Low Level Modeling of the Development of Directionally Selective Microcircuits in Cat Striate Cortex

Andrew H. Fagg
ahfagg@mensa.usc.edu

József Fiser
fiser@rana.usc.edu

Center for Neural Engineering and Computer Science Department
University of Southern California
Los Angeles, California 90089-2520

Abstract

In this paper development of a directionally selective structure in cat striate cortex is used as a paradigm for exploring issues of biologically plausible and computationally interesting learning algorithms. A compartmental model of a "canonical microcircuit" is used to represent functional units of the cortex. Initially, the thalamus projects in a random fashion to a pair of cortical microcircuits. Through the use of a sliding threshold model of LTP and LTD, these projections develop biologically plausible, directionally selective responses to randomly moving visual stimuli. Possible implications of the learning algorithm on self-organization in the developing cortex are discussed.

Introduction

The main goal of this work is to model brain structure development at the intermediate level of cortical circuits. Two primary benefits are expected from this approach. First, developing realistic models of microcircuit structures in the brain is hoped to serve as a basis for constructing biologically plausible more complex models. Second, models in this domain may be used as a vehicle by which issues involving local learning algorithms may be explored, and contrasted with biology in a principled way.

The model was developed in several stages. First, a compartmental microcircuit model of cat striate cortex was constructed and its behavior was compared with results from electrical stimulation experiments. Second, a pair of microcircuits were wired together and specific thalamic projections were added, such that the two microcircuits behaved in a directionally selective manner. Finally, the thalamic projections were replaced by random connections, and a local Hebbian-antiHebbian learning mechanism was used to develop a directionally selective behavior, similar to what has been observed in cat.

We see the importance of this work in that, to our knowledge, no full implementation of structural development in the cortex has been modeled on this low but functionally predictive level.

The microcircuit structure of the striate cortex

The behavior of the cell structures in the striate cortex is typically explained by selectivity of neurons to different aspects of the incoming information (e.g. orientation, motion direction, or color of a visual stimulus). Existing models of how such selectivity is achieved in the striate cortex can be clustered into three groups. First, it can be the result of selective neurons existing as early as in the retina e.g. (Barlow and Levick, 1965), second, it can be obtained by a specific connectivity pattern in the projection from the lateral geniculate nucleus (LGN) (Hubel and Wiesel, 1962), or it can be a phenomena emerging from the given cortical area under investigation.

In cat, it is known that the input coming from the LGN is not directionally selective, that is cells prior to the striate cortex do not respond with different frequency to stimuli moving to distinct directions (Hubel and Wiesel, 1959). Therefore models based on retinal directional selectivity e.g. (Koch et al., 1986; Borg-Graham and Grzywacz, 1992) are not applicable. According to the dominant theories, strong postsynaptic inhibitory processes in cat cortex shunt the ample non-directional excitation coming from LGN when movement in the non-preferred direction occurs. In contrast, this inhibition does not emerge in case of movement in the preferred direction (Goodwin et al., 1975; Orban, 1984; Koch and Poggio, 1985). However, careful intracellular recordings have not revealed such strong inhibition in response to non-preferred stimuli (Douglas et al. 1991). Moreover, the strongest inhibitory signal is seen when a preferred stimulus is projected to the retina (Berman et al., 1991)

Recently a new model of microcircuits in the cat visual cortex has been proposed which potentially resolves the puzzle (Douglas and Martin, 1991). The model of the basic microcircuit emerged from extended tests of the cat striate cortex with electrical stimulation. Figure 1 shows the block diagram of the proposed circuit.

The three boxes represent three populations of neurons in the cat striate cortex: (a) the pyramidal cells in the superficial layers (layer 2, 3), and the spiny stellate cells in layer 4, (b) the pyramidal cells in the deep layers (layer 5 and 6), and (c) the smooth cells in all these layers. The first two groups have excitatory outputs, the third has inhibitory output of type GABA-a (fast), and GABA-b (slow) lumped together into one type of inhibitory connection in Figure 1. All three populations receive thalamic input, though the deep layers receive only weak signals. All three populations have connections to all other populations, and every population also has recurrent connections back to itself. These connections have been justified by anatomical data (Douglas and Martin, 1991).

