

## A Learning Rule for Local Synaptic Interactions Between Excitation and Shunting Inhibition

**Chun-Hui Mo**

*mo@klab.caltech.edu*

**Ming Gu**

*gu@caltech.edu*

**Christof Koch**

*koch@klab.caltech.edu*

*Computation and Neural Systems, California Institute of Technology,  
Pasadena, CA 91125, U.S.A.*

The basic requirement for direction selectivity is a nonlinear interaction between two different inputs in space-time. In some models, the interaction is hypothesized to occur between excitation and inhibition of the shunting type in the neuron's dendritic tree. How can the required spatial specificity be acquired in an unsupervised manner? We here propose an activity-based, local learning model that can account for direction selectivity in visual cortex based on such a local veto operation and that depends on synaptically induced changes in intracellular calcium concentration. Our biophysical simulations suggest that a model cell with our learning algorithm can develop direction selectivity organically after unsupervised training. The learning rule is also applicable to a neuron with multiple-direction-selective subunits and to a pair of cells with opposite-direction selectivities and is stable under different starting conditions, delays, and velocities.

### 1 Introduction ---

The ability to distinguish the direction of movement is important. Cats reared in a stroboscopically illuminated environment developed normal orientation-selective neurons in cortex, but direction-selective neurons were virtually abolished. This effect remained after long periods of normal visual exposure (Cynader & Chernenko, 1976; Humphrey & Saul, 1998; Saul & Feidler, 2002). Therefore, direction selectivity is likely to require spatiotemporally structured synaptic input during early developmental stages. These neurons are thus excellent targets for the study of activity-dependent synaptic weight changes.

Hebb (1949) proposed his famous learning rule based on the correlation between pre- and postsynaptic activities. A computational study showed that a Hebbian learning rule performs poorly in direction-selective synapse

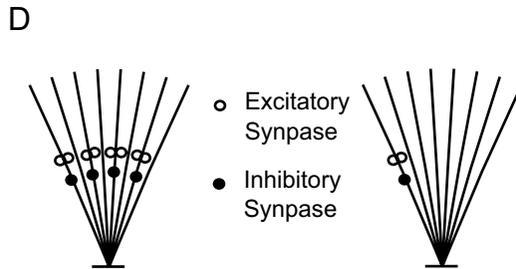
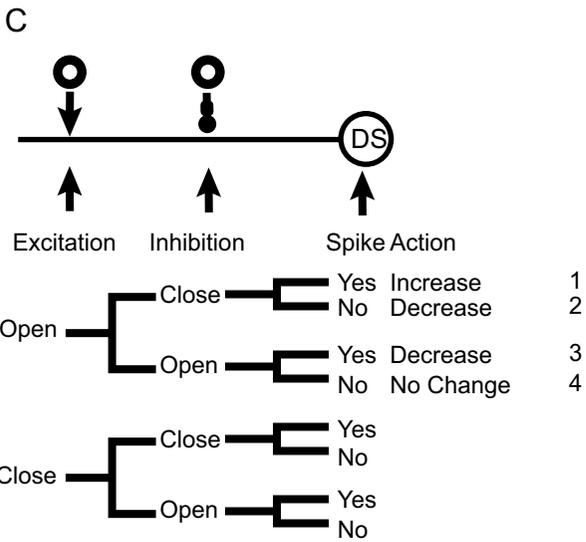
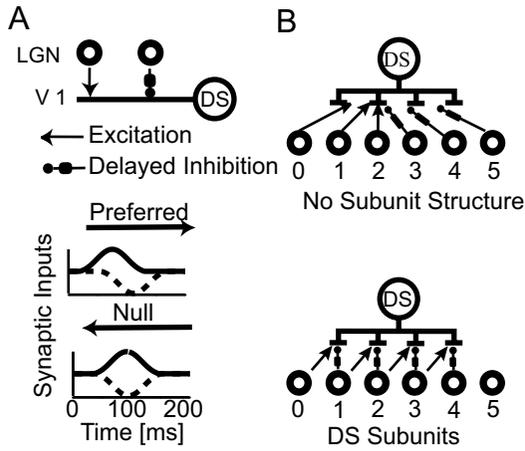
placement (Feidler, Saul, Murthy, & Humphrey, 1997). This is not surprising since the basic requirement for direction selectivity is a nonlinear interaction between two different inputs in space-time. In a natural environment, as many stimuli are expected to move in the preferred as in the null direction. While there exists a correlation between pre- and postsynaptic firing in the preferred direction, this is not the case for motion in the opposite, null, direction. There are modifications of the original Hebbian learning rule that could account for the development of direction selectivity, in particular, postsynaptic “gating” rules that link synaptic weight changes to postsynaptic activities (Feidler, et al., 1997; Blais, Cooper, & Shouval, 2000).

