Journal of Computational Neuroscience 17, 179–201, 2004
© 2004 Kluwer Academic Publishers. Manufactured in The Netherlands.

Modeling Compositionality by Dynamic Binding of Synfire Chains

MOSHE ABELES*
Department of Physiology and the Center for Neural Computation, The Hebrew University, Jerusalem, Israel; Gonda Brain Research Center, Bar-Ilan University, Ramat-Gan, Israel
abeles@vms.huji.ac.il

GABY HAYON
The Center for Neural Computation, The Hebrew University, Jerusalem, Israel
gabyh@cs.huji.ac.il

DANIEL LEHMANN
Department of Computer Science and the Center for Neural Computation, The Hebrew University, Jerusalem, Israel
lehmann@cs.huji.ac.il

Received December 2, 2002; Revised November 7, 2003; Accepted May 28, 2004

Action Editor: Misha V. Tsodyks

Abstract. This paper examines the feasibility of manifesting compositionality by a system of synfire chains. Compositionality is the ability to construct mental representations, hierarchically, in terms of parts and their relations. We show that synfire chains may synchronize their waves when a few orderly cross links are available. We propose that synchronization among synfire chains can be used for binding component into a whole. Such synchronization is shown both for detailed simulations, and by numerical analysis of the propagation of a wave along a synfire chain. We show that global inhibition may prevent spurious synchronization among synfire chains. We further show that selecting which synfire chains may synchronize to which others may be improved by including inhibitory neurons in the synfire pools. Finally we show that in a hierarchical system of synfire chains, a part-binding problem may be resolved, and that such a system readily demonstrates the property of priming. We compare the properties of our system with the general requirements for neural networks that demonstrate compositionality.

Keywords: synfire-chains, compositionality, binding-mechanism, neural-networks

1. Introduction

Neural networks are often used to model and understand the way the brain works. Most of these models assume that all semantic information is contained in the spiking rates of the neurons (Hertz et al., 1991). Nevertheless, there is evidence that the fine temporal structure of the spiking activity is also significant (Eckhorn et al., 1988; Gray and Singer, 1989; Prut et al., 1998; Villa et al., 1999). Synfire chains (Abeles, 1982, 1991) are one of the models in which the temporal structure of spike trains plays a pivotal role. A synfire chain (Fig. 1) is a feed-forward excitatory network including a large number of layers (here termed pools) of neurons, where neurons in the same pool fire in an almost synchronous

*To whom correspondence should be addressed.

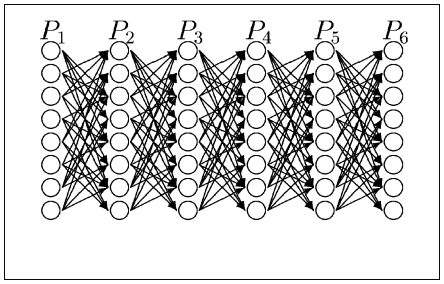

Figure 1. Schematic view of a synfire chain: Every neuron in pool i projects to m neurons in pool $i+1$. The width of the chain is the number of neurons in a pool (eight in this example), and the multiplicity (m) of a chain is the average number of cells in pool P_{i+1} to which a cell in pool P_i is connected (four in this example).

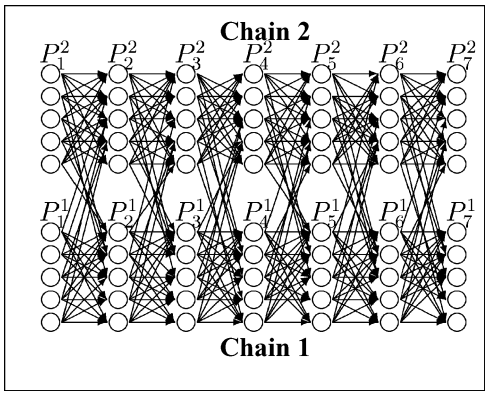

Figure 2. Schematic view of two synfire chains: The intra-chain connections are similar to Fig. 1. There are additional connections between the two chains. The number of such connections per neuron in a pool is much smaller than the number of intra-chain connections. In the above figure, each neuron is connected to four neurons in the next layer of the same chain and one neuron from the other chain.

way. Abeles et al. had reported patterns of tight locking (1–3 ms) when recording simultaneously from a number of cells in the frontal cortex of behaving monkeys (Prut et al., 1998). A simple way of explaining such findings is that these cells are part of an active synfire chain.

Synfire chain models have striking properties. Among these properties are: stability over a time span of a few hundreds of milli-seconds (Abels et al., 1993, 1994), reproducibility (Abeles, 1991; Aertsen et al., 1996a, 1996b; Diesmann et al., 1999; Postma et al., 1996; Wong, 1997), learnability by a self-organization process (Bienenstock, 1991; Doursat, 1991; Hertz and Prugel-Bennett, 1995; Horn et al., 1999; Sougne and French, 2001; Sterratt, 1999), and a large storage capacity (Bienenstock, 1995; Herrmann et al., 1995). However, the most advantageous property of the synfire chains model is their ability to account for compositionality; i.e., the ability to build complex representations out of simpler parts and to reconfigure the same parts in many different ways.

Bienenstock has suggested that compositionality may be implemented by the dynamical binding of weakly connected synfire chains (Bienenstock, 1991, 1992; Bienenstock and Geman, 1994; Doursat, 1991). Simulations by Abeles et al. have shown the feasibility of this idea (Abeles et al., 1993). According to Bienenstock's model, waves of activity in a synfire chain represent the semantic atoms (e.g. CIRCLE, RED), and synchronization of the activity in different chains serves as the dynamical binding mechanism (e.g. representation of a RED CIRCLE). Bienenstock showed, using Markov models, that two weakly connected synfire chains (as seen in Fig. 2) synchronize their activity, provided that the two chains are coactivated with appropriate initial timing (Bienenstock, 1995). Arnoldi and Brauer describe a plausible mechanism for the synchronization of waves of activity in two synfire chains with weak connections (Arnoldi and Brauer, 1996) (their architecture is similar to the one shown in Fig. 2): a leading activity wave in one of the chains accelerates the propagation of an activity wave in the other chain until both activity waves are synchronized. Using simulations, they found that this mechanism is effective as long as the initial distance between the two waves of activity is less than 14–20 ms.

The aim of this work is to show that such synchronization can be easily performed by a system of synfire chains, and that such a system may solve simple binding problems.

Since compositionality is so central to cognition, many neural network models for compositionality have been proposed. Most such models use temporal properties to bind primitive objects. von der Malsburg introduced a dynamical binding mechanism based on fast synaptic modifications (Triesch and von der Malsburg, 1996; von der Malsburg, 1981, 1987; von der Malsburg and Wiskott, 1996; Zucker, 1989). His Dynamic Link Matching mechanism has been used for computerized face identification (von der Malsburg and Wiskott, 1996). In this mechanism, synapses may be switched on very rapidly and then may relax slowly to their resting level. These synapses are used as a dynamical binding mechanism. Von der Malsburg suggests using a combination of fast and slow synaptic modifications in a model for learning and representing composition (von der Malsburg, 1987). Triesch and

von der Malsburg used this mechanism to build a model for toy vision example (SQUARE on the LEFT side, and TRIANGLE on the RIGHT side), and proposed an experiment to verify their model (Triesch and von der Malsburg, 1996).

Many neural network models for dynamic binding and compositionality have used phase locking of oscillators to express binding (Biederman and Hummel, 1992; Damasio, 1989; Horn and Opher, 1996; Hummel and Stankiewicz, 1996; Horn et al., 1991; Hummel, 2001; Shastri and Ajjanagadde, 1993). Others have used spiking neurons to carry out a variety of analog computations (Maass, 1997; Maass and Natschlager, 1997). Finally Bienenstock suggested that synfire chains may be used for such dynamic binding (Bienenstock, 1991, 1992, 1995, 1996; Bienenstock and Geman, 1994). However he did not show that this was actually possible.

In this work we introduce a neural network model for compositionality by dynamic synchronization of synfire chains. We will show that this model obeys most of the desired properties for compositionality and is consistent with biological cortical reality.

This paper proceeds as follows: In Section 2 we review the methods used for deriving the results. In Section 3 we introduce the binding mechanism, show its robustness and describe its limitations. In Section 4 we introduce a mechanism for competition among synfire chains and for limiting the amount of synchronization with in a system of synfire chains. Finally, in Section 5 we demonstrate compositionality by solving a simple part binding problem. Section 6 summarizes our results and discusses their relation to psychological and biological models.

2. Methods

In this work we used both simulations and numerical methods. The simulations were an extension of the one used by Abeles et al. (1993).

2.1. The Neuron

The main properties of the simulated neurons were:

- The dendritic and axonal topology was overlooked, such that the neurons were point neurons. However, synaptic events were modeled as currents injected into the cell body through the dendrites.
- Each neuron in the simulation had the following properties: the "cell body" exhibited passive capacity and leakage conductivities as well as active Na, K and Cl conductivities.
- The neuron dynamics had three modes:
  1. *Integrating.* The membrane potential followed Eq. (1). Slow adaptation was introduced by allowing the threshold to slowly follow the membrane potential.
  2. *Firing.* When the cell potential hit the threshold and the cell was not refractory it "fired" an action potential. The Sodium conductivity was raised for 1 milli-second.
  3. *Refractoriness.* The cell could not fire again for 3 ms (absolute refractoriness). K conductivity was turned on 1 ms after the cell fired, and decayed back with a time constant $\tau_{\mathrm{K}}^{\mathrm{ref}}$ (relative refractoriness).

  The membrane potential of the cell body (V) behaved according to:
  $$C\frac{dV}{dt} = g_{\mathrm{K}}(V - V_{\mathrm{K}}) + g_{\mathrm{Cl}}(V - V_{\mathrm{Cl}}) + g_{\mathrm{Na}}(V - V_{\mathrm{Na}}) + I \quad (1)$$
  where the resting potential was considered to be 0 mV. I is the current injected to the neuron which includes all synaptic sources as expressed by currents withdrawn from the cell body into the dendrites, and external currents injected directly into the soma, which was used to stimulate the neuron externally.
- The conductivities of the different ions were chosen such that the effective membrane time constant at rest was 5 ms. We also checked the dynamics of the system for smaller (1 ms) and larger (10 ms) membrane time constants. The threshold was set to 14 mV above the resting membrane potential.
- The excitatory synaptic current behaved like
  $$I_{\mathrm{excit}} \sim \begin{cases} te^{-\frac{t}{\tau_1}} & \text{for } t < \tau_1 \\ \frac{\tau_1}{e}e^{-\frac{t-\tau_1}{\tau_2}} & \text{for } \tau_1 \le t \end{cases}$$
  where τ_1 is the rise time, and τ_2 is the decay time constant. Note that this current did not represent the current flowing through the ionic channels at the synapse, but the current that a remote dendritic depolarization would draw from the cell body.
- The threshold level had a slight dependence on the membrane potential level. The dependence followed Eq. (2)
  $$\tau_{ac}\frac{d\theta}{dt} = -(\theta - \theta_0) + \alpha V \quad (2)$$
  where $\tau_{ac} = 20$ ms and $\alpha = 0.01$

Table 1. Values of parameters used in the simulation: The values of most parameters were similar to those used in Abeles et al. (1993), except for the membrane time constant which was increased to better fit experimental results.

Name	Description	Value
τ_1	Excitatory synaptic current, rise time	1 ms
τ_2	Excitatory synaptic current, decay time	5 ms
τ^I_{rise}	Inhibitory current, rise time	1 ms
τ^I_{Cl}	Inhibitory Cl current, decay time	5 ms
τ^I_K	Inhibitory K current, decay time	20 ms
τ_n	Noise current, decay time	5 ms
V_0	Membrane resting potential	0 mV
θ	Threshold potential	14 mV
τ_m	Membrane time constant at rest	5 ms
τ^{ref}_K	K conductivity decay time after firing	20 ms

- The noise current (I_{noise}) had a Gaussian amplitude distribution with a decay constant of τ_n ms. This noise represented the outcome of balanced random streams of excitatory and inhibitory inputs to the dendrites, derived from neurons which were not included in the simulations. The noise of the various cells was uncorrelated.
- Inhibition was represented by reducing the depolarizing currents. 17% of the inhibitory current had a decaying time constant of 20 ms, and the rest had a decay timeconstant of 5 ms.

