Dynamic linking in Stochastic Networks

Hilbert J. Kappen, Marcel J. Nijman

RWCP Novel Function SNN Laboratory

Dept. of Medical Physics and Biophysics, University of Nijmegen

Geert Grooteplein 21, NL 6525 EZ Nijmegen, The Netherlands

Tel: +31 80 614241

Fax: +31 80 541435

email bert@mbfys.kun.nl

Dynamic linking in Stochastic Networks

Hilbert J. Kappen, Marcel J. Nijman

RWCP Novel Function SNN Laboratory

Dept. of Medical Physics and Biophysics, University of Nijmegen

Geert Grooteplein 21, NL 6525 EZ Nijmegen, The Netherlands

Abstract

It is well established that cortical neurons display synchronous firing for some stimuli and not for others. The resulting synchronous subpopulation of neurons is thought to form the basis of object perception. In this paper this 'binding' problem is formulated for Boltzmann Machines. Feed-forward connections implement feature detectors and lateral connections implement memory traces or cell assemblies. We show, that dynamic linking can be solved in the Ising model where sensory input provides local evidence. The lateral connections in the hidden layer provide global correlations between features that belong to the same stimulus and no correlations between features from different stimuli.

1   Introduction

It is well established that cortical neurons display synchronous firing for some stimuli and not for others (Gray et al., 1989; Engel et al., 1990). The resulting synchronous subpopulations of neurons (cell assemblies) are thought to form the basis of segmentation and object perception (Julesz, 1971; Marr, 1982). The role of individual cells is to represent important 'atomic' visual features, such as edges, corners, velocities, colors, etc. Objects can be defined as a collection of these features: The cell assembly is a neural representation of the entire object.

There are two major complications with this idea. 1) The cells in an assembly encode different visual features. Different objects are represented as different combinations of active features. These neurons can only fire synchronously, if there exist direct or indirect synaptic connections between them. But if such a connection exists, why do these neurons not fire synchronously at all times? This problem will be referred to as dynamic linking. 2) When more than one coherent object is present in the scene, more than one cell assembly will be active. How can be determined to which cell assembly an active cell belongs? As an example consider a red square and a blue circle that are presented simultaneously. If there are separate neural pathways for shape and color, a representation for 'red', 'blue', 'square' and 'circle' will be active. It can then no longer be decided which features belong to the same object (Treismann and Schmidt, 1982).

Both problems are jointly referred to as the binding problem. A popular solution to the second problem has been proposed by several authors. If synchronous firing occurs within one cell assembly in the form of a collective oscillatory firing pattern, then different assemblies can be distinguished by either their frequency or phase (Malsburg and Schneider, 1986).

In this paper the binding problem is formulated and studied in the context of Boltzmann Machines. Feed-forward connections implement feature detectors and lateral connections implement memory traces. Boltzmann Machines provide an attractive model to study the binding problem, because the interaction of feed-forward connections and lateral connections can be analytically studied in the equilibrium distribution. In this way the 'synchronous firing' can be related to well know equilibrium properties of spin systems.

We propose a solution to dynamic linking as described above. The solution is related to the existence of long correlation lengths is spin systems near the critical temperature. This phenomenon occurs in networks that have 'sufficient' lateral connectivity. As an example we take the 2 dimensional Ising model.

2 Boltzmann Machines

Boltzmann Machines are stochastic networks. The neurons can be in two states $\sigma_i = 0, 1$ (neuron language) or ± 1 (spin language). Using so-called sequential dynamics, neurons are randomly selected one at the time at discrete time steps. The probability of firing for neuron i given the current state of the network $\vec{s}$ is given by

$$T(s'_i = 1 | \vec{s}) = \frac{1}{2}(1 + \tanh(\beta h_i)).$$

$h_i = \sum_{j \neq i} w_{ij} s_j$ ($s_0 = 1$ and w_{i0} denotes a threshold for neuron i). After long times, the probability to observe the network in a state $\vec{s}$ becomes independent of time. When the weights of the network are chosen symmetrically, this time independent equilibrium distribution is the Boltzmann distribution and is given by