The authors wish to thank Michael A. Arbib, Christoph von der Malsburg, and the members of the Brain Simulation Laboratory for their continued support and humor through the course of this work. In addition, insightful discussions with Ken Miller, Rodney Douglas and Humbert Suarez are kindly acknowledged. This work has been sponsored in part by a fellowship from the University of Southern California Graduate School (A.H.F.), School of Engineering, and Computer Science Department, and by the Orszagos Tudomanyos Kutatasi Alap (contract no:285-0813) (J. F.).

The function of the microcircuit is as follows. The thalamic excitation represents only a small percent of the total incoming excitation to each block (10-20% for superficial spiny and smooth, 1-10% for deep spiny cells). The majority of the excitatory input comes from local intracortical connections, which amplify the thalamic signals. The inhibitory control imposed by the fast GABA-a and the slow GABA-b connections prevents over-excitation in the superficial and deep spiny pyramidal and stellate populations. This control takes place at the early phase of cortical processes initiated by an input, therefore it does not need to cope with the full-blown excitation. The excitatory feedback from the spiny populations to the smooth cells assures stability of control. The less the intracortical excitation is, the less input activates the smooth population, which in turn, exhibits less inhibition. As will be seen, this delicate tandem formed by the inhibitory and excitatory signals explains well why it is so difficult to detect inhibition with non-preferred stimuli.

The microcircuit was tested with electrical stimuli by Douglas and Martin, modeling the presence of GABA-a (n-m bicuculine) and GABA-b (baclofen) blockers. The results were compared with experimental data and found to be remarkably similar. As the authors pointed out, this microcircuit can not only be a tool for reproduction of certain restricted data, but an embodiment of a general structural design, i.e. a "canonical circuit" which is potentially capable of formulating many behaviors of the cortex. It is claimed that with appropriate modifications, this microcircuit can be used for reproducing different selectivities. This structure was therefore chosen as a framework for this developmental study.

Physiological evidence for structure and development of directional selectivity

Directional selectivity has a strong relation with orientation selectivity. The evidence for this comes from structural and developmental studies. From structural point of view there is only one important aspect to be considered at the level of our model. Directional selectivity in area 17 in cats has columnar organization similar to that of orientation selectivity (Berman et al. 1987). Bicuculine (a GABA-a blocker) was shown to reduce or even to abolish directional selectivity in the cat's visual cortex (Sillito 1977), therefore inhibitory interneurons probably play a prominent role in establishment of directional selectivity in the striate cortex.

In developmental studies of selectivity, several researchers have shown a deprivation effect on directional selectivity similar to that on orientation selectivity in cats (e.g. Pasternak et al. 1985). As in the case of orientation selectivity, normal development of directional selectivity depends strongly on coherent input from the environment (Cremieux et al., 1992). This suggests that the general rules underlying activity dependent developmental processes for all kinds of selectivity are fundamentally similar. However, since the experiments depriving one type of selectivity development left other selectivities largely unaffected by the given selective deprivation (Cynader and Chernenko, 1976), these developments of different selectivities seem, at least up to a certain degree, to be driven by independent processes. Therefore, independent modeling of one given selectivity seems to be an appropriate first step in investigating developmental processes in the cortex.

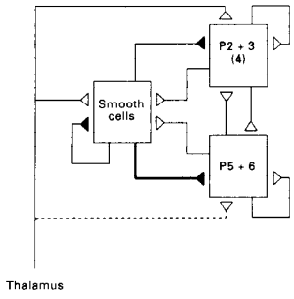

Figure 1: The principal structure of the canonical microcircuit. The three boxes stand for three different populations in the striate cortex. Hollow triangles represent excitatory, filled triangles inhibitory (GABA-a and GABA-b) connections (from Douglas & Martin 1991).

Plasticity in the visual cortex

Long lasting enhancement of the efficacy of synaptic transmission after tetanic stimulation (LTP) is widely accepted to be one of the most important ingredients of plasticity in the brain. Increasing experimental support has been given for the decrease of efficacy of mildly activated synapses (LTD) in cortex (Artola et al., 1990). Though LTP was studied mainly in the hippocampal formation, accumulating evidence shows its similarly prominent role in the cerebral cortex (Tsumoto, 1990).

