T-61.181 Special Course in Information Science

RF-SLISSOM, SNN SOM and DSOM

Jaakko Peltonen, 47052k
November 18, 2002

1 Introduction

This report is based on three articles that concern biologically inspired versions of Self-Organizing Maps. The proposed networks all use *leaky integration* to include temporal dynamics in the learning of the map, and use various strategies to control the firing of neurons to produce competitive, self-organizing behaviour.

The three models proposed in the papers are herein called RF-SLISSOM, SNN SOM and DSOM. The sections below will introduce each of these methods in order.

2 RF-SLISSOM

Choe and Miikkulainen [3] have studied computational models of the early development of the *visual cortex*. Several recent models have included *lateral interaction* between neurons, which has potential to model a wider range of phenomena than SOMs that do not account for lateral connections. Lateral connections may be important to the visual cortex: they may synchronize its spiking activity and cause feature binding and segmentation.

Lateral connections are thought to form a *topological map* of the input space, by self-organization together with *afferent connections* from the retina. A computational model of the process called Receptive Field - Laterally Interconnected Synergetically Self-Organizing Map (RF-LISSOM), based on *Hebbian learning*, can produce structures like oriented receptive fields or patterned lateral connections. Choe and Miikkulainen integrate the model with *spiking neurons*; the resulting model is called RF-SLISSOM, where the added S denotes 'Spiking'.

2.1 RF-SLISSOM architecture

The RF-SLISSOM is a two-layer network. The layers are called 'retina' and 'cortex' due to their biological inspiration. The network input is generated at active neurons of the *retinal* or *input* layer, and is sent ahead to the *cortical* neurons through so-called *afferent* connections. Each cortical neuron receives input from a fixed-size *receptive field* of retinal neurons, whose center in the retina corresponds to the location of the cortical neuron in the cortex. The left figure in Fig. 1 shows the overall organization.

Each connection between neurons has a *spike queue* that performs *leaky integration* of the inputs. The value of the current spike is added to *decayed* values of older spikes. An exponential decay of the form $1/e^{\lambda_q}$ is used, where λ_q is a decay parameter.

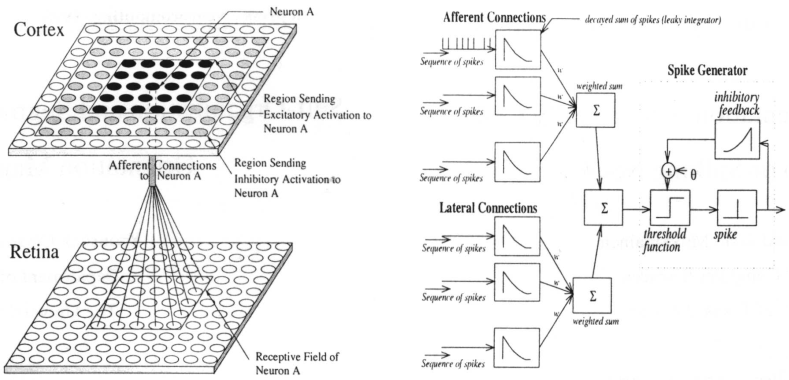

Figure 1: RF-SLISSOM structure. Left: overall struture. Right: schematic of a single cortical neuron.

Each cortical neuron calculates a weighted sum over its inputs. The sum is taken over both retinal and cortical neurons, and weighted by *afferent* and *lateral connection weights*, respectively. The lateral weights can be *excitatory* or *inhibitory*. If the sum is over a *treshold*, the neuron sends a *spike*. However, the spike generator remembers past spikes by leaky integration: active spiking increases the effective treshold. The right figure in Fig. 1 shows a schematic of a cortical neuron.

2.2 Training the model

The afferent and lateral connection weights are first initialized randomly and normalized. The RF-SLISSOM is then trained so that the input is kept constant while the cortical response settles; this reduces the redundancy of the activation. Hebbian learning is then applied to modify the connection weights $w_{ij,mn}(t)$:

$$w_{ij,mn}(t) = \frac{w_{ij,mn}(t-1) + \alpha V_{ij} X_{mn}}{\mathcal{N}} \tag{1}$$

where V_{ij} and X_{mn} are the average spiking rates of neurons (i,j) and (m,n) and $\mathcal{N}$ normalizes the weights so they do not grow without bounds.

The lateral excitation radius is reduced by a schedule during training: an ever smaller region sends neurons excitatory activation - the map then achieves fine tuning. By contrast, inhibitory and afferent learning rates *increase* the initial slow adaptation captures long-term input correlations and later fast adaptation allows quick activity modulation for segmenting input.

2.3 Experiments

The RF-SLISSOM was trained by presenting single instances of oriented, elongated 2D Gaussian bars at 8 angles. This made the afferent connections self-organize into *oriented receptive fields* like ones in the visual cortex. An orientation map (Fig. 2, left figure) of their preferred angles shows structures like *pinwheels* oriented around a point and *fractures* with sudden changes. Areas with similar orientation preference were linked by strong connections, and the connection patterns were elongated along the preferred direction.