Spike-time-dependent plasticity (STDP), a temporally asymmetric Hebbian learning rule (Montague & Sejnowski, 1994; Montague, Dayan, Person, & Sejnowski, 1995; Markram, Lubke, Frotscher, & Sakmann, 1997), in which the change in synaptic weight depends on the relative timing of the presynaptic input and the back-propagating postsynaptic spike, might be crucial here. If the postsynaptic spike arrives a few milliseconds before the presynaptic input, the synaptic weight is weakened; if it arrives a few milliseconds later, the weight is increased. A network of neocortical neurons implementing STDP developed direction selectivity after training (Rao & Sejnowski, 2000, 2001; Buchs & Senn, 2002; see also Song, Miller, & Abbott, 2002).

A great deal is known about direction-selective cortical cells. Asymmetrical delayed inhibition is likely to be one of the mechanisms that underlie direction selectivity (Koch & Poggio, 1985; Livingstone, 1998). A bar moving in the preferred direction reaches the excitatory input before the inhibitory one, which acts only after an additional delay (see Figure 1A). The excitatory input reaches the soma and causes the cell to spike because of the temporal

---

Figure 1: *Facing page*. A direction selective (DS) mechanism and basic kinds of information available to an excitatory synapse for its learning strategy. (A) A DS neuron in V1 receives two LGN inputs—one excitatory and one delayed inhibitory (via a cortical interneuron). Solid curves show excitatory synaptic conductance changes during the preferred and null direction motion stimuli. Dashed curves show inhibitory synaptic conductance changes plotted in negative. The difference in temporal alignment of these two inputs during motion stimuli forms the basis of the cell’s direction selectivity. (B) Connection diagram for two DS cells with (bottom) and without (top) DS subunits on their dendrites. Both receive inputs from the same LGN cell array. A bar needs to move across the middle line between LGN input cell 2 and 3’s receptive fields to elicit directional response in the top cell. The DS cell at the bottom has superior position-invariant direction selectivity. (C) Three types of information available to an excitatory synapse on a remote dendrite of a DS cell to determine whether its own activity is contributing to the direction selectivity of the host cell. (D) The model neuron and relative excitatory and inhibitory synapses placement. Left: multiple DS subunits. Right: single unit model.



offset between the two inputs. In the opposite, null direction, the excitatory input is “vetoed” by the inhibition if the bar’s speed is approximately matched to the delay. The wiring requirement for such a scheme is simple: excitation in visual space can reside on either side of the inhibitory zone but not on both sides, in which case the model receives symmetric input in space-time and is thus not direction selective. Large conductance changes that reverse around the cell’s resting potential have been observed in V1 during visual stimuli (Anderson, Carandini, & Ferster, 2000; Borg-Graham, Monier, & Frégnac, 1998). Given the special property of shunting inhibition, it is interesting to investigate its possible role in synaptic learning.

Both retinal and cortical direction-selective cells have subunit structures within their receptive fields (Barlow & Levick, 1965; Emerson, Citron, Vaughan, & Klein, 1987; Livingstone, Pack, & Born, 2001). A comparison of two fictitious cortical DS cells with and without subunit structures is shown in Figure 1B. Both cells receive inputs from the same group of lateral geniculate (LGN) cells, but the one with subunit structures utilizes its resources more efficiently and has superior position-invariant direction tuning.

Synaptic logic models involving complex, branch-specific synaptic placements have been proposed (Poggio & Torre, 1978; Poggio, 1982; Koch & Poggio, 1987). Recent work by Mel and colleagues suggests that the dendritic tree of cortical pyramidal neurons may function as a two-layer neural network (Poirazi, Brannon, & Mel, 2003). Standard Hebbian learning rules, describing changes in the overall connection strength between pre- and postsynaptic neurons, cannot account for dendritic learning (Mel, 2002), for they do not distinguish among postsynaptic connection locations. We demonstrate, using an idealized compartmental simulation, how a learning rule for local synaptic interactions between excitation and shunting inhibition can, in principle, account for direction selectivity and subunit learning at the dendritic level of a single neuron.