Modeling synaptic plasticity can be carried out at different levels of abstraction. The work of von der Malsburg (Malsburg, 1988) has demonstrated how a set of cells may differentiate during development into a set of feature detectors that cover a feature space. Specifically, this technique has been applied towards modeling the formation of orientation-specific cells in the primary visual cortex of cat. Key to this work is the fact that no global control is induced upon the developing system. It is only through distributed and local control that the system attains a global organization, implementing the philosophy outlined in (von der Malsburg and Singer, 1988).

One primary drawback of the algorithm suggested by von der Malsburg is its reliance upon a normalizing term that requires direct knowledge instantaneously of all incoming synapses to update a single synapse. The temporal competition mechanism for weight update (Cooper et al., 1979) is one possible alternative to this normalizing term. According to this scheme, when an incidence between an incoming signal and the firing of a cell is detected, a weight update occurs. As opposed to the standard Hebbian update, however, the weights may also be decreased in response to an incidence. The decision to increase or decrease the weights is determined by whether or not the cell's activity has reached above some threshold. Below the threshold an anti-Hebbian update (decrease of efficacy) occurs, and above the threshold, the update is a Hebbian one (increase of efficacy). This scheme, while achieving stability by normalization as the one implemented by von der Malsburg, it resembles the LTP-LTD mechanism reported from experiments in the cortex (Collingridge and Singer, 1990), and it uses inherently local information to perform the computation.

However, this scheme is not robust enough

computationally. The response of all patterns can slip under the LTP-LTD threshold and eventually decrease to zero. Conversely, in the absence of sufficient lateral inhibition, the neuron can become selective to more than one type of input pattern. A possible remedy was suggested by (Bienenstock et al., 1982). According to their scheme, the activity of the cell is kept in check by allowing the LTD-LTP threshold to change over time. In general, this threshold is controlled by the overall activity of the cell (taken over a long period of time). Thus, when a cell is overly active, the threshold will increase, causing responses to most inputs to fall below the threshold into the LTD range, resulting in a decrease in these weights. Similarly, when the cell's activity is very low, the threshold will decrease, resulting in a general increase in weights and activity (see Figure 4). Bienenstock et al. have outlined how this style of computation may be utilized to develop an array of orientation-specific cells.

To date, no explicit physiological evidence supporting this sliding threshold theory has been found. However, alternative theories closer to biology e.g. (Lisman and Goldring, 1988; Lisman, 1989) offer no distinctly articulated complete computational schemas.

Model Design

In the simplest microcircuit model the deep and superficial spiny populations may be lumped together, since from the viewpoint of directional selectivity, they do not carry out different tasks. A circuit containing two columns that are directionally selective in two different directions is depicted in Figure 2.

For the case of a rightward moving stimulus, thalamic cell T1 fires first, activating the spiny unit in microcircuit 1 and the smooth cell of microcircuit 2. When T2 fires 20 msec later, the spiny cell of microcircuit 2 does not activate because it has already been inhibited by the corresponding smooth cell, which was activated earlier by T1.

The overall principle of directional selectivity in this scheme is the same as in all of the previously proposed structures compelling the general requirement of directional selectivity (Poggio and Reichardt, 1973; Borg-Graham and Grzywacz, 1992): it is spatially asymmetric so that two adjacent units receive the excitatory information from a moving input with a delay between them, and the circuit has elements with nonlinear characteristics.

However, there is no need for strong shunting inhibition in the case of null-directional movement, as in other proposed models. Instead, in the case of null-directional movement the output of the spiny cell is reduced by early excitation of the corresponding smooth cell, which prevents full-strength spiny firing. Also, due to the structure of the scheme, the directionally selective response will increase with increasing velocity up to a given point, and then will begin to decrease. This is exactly what has been found experimentally (Orban et al., 1981).

