

Effects of Presynaptic and Postsynaptic Resource Redistribution in Hebbian Weight Adaptation ¹

Yoonsuck Choe^a Risto Miikkulainen^a Lawrence K. Cormack^b

^a*Department of Computer Sciences
The University of Texas at Austin
Austin, TX 78712*

^b*Department of Psychology
The University of Texas at Austin
Austin, TX 78712*

Abstract

The Hebbian hypothesis of activity-dependent synaptic plasticity has gained much support from experimental studies of long-term potentiation and depression. Such studies have also uncovered complex patterns of competition among the synapses. Such effects may be due to the neuron redistributing its limited synaptic resources as synaptic strengths change. In computational models this strategy is commonly known as normalized Hebbian learning. However, not much consideration is usually given to whether the weights are normalized over the presynaptic or the postsynaptic sites of the neuron. Our results show that the different loci of normalization can result in drastic differences in the model's behavior, suggesting that future experiments should investigate presynaptic factors of redistribution as well as the more widely studied postsynaptic factors.

Key words: Normalized Hebbian learning, spiking neurons, self-organizing maps.

1 Introduction

The Hebbian hypothesis of activity-dependent synaptic plasticity has recently received a lot of attention due to new neurophysiological evidence of such

¹ This research was supported in part by National Science Foundation under grant #IRI-9811478.




Fig. 1. **The SLISSOM architecture** SLISSOM is a two-dimensional map of leaky integrator neurons with afferent connections and lateral excitatory and inhibitory connections. The excitatory connections group spike activity and the inhibitory connections desynchronize groups of activity representing different objects [2].

learning in the nervous system. For example, studies of long-term potentiation (LTP) and long-term depression (LTD) in hippocampal slices and in the visual cortex have shown that synaptic efficacies are indeed modulated by synaptic activity [3, 9, 10]. However, LTP/LTD alone does not explain the more complex patterns of activity-dependent learning, including those that seem to involve competition among the synapses [14]. For example, ferret LGN cells that receive input from both ON- and OFF-center ganglia become selective to one or the other type of input, even if these inputs coincide most of the time [17]. Synapses for one type of input become weaker while those of the other type strengthen.

Such competition between synapses of the same neuron could occur through redistribution of limited synaptic resources, or neuronal regulation [7]. In recent experiments, Turrigiano et al. [16] showed that a change in a single synapse can cause the whole cell's efficacy to change. This suggests that local change in the synaptic strength scales the strength of the other synapses of the same neuron.

Synaptic resource redistribution translates to weight normalization in computational models [2, 12, 13, 14, 15]. In these models, each individual synaptic weight is normalized (i.e. divided) by the total sum of synaptic weights, so that the total sum remains constant.

In principle, normalization could take place over presynaptic sites or over postsynaptic sites. Experimental studies of normalization have concentrated on postsynaptic resources. Computational models have used either kind of normalization, and the effects of the choice has not been studied. This paper shows that the behavior of such a model can be drastically different when the site of normalization is changed. Some effects may depend on postsynap-

tic normalization, others presynaptic, suggesting that a complete model of normalization may have to take both kinds of processes into account.

2 Effect of Resource Redistribution at Different Sites of the Neuron

The effects of normalization were studied in the SLISSOM model of development and function of the primary visual cortex. SLISSOM, or Spiking Laterally Interconnected Synergetically Self-Organizing Map [1, 2] consists of two layers of spiking neurons (figure 1): the “retina” and the “cortex”. The synapses of the cortical neurons are leaky integrators that sum up the spike activities in the retina, and then generate spikes depending on the level of input activity compared to the threshold in each neuron. The neurons in the SLISSOM network become selective to particular kinds of inputs in a self-organizing process that models the development of retinotopy and orientation selectivity in the visual cortex.

An organized SLISSOM network models object perception through synchronization. Recent studies in cat and monkey visual cortex have shown that neuron firing becomes correlated when a single object (e.g. a moving bar) is presented, but uncorrelated when the input consists of two different objects [4, 5, 6]. When multiple rectangular objects are presented on the SLISSOM retina, the model will synchronize the spikes of neurons that represent one object, and desynchronize the spikes of neurons that represent different objects.

The afferent weights, the lateral excitatory weights, and the lateral inhibitory weights all adapt in the SLISSOM model based on the normalized Hebbian learning rule ²:

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

where $w_{ij,mn}(t)$ is the connection weight between neurons (i, j) and (m, n) , $w_{ij,mn}(t-1)$ is the previous weight, α is the learning rate (α_E for excitatory, and α_i for inhibitory lateral connections), V_{ij} and X_{mn} are the average spiking rates of the neurons, and N is the normalization factor $\sum_{mn} [w_{ij,mn}(t-1) + \alpha V_{ij} X_{mn}]$.

In the model, there is a single synaptic weight associated with each connection. This weight can be seen as an outgoing weight from the presynaptic cell’s point of view, or as an incoming weight from the postsynaptic cell’s point of view.

² During object perception, only the adaptation of the lateral connections is significant.




Fig. 2. Model behavior with postsynaptic and presynaptic normalization Multi-unit activities (MUAs) of the two areas in the cortex representing two different objects in the scene are plotted over time. In (a), the normalization is done over the postsynaptic weights, and in (b), over presynaptic weights. In (b), the activities in the two areas are clearly separated, while in (a), the initial uncorrelated spike activity between the two areas quickly settles to only one area and the other area remains almost silent throughout the simulation.

Weight normalization can be done from either perspective, by keeping the sum of presynaptic or postsynaptic weights constant. It turns out that depending on the site of normalization, very different behavior was obtained in the model.