Table 1 contains the values of most of the parameters used by the simulation.

2.2. The Network

- Inhibition was generated by a "pool" of neurons which had post-synaptic and pre-synaptic connections with all the excitatory neurons as well as to themselves. This pool was represented by one neuron whose output was analog (the firing rate of the pool).
- The current induced in each neuron was $I = I_{\text{den}} + I_{\text{ext}}$, where I_{den} was the dendritic current and I_{ext} was an optional external current. The dendritic current represented all the synaptic currents induced by the network as well as background noisy current produced by "other brain regions" which are not associated with the processes studied in the simulation.
- A "sensory input" to the network was represented by external current applied to a few designated neurons.
- Three types of connections between excitatory neurons were used:

  1. *Random Connections.* Each neuron had weak connections with 200 other random excitatory neurons from the network. These connections, along with the background noise, generated the background activity of the network.
  2. *Synfire chains architecture.* Synfire chains were embedded in the network by adding strong connection between pools of neurons.
  3. *Inter-chain connections.* Connectivity between synfire chains was embedded in the network by adding strong connections between pools of neurons from different synfire chains.

A second version of the simulation replaced the random connections by an effective noise.This version included only neurons from the synfire chains.

2.3. Synfire Chain Parameters

Synfire patterns may span a few hundred ms (Abeles et al., 1993). Diesmann et al. studies (1999), as well as ours have shown that it takes a wave approximately three milliseconds to propagate from one pool to the next. A synfire chain should have 100–300 pools to support such synfire patterns.The number of neurons within a pool can not be fixed by known experimental results, so a wide range of widths (30–100 neurons with in each pool) was used in the simulations. Other reports (Aertsen et al., 1996a, 1996b; Bienenstock, 1995) have used 100 neurons per pool. Each neuron in one pool may be connected to only a fraction of the neurons in the next pool. Values of around 50% have been used in many works (Aertsen et al., 1996a, 1996b; Abeles, 1991; Abeles et al., 1993) and also here. Possible effects of changing the number of inter-chain connections and their strength were studied in our work. Table 2 summarizes the synfire chain parameters which were frequently used.

In the simulation we measured the neurons' dynamic parameters, the collective parameters of pools (e.g., the number of firing neurons in a pool, correlation of activity between two synfire chains), and the activity waves (e.g., position of the wave in the chain and its velocity).

Table 2. Synfire chain parameters: The two synfire models which were used in the simulations.

Name	Description	Model 1	Model 2
L	Number of pools in the chain	100	300
W	Number of neurons in a pool	30	100
m	Intra-chain connections	15	50
m_{inter}	Inter-chain connections	2	5

2.4. *Transfer Matrix Method for Exploring the Stability of Synfire Chains*

The stability of propagation of activity along a synfire chain was analyzed (Abeles, 1991) using a transfer-function $N_{\text{out}}(N_{\text{in}})$ (The number of active neurons in the post-synaptic pool as a function of the number of active neurons in the pre-synaptic pool). This transfer function is defined solely by the synfire chain architecture. The fixed points of the synfire chain dynamics can be found through this method. This function is only deterministic when there is no noise in the system and activity (if exists) in each pool is exactly synchronous. When noise is present, responses become probabilistic and activity may be described by the pulse packet model (Aertsen et al., 1996a). Using this model the transfer function may be replaced by a transfer matrix P_{ij}, where P_{ij} is the probability that j neurons will fire in the next pool within the following τ milliseconds, given that i neurons fired in the pre-synaptic pool at the last τ milliseconds. An example of such a matrix can be seen in Fig. 3.

Figure 3 shows that the transfer function is the first moment of the transfer matrix

$$j(i) = \sum_{k=1}^{N} k P_{ik}$$

Here we use i, and j instead of N_{in} and N_{out} respectively, and N is the number of neurons in a pool.

To calculate P_{ij}, the pulse packet model (Aertsen et al., 1996a) was used. A stable isolated pulse packet was created. Then the post-synaptic firing time distribution was calculated when only i of the W neurons of the pool were active. The total activity within the first 5 milliseconds after the mean firing time of the pre-synaptic pool, yielded the mean activity level of the post-synaptic pool (μ_i). Then, using a binomial distribution, we calculated P_{ij}:

$$P_{ij} = \left(\frac{\mu_i}{W}\right)^j \left(1 - \frac{\mu_i}{W}\right)^{(W-j)} \frac{W!}{j!(W-j)!}$$

Using discrete time, a recursive formula may be used to calculate the probability $P(m, t = k+1)$ (the probability that m neurons of the $k+1$ pool are active):

$$P(m, t = k+1) = \sum_{i=0}^{W} P_{mi} P(i, t = k)$$

This method was used to study the effect of a wave in one chain on the velocity of a wave in another one, to study the effect of inhibition on waves, and the spontaneous generation of a wave in one chain by a wave propagating in another chain.

3. The Binding Mechanism

The building blocks of our compositional model are activity waves propagating along synfire chains. Dynamic binding of synfire chains is expressed by synchronizing the waves across chains. Such binding has already been introduced (Bienenstock, 1995), and implemented for synchronization between two chains (Arnoldi and Brauer, 1996). Assuming a few interconnections between two synfire chains (Fig. 2), when an activity wave is propagating along one synfire chain it will cause subliminal depolarization of neurons in the other synfire chain (Fig. 4). When there is no activity wave at the other synfire chain, this depolarization fades away. If within a short delay, an activity wave is propagating along the other synfire chain, the depolarization will speed up its velocity, until both waves synchronize.

3.1. *Speed Profile for Interaction*

For simplicity, we consider waves in two synfire chains as synchronized when activity waves in both chains exist at corresponding pools. The interaction between two such waves may be reduced to the effect of one wave on the velocity of the other. The detailed neural network simulation and the pulse packet were used for analyzing this effect. Figure 5 shows the velocity of the lagging wave as a function of its distance from the leading one, as was derived in both methods.

The neural network simulation was used to determine how waves moving along two synfire chains

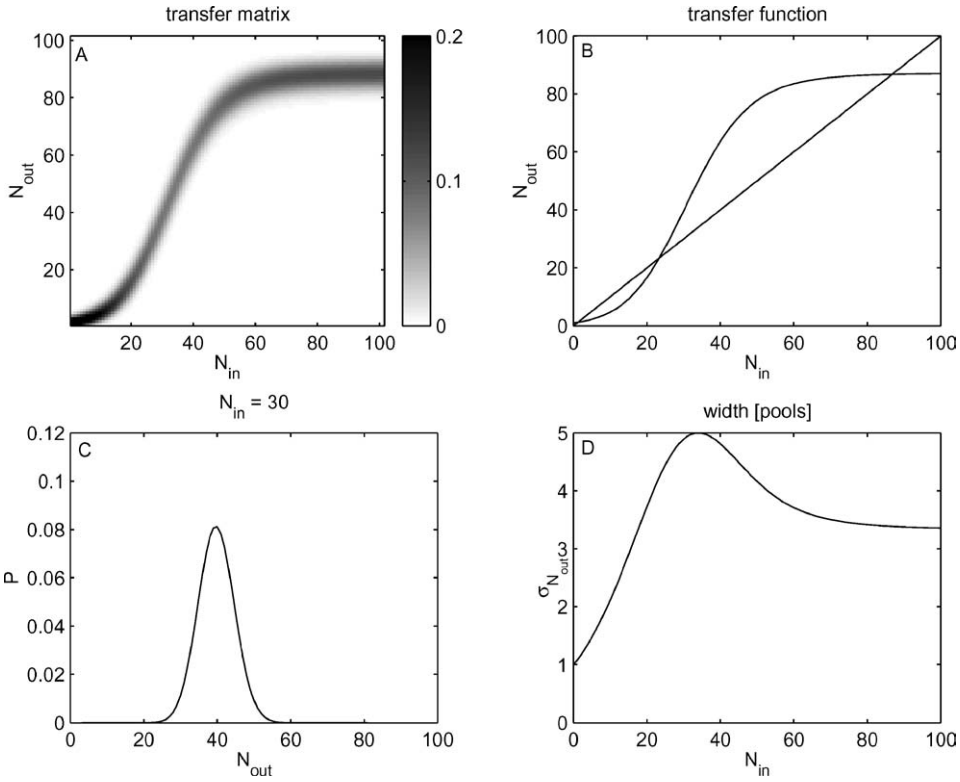

Figure 3. Transfer of activity from one pool to the next: A: The transfer matrix for one link with 100 neurons per pool. The X-axis is the number of active neurons in the pre-synaptic pool, the Y-axis is the number of active neurons in the post-synaptic pool, and the brightness of each point reflects the probability that N_{out} neurons in the post-synaptic pool would be active if N_{in} neurons were active in the pre-synaptic pool. B: shows the transfer function for the same case, that is the expected N_{out} of (A). If N_{out} of one pool serves as N_{in} of the next, as in a long synfire chain with uniform architecture, then the intersections with the line $N_{\text{out}} = N_{\text{in}}$ depict three putative fixed points ($N_{\text{out}} = N_{\text{in}}$) for propagation along the chain, as in Abeles (1991), two of which are stable. Near $N_{\text{in}} = 2$ only low spontaneous firing exists; near $N_{\text{in}} = 90$ almost all neurons in a pool fire within 5 ms and the intersection around $N_{\text{in}} = 25$ is not stable. Activity a little below it is expected to fade into the spontaneous rate, while activity a little above is expected to build into the almost full wave. C: A vertical cut through (A) at $N_{\text{in}} = 30$. This yields the probability that N_{out} neurons would become active in the post-synaptic pool when 30 neurons were active in the pre-synaptic pool. D: shows the dispersion of the number of neuronsin the post-synaptic pool relative to the number of active neurons in the pre-synaptic pool. The dispersion is maximal around the unstable fixed point and slightly above it. The above plots were derived for an isolated wave with 100 neurons within each pool, multiplicity of 50%, $\sigma_V = 4.5$ mV, $\theta = 14$ mV, $\tau_1 = 2$ ms, $\tau_2 = 5$ ms, and $\tau_m = 5$ ms.

synchronize. Figure 6 shows an example of the output of this simulation. On the plot the initial delay between the two waves is small, so, they synchronize after 60 milliseconds.

3.2. Synchronization Time and Range

We used the detailed simulation to study the relations in the initial delay between two waves and synchronization. When the initial delay is too large, the subliminal depolarization elicited by each wave in the other chain decays before the wave in the other chain arrives. To learn more about the statistics of the synchronization process we repeated each simulation, with the same parameter set and initial conditions, 40 times with different random seeds. For Model 1 chains (Table 2) we used an average connection between two neurons in successive pools of about $0.09\,\theta$, and an average connection between two neurons successive pools in two different chains of about $0.03\,\theta$. We measured the probability that the two chains would synchronize their activity before they reached the end of the chain. The results indicate that if the initial temporal delay between the leading and lagging wave is shorter than 16 ms, they will synchronize in more than 50% of the runs. For longer initial delays they will rarely synchronize (Fig. 7).