On microscopic level, we followed Douglas & Martin's approach (Douglas and Martin, 1991). Since the morphological and functional features of the neurons in each population are similar, instead of working with thousands of neurons and hundreds of compartments, each population is

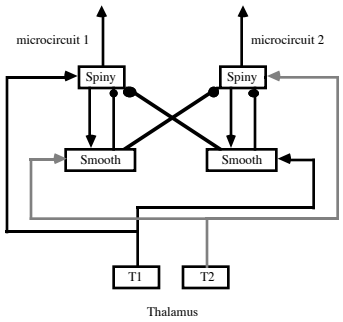

Figure 2. The pre-wired micro-circuit pair. Micro-circuit 1 (left) is sensitive to a rightward-moving stimulus, whereas micro-circuit 2 is sensitive to a left-ward moving stimulus. (Arrowhead - excitatory synapse, circle - inhibitory synapse).

modeled with a single neuron. The smooth and spiny neurons are modeled with an active soma, and two and three compartments of passive dendritic structure, respectively.

The thalamic inputs to the network are modeled as rapid, high-current spikes, which about 2 msec in duration. For each trial, the two thalamic cells fire in succession, with a delay of 20 msec. The order of firing is determined randomly.

The model is implemented using Genesis (Wilson et al., 1991), a biophysical-level simulator.

The Learning Model

Figure 3 shows the initial setup of the network. The projection from the thalamus to the cortex is complete, however the strengths of these weights are chosen to be small, random values. The goal of the learning algorithm is to develop a set of thalamo-cortical connections such that one microcircuit recognizes stimuli moving to the left, and the other recognizes movement to the right. This development occurs in response to successive presentations of moving stimuli in random order. One possible solution to this has already been shown in Figure 2.

Although learning is modeled at a more abstract level than biophysics, two important aspects from this level are integrated into the model : (1) learning is performed using only local information, and (2) learning is a continuous process in time.

The learning algorithm is based upon the correlation-based Hebbian formulation (Hebb, 1949), with the addition of a term referred to as *eligibility*. Eligibility is a function of the local membrane potential and the sliding threshold term (mthresh). This function determines the extent and direction that a synapse strength will change in response to presynaptic activity. The eligibility function is shown in Figure 4.

The weight update in the algorithm is:

$$\Delta w_{ij} = \alpha * eligibility(Vm, mthresh) * a_i * a_j$$

where α is the learning rate, and a_i and a_j are the firing rates of the presynaptic and postsynaptic neurons, respectively. The sliding threshold is updated according to:

$$\tau \frac{d\ mthresh}{dt} = -mthresh + a_j$$

where τ is the time constant for the memory of activation.

As in the standard Hebbian formulation, when activity of the presynaptic cell is correlated with activity of the postsynaptic cell, the weight connecting the two units

changes. However if the membrane potential of the postsynaptic cell does not reach above mthresh, then the weight is decreased. Likewise, when the membrane potential is very high, the weight is increased.

As discussed above, the threshold, *mthresh*, is allowed to shift, depending upon the past history of the neuron's activity. When the cell receives a large amount of input, this threshold increases, causing a reduction in the strength of the weights when the postsynaptic cell is only slightly active. On the other hand, when the cell is only slightly active (over a long period of time), *mthresh* will decay, allowing any inputs to induce a positive change in the weights.

Results

Because the initial weights from the thalamus were small, neither of the microcircuits initially responded very highly to any movement, as shown in Figure 5. As learning progresses, both the microcircuits and the weights leading into individual spiny cells compete with one another. The result is a wiring that supports recognition of each of the two different input cases. Figure 6 shows the development of this wiring over 20 seconds of simulated time (400 msec between stimulus presentations). For each spiny cell, only a single connection from the thalamus achieves a significant value. In addition, each input unit from the thalamus establishes exactly one connection to a spiny cell, yielding a connection scheme similar to the one found in Figure 2.

Figure 7 shows the post-learning microcircuit responses when a rightward moving stimulus is presented. In this case, SPINY_1 has learned to recognize this particular stimulus. SPINY_2, on the other hand, is significantly inhibited by the first pulse (thalamic input), that the second pulse cannot contribute significantly to the cell's activation even though there is a high positive connection from the thalamus. It is also important to note that SMOOTH_1 is highly active. This is consistent with experimental observations that high activity correlate with high amounts of inhibition, and the lowest inhibition is recorded when the depolarization of the excitatory cells is the smallest.