$$p(\vec{s}) = \frac{1}{Z} \exp\{\beta \sum_{i,j} w_{ij} s_i s_j\}$$

$$Z = \sum_{\vec{s}} \exp\{\beta \sum_{i,j} w_{ij} s_i s_j\}$$

The proposed Boltzmann Machine architecture is given in Fig. 1. We will assume that the network
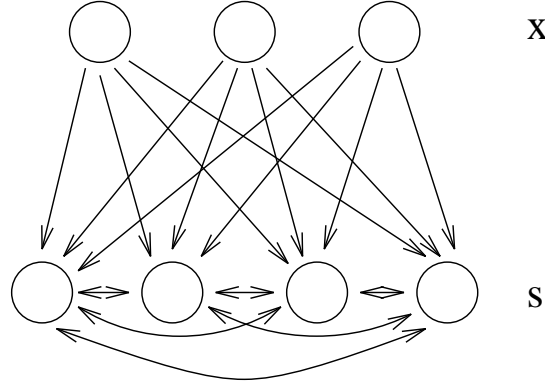

Figure 1: Boltzmann Machine architecture with feed-forward and lateral connections. Input is provided by $\vec{x} = (x_1, \ldots, x_n)$, with x_i real-valued or binary valued. Hidden units are $\vec{s} = (s_1, \ldots, s_h)$, with $s_j = \pm 1$.

at any time will receive some kind of input and study the Boltzmann distribution on the hidden layer.

2.1 Spike Interpretation

In order to study synchronous firing we need a spike interpretation of the Boltzmann Machine. This is given in figure 2. The time steps should be chosen such that the average time between updates of one neuron is of the order of the refractory period of the neuron.

3 Correlation lengths

Consider the task where the input encodes a sensory domain, such as the retina, where several stimuli are presented simultaneously. Thus the input $x_i = 1$ at the locations of the simuli and
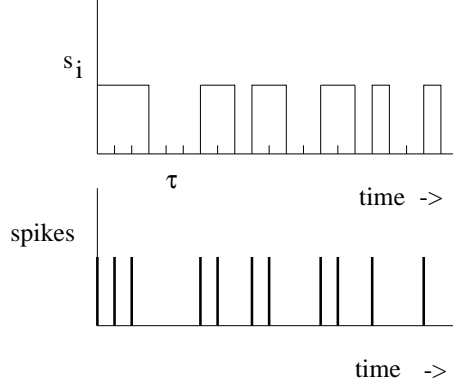
Figure 2: Spike interpretation of Boltzmann Machine. Time is discrete such that the time between two updates of one neuron is of the order of the refractory period τ of the neuron. If $s_i = 1$ in one time interval, the neuron emits a spike. If $s_i = 1$ for a long period, the neuron fires at the maximum firing rate $1/\tau$.

0 elsewhere. The task of the network is to represent these inputs in the hidden layer, such that the activity between the neurons that encode one stimulus are correlated and the activity between neurons that encode for different stimuli not. A convenient quantity that expresses correlation between neurons j and k is the correlation function:

$$\Gamma_{jk} = \langle s_j s_k \rangle - \langle s_j \rangle \langle s_k \rangle$$

The correlation function depends on the temperature β and on the connectivity of the network. For instance, in a d-dimensional Ising spin system, the connections are only between nearest neurons in a d-dimensional grid. Γ_r can be calculated in the Landau approximation and takes the form

$$\Gamma_r \propto r^{2-d} \exp(-r/\xi)$$

with r the distance in the grid. ξ depends on the temperature of the system. Around the critical temperature T_c, $\xi \propto |T - T_c|^{-\frac{1}{2}}$. In a 1-dimensional Ising system, there is no phase transition and no large range correlations exist. Thus in order to have long range correlations between neurons that encode the same stimulus, sufficient lateral connectivity must be present to display a phase transition around the critical temperature. Long correlation lengths are obtained for unfrustrated systems such as the Ising model, and for frustrated systems, such as spin glasses and Hopfield-type of attractor neural networks. The correlation lenght as a function of temperature is schematically given in Figure 3.

