world around it, builds up representations of its main constituents, defined as cohesive chunks of the perceived environment. Important questions, which we left unanswered, are how well the network generalizes, and how many patterns per cell can be discriminated. The examples shown in the figures operate in a range where the error is zero for the given parameters. The most frequent error in the final architecture occurs when too many overlapping objects (four or five) are presented simultaneously, so that the object farthest in the background is not recognized.

Finally, we wish to emphasize once more the decisive role played in our model by the phenomenon of synchronization of neural activity. For our purposes, we subdivided this synchronization into two different levels of temporal acuity: rate coherence and event coherence. This may, in fact, be an oversimplification. For instance, it is well conceivable that at the level of event coherence there is room for further subclasses of temporal synchronization (indications of this are given by e.g. Krüger and Mayer 1991; Vaadia et al. 1991; Aertsen et al. 1992, Vaadia and Aertsen 1992), or even of a whole spectrum of synchronization acuities. In this view, neural firing rate modulations, including rate-coherent modulations, might form a dynamic substrate, with an organizational hierarchy imposed on them by a corresponding hierarchy in event synchronization (e.g. Aertsen and Preissl 1991; Erb and Aertsen 1992). Further theoretical study on this issue is currently in progress.

Acknowledgements. We wish to thank Professor Valentino Braitenberg and his group for many inspiring discussions, and Volker Staiger for his expert help in making the figures. Most of the work described was performed while both authors spent a year as guests in the laboratory of Prof. Moshe Abeles at the Hebrew University in Jerusalem. It is our pleasure to thank Moshe Abeles and his group for their hospitality during our stay as well as for the highly stimulating atmosphere and intense discussions in their lab. In addition, we thank Profs. Reinhardt Eckhorn for early encouragement, Haim Sompolinsky for challenging discussions, and Gerhard Dangelmayr for many useful hints. Partial funding for this research was obtained from the Studienstiftung des deutschen Volkes (to HN) and through grants (to AA) from the Bundesministerium für Forschung und Technologie (BMFT), the Lady Davis Foundation, the Minerva Foundation, and the German-Israel Foundation for Research and Development (GIF).

References

- Abeles M (1982) Local cortical circuits. An electrophysiological study. Springer, Berlin Heidelberg New York
- Abeles M (1991) Corticonics. Neural circuits in the cerebral cortex. University Press, Cambridge
- Abeles M, Gerstein GL (1988) Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *J Neurophys* 60:909-924
- Aertsen A, Bonhoeffer T, Krüger J (1987) Coherent activity in neuronal populations: analysis and interpretation. In: Caianiello ER (ed) *Physics of cognitive processes*. World Scientific, Singapore, pp 1-34
- Aertsen AMHJ, Gerstein GL (1991) Dynamic aspects of neuronal cooperativity: fast stimulus-locked modulations of 'effective connectivity'. In: Krüger J (ed) *Neuronal cooperativity*. Springer, Berlin Heidelberg New York, pp 52-67
- Aertsen A, Preissl H (1991) Dynamics of activity and connectivity in physiological neuronal networks. In: Schuster H (ed) *Nonlinear dynamics and neuronal networks*. Verlag Chemie, Weinheim pp 281-301