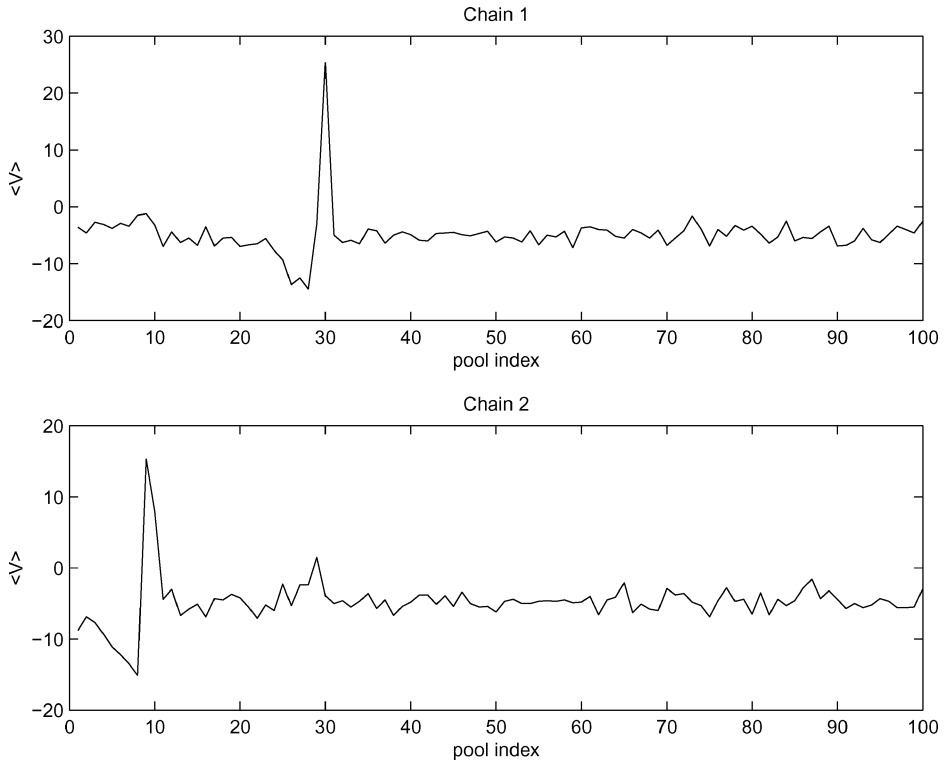

Figure 4. Intracellular voltages in two synfire chains. Each box shows the mean depolarization along 100 pools of a synfire chain not including action potentials. There is a peak at the position of the activity wave (at pool 30 in the first chain and at pool 9 in the second chain) followed by valley (relative refractoriness effect). There is a smaller peak at the appropriate position on the other chain (pools 8–10 in chain 1 and pools 28–30 in chain 2). This peak decays to baseline within a few pools.

A similar analysis with long synfire chains (*Model 2* with 300 pools) shows that the maximal initial delay for synchronization may be larger (23–27 milli-seconds), yet it takes a long time to synchronize waves that start with such delays (about 500–600 ms).

To study the relations between the initial delay and the synchronization time, we again used the statistics from the simulation, and found (Fig. 8) that the synchronization time increases almost exponentially with the initial delay.

Reaction time in psychophysical experiments which require simple binding is 400–600 ms (Cohen and Shoup, 2000). Since such tasks include processes beyond simply binding, the binding time must be less than 200 ms. This corresponds to a maximal initial delay between waves of 15–20 ms. Arnoldi and Brauer obtained similar values for the maximal initial delay (Arnoldi and Brauer, 1996). With other architectures and/or synaptic time profiles, there may be longer maximal delays.

3.3. Bounds on the Binding Mechanism

The synchronization time and maximal initial delay for which waves become synchronized depend on the strength of the connections between neurons in the same chain (intra-chain) and neurons from different chains (inter-chain). Strength is defined as the product of the multiplicity of connections by the strength of each synaptic contact. In the above example we used a ratio of 1:15–1:20 between the strength of intra-chain connections and inter-chain connections. Strong connections between chains enable synchronization when the initial delay is larger. However with strong inter-chain connections a wave propagating along one chain may spontaneously elicit a wave in the other chain. We denote this phenomenon 'spontaneous synchronization', because the second wave will arise in synchrony with the first wave.

We studied the process of spontaneous synchronization using the transfer matrix method (Section 2.4).

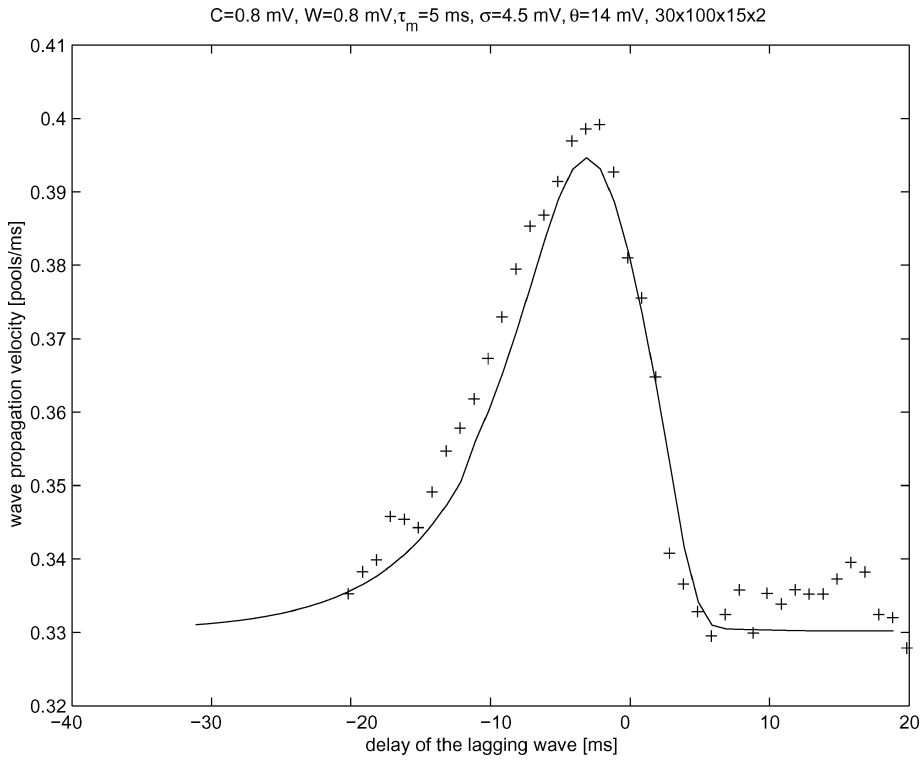

Figure 5. Velocity profile during synchronization: The X-axis is the time difference between the leading and lagging waves. When this difference is above 25 ms, the effect of the leading wave is negligible and the two waves progress at the speed of one pool every 3 ms. When the delay to the lagging wave is below -5 ms (i.e. the lagging wave is actually in front of the leading one) there is no effect either. In between, the lagging wave is being speeded up, reaching a maximal speed at approximately 3 ms. The position of the peak depends on the cross connection architecture and the shape of the EPSP. Here the cross connections are from pool i of one chain to pool $i+1$ in the other, and the EPSP builds gradually within 3 ms. Had the cross connection been from pool i in one chain to pool $i+2$ in the other the peak would be around 0 delay. The solid line results from numerical analysis using the transfer matrix. The results are similar to those obtained by simulation (crosses). Simulation parameters as in Table 1 and Model 1 of Table 2. The synaptic strength within the chain was 0.8 mV, and between chains 0.8 mV.

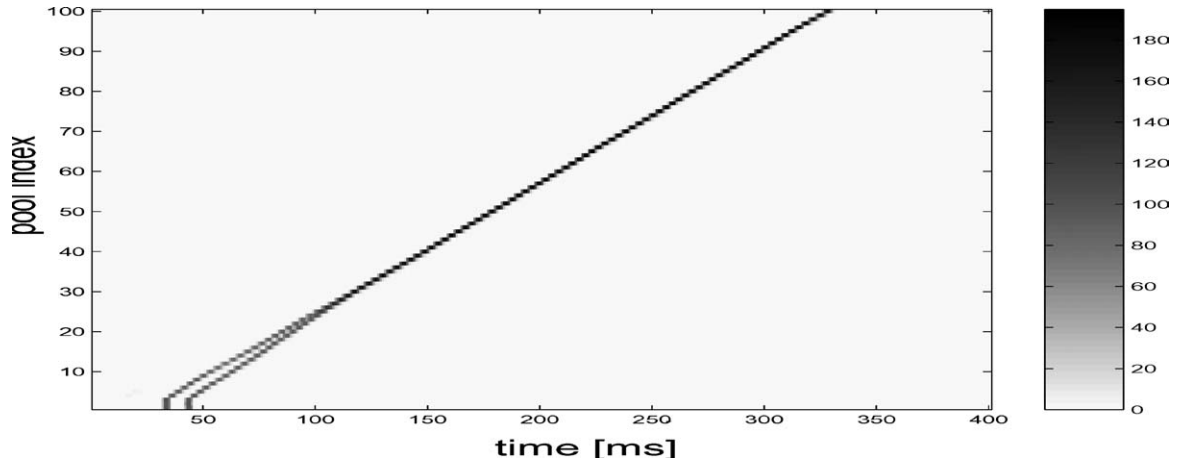

Figure 6. Evolution of activity in two synfire chains: The simulation results for a system of two synfire chains. The X-axis is time in milliseconds, the Y-axis is the position along the chains, and the gray scale of the pixel is the number of active neurons at a pool at a given time point. In each chain an activity wave was initiated by injecting strong current to the neurons of the first three pools. The initial delay was 11 ms and the waves become synchronized after a short while.

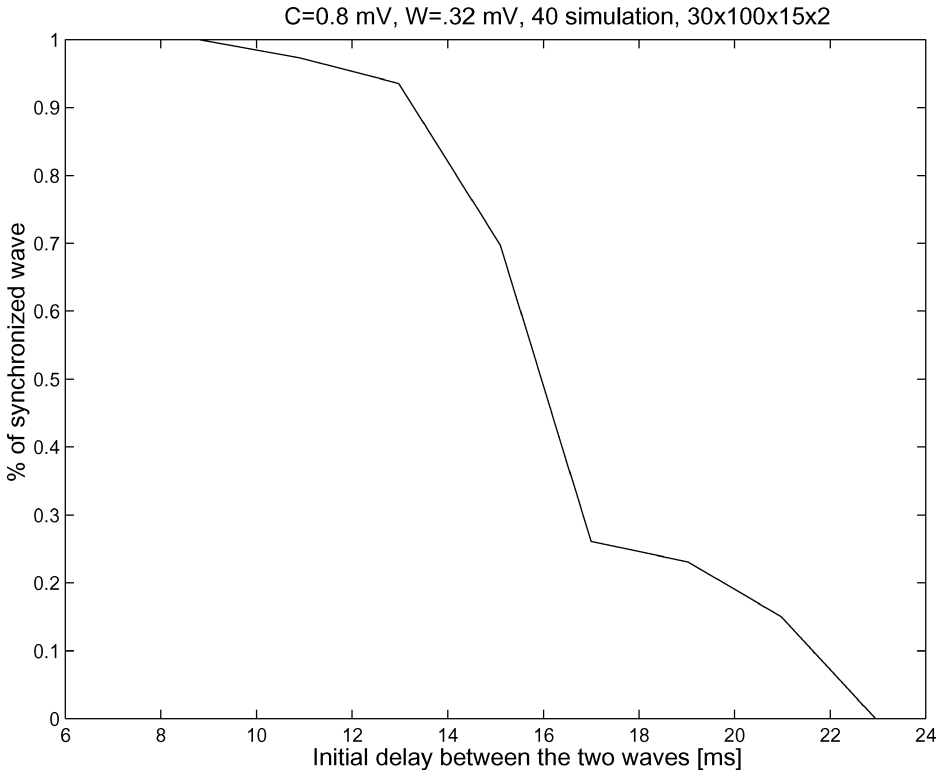

Figure 7. Probability of synchronization. The probability that waves of activity in two synfire chains will become synchronized as a function of the initial delay between them. When the initial delay between the two waves is less than 13 ms they almost always synchronize their activities within 100 pools. When the initial delay is larger than 18 ms they rarely become synchronized. Using 50% as criterion, the maximal initial delay for synchronization is around 16 ms.

The architecture included two synfire chains with some connections between them (Fig. 2). The effect of the following parameters were studied: the number and strength of the inter-chain connections; the intra-chain connections; and the level of background noise.The latter determines the spontaneous firing rate of the neurons (Fig. 9).

We first calculated a transfer matrix for the activity in chain no. 2, given a wave in chain no. 1 (Fig. 10(A)). Then we used this matrix to analyze the effect of a traveling wave in chain no. 1 on the activity in the corresponding pool of chain no. 2. This was done by calculating the probability that N neurons within a pool of the second chain would fire when the wave in the first synfire chain reached the corresponding pool (Fig. 10(B)).

