Segmentation, Binding, and Illusory Conjunctions

D. Horn
School of Physics and Astronomy,
Raymond and Beverly Sackler Faculty of Exact Sciences,
Tel Aviv University, Tel Aviv 69978, Israel

D. Sagi
M. Usher
Department of Applied Mathematics and Computer Science,
Weizmann Institute of Science, Rehovot 76100, Israel

We investigate binding within the framework of a model of excitatory and inhibitory cell assemblies that form an oscillating neural network. Our model is composed of two such networks that are connected through their inhibitory neurons. The excitatory cell assemblies represent memory patterns. The latter have different meanings in the two networks, representing two different attributes of an object, such as shape and color. The networks segment an input that contains mixtures of such pairs into staggered oscillations of the relevant activities. Moreover, the phases of the oscillating activities representing the two attributes in each pair lock with each other to demonstrate binding. The system works very well for two inputs, but displays faulty correlations when the number of objects is larger than two. In other words, the network conjoins attributes of different objects, thus showing the phenomenon of "illusory conjunctions," as in human vision.

1 Introduction

Recent observations of synchronous oscillatory behavior of neural firings (Eckhorn *et al.* 1988; Gray *et al.* 1989) have strengthened the idea that temporal correlations are the means by which binding is achieved (von der Malsburg and Schneider 1986).

The binding problem may be viewed as the quest for a mechanism uniting parts of incoming sensory information into coherent activation patterns representing objects or situations in the external world. In the case when the assembled parts are essential for the object's identity (as when the object is defined by specific relations between its parts), this mechanism could be provided by an underlying synaptic connectivity reflecting prior knowledge. A theoretical attempt in this direction was made by several researchers (Sompolinsky *et al.* 1989; Kammen *et al.*

1989) who showed that synchronized oscillations between distant neural populations may be obtained, once an explicit connectivity among the neural populations is assumed. The binding problem is more acute in the case when the relation among the parts of the integrated objects is of contingent nature, that is, when the parts do not bear any relations essential to the identity of the objects. In this case a mechanism that does not rely on a priori connectivity should be provided.

We will limit our discussion to the binding of intramodality information. To illustrate the problem consider the case studied in psychophysical experiments (Treisman and Schmidt 1982), in which an observer is presented with a display consisting of three colored shapes, for example, a red diamond, a blue square, and a green circle. If we suppose that shapes and colors are stored in different cortical modules (networks), we are faced with the double problem of segmentation and binding. That is, the "shape" module should recognize and segment the shapes, while the "color" module should recognize and segment the colors. The binding problem, then, is to provide the correct matching between the shapes and their corresponding colors. Treisman and Schmidt suggested that correct matching can be obtained by the human visual system only when focusing attention on each of the objects separately, otherwise illusory conjunctions may occur.

We wish to study this problem within a model of coupled oscillatory networks that receive a mixed input, as illustrated schematically in Figure 1. Such a system is then required to perform simultaneously both segmentation and binding.

Segmentation is the task of parallel retrieval of the individual memorized patterns composing the input. This can be achieved in oscillating networks as was demonstrated by Wang *et al.* (1990) and by Horn and Usher (1991). What happens is that the activities of the different memory patterns that are turned on by the input oscillate in a staggered fashion.

Binding is modeled by assuming that patterns corresponding to the related attributes oscillate in phase (e.g., the activity of the pattern representing the shape "diamond" should oscillate in phase with the activity of the color "red"). Modeling the binding process is especially challenging since no a priori stored synaptic structure relating the corresponding patterns is allowed, due to the fact that their relation is contingent. We will show how a solution to the binding problem is achieved by using a mechanism based on enhancement of noise correlations. Moreover we will show that for more than two input patterns synchronization faults occur. These faults may provide a natural explanation for perceptual errors of the illusory conjunction type.

The neural networks that we study are based on coupled formal neurons that possess dynamic thresholds that exhibit adaptation: they vary as a function of the activity of the neurons to which they are attached. As such they introduce time dependence, which can turn a neural network from a dissipating system that converges onto fixed points into one that
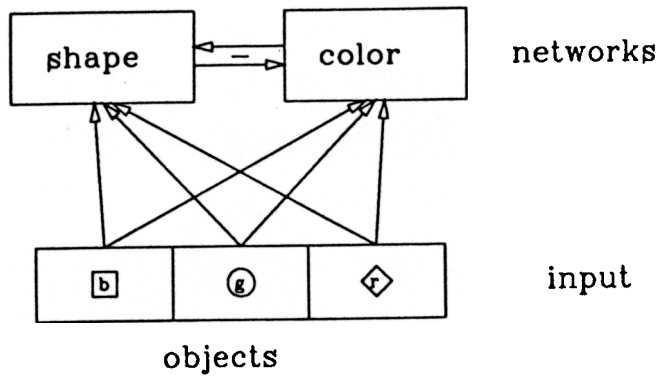
Figure 1: Schematic drawing of a problem of joint segmentation and binding. The example of objects of different shapes and colors will be discussed throughout this paper. The inhibitory connection between the two networks is an important element of our solution.