- Aertsen A, Vaadia E, Abeles M, Ahissar E, Bergman H, Karmon B, Lavner Y, Margalit E, Nelken I, Rotter S (1992) Neural interactions in the frontal cortex of a behaving monkey: signs of dependence on stimulus context and behavioral state. *J f Hirnforschung* 32:735-743
- Amit DJ (1989) Modeling brain function. The world of attractor neural networks. University Press, Cambridge
- Barnard E, Casasent D (1990) Shift invariance and the neocognitron. *Neural Networks* 3:403-410
- Bauer M, Martienssen (1991) Coupled circle maps as a tool to model synchronisation in neural networks. *Network* 2:345-351
- Bonhoeffer T, Grinvald A (1991) Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature* 353:429-431
- Boven K-H, Aertsen A (1990) Dynamics of activity give rise to fast modulations of functional connectivity. In: Eckmiller R et al. (eds) *Parallel processing in neural systems and computers*. Elsevier, New York Amsterdam Oxford, pp 53-56
- Braitenberg V (1978) Cell assemblies in the cerebral cortex. In: Heim R, Palm G (eds) *Theoretical approaches to complex systems*. Springer, Berlin Heidelberg New York, pp 171-188
- Braitenberg V (1985) Charting the visual cortex. In: Peters A, Jones EG (eds) *Cerebral cortex*, Vol. 3. Plenum Press, New York London, pp 379-414
- Braitenberg B (1986) Two views of the cerebral cortex. In: Palm G, Aertsen A (eds) *Brain theory*. Springer, Berlin Heidelberg New York, pp 81-96
- Braitenberg V, Schüz A (1991) *Anatomy of the cortex. Statistics and geometry*. Springer, Berlin Heidelberg New York
- Eckhorn R, Bauer R, Jordan W, Kruse W, Munk M, Reitboeck HJ (1988) Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analysis in the cat. *Biol Cybern* 60:121-130
- Eckhorn R, Reitboeck HJ, Arndt M, Dicke P (1990) Feature linking via synchronization among distributed assemblies: simulations of results from cat visual cortex. *Neural Comp* 2:293-307
- Eckhorn R (1991) Stimulus-specific synchronisations in visual cortex: linking features into global figures? In: Krüger J (ed) *Neural cooperativity*. Springer, Berlin Heidelberg New York, pp 184-219
- Erb M, Aertsen A (1992) Dynamics of activity in biology-oriented neural network models: stability at low firing rates. In: Aertsen A, Braitenberg V (eds) *Information processing in the cortex: experiments and theory*. Springer, Berlin Heidelberg New York, pp 201-223
- Frégnac Y, Imbert M (1984) Development of neural selectivity in primary visual cortex of cat. *Physiol Rev* 64:325-434
- Fukushima K (1988) Neocognitron: A hierarchical neural network capable of visual pattern recognition. *Neural Networks* 1:119-130
- Gerstein GL (1970) Functional association of neurons: detection and interpretation. In: Schmitt FO (ed) *The neurosciences: second study program*. Rockefeller University Press, New York, pp 648-661
- Gerstein GL, Bloom MJ, Espinosa IE, Evanczuk S, Turner MR (1988) Design of a laboratory for multi-neuron studies. *IEEE Trans Systems Man Cyb SMC-13:668-676*
- Gerstein GL (1988) Information flow and state in cortical neural networks: Interpreting multi-neuron experiments. In: von Seelen W, Shaw G, Leinhos UM (eds) *Organization of neural networks*, Verlag Chemie, Weinheim, pp 53-75
- Gerstein GL, Bedenbaugh P, Aertsen AMHJ (1989) Neuronal assemblies. *IEEE Trans Biomed Engin BME-36:4-14*
- Gilbert CD, Wiesel TN (1990) The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Res*. 30:1689-1701
- Gilbert CD, Wiesel TN (1991) Short and long term changes in receptive field size and position following focal retinal lesions. *Neurosci Abstr* 17:1090
- Gray CM, Singer W (1989) Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc Natl Acad Sci USA* 86:1698-1702