The transfer matrix of the second chain (Fig. 10(A)) is biased toward the creation of new waves (e.g., when 40 neurons fire in pool i, 60 to 80 are expected to fire in pool $i + 1$), so it is expected that activity should build up until a full blown activity wave will propagate along the second chain. Figure 10(B) shows that at each pool the activity along the second chain builds up, until after10–15 pools it becomes a full blown activity wave (at least 80 of the 100 neurons of the pool are firing).

We repeated the above analysis for different values of background noise and inter chain connectivity. For each set of parameters, we estimated the probability of generating a spurious wave as the probability that more than half of the neurons in pool 70 of the second chain fired when the activity wave in the first chain reached pool 70.

Figure 11 shows that spontaneous synchronizations are very rare for low background noise ($\sigma_v < 4.2$ mV, which means a background firing rate of 1 spike per second or less). For stronger noise levels the need to avoid spontaneous synchronizations sets a limit on the maximal possible inter-chain connection strength and therefore shortens the maximal synchronization distance. Similar results were found in detailed simulations.

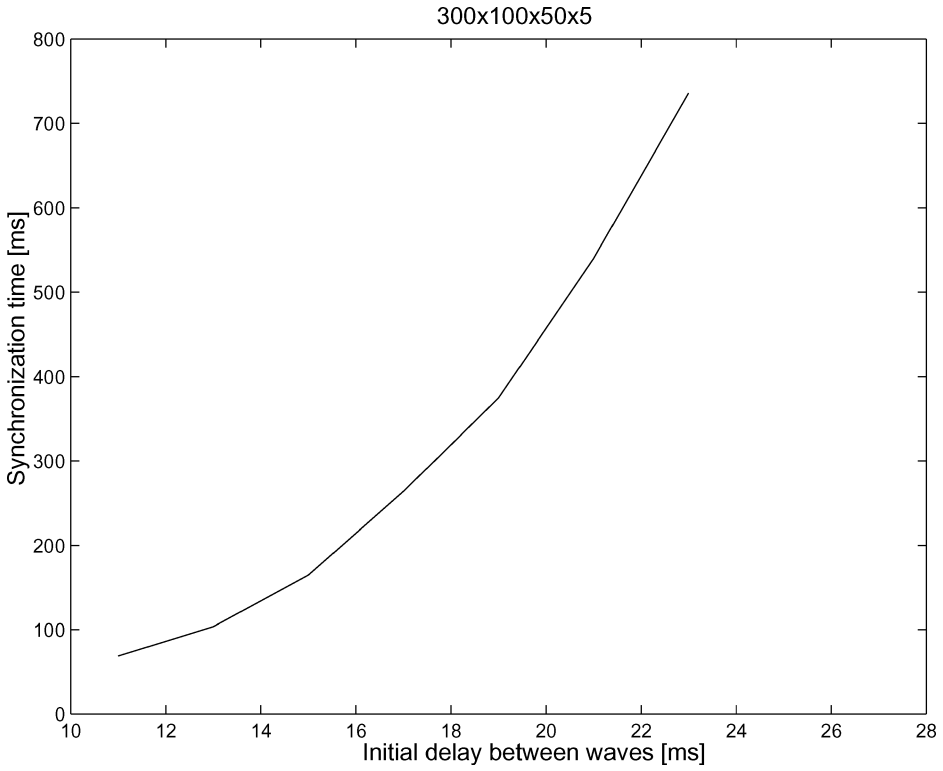

Figure 8. Effect of the initial delay on the synchronization time.

4. Competition Among Synfire Waves

Composite objects may be represented by binding simpler components. In our model, each component is represented by an activity wave propagating along a synfire chain, and binding between components is achieved by synchronization of these activity waves. Components may bind in several ways to form completely different composite objects (like words which are composed of characters). In such a scenario almost all synfire chains will be connected to each other. When two distinct objects are presented simultaneously, there is a danger that all the component-waves will synchronize to represent one mega object.

To avoid such situations some form of competition among synfire chains is needed. This may be achieved by inhibitory neurons that increase their activity when waves become synchronized. To test the feasibility of such a mechanism we first evaluated the stability of bound and unbound waves (Section 4.1), and then introduced a modified synfire chain which increases the sensitivity of inhibition to synchronization (Section 4.2).

4.1. The Effect of Inhibition on the Stability of Waves

Representing composite objects by synchronized waves requires that these waves be more robust to inhibition, such that when there is a competition between a group of synchronized waves and a solitary wave the latter will fade away, while the synchronized waves will remain. To examine the effect of inhibition on the stability of synchronized and unsynchronized waves, we simulated such waves in synfire chains subject to constant inhibition. When less than 50% of the neurons in a pool were active we considered the wave to be extinguished. Using the transfer matrix method we repeated this experiment for different values of inhibition level. Where, inhibition level is the amount of inhibitory current injected into the neuron. Figure 12 shows the range of inhibition where synchronized waves are stable and unsynchronized waves fade away. The width of this region depends on the amount of synchronization between the chains. Stronger connections between synfire chains broaden this region. The stability of the waves is almost binary: they are either stable for the entire chain length, or fade

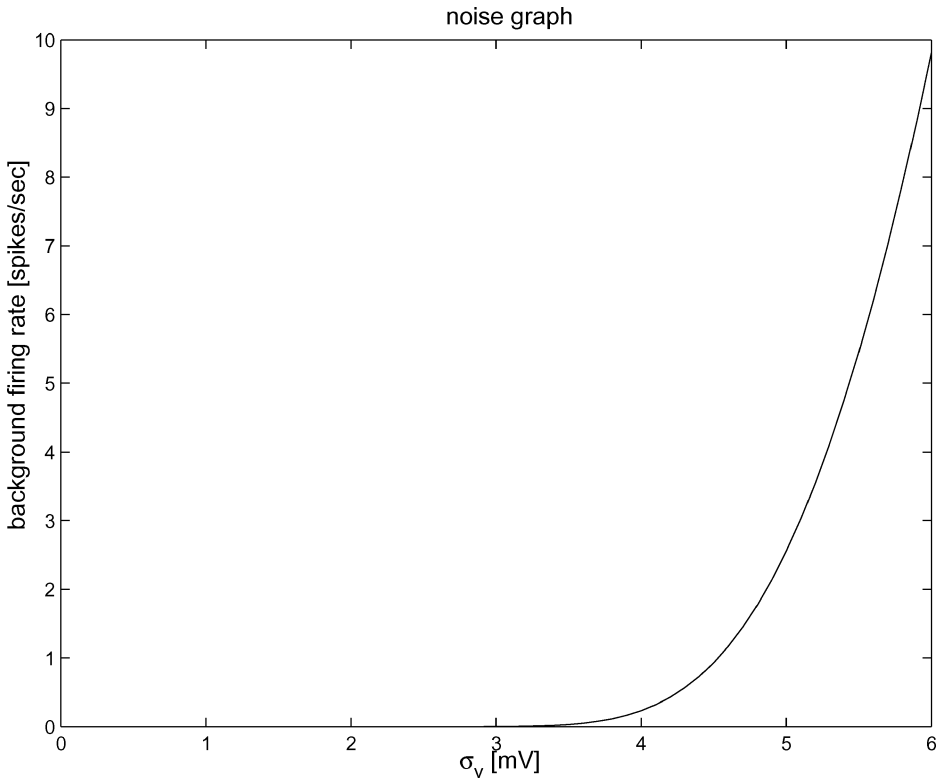

Figure 9. The background firing rate relative to the noise level σ_v.

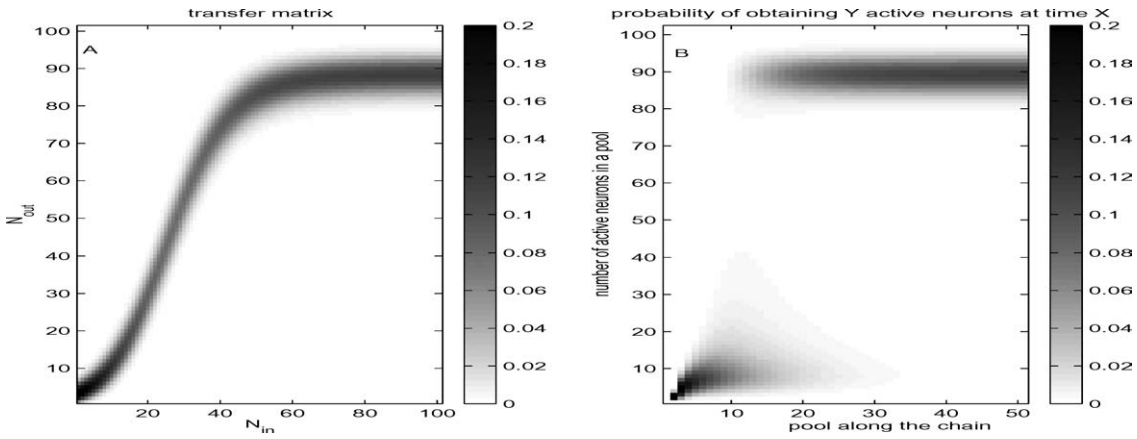

Figure 10. Example of transfer between two chains: Chains as for Model 2 of Table 2. In these conditions spontaneous synchronization occurred. A: The transfer matrix for the second chain when biased by a wave traveling along the first chain. B: The probability that N neurons would fire at pool t of the second chain when an activity wave along the first chain reached pool t. Intra-chain EPSP equals inter-chain EPSP of 0.5 mV, and background noise of 4.5 mV. The inter chain connectivity was about $\frac{1}{10}$ of the intra chain connectivity.

away within 10–20 pools, which is concordant with results that propagation along synfire chains has two stable fixed points stable propagation (almost all neurons fire in synchrony), or decay. All waves converge to one of 15 these two modes within a few pools.

Inhibition can be effective in selecting between synchronized and unsynchronized waves if:

1. When none of the waves are synchronized, unsynchronized waves propagate freely.

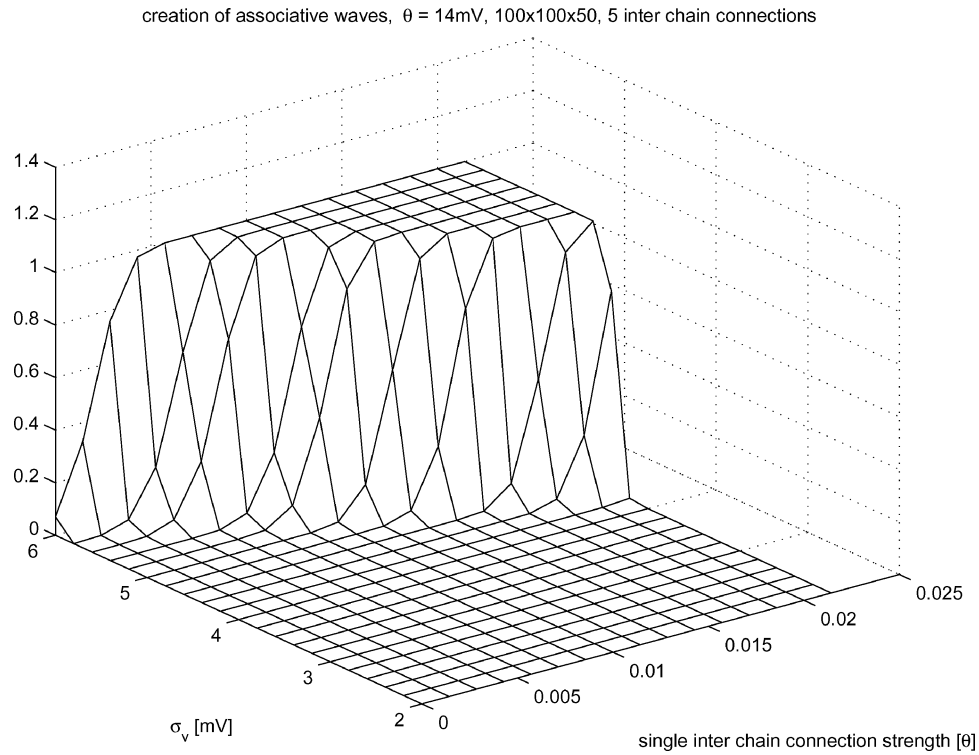

Figure 11. The probability of spontaneous synchronization: The probability of spontaneous synchronization versus the background noise level (σ_v) and the strength of an inter-chain connection. The noise level is given in mV, the EPSP of the inter-chain synapse is given as a fraction of the threshold. The graph shows that for a low noise level ($\sigma_v < 4.2$ mV) there is no spontaneous synchronization. Yet, for noise levels that can account for spontaneous firing rates of more than 2 spikes per second, spontaneous synchronization occurs at moderate levels of inter-chain connections.