moves from one center of attraction to another (Horn and Usher 1989). Here we will use a variant of these models (Horn and Usher 1990) that is based on a model of excitatory and inhibitory neurons. This model is explained briefly in the next section. It can be expressed in terms of a set of differential equations that involves the activities of the excitatory cell assemblies that represent the memories of this model. In the following section we describe how two such networks can be coupled through their inhibitory neurons in a nonsemantic fashion, that is, no explicit connections exist between the patterns that are to be bound. We show that this coupling leads to matching their periods and phases. The next section is devoted to an explanation of how binding is achieved. Afterward we turn to an analysis of the performance of our model and dwell on the illusory conjunctions that it exhibits when the number of inputs is larger than two.

2 The E-I Model

The system that we will study is based on a model of excitatory and inhibitory neurons with dynamic thresholds. These two kinds of neurons are assumed to have excitatory and inhibitory synapses exclusively. Memory patterns are carried by the excitatory neurons only. Furthermore, we make the simplifying assumption that the patterns do not overlap with one another, that is, the model is composed of disjoint Hebbian

cell assemblies of excitatory neurons that affect one another only through their interaction with a group of inhibitory neurons common to all of them. We refer to a previous paper (Horn and Usher 1990) for details of the microstructure of this model. Here we will limit ourselves to its description in terms of differential equations for the activities of the cell assemblies.

To start with let us consider the case of static thresholds. We denote by $m^\mu(t)$ the fraction of cell assembly number μ that fires at time t, and by $m^I(t)$ the fraction of active inhibitory neurons. We will refer to m^μ as the activity of the μth memory pattern. There are p different memories in the model, and their activities obey the following differential equations

$$
\begin{aligned}
dm^\mu/dt &= -m^\mu + F_T(Am^\mu - Bm^I - \theta^E) \\
dm^I/dt &= -m^I + F_T(CM - Dm^I - \theta^I)
\end{aligned}
\tag{2.1}
$$

where

$$
M = \sum_\mu m^\mu \qquad F_T(x) = (1 + e^{-x/T})^{-1} \tag{2.2}
$$

θ^E and θ^I are the (constant) thresholds of all excitatory and inhibitory neurons, respectively. The four parameters A, B, C, and D are all positive and represent the different couplings between the neurons. This system is an attractor neural network. It is a dissipative system that flows into fixed points determined by the memories.

This system is a generalization of the E-I model of Wilson and Cowan (1972) in which we have introduced competing memory patterns. The latter make it into an attractor neural network. Wilson and Cowan have shown that a pair of excitatory and inhibitory assemblies, when properly connected, will form an oscillator. We induce oscillations in a different way, keeping the option of having the network behave either as an attractor neural network or as an oscillating one: we turn the thresholds of the excitatory neurons into dynamic variables. For this purpose we introduce new variables r^μ that represent the average alternating behavior of the thresholds of the excitatory neurons in cell assembly μ, and change the p equations of the excitatory neurons to the following $2p$ equations:

$$
\begin{aligned}
dm^\mu/dt &= -m^\mu + F_T(Am^\mu - Bm^I - \theta_0^E - br^\mu) \\
dr^\mu/dt &= (1/c - 1)r^\mu + m^\mu
\end{aligned}
\tag{2.3}
$$

For $c > 1$ and appropriate values of $g = bc/(c-1)$, this system exhibits local fatigue effects. The reason is simple. Imagine a situation in which the system would move into a fixed point $m^\mu = 1$. r^μ will then increase until it reaches the value $c/(c-1)$. This means that the argument of the F_T function in the equation for m^μ decreases by g. If this overcomes the effect of the other terms the amplitude m^μ decreases and the system moves out of the attractor and falls into the basin of a different center of attraction. This process can continue indefinitely.