The results are similar to RF-LISSOM results. That networks with spiking neurons attain the same results helps understand visual cortex development.

Next, the trained map was tested in segmenting multiple objects. Activity in part of the cortex shows it is representing the corresponding input part. The total spike amount in an area per time tells the representation strength; this can be plotted over time.

Neurons for the same object were active or silent together, whereas neurons for different objects were not active together. The activation rotates through the different areas. The network was even able to bind together close-by bars with the same orientation, belonging to the same texture.

Quick lateral adaptation and the leaky integration treshold were crucial: the first helps adapt to activity correlations, and the second helps desynchronization.

2.4 Discussion

Fast adaptation is necessary to bind features by temporal coding, and fast lateral adaptation was necessary in RF-SLISSOM experiments. Yet slow adaptation is needed to learn long-term correlations. To solve this, one could use two lateral connection sets, or vary the learning rate. The latter is simpler and is used here: learning rate grows during training.

RF-SLISSOM, like RF-LISSOM, finds a sparse coding for visual input, where lateral connections represent input correlation. Redundant activity is reduced, and there are minimal amounts of active neurons for input.

Currently RF-SLISSOM is limited by network size, since memory and time requirements grow by $O(n^4)$. Future work will include more computationally intensive simulations.

How large objects can be identified depends on the extent of lateral connections. Object activity wider than the lateral excitation radius can form multiple areas. The inhibitory connection extent affects how far apart objects are segmented. Future work will then include experiments with hierarchical RF-SLISSOMs.

3 SNN SOM

Ruf and Schmitt [1] have studied unsupervised learning in spiking neuron networks (SNNs) with *action potential timing*. Topology-preserving maps have been found in e.g. the visual, auditory or somatosensory cortex, and the SOM gives a possible explanation for their development.

Previous SOMs do not characterize neuron output by single-event timing. With lateral connections the winner is found by a recurrent network, after the network has stabilized. Temporal coding can benefit fast information processing in SNNs. Models like RF-SLISSOM also learn only after stabilization. Ruf and Schmitt propose an unsupervised learning mechanism based on single-event timing.

3.1 Model architecture and learning

A leaky integrate-and-fire neuron v receives m excitatory inputs s_i from neurons $u_1, \ldots, u_m$ with weights $w_1, \ldots, w_m$. Each neuron fires once in an interval. Under certain assumptions, v will fire at a time determined from $T_{out} - \sum w_i s_i$ where T_{out} is a constant.

Self-organization is done by, at each cycle, choosing an input $\mathbf{s}^l = (s_1^l, \ldots, s_m^l)$ and firing input neurons to encode it. Each neuron v_j then computes a *vector similarity* $\sum w_i j s_i^l$, and the most similar neuron will thus fire first. If the lateral connections are strongly inhibitory, the winner will inhibit competitors from firing, and one can use competitive learning by the standard rule

$$\Delta w_{ik} = \eta(s_i^l - w_{ik}) \tag{2}$$

where η is the learning rate. For the SOM, one must also implement a *neighborhood* from local information. Ruf and Schmitt use a monotonously increasing function from neuron distance m_{kj} to connection weight w_{kj}: lateral connections are excitatory for close-by neurons and inhibitory for remote neurons. The winner drives the firing times of neighbors towards its own firing time. Ruf and Schmitt suggest the learning rule

$$\Delta_{wij} = \eta \frac{T_{out} - t_j}{T_{out}} (s_i^l - w_{ij}) \tag{3}$$

where t_j is the firing time of competitive neuron v_j. The lateral weights decrease during learning, which shrinks the neighborhood.

3.2 Experiments

Ruf and Schmitt test their method by choosing 10 uniformly distributed patterns and presenting them to 10 competitive units with random initial weights. The

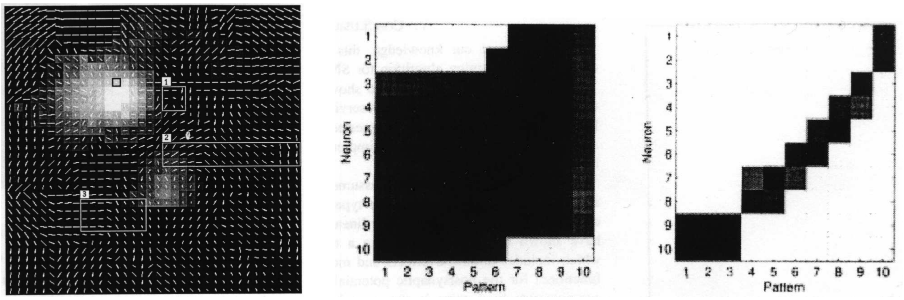

Figure 2: Left: RF-SLISSOM orientation map. Center and right: SNN SOM experiment with 10 random patterns, activation initially (center) and after training (right).

goal was to form a linear map. At first all neurons fire early, but after learning only topologically close neurons react for each input. The right figure in Fig. 2 shows the initial and final activations.