2. When some waves are synchronized, the activity of the inhibitory neurons increases to the level where synchronized waves are stable and unsynchronized waves fade after a short period of time.

Thus, the inhibitory level before synchronization should be below 1.2 pA, and synchronization of waves should increase the inhibition to the level of 1.2–1.5 pA, where synchronized waves are stable while unsynchronized are not. Synchronization between activity waves increases the propagation velocity of these waves and therefore the number of units firing per unit time increases. Thus, even if the inhibition is only affected by the global firing rate, synchronization will lead to enhanced inhibition.

4.2. Inhibition Which is Sensitive to Synchronization

By adding a few inhibitory neurons to each pool in the synfire chain (Fig. 13) one can further increase the amount of inhibitory current when waves are synchronized.

Each inhibitory neuron receives synaptic connections from the excitatory neurons of the previous pool, but inhibits the entire network. Preferably, the strength of the excitatory connections to the inhibitory neuron are lower than the strength of connection between two excitatory neurons. In this way, when unsynchronized waves propagate along a synfire chain only a small fraction of the inhibitory neurons fire, while when waves are synchronized most of the inhibitory neurons will. Note that each inhibitory neuron may take part in many synfire pools, so that the total number of inhibitory neurons may be considerably smaller than the excitatory ones. Each inhibitory neuron in one chain also receives excitatory connections from corresponding pools in synfire chains which may become synchronized with its chain. Such inhibitory neurons are expected to act in one of the following modes:

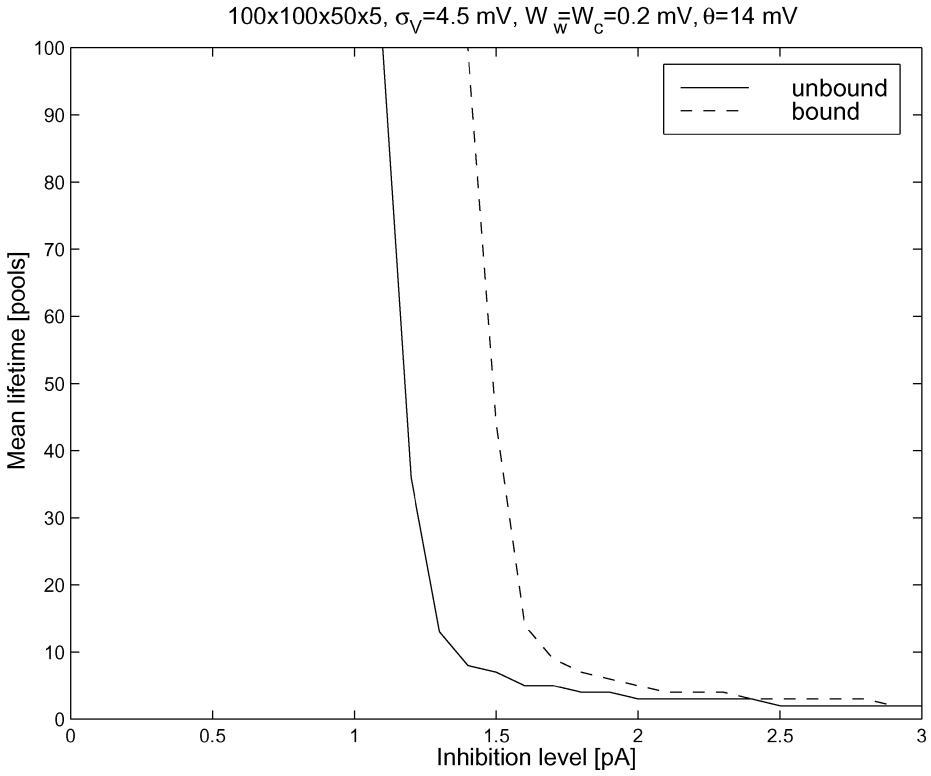

Figure 12. Effect of inhibition on synfire waves: Mean lifetime of unbound and bound waves in synfire chains as a function of the inhibition level (total inhibitory current injected into the neuron). The mean lifetime of a wave is the number of pools it traversed until the probability of its decay is 50%. The synfire chains were as described by Model 2 in Table 2. The background noise had $\sigma = 4.5$ mv, the strength of each connection (both inter-chain and intra-chain connections) was 0.2 mV. The bound wave was more stable for all inhibitory levels. The maximal difference was in the 1.2–1.5 region, where unbound waves fade away almost immediately, whereas bound waves are very stable.

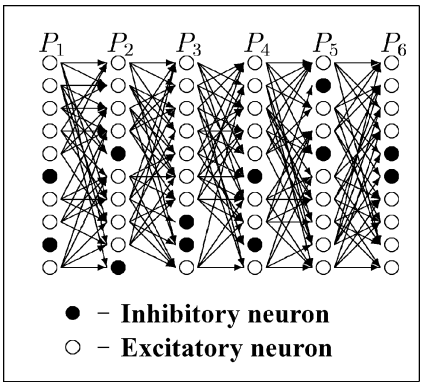

Figure 13. Synchronization sensitive synfire chain: Each inhibitory neuron is activated by excitation from neurons in the pre-synaptic pool and inhibits the whole network. We used synfire chains with 100 excitatory neurons and 20 inhibitory ones in each pool. Each excitatory neuron provides 50 connections to excitatory neurons in the next pool and 10–25 connections to inhibitory neurons in the next pool. All the connections between excitatory and excitatory neurons and excitatory to inhibitory neurons have the same strength.

1. When there is no activity wave traveling along the synfire chain, the inhibitory neurons fire at a low background activity level, and the inhibition is very weak.
2. When an unsynchronized activity wave is propagating along the chain, the tendency of the inhibitory neurons to fire increases and a partially saturated wave of inhibitory activity propagates along with the excitatory wave. This causes some increase in the inhibitory level.
3. When synchronized waves propagate along two chains, inhibitory neurons acquire more pre-synaptic inputs, increasing the saturation of the inhibitory wave. Synchronization amongst several waves will result in an even higher inhibitory level.

In this way, once some waves bind to each other the increased inhibition slows down the other waves and lowers the likelihood of their becoming synchronized too (preventing the formation of one mega object), or

even causes the unsynchronized waves to decay and disappear (separating between object and background).

4.3. The Effect of Synchronization on the Inhibitory Level

We explored the effect of unsynchronized as well as synchronized waves on the inhibitory level by the pulse packet model (Aertsen et al., 1996a). Graphs of the firing probability of the inhibitory neuron versus the number of pre-synaptic excitatory to inhibitory connections were plotted for unsynchronized and synchronized waves (Fig. 14). Extensive excitation from the previous pool will increase the firing probability of the inhibitory neuron. Hence, additional excitatory input from other waves will only have a marginal effect on its firing probability and therefore the synchronization effect on the inhibition will be weak. In contrast when the number of connections from the previous pool is small, the firing probability due to propagation of unsynchronized waves is low, and the effect of synchronization with other waves is much stronger. Thus, effective sensitivity to synchronization is achieved when the vertical distance between the two graphs of Fig. 14 is large.

From Fig. 14 we expect the regulating mechanism to work well when the firing probability in the presence of unsynchronized waves is less than 50% (In our architecture there were less than 25 synaptic connections between excitatory neurons and an inhibitory neuron in a subsequent pool of the same chain). Adding more inter-chain connections will increase the difference between the firing probability for synchronized and unsynchronized states, but also increase the probability of eliciting spurious waves.

5. Hierarchy

We use the simple part binding problem to illustrate how synchronization among synfire chains may resolve a complex scene as shown in Fig. 15.

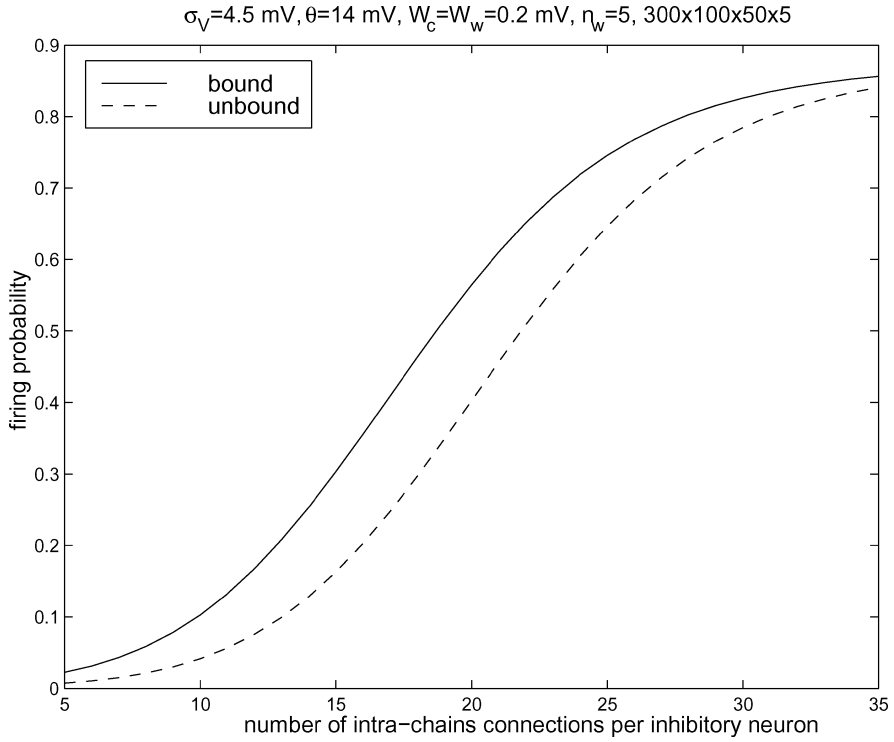

Figure 14. The effect of excitatory inputs on the inhibition: The firing probability of an inhibitory neuron versus the number of its excitatory inputs is plotted when an unbound (dashed line) or bound (continuous line) waves are propagating along the chain. The number of connections to inhibitory neurons from other chains per pool is fixed (5 connections per neuron). It is apparent that binding increases the inhibitory level in all cases. When the firing probability for an unbound wave is smaller than 50% (there are less than 25 synaptic connections between excitatory neurons and an inhibitory neuron in the subsequent pool), binding of the wave increases the inhibitory level by more than a third of its original level. Each chain was as described by Model 2 in Table 2. The background noise level was $\sigma_V = 4.5$ mV.

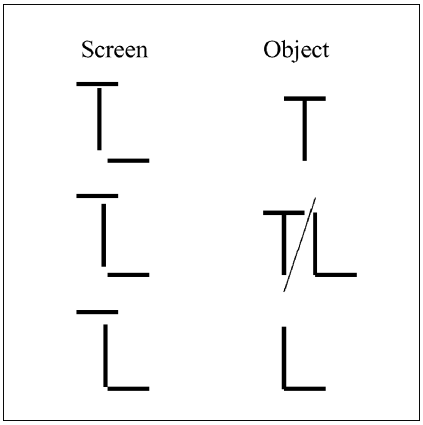

Figure 15. A simple part binding: On the left three possible images are shown. In the upper case, the upper two lines combine and form the letter 'T'. In the middle, either letter ('T' or 'L') is possible, but not both of them simultanously. In the lower case, the lower two lines combine and form an 'L'.

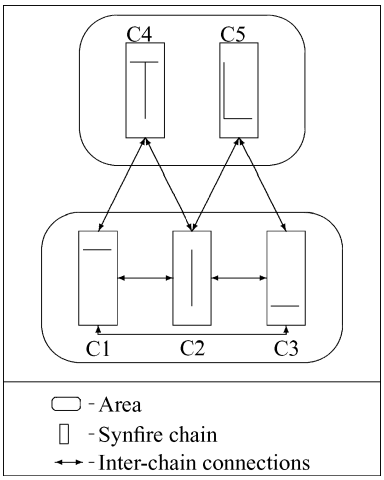

Figure 16. Network architecture: Synfire chains 1,2,3 represent lines, whereas chains 4 and 5 represent symbols. All the chains within an area share inhibition. There are weak connections between the first three chains. Chain 4 is connected to chains 1 and 2 in the lower area, and chain 4 is connected to chains 2 and 3. All the chains include excitatory and inhibitory neurons.