3 The Model of Coupled Networks

The system we study is presented diagrammatically in Figure 1. We realize it by using two E–I networks that are coupled through their I neurons. This type of coupling is chosen to avoid any a priori specific connection between memory patterns of the two different networks. They may still affect each other through their couplings to the connected sets of I neurons, but there is no explicit relation between the two sets of memories. Let us introduce also external inputs to the E neurons designated by i^μ. The system of the two coupled networks takes then the form

$$dm_{1,2}^\mu/dt = -m_{1,2}^\mu + F_T(Am_{1,2}^\mu - Bm_{1,2}^I - \theta_0^E - br_{1,2}^\mu + i_{1,2}^\mu) \quad (3.1)$$

$$dr_{1,2}^\mu/dt = (1/c - 1)r_{1,2}^\mu + m_{1,2}^\mu$$

$$dm_{1,2}^I/dt = -m_{1,2}^I + F_T(CM_{1,2} - Dm_{1,2}^I - \theta^I - \lambda m_{2,1}^I)$$

The subindices refer to the two different networks, whose only connection is through the term λm^I representing the coupling between the two sets of inhibitory neurons. We present in Figure 2 a schematic drawing of the relation between the variables that appear in this set of differential equations. Drawn here is one of the two networks with three memories and an input that feeds into two of them.

Let us start our discussion of this system of differential equations by limiting ourselves to the case of a single excitatory cell assembly in each
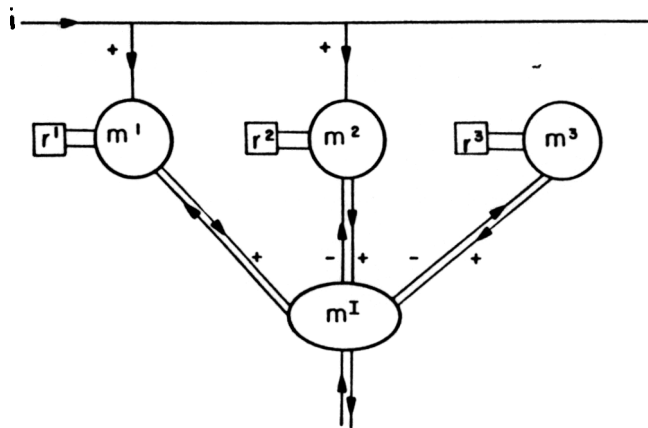

Figure 2: Schematic drawing of the relations between the variables of one of the two networks described by equations 3.1. Shown here is the case of three memories and two inputs.

network. Assuming first no input and no coupling between the two networks we obtain the results shown in Figure 3. We have chosen the b parameters to be different in the two networks, therefore we obtain oscillations with different frequencies. Figure 4 shows how the situation changes when the coupling between the two I assemblies is turned on, $\lambda = 1$. It is quite evident that this coupling forces the two networks to move in tandem. The common frequency is lower than the two frequencies of the free networks. We observe a difference in the shape and phase of the activities of the two networks, which is the remnant effect of the two different frequencies. The phase shift is particularly strong between the two I activities because they inhibit each other. The regular shape of the average I activity in the coupled case justifies a posteriori approximating its equation of motion by

$$dm^I/dt = -m^I + F_T(CM - Dm^I - \theta^I - \lambda m^I)$$

meaning that its effective autoinhibition increased from D to $D + \lambda$. This seems to be the reason for the lower overall frequency.

Let us turn now to the general case of p excitatory cell assemblies, n of which receive a common input:

$$i_{1,2}^\mu = i \quad \mu = 1, \cdots, n \quad (3.2)$$

$$i_{1,2}^\mu = 0 \quad \mu = n + 1, \cdots, p$$

For $\lambda = 0$ we find the phenomenon of temporal segmentation discussed by Wang et al. (1990) and by Horn and Usher (1991). This means that different memories oscillate in a staggered manner, each one peaking at different times, thus leading to segmentation of the mixed input. This scenario works as long as n is small. Once we couple the two inhibitory assemblies we may expect the oscillations of the two networks to match one another in period and phase. However, this matching will be random, since there is no reason for a particular cell assembly of one network to oscillate in phase with a particular one of the other network. How to achieve such binding will be discussed in the next section.