4   Binding

First we give the explicit relation between $0, 1$ coding and ± 1 coding for the Ising model. The Ising model Hamiltonian is given by:

$$
\begin{aligned}
H &= -W \sum_{\langle ij \rangle} s_i s_j - H \sum_i s_i \\
&= -w \sum_{\langle ij \rangle} \sigma_i \sigma_j - h \sum_i \sigma_i + const.
\end{aligned}
$$

W, H and w, h are the lateral connectivity and external magnetic field in the ± 1 and $0, 1$ representation, respectively. $\sum_{\langle ij \rangle}$ denotes the sum over all pairs of neighbors in the Ising lattice. It is
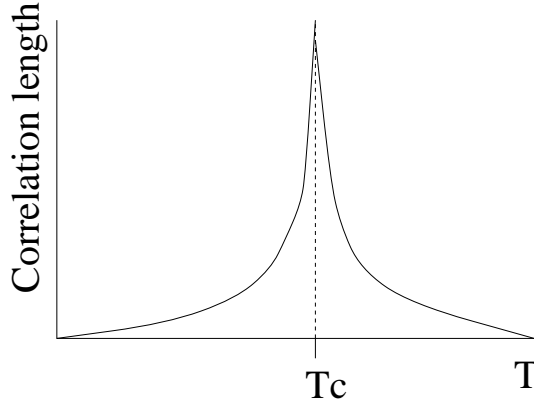
Figure 3: Correlation lenght as a function of temperature. Large scale correlations are necessary for binding and develop when $T \approx T_c$. For $T = T_c$, $\langle s_j \rangle = 0$, which means that the neurons fire at 50 % of their maximal firing rate.

easy to show that $w = 4W$ and $h = 2(H - \alpha W)$. α is the number of neighbors of each Ising spin ($\alpha = 4$ for the 2D Ising model).

We will apply the Ising model to simulate the binding of features in the presence of a stimulus. W will be assumed constant for all neighbor connections. H will be different depending on whether neurons receive stimulus input or not.

First of all, we must ensure that no large scale correlations and low firing rates develop outside the cell assemblies. Therefore, we assume that the total external field for neuron i $H_i = H^- + H_i^+$, with H_i^+ the stimulus contribution, and H^- a constant inhibitory external field all hidden neurons. This global inhibition could be due to a lateral inhibition mechanism but will be left unspecified at this point. Outside the cell assemblies the neural activity will be low and no large scale correlations will develop.

In order to assure that long correlation lenghts occurs in the presence of stimuli for subsets of neurons that encode the different objects (cell assemblies), we need $H = 0$ and $T = T_c$. Thus the stimulus strength should cancel the inhibitory field $H_i^+ = -H^-$. The architecture is given in Figure 4

The result of a small simulation with 2 objects in an 11×11 Ising lattice is shown in Figure 5. The values of H and W translate in neural connections $w = 4$ and stimuli: $h = -2W\alpha = -8 = h^+ + h^-$ (stimulus); $h = 2(H^- - W\alpha) = -16 = h^-$ (no stimulus). Thus $h^+ = 8$ and $h^- = -16$. At the stimulus locations, the neurons have $\langle \sigma_j \rangle = 0.5$. Away from the stimuli, $\langle \sigma_j \rangle \approx 0$

As is clear from the figure, all cells belonging to the same object are highly correlated, whereas cells belonging to different objects are not correlated. So, the network concludes the existance of 2 objects.

5 Discussion

We have shown how long range correlations in spin models can be used to solve a part of the binding problem, the dynamic linking problem. We have shown, that dynamic linking can be solved in the Ising model. This model contains a sensory layer and a hidden layer. The hidden layer represents an interpretation of the sensory input. Sensory input provides local evidence. The lateral connections in the hidden layer provide global correlations between features that belong to the same stimulus and no correlations between features from different stimuli.

Several things are still unclear. First of all, two cell assemblies should merge into one if the distance between them is very small. In that case the sensory input coming from the small area
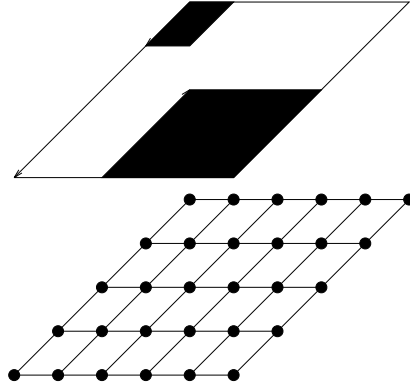
Figure 4: Boltzmann Machine architecture used for simulation of binding phenomenon. All neurons in the sensory layer (top) are only connected to the corresponding neuron in the hidden layer (bottom). Neurons in the hidden layer receive sensory input of strength 1 when a stimulus in the sensory layer is present (dark areas).

between the two cell assemblies is overruled by the urge to create large coherent objects. In our model, this means that the correlated activity must be able to spread somewhat outside the edges of an object. From Figure 5 can be seen that this does happen.