- Gray CM, König P, Engel AK, Singer W (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronisation which reflects global stimulus properties. *Nature* 338:334-337
- Hebb DO (1949) *The organisation of behavior*. Wiley, New York
- Hubel DH, Wiesel TN (1965) Receptive fields and functional architecture in two nonstriate visual areas (18 and 19) of the cat. *J Neurophysiol* 28:229-289
- Hubel DH, Wiesel TN (1977) Functional architecture of macaque monkey visual cortex (Ferrier Lecture). *Proc Roy Soc London B* 198:1-59
- Johannesma P, Aertsen A, van den Boogaard H, Eggermont J, Epping W (1986) From synchrony to harmony: Ideas on the function of neural assemblies and on the interpretation of neural synchrony. In: Palm G, Aertsen A (eds) *Brain theory*. Springer, Berlin Heidelberg New York, pp 25-47

- Kandel ER, Schwartz JH (1985) Principles in neural science. Elsevier, New York Amsterdam Oxford
- Kaneko K (1990) Clustering, coding, switching, hierarchical ordering and control in a network of chaotic elements. *Physica D* 41:137–172
- Kanizsa G (1979) Organization in vision. Essays on Gestalt perception. Praeger, New York
- Klein F (1921) Zum Erlanger Programm (1872). In: Fricke and Ostrowski (eds) Gesammelte mathematische Abhandlungen, vol. 1. Springer, Berlin Heidelberg New York
- Koenderink JJ (1984a) The concept of local sign. In: van Doorn AJ, van der Grind WA, Koenderink JJ (eds) Limits in perception. VNU Science Press, Utrecht, pp 495–547
- Koenderink JJ (1984b) Simultaneous order in nervous nets from a functional standpoint. *Biol Cybern* 50:35–41
- Koenderink JJ (1984c) Geometrical structures determined by the functional order in nervous nets. *Biol Cybern* 50:43–50
- König P, Schillen TB (1991) Stimulus-dependent assembly formation of oscillatory responses. 1: Synchronisation. *Neural Comp* 3:155–167
- Kohonen T (1982) Self-organized formation of topologically correct feature maps. *Biol Cybern* 43:59–69
- Krüger J (1983) Simultaneous individual recordings from many cerebral neurons: techniques and results. *Rev Physiol Biochem Pharmacol* 98:177–233
- Krüger J, Mayer M (1991) Two types of neural synchrony in monkey striate cortex. *Biol Cybern* 64:135–140
- Kuramoto Y (1984) Cooperative dynamics of oscillator community. *Progress of theoretical physics*, [Suppl 79], 223–240
- Lampert L (1978) Time, clocks and the ordering of events in a distributed system. *Comm ACM* 21:558–565
- Mattern F (1989) Virtual time and global states of distributed systems. In: Cosnard M, Robert I, Quinton P, Raynal H (eds) Parallel and distributed algorithms. Elsevier, New York Amsterdam Oxford, pp 215–226
- Marr D (1982) Vision. Freeman, San Francisco
- Marr D, Ullmann S (1981) Directional selectivity and its use in early visual processing. *Proc R Soc London B* 211:151–180
- Menon MM, Heinemann KG (1988) Classification of patterns using a self-organising neural network. *Neural Networks* 1:201–215
- Minsky M, Papert S (1988) Perceptrons (expanded edition). MIT Press, Cambridge
- Oja E (1989) Neural networks: principal components and subspaces. *Int J Neural Systems* 1:61–68
- Orban GA (1984) Neural operations in the visual cortex. Springer, Berlin Heidelberg New York
- Palm G (1980) On associative memory. *Biol Cybern* 36:19–31
- Pettet MW, Gilbert CD (1991) Contextual stimuli influence receptive field size of single neurons in cat primary visual cortex. *Neurosci Abstr* 17:1090
- Rumelhart DE, Zipser D (1986) Feature discovery by competitive learning. In: Rumelhart D, McClelland J (eds) and the PDP research group. Parallel distributed processing. MIT Press, Cambridge, pp 151–193
- Sagi D, Julesz B (1985) 'Where' and 'what' in vision. *Science* 228:1217–1219