4 Binding by Correlated Fluctuations

Our problem, which is symbolically presented in Figure 1, assumes that the two networks describe two attributes of objects that appear in a mixed form in the input. We expect our combined network to be able to segment this information and, moreover, to order the staggered oscillations in such a form that the activities of the two attributes of the same object have the same phase. To achieve the latter we make use of noisy inputs. For two attributes of the same object we assume that both are affected by some common random activity fluctuation. However, the noises affecting two different objects are assumed to be uncorrelated. The noise is transmitted
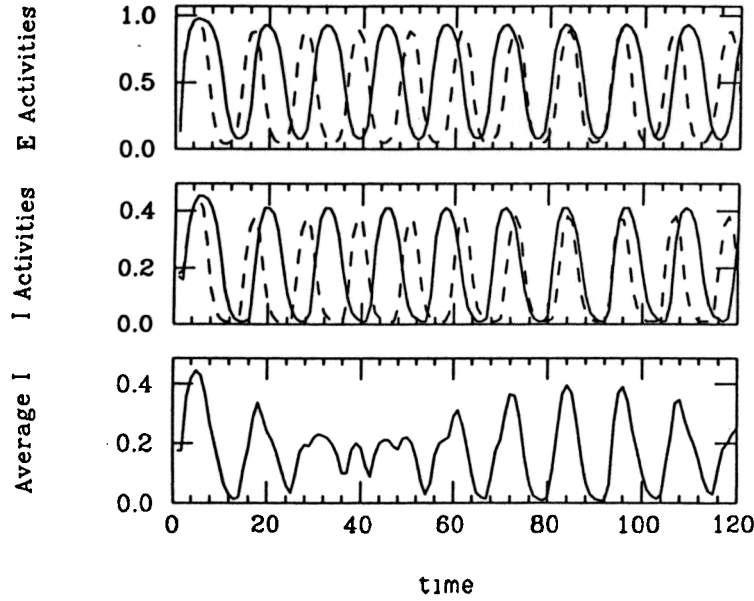
Figure 3: Activities of two networks with one E assembly each, no input, and no coupling. The activities of the first and the second network are represented by full and dashed curves, respectively. The results are numerical solutions of the differential equations 3.1 using time steps of $dt = 0.1$ and parameters $A = 1$, $B = 0.7$, $C = 1$, $D = 1$, $T = 0.1$, $c = 1.2$, $\theta^E = 0.1$, $\theta^I = 0.55$. The parameters b are chosen differently for the two networks, $b_1 = 0.15$ and $b_2 = 0.2$, hence the different frequencies of oscillation.

together with the constant input to the relevant cell assemblies of the two networks.

The inputs we use take the form

$$i^\mu(t) = 0.1 + 0.1[\rho^\mu(t) - 0.5] \tag{4.1}$$

where ρ^μ is a random variable distributed between 0 and 1. The same input is used for both attributes which refer to the same object $i^\mu_{1,2} = i^\mu$, yet different pairs of attributes are driven by different and uncorrelated random noises ρ^μ. We solve numerically the differential equations 3.1 using small time steps of $dt = 0.1$ for each iteration. We assume that the inputs are updated on a time scale τ that is an order of magnitude larger, either $\tau = 1$ or 2. Correspondingly we represent the time scale in the following figures by integers.
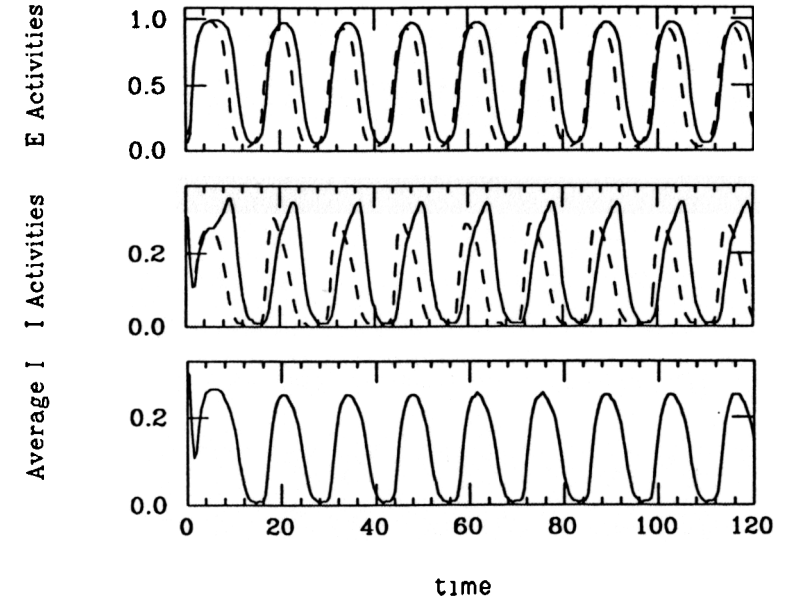

Figure 4: The result of turning on the coupling $\lambda = 1$ between the two networks of the previous figure.

Figure 5 describes results when two pairs of input of the type of 4.1 were used. Starting from random initial conditions we observe correct phase correlations after 10 time units, turning into almost perfect binding after 30 time units. Binding occurs almost instantly if one starts from zero (instead of random) initial conditions for the activities. In this figure we show in addition to the activities of the two different cell assemblies in the two networks also the random noises used for the two pairs of inputs. Note that the time scale of phase-locked oscillations is much larger than that of the autocorrelations of the fluctuating noise (which is $\tau = 1$).

In the case of three objects, shown in Figure 6, it takes longer time to achieve correct binding. Moreover, we have noticed that the system can move out of correct binding into erroneous phase correlations, of the type shown here from $t = 30$ to 90.