Discussion

It was found that the temporal competition between synapses on a single cell worked well as a mechanism for normalization for a wide range of parameter values. Sometimes , however, both of the spiny cells would learn to recognize the same direction (15% of the time). This problem was minimized by adjusting the strength of the cross-connections between the two microcircuits, to the point that a slight advantage of one spiny cell would cause the other to be shunted, indirectly through the smooth cell. In effect, this implements a contrast enhancement operation, similar to that used by (von der Malsburg and Singer, 1988). This contrast enhancement is key because during a single learning trial, it forces the weight updates to be concentrated onto a small number of weights, rather than affecting all weights in the network.

We also found a similar problem in the simultaneous adjustment of the thalamic-to-spiny projections and the thalamic-to-smooth projections. Not only is it possible for the smooth cells to fall into the situation where they both respond

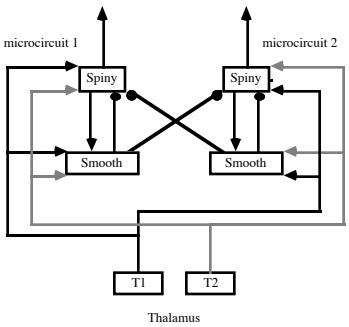

Figure 3. The learning micro-circuit model. The connections from the thalamus to the cortical cells adjust their strengths towards the development of directionally-selective cells.

to stimuli moving in the same direction, they very often chose to respond inconsistently with their corresponding spiny cells. The learning parameters were tuned such that the thalamus-to-spiny projections adjusted quicker than the thalamus-to-smooth projections. This strategy helped to reduce the incidence of inconsistent wiring, because once the thalamus-to-spiny projections began to commit towards a solution, the thalamus-to-smooth projections would follow their lead. More work along these lines is necessary in the future. It would be interesting to test experimentally whether any difference in learning rate between projections to the excitatory and the inhibitory cells can be detected in the cortex of cat.

The sliding threshold mechanism appears to satisfy the goal of normalization through temporal competition. It is necessary, however, to ask about the generality of the algorithm as it is applied to other problems. The primary concern at this point is the apparent dependence of the sliding threshold memory time constant on how often a particular stimulus is presented to the system. As additional stimuli are added to the presentation list, or the frequency of presentation for an individual stimulus is changed, it may be necessary to further adjust the parameters of the model. Further work is necessary to explore this issue.

In addition to the benefits discussed above for a local learning algorithm, we posit that the sliding threshold mechanism might possess further computational ability above (or at least different than) the standard normalization that is so often used in neural algorithms. This mechanism implements temporal competition, as opposed to spatial competition. In other words, the competition limits the overall time that the cell participates in computations, rather than the total "possible" contributions of all of the weights.

Suppose that we wish to train a system to recognize features in the environment and respond with some action. If two features happen to map to the same action, the weights

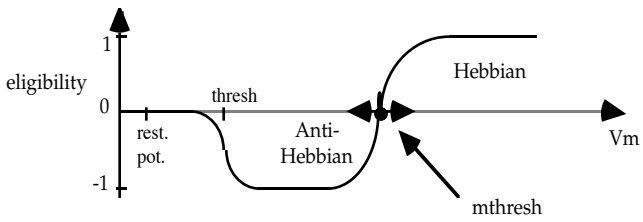

Figure 4. Eligibility as a function of the local membrane potential, a fixed threshold *thresh*, and a sliding threshold *mthresh*.