## 2 Methods

---

All compartmental simulations are carried out using the program NEURON (Hines & Carnevale, 1997). The idealized cell morphology of a direction-selective neuron includes eight dendrites (width  $0.5 \mu\text{m}$ , length  $100 \mu\text{m}$ ) that are directly connected to the soma (width  $16 \mu\text{m}$ , length  $16 \mu\text{m}$ ). Each dendrite is unbranched and has 20 compartments, for a total of 161 compartments (the soma is modeled as a single compartment). The dendrites are passive except for an N-type calcium conductance, while the cell body contains sodium and potassium conductances that give rise to fast Hodgkin-Huxley-like action potentials without spike adaptation. The biophysical parameters are:  $R_a = 250 \Omega \bullet \text{cm}$ ,  $C_m = 0.5 \mu\text{F}/\text{cm}^2$ ,  $E_{\text{leak}} = -60 \text{ mV}$ ,  $R_m = 10\text{k} \Omega \bullet \text{cm}^2$ ,  $g_{\text{Na}} = 0.030 \text{ S}/\text{cm}^2$ ,  $g_{\text{K}} = 0.028 \text{ S}/\text{cm}^2$ ,  $E_{\text{NMDA}} = 0 \text{ mV}$ ,  $g_{\text{NMDA}} = 0 - 2 \text{ nS}$ ,  $\tau_{\text{NMDAon}} = 0.1 \text{ ms}$ ,  $\tau_{\text{NMDAoff}} = 80 \text{ ms}$ ,  $E_{\text{AMPA}} = 0 \text{ mV}$ ,  $g_{\text{AMPA}} = 0 - 2 \text{ nS}$ ,  $\tau_{\text{AMPAon}} = 0.1 \text{ ms}$ ,  $\tau_{\text{AMPAoff}} = 2 \text{ ms}$ ,  $E_{\text{GABA}} = -60 \text{ mV}$ ,  $g_{\text{GABA}} = 5.0 \text{ nS}$ ,

$\tau_{GABA_{on}} = 1$  ms, and  $\tau_{GABA_{off}} = 80$  ms. Synaptic input is modeled using the point process in NEURON (adopted from Archie & Mel, 2000).

The N-type voltage-gated calcium conductance is taken from Benison, Keizer, Chalupa, and Robinson (2001) and mapped to dendrites with a density of 1 mS/cm<sup>2</sup>. We assume that excitatory synapses are located to dendritic spines and that the rapid rise in intracellular calcium concentration at the postsynaptic site inside the spine following synaptic activation has two additive sources (see Figure 2A): the calcium current through the NMDA synapse,  $I_{Ca\_NMDA}$ , and the calcium current,  $I_{CaN}$ , located in the dendrite and a certain amount of which diffuses up into the spine.  $I_{Ca\_NMDA}$  is calculated as one-third of the total current through the NMDA synapse with a reversal potential  $E_{Ca} = 130$  mV. Of the  $I_{CaN}$  entering the dendritic compartment associated with the spine, 5% is assumed to contribute instantaneously to the calcium concentration in the spine head compartment (Koch, 1999). These scaling factors are chosen to bring the amounts of calcium at the spine entering through NMDA channels and through N-type voltage-gated calcium channels within the same order of magnitude. The final concentration of free, intracellular calcium at the spine,  $[Ca^{2+}]$ , is given by a simple decay equation,

$$\frac{d[Ca^{2+}]}{dt} = I_{Ca\_NMDA} + I_{CaN} - \frac{[Ca^{2+}]}{\tau_{Ca^{2+}}},$$