In order to quantify the binding quality we measure the fraction of correct activity correlations:

$$B = \frac{\int dt \sum_{\mu=1}^n m_1^\mu(t) m_2^\mu(t)}{\int dt \sum_{\mu=1}^n m_1^\mu(t) \sum_{\nu=1}^n m_2^\nu(t)} \tag{4.2}$$
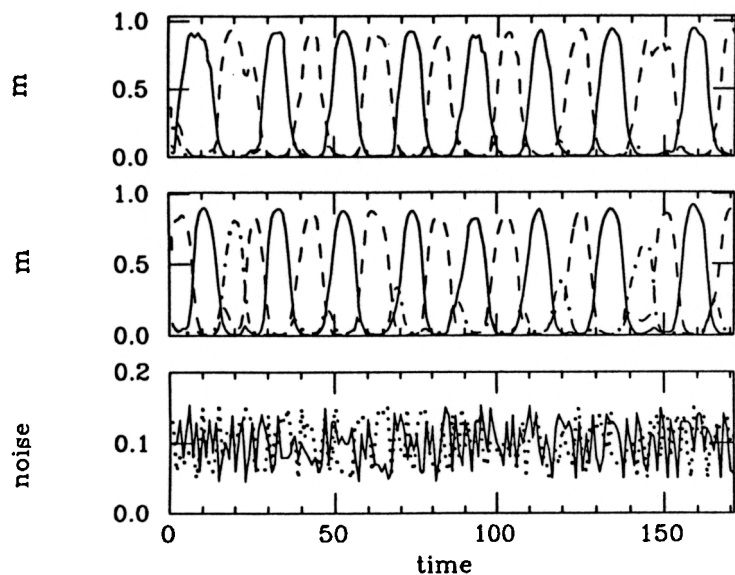
Figure 5: The first two frames exhibit activities of memory patterns (excitatory cell assemblies) of two coupled networks. The parameters b are 0.1 and 0.12. The other parameters used here and in all following figures are $A = 1$, $B = 1.1$, $C = 1.2$, $D = 1$, $T = 0.1$, $c = 1.2$, $\theta^E = 0.1$, $\theta^I = 0.55$, $\lambda = 1.2$. The first network has five memory patterns and the second has three. Two cell assemblies receive inputs of the type of equation 4.1. The activities of these two memories are shown by the full line and the dashed line. The dot–dashed curve represents an activated memory that does not receive an input. We observe both segmentation and binding. Segmentation means that the two different patterns in the two networks oscillate in a staggered fashion, and binding means that the patterns that are associated with one another oscillate in phase. The association is brought about by the common noise, which is shown in the third frame.

For the case of $n = 2$ in networks with different parameters ($b_1 = 0.1$, $b_2 = 0.15$) we find high correct correlations, $B = 0.83$. In general binding is best when the frequencies of the two coupled networks are identical. Nonetheless, when we turn to $n = 3$ in networks with identical parameters as shown in Figure 6, we find that B reduces to $0.41 \pm .02$. Better performance is obtained if we allow the noise correlation time to be longer, for example, we change ρ only every two time units ($\tau = 2$). This leads to $B = 0.51 \pm .02$.
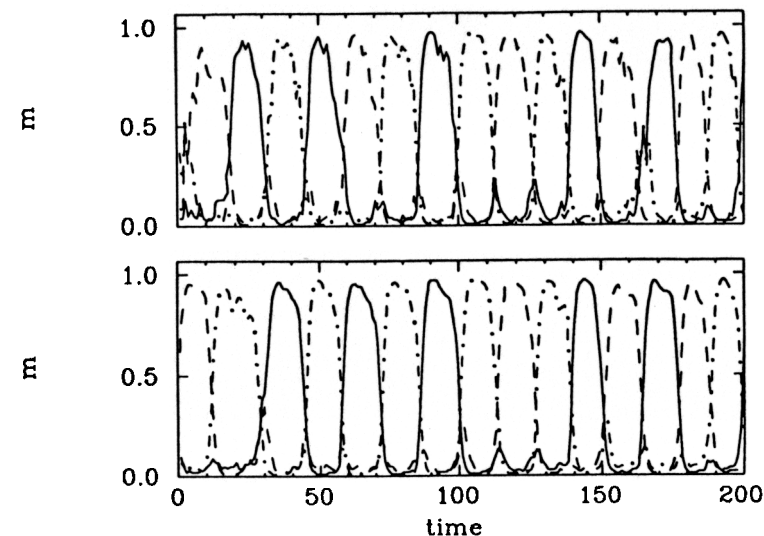
Figure 6: The behavior of the joint networks in the case of three inputs. Using equal frequencies, $b_1 = b_2 = 0.1$, we find that temporal segmentation works very well but binding is less successful than in the case of two patterns. Associated patterns are represented by the same type of lines in the two different networks.