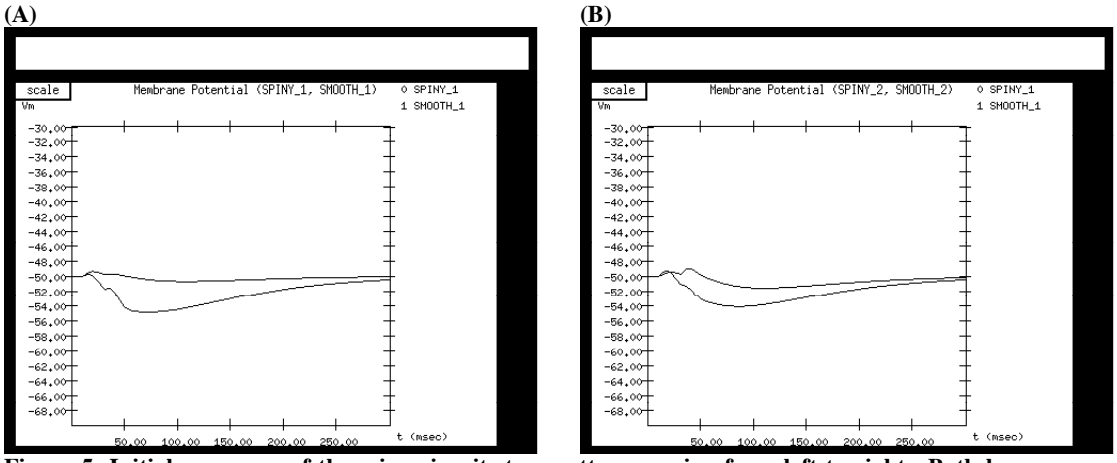

Figure 5: Initial responses of the microcircuits to a pattern moving from left to right. Both lower curves correspond to the membrane potential of (A) spiny cell of microcircuit 1, and (B) spiny cell of microcircuit 2 (the upper curves are the membrane potential of the corresponding smooth cells). The activity of neither the excitatory nor the inhibitory units achieves a high state. Note also that the hyperpolarization of the spiny cells (due to smooth cell inhibition) is only 5 mV below the resting potential.

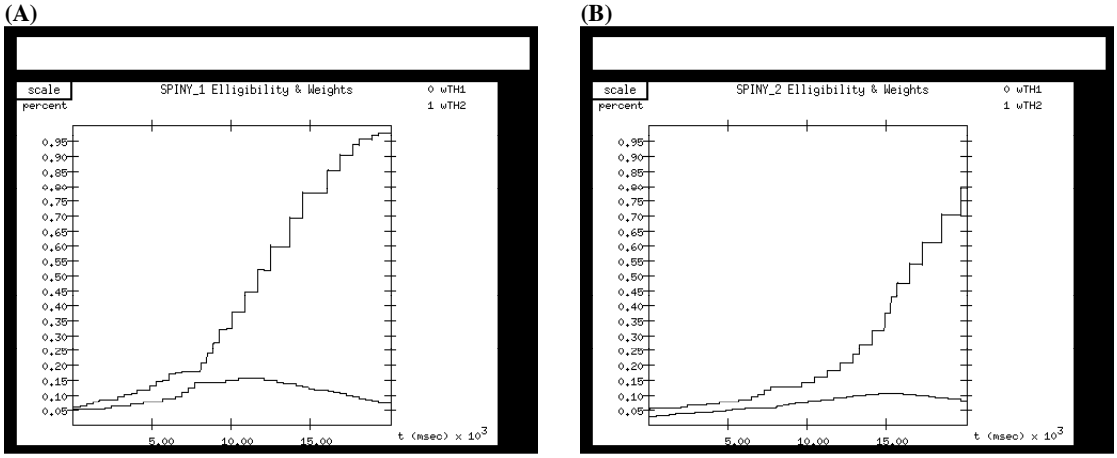

Figure 6: Change of thalamic-spiny weights over the course of an entire experiment (50 stimulus presentations). For spiny cell 1 (A), the effects of competition between the rightward projection (upper curve) and the leftward projection (lower) begin to be very visible at about 7.5 sec. (B) Corresponding weights for micro-circuit 2.

for one feature compete with the weights of the other feature, even if the corresponding inputs do not occur at the same time. The result in the standard normalization case is that the cell is restricted to respond to half the degree to each of the features than it would if it only learned to respond to a single feature. With the sliding threshold mechanism, on the other hand, the cell may be allowed to respond to both features with the original strength, as long as the features did not temporally overlap.