The three lines in Fig. 15 do not make much sense. However, pairs of lines do (either the letter 'L' or 'T'). We extend our model by assuming that on a higher level the abstract notions of T and L are represented, such that when the first two chains are synchronized they initiate an activity wave in a synfire chain representing the letter 'T'. Binding the second and third chains initiates an activity wave at a chain representing the letter 'L'. The system architecture is shownin Fig. 16, where the five chains include both excitatory and inhibitory neurons (Fig. 13). The inhibitory connections are local to their level (Braitenberg, 1986); that is, the inhibitory neurons from the lower level inhibit all the neurons in the lower level, whereas the inhibitory neurons from the upper level inhibit all the neurons in the upper level. The relative positioning of the lines is represented by the initial times of the activity waves in the appropriate synfire chains.

To investigate under which conditions chain 4 (or chain 5) became active, we fixed the temporal distance between activation of chains 1 and 3 at 26 ms and studied the effect of the time of activation of a wave in chain 2 (in between 1 and 3) on the system dynamics.

5.1. Part Binding Problem

Simulation results for such a system are shown in Fig. 17. The top panel shows the activity of all three synfire chains at the lower level. Activity is started in chain 1 at 60 ms by strongly stimulating its first 3 pools, 13 ms later chain 2 is activated and 13 ms later chain 3 is activated. At the beginning all 3 chains are active and weak activity arises in chains 4 and 5 of the higher level. The closer waves 1 and 2, the stronger the activity in chain 4, and shortly after waves 1 and 2 synchronize, a full blown activity wave is seen in chain 4. The competition between waves from the same level is apparent by the fact that synchronization of the first two waves stops the activity of the third wave, and the fact that the stronger the activity in chain 4 the weaker the activity in chain 5.

The confusing initial scene was resolved into a 'T', and the components of this letter are recognizable by the fact that the waves representing these components are synchronized with the wave representing the T.

By repeating simulations such as in Fig. 17 for 48 times we estimated the probability of retrieving 'T' (or 'L'). Figure 18 shows that the system succeeded in retrieving a representation under all conditions. When two of the lower level waves are close, it chose the same representation on every trial. For ambiguous scenes both representation are possible, yet only one of them was retrieved at any given simulation run.

Note the slight asymmetry in the crossover time (57% of the delay between 1 and 3). Thus, synchronization of the leading two waves is preferred upon the synchronization of the lagging two waves; that is, binding of early shown components is more probable.

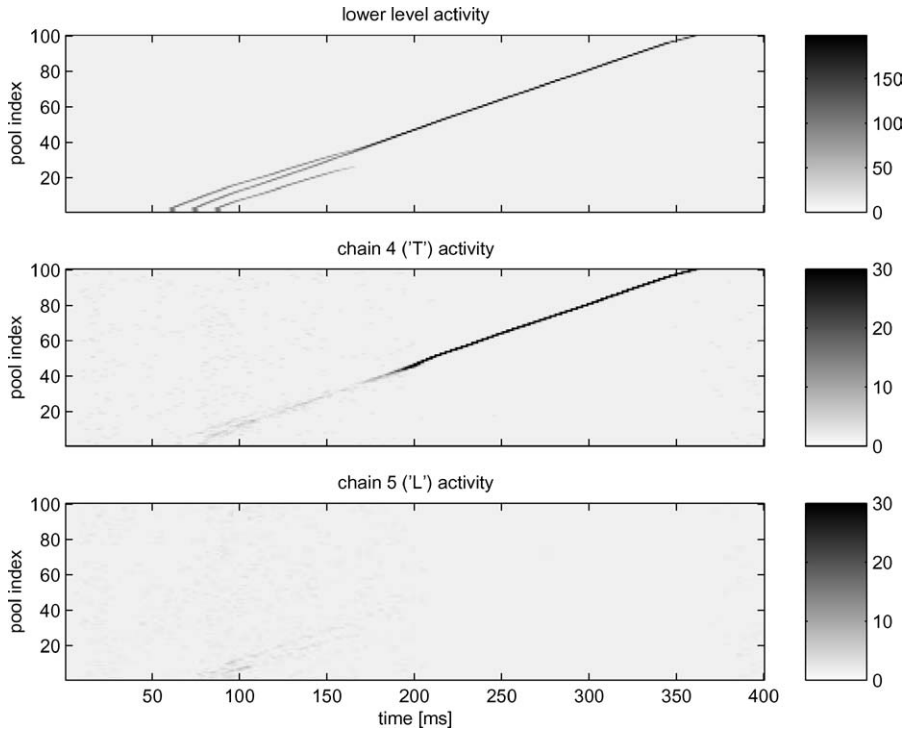

Figure 17. Example of activities during binding: The activity of the five synfire chains during a simulation with $d_{12} = 13$ ms and $d_{13} = 26$ ms. The X-axis of each box is the simulation time, the Y-axis is the pool index, and the color of the pixel describes the number of active neurons in a pool at a certain time. The upper box describes the activity of the three synfire chains in the lower level (the lines). The middle box describes the activity of chain 4 (the representation of the letter 'T'). The lower box describes the activity of chain 5 (the representation of the letter 'L'). The activity of the three chains from the lower level initiates some activity at the two chains from the higher level. When the leading two waves become synchronized, an activity wave starts propagating along chain 4.

5.2. Priming

When a reader sees the characters 'SHAL' he expects to see another 'L'. If the next character is ambiguous, the probability of interpreting it as 'L' increases. Such effects are called priming in the psychological literature. We tried to model such priming by increasing the background firing of the neurons in the chains representing 'L', without creating activity waves in this chain. This was done by decreasing the threshold of the excitatory neurons in the appropriate synfire chain by 1.1 mV.

Figure 19 shows an example of the priming effect on the symmetric initial conditions (chain 2 is excited exactly in the mid-time between chains 1 and 3). While in the unprimed situation the system retrieved 'T', it now retrieved 'L', because partial synchronization between waves 2 and 3 at the lower area was enough to initiate an activity wave at the higher area, which in turn facilitated synchronization between the component waves.

Priming was tested for the same range of initial conditions as in Fig. 18. The priming effect (Fig. 20) on the two probabilities graphs is different: (1) The additional excitation makes the retrieval of 'L' easier almost without changing the behavior of its probability shape, (2) for the other chain ('T') there are two effects: (a) the fact that it is easier to retrieve 'L' makes it harder to retrieve 'T' and lower its retrieval probabilities, (b) increasing the background firing of some neurons in the area (those of chain representing 'L') caused an increase in the inhibition level which reduced the probabilities of retrieving 'T' even more. The first two properties caused a left shift of the range of confusion. The last one caused some asymmetry in the graph because the additional inhibition made retrieval of any solution harder.

The above example illustrates an important property of the hierarchical system of synfire chains; mainly, its ability to use a bidirectional stream of effects in resolving an ambiguous scene. The bottom-up stream is clear: activity waves in the lower level of the system may

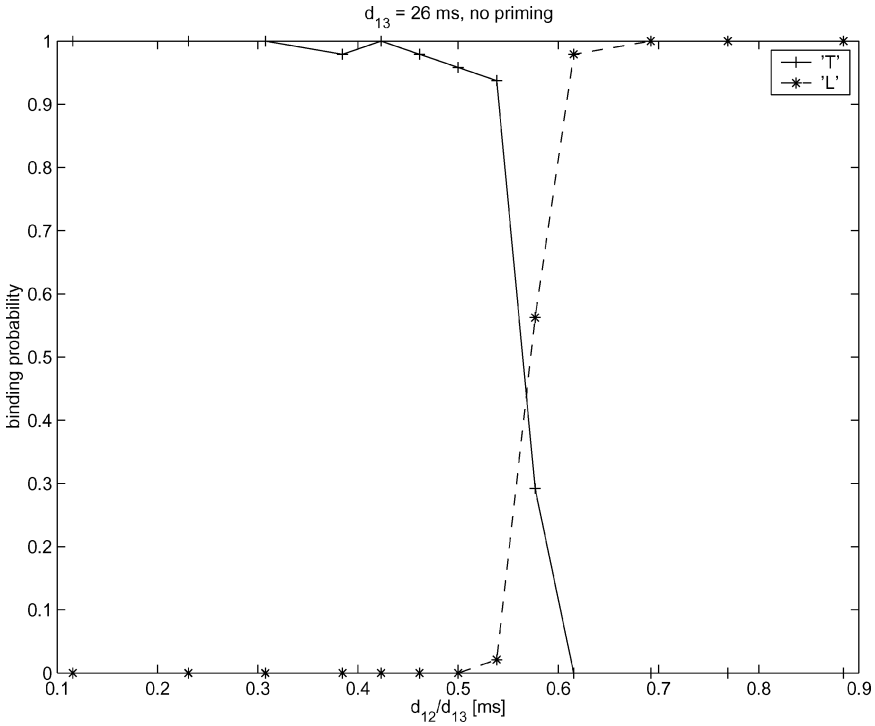

Figure 18. Retrieving 'L' or 'T': Retrieval probability of the letters 'L' and 'T' versus the initial temporal distance between the leading two waves. When the two parts of the letter 'T' are very close (small d_{12}) the system always retrieved this representation. When the two parts of the letter 'L' are close, it retrieved this letter representation, while at an intermediate range both representation were possible. The model almost always (in more than 88% of the simulations) resolved the scene and selected either 'T' or 'L'.

initiate activity at the higher level. The top-down stream is demonstrated in Fig. 19, when around 100 ms partial synchronization of waves 2 and 3 initializes some activity at chain 5. This stabilizes the activity waves at chains 2 and 3 in the lower level and helps them to overcome the possible synchronization between chains 1 and 2. The competition at both levels is shown both by the fact that the leading wave fades due to synchronization of the lagging two waves, and the negative correlation between the activity in chains 4 and 5.

6. Discussion

The issue of defining the neural mechanisms for binding is being extensively debated (e.g. Gray, 1999; Reynolds and Desimone, 1999; Riesenhuber and Poggio, 1999; Roskies, 1999; Shadlen and Movshon, 1999; Singer, 1999; Treisman, 1999; von der Malsburg, 1999). Shadlen and Movshon (1999) argued that unless connectivity is very low, the network is small, and the firing-rates are low, synchronization cannot be used to mark the synchronized component, nor can it be used to elicit a higher level representation of a compound object. Our simulations demonstrate, by way of example, that these claims are false. Our simulations of a part-binding problem contained 50,000 excitatory neurons, and the background noise in these simulations represented the uncorrelated activity of approximately 100,000 spikes per neuron per second. Nevertheless the system could easily resolve the two different scenarios (two out of three lines represent either the character 'L', or 'T'). Synchronization here was not just a signal of what is synchronized to what (as Shadlen and Movshon claimed), but also the mechanism by which the components selected the higher level representation. The main difference between our simulations and their argument is that they assumed all neurons of a component (or object) should fire together, whereas our synfire representations are based on orderly sequences of co-firing neurons.