5 Binding Errors and Illusory Conjunctions

We saw that the binding obtained by our model is not perfect and that some degree of erroneous matching of oscillation occurs. The frequency of these matching errors increases with the increase in the number of displayed objects from two to three. We propose that this could be the mechanism responsible for the phenomenon of illusory conjunctions.

Let us shortly describe the outcome of a typical experiment in the illusory conjunction paradigm. When an observer is presented with a display containing several visual shape–color patterns for a short exposure time, and when due to experimental set-ups his attention is spread over the whole display, perceptual errors (such as reporting a green diamond when presented with a red diamond and a green circle) occur. As we previously mentioned, Treisman and Schmidt (1982) suggest that integrating shape and color information related to one object requires focusing attention on the object. Thus when attention is distributed across several

objects incorrect matching occurs. However, in all illusory conjunction experiments a considerable amount of correct responses is obtained, even when the experiment is designed to maximize the size of the attention window. As we will show our model can provide the explanation.

In the previous section we have defined the fraction of correct binding B and have given several numerical examples. Thinking of our model as a binding predictor we should however take into account that an observer randomly conjoining n objects' shapes and colors will obtain $B = 1/n$ by purely guessing the result. Trying to correct for this trivial baseline we define the significance of the binding probability by

$$S = \frac{B - 1/n}{1 - 1/n} \qquad (5.1)$$

The denominator serves as a normalization factor, allowing S to vary between 0 and 1. To demonstrate the systematic trend of our model we display in Figure 7 both B and S as a function of n; the number of objects, for two coupled networks of the same frequency. Clearly $n = 3$ is the worst case. Higher n seem to lie on a plateau of $S \approx 0.3$. Explaining illusory conjunctions by our model, we expect their number to increase strongly when the number of objects increases from 2 to 3, but to level off at a rate that is significantly different from pure chance after that.

Until now we have discussed the case in which all color and shape patterns were different. It is interesting to examine the model in the case in which one of the attributes is shared by several objects. Consider the case in which the display consists of three objects, two of which share the same color, for example, using a green square instead of a blue one in the example of Figure 1. The repeated color (green) is represented just once in the color network, but it will receive a large input, which is the linear sum of the inputs of the two objects that share the same color. The result of such a simulation is illustrated in Figure 8. We observe that the repeated color is indeed strongly enhanced. The amount of conjunctions between the repeated color (green represented by the dashed curve) and the unrelated shape (diamond represented by a full line) is higher than that for the unrepeated color (red) and an unrelated shape (square or circle). In this asymmetric case it is advantageous to consider a correlation matrix

$$C_{\alpha\beta} = \frac{\int dt\, m_1^{\alpha}(t) m_2^{\beta}(t)}{\int dt \sum_{\mu=1}^{n_1} m_1^{\mu}(t) \sum_{\nu=1}^{n_2} m_2^{\nu}(t)} \qquad (5.2)$$

which describes the probability of binding shape α with color β. Running the system of Figure 8 with $n_1 = 3$ $n_2 = 2$ for a long time we obtain

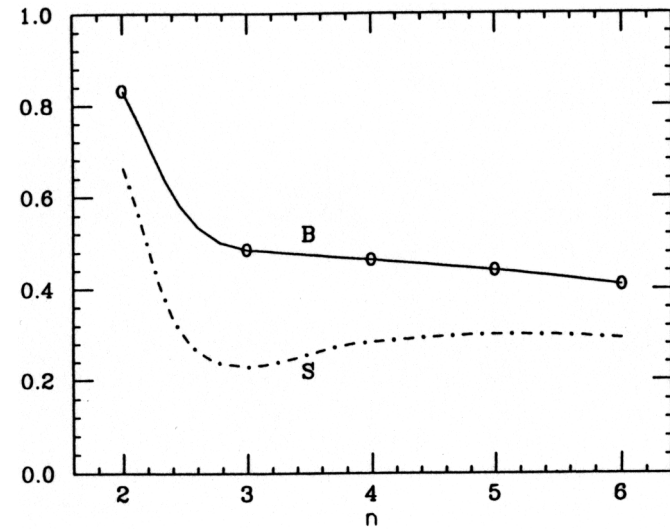$$C = \begin{pmatrix} .148 & .183 \\ .088 & .248 \\ .087 & .244 \end{pmatrix}$$

Figure 7: Typical variation of the binding and significance parameters as a function of the number of objects. The two networks have $b_1 = b_2 = 0.1$. The noisy input has correlations of $\tau = 2$.

This corroborates the statement we made before that there are more illusory conjunctions with the repeating attribute. Note, however, that for a given color the strongest correlations are with the correct shapes.