with  $\tau_{Ca^{2+}} = 15$  ms. We model all internal calcium buffers and calcium pumps using a single decay constant. The stimulus is a one-dimensional bar moving at 10 degrees per second across the receptive fields of six LGN cells (see Figure 1B; for details, see Mo & Koch, 2003). There are eight excitatory and four inhibitory, deterministic synapses in the model (each corresponding to a cluster of probabilistic synapses; see Figure 1D). Each geniculate input is assumed to directly excite its appropriate dendrite at a single synaptic cluster and, via a local interneuron, to inhibit a dendrite. This is modeled by a single inhibitory synapse that is delayed by 10 ms with respect to excitation. Excitatory synapses are mapped to the dendritic compartment 60  $\mu$ m away from the soma; inhibitory synapses are located 50  $\mu$ m away from the cell body. In the main model, inhibition is assumed to be of the shunting (or silent) type, with a reversal potential  $E_{GABA} = -60$  mV (the cell's resting potential is  $-60$  mV). The temporal dynamics of the excitatory (NMDA) and inhibitory (GABA) synaptically induced conductance changes are as described by Destexhe, Mainen, and Sejnowski (1994).

The change in the weight of the excitatory synapse is a function of the intracellular calcium concentration in the spine head compartment following synaptic activation:  $Weight\ Change = a * \sqrt{c * (-[Ca^{2+}] - d) - \exp(c * (-[Ca^{2+}] - d))} + b$ , with  $a = -3.3$ ,  $b = 1$ , and  $c = 13$ .  $d$  is a constant that varies continuously between  $-0.10$  and  $-0.22$  as  $g_{NMDA}$  varies from 0 to 2 nS. The learning curve is chosen to give a negative output at a medium calcium concentration and a positive one at high calcium con-



centration (see Figure 2C). The parameters  $a$  and  $b$  are chosen to restrict the function's output to between  $-1$  and  $1$ . The parameter  $c$  is a scaling factor that determines the width of the curve, and parameter  $d$  is a sliding threshold that linearly shifts the curve according to different values of  $g_{NMDA}$  (see Figure 2D).  $g_{AMPA}$  is always set to equal  $g_{NMDA}$ . Other types of equations that decrease synaptic weight at medium calcium concentrations and increase it at high calcium concentrations can also be used to model the learning curve (personal communication with Ming Tao, 2003).

### 3 Results

---

We consider the local information available to an excitatory geniculate synapse on a dendrite of a cortical neuron to correctly judge whether its activity contributed to or degraded the direction selectivity (DS) of the host cell, assuming the inhibitory input has already been connected and fixed. There are three major pieces of information accessible to local mechanisms: the state of the excitatory input, the state of the inhibitory input, and whether the host cell generated a somatic action potential within a small time window and this spike propagated back into the dendrite to the postsynaptic site of excitation. Assuming binary states (e.g., excitatory input is either on or off), this gives rise to eight possible scenarios (for instance, both excitation and inhibition are active and the host cell spikes). We assume the excitatory synapse can be modified only when it is active; this reduces the combinations to four scenarios (see Figure 1C). In scenario 1, there is no inhibition, and the cell spikes after the excitatory synapse opens. The assumption is that the excitatory synapse directly contributes to the cell's direction

---

Figure 2: *Facing page*. Calcium dynamics at a spine and the BCM learning curve based on local calcium concentration change. (A) A schematic drawing of calcium entrance points at spines and the local dendritic branch. Calcium can enter spines directly through NMDA channels or indirectly through N-type voltage-gated calcium channels activated by backpropagating somatic action potentials. Shunting inhibition is located between the spine and soma; it can block backpropagating spikes and clamp the membrane voltage to reduce calcium entry. (B) Local calcium concentration changes (solid curves) following synaptic activation for four different learning scenarios. Dashed (resp. dotted) curves show the calcium concentration changes due to calcium entering through N-type voltage-gated calcium channels (resp. local NMDA channels). Peak calcium exposures within 30 ms are marked. (C) A BCM type learning curve ( $g_{NMDA} = 1$  nS). Peak calcium concentration reached in each of the four scenarios described. The corresponding synaptic weight changes are marked on the curve. Calcium concentrations are in an arbitrary unit. The synaptic weight change is relative to the maximum allowed learning step size in each trial. (D) A linear sliding threshold is chosen for the learning curve. The three learning curves shown are calculated at  $g_{NMDA} = 0$  nS (dashed),  $g_{NMDA} = 1$  nS (solid),  $g_{NMDA} = 2$  nS (dotted).