- Schillen TB, König P (1991) Stimulus-dependent assembly formation of oscillatory responses. 2: Desynchronisation. *Neural Comp* 3:167–178
- Schuster HG (1991) Nonlinear dynamics and neural oscillations. In: Schuster HG (ed) Nonlinear dynamics and neural networks. Verlag Chemie, Weinheim, pp 131–151
- Schwartz EL (1977) Spatial mapping in primate visual sensory projection and relevance to perception. *Biol Cybern* 25:181–195
- Schwarz C, Bolz J (1991) Functional specificity of the long-range horizontal connections in cat visual cortex: a cross-correlation study. *J Neurosci* 11:2995–3007
- Sompolinsky H, Golomb D, Kleinfeld D (1990) Global processing of visual stimuli in a network of coupled oscillators. *Proc Natl Acad Sci USA* 87:7200–7204
- Sompolinsky H, Golomb D, Kleinfeld D (1991) Cooperative dynamics in visual processing. *Physical Rev A* 43:6990–7011
- Spelke ES (1990) Origins of visual knowledge. In: Osherson DN, Kosslyn SM, Hollerbach JM (eds) Visual cognition and action. An invitation to cognitive science, vol. 2. MIT Press, Cambridge, pp 99–127
- Sporns O, Gally JA, Reeke GN Jr, Edelman GM (1989) Reentrant signaling among simulated neuronal groups leads to coherency in their oscillatory activity. *Proc Natl Acad Sci USA* 86:7265–7269
- Treisman A (1986) Features and objects in visual processing. *Sci Am* 255:106–115
- Ts'o DY, Gilbert CD, Wiesel TN (1986) Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J Neurosci* 6:1160–1170
- Ts'o DY (1991) Connectivity and functional organisation in the mammalian visual cortex. In: Krüger J (ed) Neural cooperativity. Springer, Berlin Heidelberg New York, pp 133–164
- Tusa RJ, Palmer LA, Rosenquist AC (1981) Multiple cortical visual areas: visual field topography in the cat. In: Woolsey CN (ed) Cortical sensory organisation vol. 2. Humana Press, Clifton, pp 1–29
- Vaadia E, Ahissar E, Bergman H, Lavner Y (1991) Correlated activity of neurons: a neural code for higher brain functions? In: Krüger J (ed) Neuronal cooperativity. Springer, Berlin Heidelberg New York, pp 249–279
- Vaadia E, Aertsen A (1992) Coding and computation in the cortex: single neuron activity and cooperative phenomena. In: Aertsen A, Braitenberg V (eds) Information processing in the cortex: experiments and theory. Springer, Berlin Heidelberg New York, pp 81–121
- van Essen DC (1985) Functional organisation in primate visual cortex. In: Peters A, Jones EG (eds) Cerebral cortex vol. 3. Plenum Press, New York, pp 259–320
- von der Heydt R, Peterhans E, Baumgartner G (1984) Illusory contours and cortical neuron responses. *Science* 224:1260–1262
- von der Malsburg C (1981) The correlation theory of brain function. Internal report 81-2. MPI for Biophys Chem, Göttingen
- von der Malsburg C (1986) Am I thinking assemblies? In: Palm G, Aertsen A (eds) Brain theory. Springer, Berlin Heidelberg New York, pp 161–176
- von der Malsburg C, Schneider W (1986) A neural cocktail-party processor. *Biol Cybern* 54:29–40
- von der Malsburg C, Singer W (1988) Principles of cortical network organization. In: von Seelen W, Shaw G, Leinhos UM (eds) Organization of neural networks. Verlag Chemie, Weinheim, pp 109–126
- Wertheimer M (1923) Principles of perceptual organization. *Psychologische Forschungen* 4:301–350 (Edited and translated by Wertheimer M) in: Beardslee DC and Wertheimer M (eds) Readings in perception. Van Nostrand, Princeton
- Wiesel TN (1982) Postnatal development of the visual cortex and the influence of environment. (Nobel lecture) *Nature* 299:583–591
- Willshaw DJ, Buneman OP, Longuet-Higgins HC (1969) Non-holographic associative memory. *Nature* 222:960–962