The fact that relative duration of green is roughly twice that of red in Figure 8 is a consequence of using a linear sum of the two inputs that contribute to green. Changing the parameters of our model we can change the relative strengths of these signals, but usually in the direction of further amplifying the effect of the stronger amplitude. One can conceive of a different situation, in which the two colors appear with about the same strength. This calls, however, for a modification of our model: it necessitates a nonlinear interaction at the sensory input stage.

6 Discussion

The model that we propose is based on two oscillatory networks, in which excitatory neural assemblies represent attributes of objects, such as shape and color. The networks are coupled through their inhibitory neurons, whose role is to mediate competition among the excitatory assemblies in each network, and also to phase lock the oscillations of the
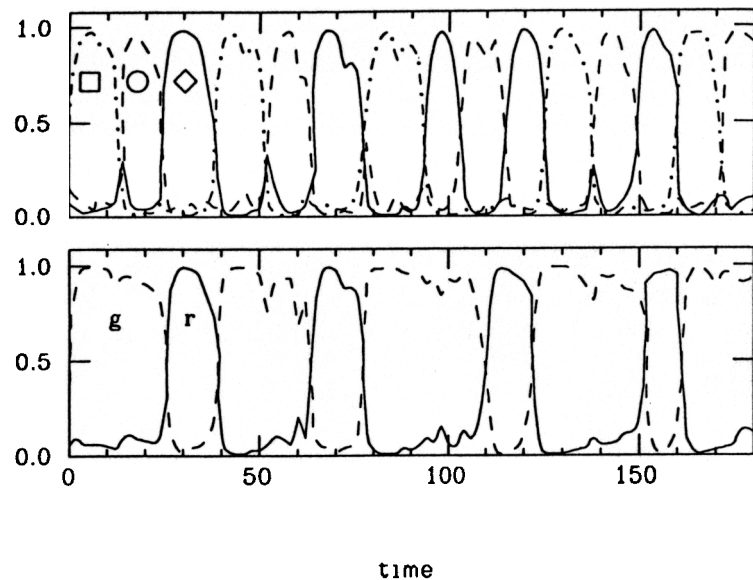
Figure 8: An example of three objects, two of which share one attribute (green color). Most conjunctions are correct. The shape and the color networks have the parameters $b_1 = b_2 = 0.1$, $p_1 = 5$, $p_2 = 3$.

shape and color networks. The problem of binding the correct assemblies in the shape and color networks (i.e., assemblies representing attributes of the same object) is solved by introducing correlated noise fluctuations into the corresponding assemblies. Thus our model provides the means by which binding via amplification of noise correlations can be obtained.

This model can serve as an example of intramodality binding, in which we can assume that the input carries some information regarding the connection between the two attributes that are to be bound. This cannot be applied in the same form to the interesting question of cross-modality binding (e.g., connecting visual and acoustical attributes) where the common input layer does not exist. A model for cross-modality binding needs a different approach, which may have to rely on prior knowledge that introduces explicit synaptic connectivity between memory patterns, an element that we have successfully avoided in our model.

Two characteristics of noisy inputs are worthwhile stressing in the context of our model. First, we wish to point out that noisy input increases the segmentation power of the network. Running the network with a constant input we find that it cannot segment successfully more

than about five objects. This is in complete agreement with other oscillatory networks performing segmentation (Wang et al. 1990; Horn and Usher 1991). If more than five excitatory assemblies receive a constant input many activities try to rise simultaneously leading to the collapse of all of them. When noise is added to the constant inputs the network can continue its staggered oscillations for very large numbers of objects. The reason for this is that noise fluctuations will always enhance momentarily the input of one of the assemblies, enabling it to overtake the other ones. It seems that this has to do with the fact that we run the network in a chaotic phase, which is the second point we wish to stress. When we use $n = 3$ input patterns the networks segment the input into a well ordered sequence of staggered oscillations. This is no longer true for $n = 4$ or 5. The order of the staggered oscillations is quite random, indicating chaotic behavior. Therefore, if synchronicity between the activities of the two connected networks fails, it is easier to amend it when $n > 3$. In other words, sensitivity to noise correlations is enhanced when the network is in its chaotic phase, leading to an increase in the value of S beyond three displayed objects. The importance of chaos in the sensory processing of information by the brain was discussed by Skarda and Freeman (1987), who found that neural activity in the olfactory bulb shows chaotic characteristics when the animal is engaged in odor recognition. They suggested that the advantage of chaos for the processing of sensory information is that a chaotic state is more sensitive to changes in the incoming input. It seems that this characteristic is also demonstrated by our model.