Future Directions

We see several important directions that this model may be taken. First, the model should be extended from the simple two microcircuit model to an array or matrix of connected microcircuits. Second, more biological aspects of the model may be developed further, including adaptation of neural firing (due to calcium), realistic thalamic input, use of three units within an individual microcircuit, molecular encoding of learning, and directional reversal between superficial and deep layers. Third, an important step in this modeling work will be in the learning of additional selectivities, such as orientation, within the same cortical structure. Finally, it will be interesting to investigate how the form of normalization discussed in this work relates to the more general framework of subtractive and multiplicative normalization methods put forth by (Miller and MacKay, 1992).

Conclusions

This paper has presented a self-organizing model of directional selectivity in cat striate cortex. The neural units were modeled at a compartmental level, and were used to construct complex microcircuit building blocks to form the model of the cortex . The self-organization was performed by a correlation-based sliding threshold algorithm. This algorithm is inherently local in the computation of weight updates, lending credence to its biological plausibility, and making it a candidate algorithm for implementation using VSLI technology.

More detailed information regarding this work may be

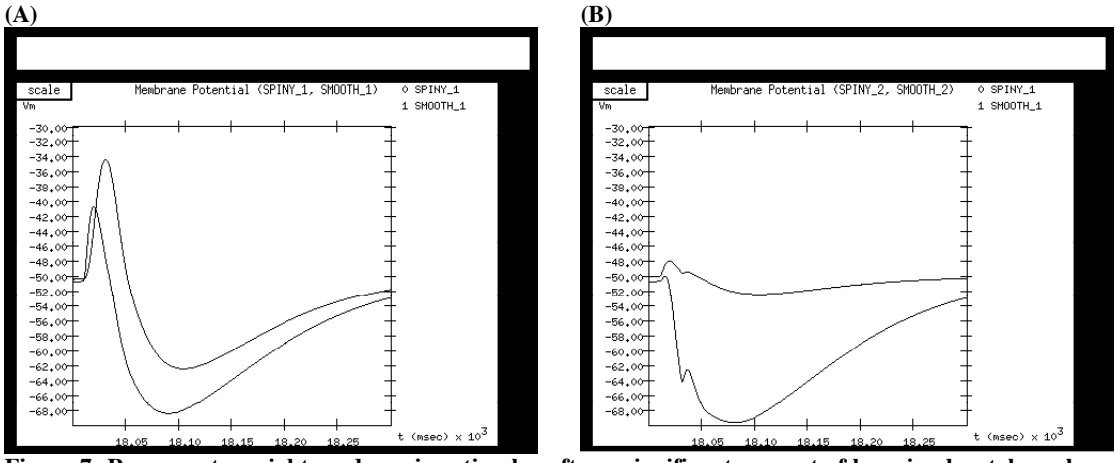

Figure 7: Response to a rightward-moving stimulus after a significant amount of learning has taken place. Spiny cell 1 (A) responds very highly to the stimulus (lower curve). Note also, that smooth cell 1 also responds highly, due to the excitation fed from the spiny cell. Spiny cell 2 (B, lower curve) is already significantly shunted by the time the second pulse (from TH1) arrives from the thalamus, therefore it does not respond significantly but becomes strongly hyperpol,arized.

found in (Fagg and Fiser, 1992).

References

Artola, A., S. Broecher and W. Singer (1990). "Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex." *Nature* **347**: 67-72.

Barlow, H. B. and W. R. Levick (1965). "The mechanism of directionally selective units in rabbit's retina." *J. Physiol*. **178**: 477-504.

Berman, N. J., R. J. Douglas, K. A. C. Martin and D. Whitteridge (1991). "Mechanisms of Inhibition in Cat Visual Cortex." *Journal of Physiology* **440**: 697 - 722.

Bienenstock, E. L., L. N. Cooper and P. W. Munro (1982). "Theory for the Development of Neuron Selectivity: Orientation Specificity and Binocular Interaction in Visual Cortex." *J. Neurosci*. **2**(1): 32-48.

Borg-Graham, L. J. and N. Grzywacz (1992). A model of the direction selectivity circuit in retina: Transformation by neurons singly and in concert. In *Single neuron computation*. Academic Press Inc.