A standard SOM experiment was also run, with input chosen randomly from a square; the competitive neurons should organize in a topology-preserving grid. Neighborhood preservation was measured by *metric multidimensional scaling*:

$$E_{MDS} = \sum_{i=1}^{N} \sum_{j<i} [F(i,j) - F(M(i), M(j))]^2 \tag{4}$$

where N is the number of input patterns, $M(i)$ is the winner index for input $\mathbf{s}^i$, G is the neighborhood matrix, and F is the Euclidean dissimilarity matrix between input pairs. A rectangular grid and Manhattan distance were used.

A reescaled version of E_{MDS} stabilizes at a small value. A standard SOM converges about 10 times faster and has qualitatively similar behavior.

3.3 Discussion

Ruf and Schmitt's model is a first implementation of Kohonen's learning algorithm for SNNs by temporal coding, and helps show biological neurons can preserve topology. It may lead to fast hardware implementation of self-organization.

An assumption was made that neurons are a leaky integrate-and-fire type, but the simplifying assumptions can be dropped.

4 DSOM

Privitera and Shastri [2] have studied *compositional processes* that extract progressively complex and structured representations hierarchically. E.g. visual complex object recognition needs quick compiled information about object *subcomponents.*

According to one hypothesis, information about an entity is bound together by synchronous firing of nodes that encode its information. This relates binding, temporal processing and compositionality.

Privitera and Shastri seek compositionality with Dynamic SOM (DSOM), a hierarchical SOM where temporal processing comes through *short-term memory*, and unification among maps in the hierarchy through *synchronized activity.*

4.1 DSOM architecture and training

The DSOM model represents map element activations as a *membrane potential* $P_i(t)$, a function of the input activation $u_i(\epsilon(t))$ that measures match between weights and input $\epsilon(t)$. The input activation depends on time, so membrane potential correlates with temporal input change and can (leakily) integrate activation over time. An element is then a recognizer for a complex event.

Each element recognizes m temporal features with m *taps*, each with its own weight vector. The mebrane potential for neuron i is evaluated by

$$\frac{dP_i}{dt} = -\frac{P_i(t)}{\tau_i} + \sum_{j \in Act_i(t)} u_{ij}(\epsilon)t)) \tag{5}$$

where u_{ij} is the activation of the jth tap, $Act_i(t)$ is a set of *active* taps, and τ_i is a constant. Taps whose activation is over a treshold are inhibited and excluded from $Act_i(t)$.

There may be different activation strategies: taps may become active in a sequence and respond to features in a specific order. The taps have a common input and are integrated into a continuous potential by neuron n_i which fires if all taps are activated. Tap weights are sought by Hebbian learning.

4.2 Experiments and discussion

Privitera and Shastri test the DSOM by analyzing two-dimensional moving points for coordinated movement patterns. This involves three representation levels: point positions, the relationship of their trajectories, and its change over time. There are then three DSOM levels.

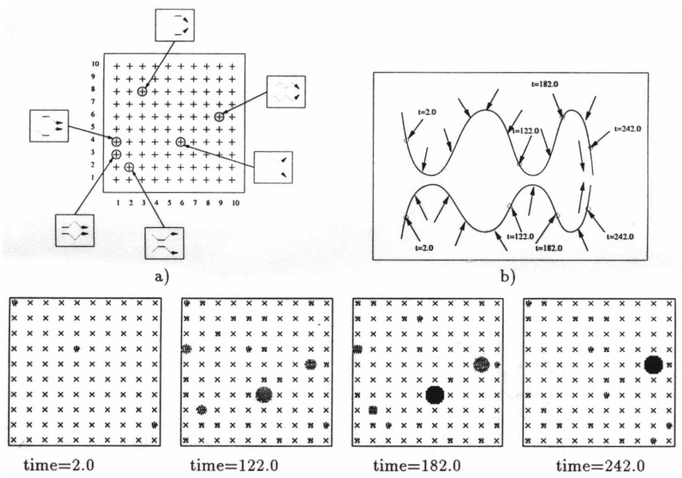

c)

Figure 3: DSOM experiment. a) DSOM after learning. b) Test pattern. c) Snapshots of third level activity.

The trained DSOM fires neurons on its third level when the system has observed appropriate patterns in the test sequence (a zig-zag motion).

DSOM implements compositional, hierarchical processing of time-varying situations into progressively abstract and structured descriptions. Information is represented by an impulse sequence, and output impulses represent pattern occurrence in the input. Temporal capabilities of map elements let DSOM keep input and output correctly bound. The tap firing is synchronous with sub-sevents in the input. Future work will investigate DSOM in language, handwriting recognition and rapid reasoning.

References

[1] Ruf, B., and Schmitt, M. Self-Organization of Spiking Neurons Using Action Potential Timing. *Ieee Transactions on Neural Networks* **9**(3), May 1998.

[2] Privitera, C. M., and Shastri, L. Temporal Compositional Processing by a DSOM Hierarchical Model.

[3] Choe, Y., and Miikkulainen, R. Self-Organization and Segmentation in a Laterally Connected Orientation Map of Spiking Neurons. *Neurocomputing*, 1998. In press.