Although we have not attempted to model the physiological observations in the visual cortex (Eckhorn et al. 1988; Gray et al. 1989), we should be aware of an interesting qualitative difference: binding in our model takes some time to develop, as seen in Figures 5 and 6, whereas in the experimental results phase locking develops rapidly. The delay in our model comes about because we start from random initial conditions that the input has to overcome. It is quite possible that the physiological process is also assisted by auxiliary mechanisms. For example, "spotlight" attention (Koch and Ullman 1985) can eliminate all but one object and, therefore, lead to fast binding. Our model shows that even when the attention spotlight is spread, as in illusory conjunction experiments, significant amount of binding can be obtained by making use of noise effects.

Within the context of the psychological phenomenon of illusory conjunctions, our model differs from the Feature Integration Theory (Treisman and Schmidt 1982). While according to this approach, feature representations (e.g., shape and color) are completely separate, according to our model some early mixed representation of shape and color information exists in the input layer. Only at a higher order memory level shape and color information are separated.

The main prediction of our model is that binding performance depends on the number of displayed objects (Fig. 7). In particular, we find

a strong increase in the rate of illusory conjunctions when the number of objects is increased from two to three. We predict, however, quite uniform behavior when the number of objects is larger than three. Due to the linear dependence on the input we expect predominance of repeated attributes both in correct and illusory conjunctions. In this context we wish to stress that we represented the different objects with the same weights, that is, by equal numbers of neurons. If this is modified, we expect oscillations of the excitatory cell assemblies to be ordered according to input strength, thus producing biased errors (e.g., if "red" is stronger than "green" and "circle" is stronger than "square," then circles may be always red regardless of spatial coincidence). An experimental examination of this issue is needed in spite of its difficulty. It may call for additional mechanisms to rescale the representation on the input level.

Finally we wish to address the issue of temporal versus spatial coincidence. We assume that a shape and a color appearing in synchrony will be matched by some higher level process. To test this assumption is rather difficult since it involves rapid presentation (at the oscillation rate, probably higher than 40 Hz, which is probably higher than sensory integration rate) of isolated object attributes at different locations. Thus it is not surprising that Keele *et al.* (1988) failed to find direct support for temporal binding, and concluded that spatial coincidence is the preferred mechanism for binding that is revealed by psychophysical experiments. Note, however, that our model is making use of spatial coincidence as a binding clue (by local noise) and thus is not in disagreement with these results. It is also possible to introduce spatial location explicitly into the model by adding a network encoding relative or absolute location as an attribute that can oscillate synchronously with all other attributes and enhance the role of spatial coincidence.

In conclusion, we have shown that noise correlations in the input layer can provide the mechanism by which binding via matching of oscillations is achieved. It remains to be seen whether this mechanism is used by the brain to conjoin sensory attributes when attention is distributed across several objects.

Acknowledgment

References

Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., and Reitboeck, H. J. 1988. Coherent oscillations: a mechanism of feature linking in the visual cortex? *Biol. Cybern.* **60**, 121–130.

Gray, C. M., Konig, P., Engel, A. K., and Singer, W. 1989. Oscillatory response in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus attributes. *Nature (London)* **338**, 334–337.

Horn, D., and Usher, M. 1989. Neural networks with dynamical thresholds. *Phys. Rev. A* **40**, 1036–1044.

Horn, D., and Usher, M., 1990. Excitatory-inhibitory networks with dynamical thresholds. *Int. J. Neural Syst.* **1**, 249–257.

Horn, D., and Usher, M. 1991. Parallel activation of memories is an oscillatory neural network. *Neural Comp.* **3**, 31–43.

Kammen, D., Koch, C., and Holmes, P. J. 1989. Collective oscillations in the visual cortex. *Proceedings of the NIPS Conference*, pp. 76–83.

Koch, C., and Ullman, S. 1985. Shifts in selective attention: Towards the underlying neural circuitry. *Human Neurobiol.* **4**, 219–227.

Keele, S. W., Cohen, A., Ivry, R., Liotti, M., and Yee, P. 1988. Tests of a temporal theory of attentional binding. *J. Exp. Psychol: Human Percept. Perform.* **14**, 444–452.

Skarda, C. A., and Freeman, W. J. 1987. How brains make chaos in order to make sense of the world. *Behav. Brain Sci.* **10**, 161–195.

Sompolinsky, H., Golomb, D., and Kleinfeld, D. 1989. Global processing of visual stimuli in a neural network of coupled oscillators. *Proc. Natl. Acad. Sci. U.S.A.* **87**, 7200–7204.

Treisman, A., and Schmidt, H., 1982. Illusory conjunctions in the perception of objects. *Cognit. Psychol.* **14**, 107–141.

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

Wang, D., Buhmann, J. and von der Malsburg, C. 1990. Pattern segmentation in associative memory. *Neural Comp.* **2**, 94–106.

Wilson, H. R., and Cowan, J. D. 1972. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys. J.* **12**, 1–24.