Our results were obtained in a small network. More robust correlations will develop in larger networks, which we are currently investigating.

Animal brain studies have shown the existance of 'oscillations' of activity with a frequency of 40 Hz. In our model binding occurs when the neurons fire at half of the maximum firing rate, leading to oscilations of 167 Hz, if one assumes a refractory period of 3 ms. These frequencies are not very dissimilar, given the crude assumptions of our model.

In this paper, we have only studied correlations at 0 time delay. In (Ginzburg and Sompolinsky, 1994), delayed correlations were studied in networks composed of several sub-populations. Each sub-population is fully connected, and connections between sub-populations only depend on the sub-populations and not on the individual neurons. The issue of dynamic linking was not addressed there.

Clearly, we are not proposing the Ising model as a serious computational model for the cortex and it should be investigated whether and how this mechanism can be extended to other network architectures. In a more realistic network, the lateral connectivity would arise from learning. As a result, an inhomogeneous network would result with both inhibitory and excitatory connections. The resulting network will be more like a spin glass than an Ising model. It is known, that long range correlations also exist in these frustrated systems.

There are several ways to learn the lateral connections strengths from an environment. For instance, a Hebb learning rule based on 1) correlated activity $\langle s_i s_j \rangle$ or 2) on non-correlated activity $\langle s_i \rangle \langle s_j \rangle$. If correlation plays a functional role, the first rule seems most promising. However, since correlations arise naturally near criticality, also the second rule may be able to yield correlated patterns of activity, whereafter the effects of rule 1) and 2) become identical.

References

Engel, A., König, P., Gray, C., and Singer, W. (1990). Stimulus-dependent neuronal oscillations in cat visual cortex: Inter-columnar interaction as determined by cross-correlation analysis. *Eur. J. Neurosci.*, 2:588.
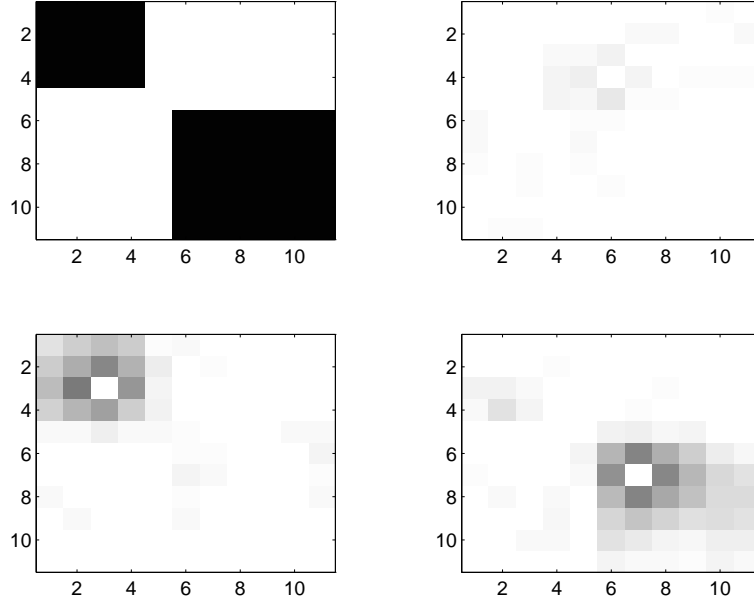
Figure 5: Top Left: sensory input containing 2 black objects ($H = H^+ + H^-$) against a white background ($H = H^-$). $H^+ = -H^- = 4$, $W = 1$, $T \approx T_c$. Top Right: Correlation Γ_{jk} with the neuron located at lattice site (6,4). Since $W > 0$, $\Gamma \geq 0$. White (black) encodes $\Gamma = 0, 1$, respectively. Bottom Left: Correlation with point (3,3). Bottom Right: Correlation with point (7,7).

Ginzburg, I. and Sompolinsky, H. (1994). Theory of correlations in stochastic neural networks. *Physical Review E*, 50:3171–3191.

Gray, C., König, P., Engel, A., and Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338:334.

Julesz, B. (1971). *Foundations of cyclopean perception*. University of Chicago Press, Illinois.

Malsburg, C. and Schneider, W. (1986). A neural cocktail-party processor. *Biological Cybernetics*, 54:29–40.

Marr, D. (1982). *Vision: A computational investigation into the human representation and procesing of visual information*. Freeman, San Francisco.

Treismann, A. and Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology*, 14:107–141.