The main features of the present neural-network model for compositionality may be summarized as follows:

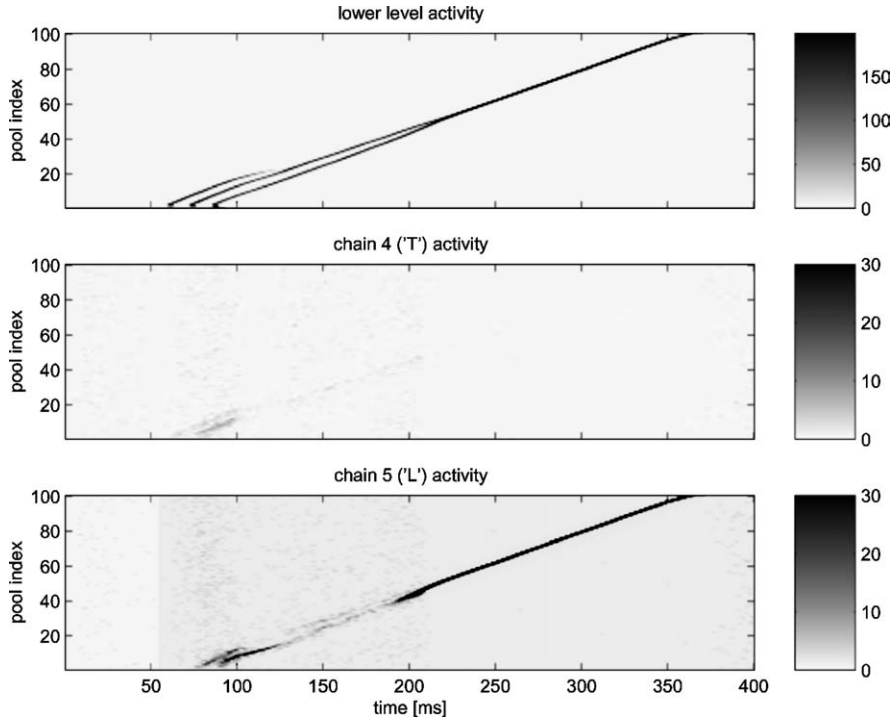

Figure 19. Example of activity during binding with priming: The activity of the five synfire chains along the simulation with priming for the letter 'L' (chain 5). $d_{12} = 13$ ms and $d_{13} = 26$ ms. The X-axis of each box is the simulation time, the Y-axis is the pool index, and the color of the pixel describes the number of active neurons in a pool at a certain time. The upper box describes the activity of the three synfire chains in the lower level (the lines). The middle box describes the activity of chain 4 (the representation of the letter 'T'). The lower box describes the activity of chain 5 (the representation of the letter 'L'). The activity of the three chains from the lower level starts some activity at the two chains from the higher level. The priming effect can be seen as some low level of activity in all the pools of chain 5 (Starting at 60 ms). Chain 5 activity strengthens the activity waves at chains 2 and 3. After a short time the activity wave in chain 1 dies out, and the lagging two waves are synchronized. Their synchronization strengthen the activity wave in chain 5 and stop the activity at chain 4.

1. *Components are represented by activity waves propagating along synfire chains*. The relation between the different components is mapped to relative delay between these activity waves. Each synfire chain represents one element ('a line', 'a junction', 'a letter', 'some part of an object' , . . .).
2. *Binding is expressed by synchronization among activity waves at different synfire chains*. The relative delay between activity waves in different synfire chains provides the binding information (when the relative delay is large they are detached, whereas when it is short they are bound). A few cross links between synfire chains allows synchronization when the relative initial delay is not too large.
3. *Global inhibition can regulate the total amount of synchronization within each group of synfire chains*. The sensitivity of the global inhibition to the amount of synchronized waves may be increased by adding inhibitory neurons to the synfire chains. These are activated in the same manner as the excitatory neurons in the chain, but inhibit all the neurons in the region.
4. *Solutions to binding problems are facilitated by introducing a hierarchy of synfire chains*. Each item is represented by a synfire chain. Objects of similar complexity reside in one 'area' and share inhibition. Reciprocal links between synfire chains in different areas represent which components might comprise a compound object. The connections between synfire chains from different areas and the existence of competition only within an area (by the global inhibition) constitute the 'grammar' of the synchronization process.
5. *Reciprocal links between areas provide the means for simultaneous bottom-up and top-down binding*. When two objects from different hierarchy levels are connected, they mutually affect each other. Synchronization of chains from lower levels initiates

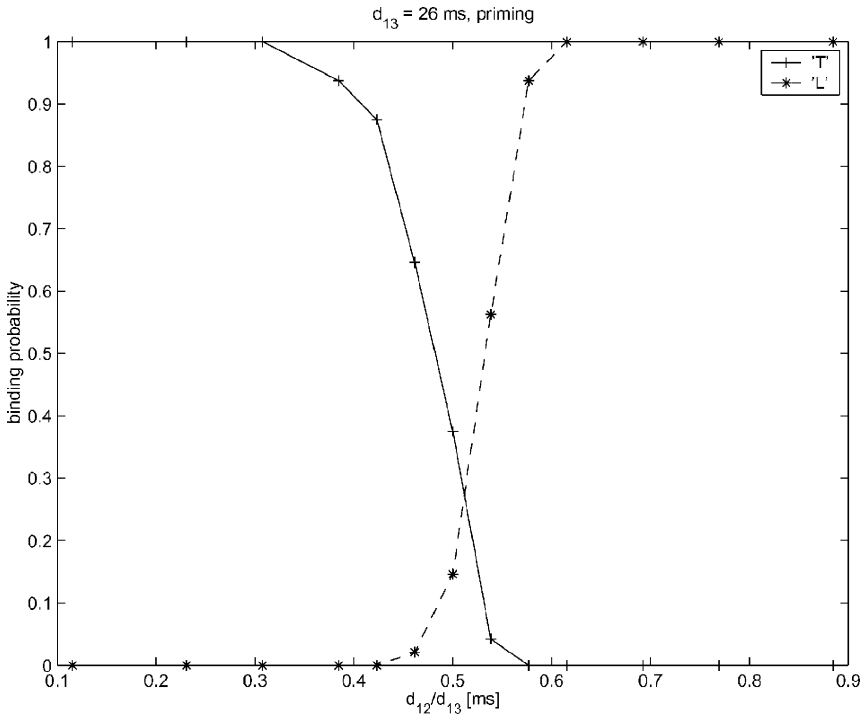

Figure 20. Retrieving 'L' or 'T' with priming: Retrieval probability of the letters 'L' and 'T' versus the initial temporal distance between the leading two waves. When the two parts of the letter 'T' are very close (short d_{12}) the system always retrieves this representation. When the two parts of the letter 'L' are close, it retrieves this letter representation, while at an intermediate range both representations are possible. The main effect of priming is seen as a shift of the intermediate region to the left.

activity at higher levels, and bound activity from higher level stabilizes the activity of synfire chains from lower levels.

Several of the points above were suggested and some were studied in the past. Representation by waves in synfire chains were explicitly or implicitly assumed by most studies on synfire chains. The process of synchronization between two synfire chains was studied in simulation by Arnoldi and Brauer (1996). The suggestion that such synchronization is the basis for binding was first made by Bienenstock (1991, 1996). We have extended these studies by more quantitative studies and by introducing the transfer matrix method, which allows for numerical evaluation (rather than detailed simulations) of the behavior of waves in a synfire chain.

To the best of our knowledge, the issue of competition among synfire chains has only been raised and studied quantitatively here. This is also the first time the advantage of including inhibitory neurons in synfire chains, and the quantitative study of its effects have been investigated. A somewhat similar structure of a multi-layered feed-forward network with inhibitory neurons was suggested by Shadlen and Newsome (1998). In their case the inhibition was fed only to the next layer and was designed to prevent synchrony from being built up. In our case the output of inhibitory neurons is fed to the entire region and not just to the subsequent layer. As a side issue, we note that even the strict feed forward inhibition between layers could not prevent the buildup of synchrony (Litvak et al., 2002).

It is well accepted that in the brain, different levels of complex representations are instantiated in different cyto-architectonic areas, and that the connections between such areas are reciprocal (Van Essen et al., 1991). Many have suggested that these connections are used in a top-down manner to speedup resolution of ambiguous scenes. To the best of our knowledge, we are the first to demonstrate how this is actually done by way of a detailed neural network model. In our simulations each neuron fired one spike during half a second. It is not known whether attractor neural network can work at such low firing rates. Activity in a network of synfire chains may be 'read out' very quickly, a couple of synchronously firing groups can be used to resolve which synfire chain is active within 6 msec. Fix point

attractors with low firing rates may take much longer to identify. Furthermore, since each neuron may take part in multiple synfire chains (the number of synfire pools is of the same order of magnitude as the number of neurons (Bienenstock, 1995; Herrmann et al., 1995), the same network may be used to solve a large number of problems. Thus we believe that synfire chains can accomplish this task faster and with lower firing rates than in a system of cardinal cells (Barlow, 1995) based on firing rate codes. In the examples given here, we showed two types of binding: conjunction of two parts of a novel object, and the binding between parts and retrieval of a known complex object.

Prior expectation for a given solution biases the resolution of an ambiguous scenario towards that solution. The effect was small. The crossover point was shifted from 57 to 50% (Figs. 18 and 20), but so is the priming effect in real life (For example very strong priming reduced reaction time in the visual system from 295 to 280 ms (Carreiro et al., 2003)).

Typically, neural-network models of compositionality are expected to have the following properties:

1. *Stable representations.* To resolve a complex scene, a neural network must have the ability to maintain the representation of the elements of this scene for a duration on the order of magnitude of a second. Furthermore, a neural network model for compositionality must have a representation for the different levels of processing. Such representations should also be stable for a time span on the order of a second. The synfire chain model introduced here has these properties.
2. *Dynamic binding mechanism.* Representing a complex scene by the composition of its components calls for a mechanism for binding the above components. Hebb introduced the cell assembly model (Hebb, 1949) as a model for composition. Such a model assumes that the complete representation can be decomposed into smaller units, or semantic atoms. Co-activation of a number of semantic atoms represents a complex semantic entity (in the toy world, the semantic atoms are the components, and co-activating CIRCLE and RED will represent a RED CIRCLE). This family of models was criticized by C. von der Malsburg for having no internal structure (von der Malsburg, 1981, 1987). In such models, one cannot distinguish between a situation in which one sees simultaneously a RED CIRCLE and a GREEN SQUARE, and the one in which one sees a RED SQUARE and a GREEN CIRCLE. This is the *superposition catastrophe* (von der Malsburg, 1987). The question of how to prevent such a confusions is known as the binding problem (Roskies, 1999; Treisman, 1996, 1999). We showed here that synchronization among waves in synfire chain does tag the activities in the appropriate way.
3. *Regulation of the binding level.* The above dynamic binding mechanism cannot always prevent the *superposition catastrophe* (e.g. when all the presented entities become bound). Therefore, a mechanism for regulating the number of components which are bound to each other is needed (von der Malsburg, 1995). We showed that by including inhibitory neurons in synfire chains such regulation may be achieved.
4. *Hierarchy.* The need for hierarchy comes both from the hierarchical nature of compositionality (e.g., "T-junction" is formed out of lines, "words" are formed out of "letters") and from the hierarchical organization of the cortex (Van Essen et al., 1991). The hierarchy must contain a kind of rules which defines the relations between objects from different levels (The word 'CAT' is formed from the letters 'C', 'A' and 'T' at certain locations). In the example used here (Fig. 16), a second hierarchy level was added, where the letters 'T' and 'L' are represented. These rules (or grammar) is embodied in the connections between 'T' and its parts (horizontal and vertical lines with some relation between both lines) and connections between 'L' and its parts (horizontal and vertical lines with a different relation).
5. *Bidirectional stream of information.* Resolving a complex scene involves parallel processing at the different levels of hierarchy. An efficient usage of such parallel processing calls for bidirectional streams of influences (bottom-up and top-down). In computer vision, the integration of bottom-up with top-down processing has been a major concern (Ullman, 1989, 1996). Our simulations demonstrated how this is done with a hierarchical system of synfire chains.
6. *Invariance.* Relational descriptions are invariant. This is the basis for many object recognition algorithms. Dickinson et al. (1992) (Example taken from Bienenstock and Geman, 1994) for instance introduced such algorithm: defining the objects of interest relationally, and often hierarchically, in terms of the relative positioning of identifiable subparts. The parts, furthermore, may themselves be defined as

relational compositions of still more primitive elements. We have not demonstrated invariance here, but at least position invariance can be done with synfire chains.

7. *MDL prior.* Images are often ambiguous but our eyes and brain tend to resolve the ambiguity. Sometimes, complex scene can be resolved as a combination of a few simple objects, or by one more complex object. Gestalt theory (Ellis, 1938) summarizes the competition between a combination of several simpler representations and one composite representation by an economical rule, that involves using as few representations as possible. Rissanen's Minimal Description Length (MDL) principle (Rissanen, 1989) was used to formulate a probabilistic model for compositionality (Bienenstock et al., 1997; Geman et al., 1998; Potter, 1999). We have shown elsewhere that this may be achieved with synfire chains (Hayon, 2004).