selectivity, and thus its connection strength should increase. In scenario 2, there is no inhibition and no spike when the excitatory synapse opens. The fact that the DS cell is not spiking suggests that the stimulus moved in the null direction, yet the excitatory input is not gated by the inhibition and counteracts the cell's direction selectivity. Its connection strength should decrease. In scenario 3, both excitation and inhibition are active, yet the cell still fires an action potential. The assumption is that this scenario corresponds to null direction of motion and that the excitatory synapse landed on a spot with incorrect matching of inhibition. Its connection strength should therefore decrease. In scenario 4, the excitation is successfully blocked by inhibition, suggesting a null direction movement, in which case the blocking by inhibition is legitimate. On the other hand, this could also correspond to a preferred direction movement. The cell generates an action potential in the presence of inhibition, yet this spike, propagating back from the soma into the dendrite, is blocked by inhibition from reaching the site of excitation. Given this ambiguity, the best possible action is to do nothing and keep the excitatory weight constant.

An excitatory synapse can adjust its weight, via the local calcium concentration change, provided it can distinguish these four scenarios. This biophysical variable then determines whether the amplitude of the excitatory input is increased, remains the same, or is decreased.

**3.1 Calcium Dynamics at the Spine and the Local Learning Rule.** Figure 2A illustrates a dendritic branch with two spines, each with an independent excitatory input. Inhibition is mapped to a dendritic compartment between excitation and the soma, fulfilling the "on-the-path" requirement (Koch, Poggio, & Torre, 1982). We mapped the two excitatory synapses onto one electrically equivalent, dendritic compartment. While calcium can enter the spine from the dendrite, there is likely to be a severe calcium concentration gradient from the spine to the dendrite (given the large volume difference between the spine head and the dendrite and calcium pumps along the thin neck of the spine). Therefore, two nearby spines may be chemically independent, although they are electrically equivalent (Zador, Koch, & Brown, 1990).

The excitatory and inhibitory inputs and the backpropagating spike from the soma affect the intracellular calcium concentration at the spine in their own way. Excitation directly correlates with the calcium current entering through NMDA channels, which is spine and synapse specific. Its time course is mainly determined by the conductance change of the local NMDA synapse. Given the NMDA synapse's reversal potential (set to zero), however, the synaptic input current by itself cannot elevate the membrane potential high enough to cause significant activation of the N-type voltage-gated calcium channel, which are mainly activated when there is a backpropagating spike signaling a global activation state of the DS cell. On-the-path shunting inhibition affects the time course of both calcium currents. It clamps the

membrane voltage to the resting potential when activated, thereby reducing the amount of calcium current entering through NMDA channels. Furthermore, it blocks backpropagating spikes, which significantly reduces calcium entrance into the spine through N-type voltage-gated calcium channels. The clamping and blocking effects are branch specific due to the local action of shunting inhibition.

We implemented the above calcium scheme and computed the resultant changes in free, intracellular calcium concentration at the spine for the four scenarios considered earlier (see Figure 2B). In scenario 1, there is no local inhibition, and the cell spikes after the excitatory synapse opens. Calcium enters through both channel populations. Changes of  $[Ca^{2+}]$  are high. In scenario 2, there is neither local inhibition nor a backpropagating spike immediately after the excitatory synapse opens. Calcium mainly enters through NMDA channels. Changes of  $[Ca^{2+}]$  fall into an intermediate range. In scenario 3, the excitatory input is blocked by the inhibition, but the cell spikes. The amount of calcium entering through NMDA channels is reduced by inhibition. Meanwhile, residual calcium enters through N-type voltage-gated calcium channels due to the backpropagating action potential immediately before the synapse opens, adding to the total calcium concentration. Changes of  $[Ca^{2+}]$  fall, again, into an intermediate range. In scenario 4, excitation is blocked by inhibition in the absence of any action potential. Only a limited amount of calcium enters through the NMDA synapse due to the clamping effect of shunting inhibition. Changes of  $[Ca^{2+}]$  are low.