Collingridge, G. L. and W. Singer (1990). "Excitatory amino acid receptors and synaptic plasticity." *Trends in Pharmachological Sciences* **11**: 290-296.

Cooper, L. N., F. Lieberman and E. Oja (1979). "A theory for the acquisition and loss of neuron specificity in visual cortex." *Biol. Cybernetics* **33**: 9-28.

Cremieux, J., P. Buisseret and E. Gary-Bobo (1992). "Experimental evidence that Rearing Kittens in stroboscopic light retards maturation of the visual cortex: a nev tool for studying critical periods." *Vision Res*. **32**(1): 41-45.

Cynader, M. and G. Chernenko (1976). "Abolition of directional Selectivity in the Visual Cortex of the Cat." *Science* **193**: 504-505.

Douglas, R. J. and K. A. C. Martin (1991). "A Functional Microcircuit for Cat Visual Cortex." *J. Physiol*. **440**: 735-769.

Fagg, A. H. and J. Fiser (1992). *Low Level Modeling of the Development of Directionally Selective Microcircuits in Cat Striate Cortex* No. TR #CNE-92-04). Center for Neural Engineering.

Goodwin, A. W., G. H. Henry and P. O. Bishop (1975). "Direction selectivity of simple striate cells: properties and mechanism." *J. Neurophysiol*. **38**: 1524-1540.

Hebb, D. O. (1949). *The Organization of Behavior*. Wiley.

Hubel, D. H. and T. N. Wiesel (1959). "Receptive fields of single neurones in the cat's striate cortex." *J. Physiol*. **148**: 574-591.

Hubel, D. H. and T. N. Wiesel (1962). "Receptive fields, binocular interaction and functional architecture in the cat's visual cortex." *J. Physiol*. **160**: 106-154.

Koch, C. and T. Poggio (1985). The synaptic veto mechanism: does it underlie direction and orientation selectivity in the visual cortex? In D. R. Rose and V. G. Dobson (Eds.), *Models of the Visual Cortex*. Chichester, New York, John Wiley & Sons. 408-419.

Koch, C., V. Torre and T. Poggio (1986). "Commputations in the vertebrate retina: gain enhancement, differentiation and motion discrimination." *Trends in Neurosci*. **9**(5): 204-211.

Lisman, J. (1989). "A Mechanism for the Hebb and the Anti-Hebb Proceses Underlying Learning and Memory." *Proc. Natl. Acad. Sci* **86**: 9574-9578.

Lisman, J. E. and M. A. Goldring (1988). "Feasibility of Long-term Storage of Graded Information by the CA2+/Calmodulin-dependent Protein Kinase Molecules of the Postsynaptic Density." *Proc. Natl. Acad. Sci* **85**: 5320-5324.

Miller, K. D. and D. J. C. MacKay (1992). *The Role of Constraints in Hebbian Learning* No. CNS Memo 19). California Institute of Technology, Computaion and Neural Systems Program.

Orban, G. A. (1984). *Neural operations in the visual cortex*. Berlin, Heidelberg, Springer Verlag.

Orban, G. A., H. Kennedy and H. Maes (1981). "Response to movement of neurons in areas 17 and 18 of the cat: Directional selectivity." *J. Neurophysiol*. **45**: 1059-1072.

Poggio, T. and W. E. Reichardt (1973). "Consideration on models of movement detection." *Kybernetics* **13**: 223-227.

Tsumoto, T. (1990). "Long-Term Potentiation and Depression in the Cerebral Neocortex." *Japanese Journal of Physiology* **40**: 573 - 593.

von der Malsburg, v. d. (1988). Synaptic Plasticity as Basis of Brain Organization. In J.-P. Changeux and M. Konishi (Eds.), *The Neural and Molecular Bases of Learning*. John Wiley & Sons Limited. 411-432.

von der Malsburg, C. and W. Singer (1988). Principles of cortical network organization. In P. Rakic and W. Singer (Eds.), *Neurobiology of Neocortex*. John Wiley & Sons Limited. 69-99.

Wilson, M., J. M. Bower and e. al (1991). *Genesis and Xodus Documentation* No. Caltech.