Normalizing the outgoing weights (i.e., presynaptic weights) resulted in a clear segmentation of two objects (figure 2b). The multi-unit activity (MUA, i.e. the number of neurons spiking in unit time) of the two areas representing two different objects are desynchronized (their hills and valleys occur at different times), whereas the spikes within the same area are synchronized (i.e. the neurons fire at the same time). When the normalization is done on the incoming weights (i.e. by redistributing the postsynaptic resources), segmentation does not occur. Only one area shows dominant activity and the other area is mostly silent, with only occasional spikes. This way, the site of normalization makes a big difference in the behavior of the model.

3 Discussion

Why does such a small change cause such a drastic difference? In presynaptic normalization, the postsynaptic cell receives inputs through weights that are

each scaled differently (according to the outgoing weights of each presynaptic cell). Therefore, even relatively low activity can result in a large weight, and the postsynaptic cell can be more sensitive to small changes in the input. Therefore the response becomes more dynamic. In postsynaptic normalization, all incoming weights are scaled by the same value. The inputs are treated more equally, and the behavior of the neuron becomes more stable.

In the segmentation task, small differences in the activation levels must be magnified to establish desynchronization, and therefore presynaptic normalization will work better. In other tasks where a whole pattern needs to be captured in the weights [8], postsynaptic normalization is more appropriate. The question then arises, are both processes active in real neurons? It is possible that the connection weights should be modeled as a product of two factors, the postsynaptic and the presynaptic weight, each normalized separately [11]. The different effects may interact, and depending on the input, one or the other may dominate. Another interesting question is if the effect is specific to only the excitatory or inhibitory synapses. Future research, both experimental and computational, is necessary to verify these things.

4 Conclusion

Hebbian hypothesis of activity-dependant synaptic plasticity is gaining support from experimental data, and synaptic resource redistribution may play an important role in explaining the competitive effects observed. In this paper, synaptic resource redistribution at different sites (presynaptic or postsynaptic) were shown to cause drastic differences in segmentation behavior of the visual cortex model. Future studies of normalization should therefore consider both types of redistribution, and a more complex model including both presynaptic and postsynaptic normalization of each synaptic strength may be necessary to explain the different effects.

References

- [1] Yoonsuck Choe and Risto Miikkulainen. Self-organization and segmentation with laterally connected spiking neurons. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, pages 1120–1125, San Francisco, CA, 1997. Morgan Kaufmann.
- [2] Yoonsuck Choe and Risto Miikkulainen. Self-organization and segmentation in a laterally connected orientation map of spiking neurons. *Neurocomputing*, 21:139–157, 1998.
- [3] Serena M. Dudek and Mark F. Bear. Homosynaptic long-term depression in area ca1 of hippocampus and effects of n-methyl-d-aspartate receptor blockade. *Proceedings of the National Academy of Sciences, USA*, 89:4363–4367, 1992.

- [4] R. Eckhorn, R. Bauer, W. Jordan, M. Kruse, W. Munk, and H. J. Reitboeck. Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biological Cybernetics*, 60:121–130, 1988.
- [5] C. M. Gray, P. Konig, A. Engel, and W. Singer. Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338:334–337, 1989.
- [6] C. M. Gray and W. Singer. Stimulus specific neuronal oscillations in the cat visual cortex: A cortical functional unit. In *Society of Neuroscience Abstracts*, volume 13, page 404.3, 1987.
- [7] David Horn, Nir Levy, and Eytan Ruppin. Memory maintenance via neuronal regulation. *Neural Computation*, 10:1–18, 1998.
- [8] George J. Kalarickal and Jonathan A. Marshall. Visual classical rearing and synaptic plasticity: Comparison of exin and bcm learning rules. In *Proceedings of Vision, Recognition, Action: Neural Models of Mind and Machine*, page 28, 1997.
- [9] Alfredo Kirkwood, Marc G. Rioult, and Mark F. Bear. Experience-dependent modification of synaptic plasticity in visual cortex. *Nature*, 381:526–538, 1996.
- [10] C. Koch. Computation and the single neuron. *Nature*, 385:207–210, 1997.
- [11] Wee Kheng Leow. *VISOR: Learning Visual Schemas in Neural Networks for Object Recognition and Scene Analysis*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin, 1994. Technical Report AI94-219.
- [12] Risto Miikkulainen. Self-organizing process based on lateral inhibition and synaptic resource redistribution. In Teuvo Kohonen, Kai Mäkisara, Olli Simula, and Jari Kangas, editors, *Proceedings of the International Conference on Artificial Neural Networks 1991* (Espoo, Finland), pages 415–420, Amsterdam; New York, 1991. North-Holland.
- [13] Risto Miikkulainen, James A. Bednar, Yoonsuck Choe, and Joseph Sirosh. Self-organization, plasticity, and low-level visual phenomena in a laterally connected map model of the primary visual cortex. In R. L. Goldstone, P. G. Schyns, and D. L. Medin, editors, *Perceptual Learning*, volume 36 of *Psychology of Learning and Motivation*, pages 257–308. Academic Press, San Diego, CA, 1997.
- [14] Kenneth D. Miller. Synaptic economics: Competition and cooperation in correlation-based synaptic plasticity. *Neuron*, 17:371–374, 1996.
- [15] Joseph Sirosh and Risto Miikkulainen. Topographic receptive fields and patterned lateral interaction in a self-organizing model of the primary visual cortex. *Neural Computation*, 9:577–594, 1997.
- [16] Gina G. Turrigiano, Kenneth R. Leslie, Niraj S. Desai, Lana C. Rutherford, and Sacha B. Nelson. Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, 391:845–846, February 1998.
- [17] R. O. Wong and D. M. Oakley. Changing patterns of spontaneous bursting activity of on and off retinal ganglion cells during development. *Neuron*, 16:1087–1095, 1996.