As we mentioned earlier, the proper action for scenario 1 is to increase the weight of the excitatory synapse. Changes of  $[Ca^{2+}]$  in this case are high. The proper action for scenarios 2 and 3 is to decrease the synaptic weight. Changes of  $[Ca^{2+}]$  in these cases are medium. The proper action for scenario 4 is to keep the synaptic weight unchanged. Changes of  $[Ca^{2+}]$  are low. In order to link the synaptic weight change with maximum calcium exposure, the amplitude of each excitatory synapse from the geniculate input to the target cell is changed in accordance with the maximum calcium concentration change at the synapse within 30 ms of synaptic activation. Such a change is computed using the BCM rule (see Figure 2C), a well-known learning rule for cortical plasticity, proposed by Bienenstock, Cooper, and Munro (1982).

As training progresses and the synaptic weight increases, the threshold for learning needs to increase to prevent runaway excitation and to stabilize the synapse (Bienenstock et al., 1982; Abbott & Nelson, 2000). We use a linear sliding threshold to shift the learning curve without changing its shape, as shown in Figure 2D. We link the sliding threshold directly to the excitatory synaptic connection strength (see section 2 for details): the larger the connection strength (i.e., its postsynaptic conductance), the larger the increase in intracellular calcium concentration necessary to trigger further increase in synaptic weight.

**3.2 Direction-Selective Single Unit Learning.** We first tested our learning rule in a model cell without subunits (see Figure 3A). The model initially receives balanced excitatory inputs from both left and right LGN neurons. Two excitatory inputs are mapped onto the same compartment of the dendrite. The initial connection strength is 1 nS each. Delayed inhibition is fixed at 5 nS and is mapped to a compartment between excitation and the soma. During each trial, a bright bar randomly moves to either the left or to the right, and the maximum change of  $[Ca^{2+}]$  at each excitatory synapse within 30 ms following activation is recorded. After each trial, the synaptic weight change is calculated based on the learning curve shown in Figure 2D.

A synapse can be rewarded only if the host cell spikes and the spike successfully invades the dendrite. Our learning model cannot become direction selective if none of the excitatory inputs is strong enough to generate a spike. Therefore, we need to impose a “competition” rule to control the model’s excitability. We do this by holding the total excitatory connection strength over each local dendrite constant during simulations. This rule is of the “subtraction” type (Abbott & Nelson, 2000); that is, after each trial, half of the value of the synaptic weight above or below the total connection strength is subtracted from or added to both excitatory synapses. If the two excitatory synaptic weights are  $g_{e1}$  and  $g_{e2}$ , then the rule specifies  $g_{e1\ new} = g_{e1\ old} - (g_{e1\ old} + g_{e2\ old} - 2\ nS)/2$ . This competition rule prevents both inputs from slipping to zero. We assume there is no initial bias, and the model cell initially receives balanced input from the left and the right. The outcome is dependent on the training sequence.

The weight changes for both synapses during a simulation run that lasted 300 trials are shown in Figure 3B. Before training, the model cell responds equally for motion in either direction (see Figure 3D). During the initial training period, if a rightward-moving bar is present, the left excitatory synapse opens first and causes the host cell to spike (scenario 1 for the left synapse). Then the inhibitory synapse opens and blocks the excitatory input from the right excitatory synapse (scenario 3 for this synapse). After the trial, the left connection is strengthened and the right one weakened. Similarly, if a leftward-moving bar is shown to the model, the right connection is strengthened and the left weakened after the trial. If the training regime consisted of alternating left and right stimuli, we would expect the synaptic strengths of both sides to oscillate within a range close to the learning step size but never converge. However, a random training sequence contains consecutive left or right trials and thus causes the oscillation to be larger than one learning step. The longer the training sequence of the same direction, the bigger the expected oscillation amplitude is. Once the oscillation reaches a large enough value such that the connection strength of one excitatory input, say, the right input, drops below a value sufficient to elicit a somatic spike, the oscillation stops. Now, during its preferred direction motion (a bar moving from the right to left), its weight is decreased according to scenario 2 instead of being increased according to scenario 1. The right