Our simulated neurons are quite simple. However they come closer to biological reality than the popular integrate and fire point neuron models. Here synaptic effects were modeled by currents rather than voltage. The duration and shape of these synaptic currents mimic the dendritic currents withdrawn from the soma, rather than the ionic channel currents. After an action potential, the membrane is reset by an increase of potassium conductivity, but the synaptic currents continue to flow to the soma. This prevents the very deep after hyper-polarization seen in integrate and fire models. Our model exhibits a relative refractory period and threshold adaptation, which is not included in typical integrate and fire models. Although we did not study systematically how crucial these features are to the performance of the network, we are certain that the behavior seen here could well conform to the biological reality of brain activity. The ability to resolve ambiguous scenarios in a hierarchical system of synfire chains depends on the architecture assumed here. In each area there are both excitatory and inhibitory connections, but between areas the connections are only reciprocal and excitatory. This matches the anatomical reality in the cortex.

In order to synchronize two synfire chains within a reasonable time (up to 200 ms) the initial delay between the two waves should not exceed 20 ms. This timing might be too tight for some sensory systems. We suggest that this limit might be relaxed if the cross links have prolonged EPSPs with slow rise times. This is the case for EPSPs mediated by NMDA receptors that are more prevalent in excitatory connections between different cortical areas.

This work did not touch upon several issues. How are synfire chains generated in the first place? Several works attempted to tackle this issue, but the results are controversial (Bienenstock, 1991; Hertz and Prugel-Bennett, 1995; Horn et al., 1999). Can the desired cross-links between synfire chains be learned? This was shown to be feasible (Abeles et al., 1993), but there is no quantitative analysis of this process. When the same neurons participate in several pools of the same synfire chain, activity waves may reverberate within the chain for a long time (Abeles et al., 1993). However, it is not clear whether synchronized waves will reverberate in a synchronized manner under these conditions.

The simulations shown here treat only 'toy' problems of binding, yet they required a simulation of 50,000 neurons. Any real world problem will require much larger simulations. This renders the study of such problems impractical. In a subsequent paper we show how this situation may be alleviated (Hayon et al., 2004).

Acknowledgment

This study was supported in part by grants from the GIF and the Paula Rich Foundation. M. Abeles is a Shamoon Prof. at Bar-Ilan University.

References

Abeles M (1982) Local Cortical Circuits—An Electrophysiological Study. Springer-Verlag, Berlin.

Abeles M (1991) Corticonics, Neural Circuits of the Cerebral Cortex. Cambridge University Press.

Abeles M, Prut Y, Vaadia E, Bergman H (1994) Synchronization in neuronal transmissionand its importance for information processing. In: G. Buzsaki, et al., eds. Temporal Coding in the Brain, pp. 39–50.

Abeles M, Vaadia E, Bergman H, Prut Y, Haalman I, Slovin H (1993) Dynamics of neural interaction in the frontal cortex of behaving monkeys. Concepts in Neuroscience 4(2): 131–158.

Aertsen A, Diesmann M, Gewaltig MO (1996) Characterization of synfire activity by propagating pulse packets. In: JM Bower, ed. Advances in Computational Neurosc. Plenum.

Aertsen A, Diesmann M, Gewaltig MO (1996) Propagation of synchronous spiking activityin feedforward neural networks. J. Physiology 90(3/4): 243–247.

Arnoldi R, Brauer W (1996) Synchronization without oscillatory neurons. Biol. Cybernet. 74: 209–223.

Barlow H (1995) The neuron doctrine in perception. In: M. Gazzaniga, ed., The Cognitive Neuronsciences pp. 415–435. MIT Press.

Biederman I, Hummel JE (1992) Dynamical binding in a neural network for shape recognition. Psychol. Review 99(3): 480–517.

Bienenstock E (1991) Notes on the growth of a composition machine. In: D. Adler, E. Bienenstock. Laks, eds., Proceeding of the Royaumont Interdisciplinary Workshop on Compositionality in Cognition and Neural Network-I.

Bienenstock E (1992) Suggestions for a Neurobiological Approach to Syntax. Private Communication.

Bienenstock E (1995) A model of neocortex. Network: Comput. Neural Syst. 6: 179–224.

Bienenstock E (1996) Composition. In: A. Aertsen and V. Braitenberg, eds., Brain Theory: Biological Basis and Computational Theory of Vision. Elsevier.

Bienenstock E, Geman S (1994) Compositionality. In: Michael Arbib, ed., The Handbookof Brain Theory and Neural Networks. MIT Press.

Bienenstock E, Geman S, Potter D (1997) Compositionality, mdl priors, and object recognition.In: MI Jordan, MC Mozer, T Petsche, eds., Advances in Neural Information Processing Systems, vol. 9. MIT Press.

Braitenberg V (1986) Two views of the cerebral cortex. In: G Palm, A Aertsen, eds., Brain Theory Springer, Berlin, pp. 81–96.

Carreiro LRR, Haddad H Jr, Baldo MVC (2003) The modulation of simple reaction time by the spatial probability of a visual stimulus. Braz. J. Med. Biol. Res. 36: 907–911.

Cohen A, Shoup R (2000) Response selection target for conjunctive targets. J. Experim. Psych.: Human Percep. Perform. 26: 391–411.

Damasio AR (1989) Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recalll and cognition. Cognition 33: 25–62.

Dickinson SJ, Pentland AP, Rosenfeld A (1992) From volumes to views: An approach to 3-d object recognition. Image Understanding 55: 130–154.

Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. Nature 402: 529–533.

Doursat R (1991) Contribution à l'étude des représentations dans le système nerveux et dans les réseaux de neurones formels. PhD thesis, Université Paris, 6.

Eckhorn R, Brauer R, Jordan W, Brosch M, Kruse W, Munk M, Reitboeck HJ (1988) Coherent oscillations: A mechanism for feature linking in the visual cortex? Biol. Cyber. 60: 121–130.

Ellis W (1938) A Source Book of Gesstalt Psychology. Humanities Press.

Van Essen DC, Felleman DJ, De Yoe EA, Olavarria J, Knierim J (1991) Modular and hierarchical organization of extrastriate visual cortex in macaque monkey. In: Cold Spring Harbor Symposiom of Quantitative Biology vol. 55, pp. 679–696.

Geman S, Potter DF, Chi Z (1998) Composition systems. Technical report, Division of Applied Mathematics, Brown University.

Gray CM (1999) The temporal correlation hypothesis of visual feature integration: Still alive and well. Neuron 24: 31–47.

Gray CM, Singer W (1989) Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. Proc. Natl. Acad. Sci USA 92: 6655–6662.

Hayon G (2004) Modeling compositionality in biological neural networks by dynamic binding of synfire chains. PhD thesis, the Hebrew University in Jerusalem, 2003.

Hayon G, Abeles M, Lehmann D (2004) A model for reresenting the dynamics of a system of synfire chains. J. Comp. Neurosci., submitted.

Hebb DO (1949) Organization of Behaviour. Wiley Press, New York.

Herrmann M, Hertz J, Prugel-Bennett A (1995) Analysis of synfire chains. Network: Computation in Neural Systems 6: 403–414.

Hertz J, Krogh A, Palmer R (1991) Introduction to the Theory of Neural Computation. Addison-Wesley Publishing Company.

Hertz J, Prugel-Bennett A (1995) Learning short synfire chains by self-organization. Network: Computation in Neural Systems.

Horn D, Levy N, Meilijson I, Ruppin E (1999) Distributed synchrony of spiking neurons in a hebbian cell assembly. In: NIPS99.

Horn D, Opher I (1996) The importance of noise for segmentation and binding in dynamical neural systems. International Journal of Neural Systems 7(4): 529–535.

Horn D, Sagi D, Usher M (1991) Segmentation, binding and illusory conjunctions. Neural Comput. 3: 510–525.

Hummel JE (2001) Complementary solution to the binding problem in vision: Implications for shape perception and object recognition. Visual Cognition 8: 489–517.

Hummel JE, Stankiewicz BJ (1996) An architecture for rapid hierarchical structural description. In: T Inui, J McClelland, eds., Information Integration in Perception and Communication vol. XVI of Attention and Performance, MIT Press, Cambridge, MA, pp. 93–121.

Litvak V, Sompolinsky H, Segev I, Abeles M (2002) On the transmission of rate code in long feed-forward networks with excitatory-inhibitory balance. J. Neurosci. 23: 3006–3015.

Maass W (1997) Fast sigmoidal networks via spiking neurons. Neural Computation 9: 279–304.

Maass W, Natschlager T (1997) Network of spiking neurons can emulate arbitrary hopfield nets in temporal coding. Network: Computation in Neural Systems.

Postma EO, van der Herik HJ, Hudson PTW (1996) Robust feedforward processing in synfire chains. International Journal of Neural Systems 7(4): 537–542.

Potter DF (1999) Compositional Pattern Recognition. PhD thesis, Division of Applied Mathematics, Brown University.

Prut Y, Vaadia E, Bergman H, Haalman I, Slovin H, Abeles M (1998) Spatiotemporal structure of cortical activity: Properties and behavioral relevance. J. Neurophysiol. 79(6): 2857–2874.

Reynolds JH, Desimone R (1999) The role of neural mechanisms of attention in solving the binding problem. Neuron 24: 19–29.

Riesenhuber M, Poggio T (1999) Are cortical models realy bound by the "binding problem". Neuron 24: 87–93.

Rissanen J (1989) Stochastic Complexity in Statistical Inquiry. World Scientific Press.

Roskies AL (1999) The binding problem. Neuron 24: 7–9.

Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. J. Neurosci. 18(10): 3870–3896.

Shadlen MN, Movshon A (1999) Synchrony unbound: A critical evaluation of the temporal binding hypothesis. Neuron 24: 67–77.

Shastri L, Ajjanagadde V (1993) From simple associations to systematic reasoning: A connectionist representation of rules, variables and dynamic binding using temporal synchrony. Behav. Br. Sci. 16: 417–494.

Singer W (1999) Neural synchrony: A versatile code for the definition of relations? Neuron 24: 49–65.

Sougne JP, French RM (2001) Synfire chains and catastrophic interference. In: Proceedings of the 23rd Annual Conferences of the Cognitive Science Society. Lawrence Erbaum Ass., Mahwah, NJ.

Sterratt DC (1999) Is a biological temporal learning rule compatible with learning synfire chains? In: Proceedings of ICANN 99 can be found at http://www.cogsci.ed.ac.uk/ dcs/research.html.

Treisman A (1996) The binding problem. Current Opinion in Neurobiology 6: 171–178.

Treisman A (1999) Solutions to the binding problem: Progress through controversy and convergence. Neuron 24: 105–110.

Triesch J, von der Malsburg C (1996) Binding—A proposed experiment and a model. In: Proceedings of the International Conference on Artificial Neural Networks Bochum.

Ullman S (1989) Aligning pictorial descriptions: An approach to object recognition. Cognition 32(3): 193–254.

Ullman S (1996) High-level Vision. The MIT Press.

Villa AEP, Tetko IV, Holand B, Najem A (1999) Spatiotemporal activity patterns of rat cortical neurons predict responses in a conditioned task. In: Proceedings of the National Academy of Science USA 96: 1106–1111.

von der Malsburg C (1981) The Correlation Theory of Brain Function, Springer, Berlin, vol. Models of Neural Networks II, Chap. 2, pp. 95–119.

von der Malsburg C (1987) Synaptic plasticity as basis of brain organization. In: JP Changeuxand, M Konishi, eds., The Neural and Molecular Bases of Learning Wiley, New York.

von der Malsburg C (1995) Binding in models of perceptions and brain function. Current Opinion in Neurobiology 5: 520–526.

von der Malsburg C (1999) The what and why of binding: The modeler's perspective. Neuron 24: 95–104.

von der Malsburg C, Wiskott L (1996) Recognizing faces by dynamic link matching. In: Proceedings of the International Conference on Artificial Neural Networks, Bochum, pp. 347– 352.

Wong KYM (1997) Exact dynamics in feedforward neural networks. Europhysics Letters 38(8): 631–636.

Zucker RS (1989) Short-term synaptic plasticity. Ann. Rev. Neuroscience 12: 13–31.