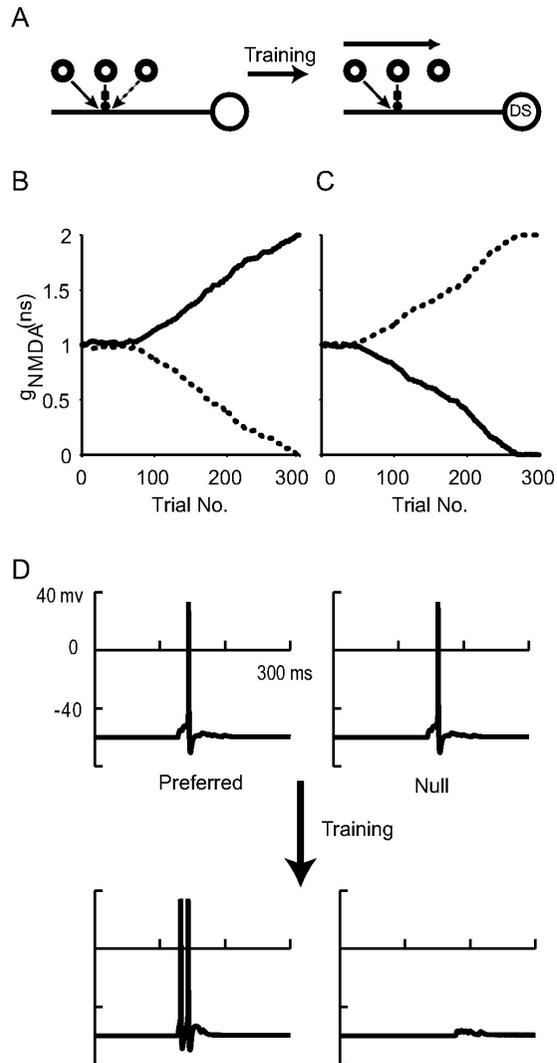


Figure 3: Model cell with a single subunit implementing these rules acquires DS following unsupervised training. (A) The model cell initially receives balanced inputs from the left and right and is not direction selective. After training with random bar movements, the input connection from one side is strengthened, while the input connection from the other side is weakened and the model cell becomes direction selective. (B) Synaptic weight changes for motion for the left (solid curve) or right (dashed curve) input during one simulation run. Learning step size is 0.01 nS. (C) A different simulation run. After learning, the cell responds only to leftward motion. (D) The cell's response to a bright bar moving at 10 degrees per second across its receptive field before and after training.

input thus enters a downward spiral and gradually decreases its weight to zero, while its counterpart, the left input, gradually increases its weight to the maximum allowed value. In the simulation shown in Figure 3B, this transition occurs around 80 trials. After training, only the excitatory input to the left side of inhibition remains, and the cell fires in a direction-selective manner (see Figure 3D).

Figure 3C shows another simulation run during which the right excitatory input cell won the competition and the preferred direction of motion was reversed. We ran 100 simulations with a 0.032 nS learning step size; during all of this, the model cell converged to a DS cell within 200 trials (in 52 out of these 100 simulations, the preferred direction was rightward). An index of direction selectivity (DI) is computed as (preferred direction response – null direction response) / (preferred direction response + null direction response). DI values close to zero indicate a lack of direction selectivity, while the maximal extent of selectivity yields  $DI = 1$ . In all above cases, the model cell reached  $DI = 1$  after training.

**3.3 Learning Multiple, Direction-Selective, Subunits.** How can our learning rule ensure that direction selectivity in different dendritic subunits of the host neuron is the same? To answer this question, we tested our DS learning rule in a model cell with four direction-selective subunits on four of its eight dendrites (see Figure 4A). Each of the middle four LGN cells (1–4) provides delayed inhibitory input to one dendrite (1–4 from the left) of the model cell. LGN cells 0–3 each provides a left excitatory input to dendrites 1–4 respectively. We refer to this group as the left input connection group. LGN cells 2–5 each provides a right excitatory input to dendrites 1–4, respectively (right input connection group). We refer to the connections within a group as “friends” and the connections between groups as “competitors.” The learning goal is to have all members within one group outcompete their competitors after training.

If the four subunits were completely independent and all received the exact same sequence of visual stimuli, we would expect them to converge to the same direction selectivity. Unfortunately, neither of these two conditions is true. Although at each trial the same moving bar is presented to each subunit, the exact timing of the bar reaching the receptive field of each geniculate cell is different. Both NMDA and GABA synapses have long off-ramps. Their late currents can cause differences in the status of subunits. The shunting inhibition mostly affects local connections, but it also has a more global effect. Because the same learning curve is used for all synapses, these differences can cause different branches to learn to respond to opposite directions of motion. The model cell may thus not become direction selective. The problem can be solved if there are internal links between group members and competition between the groups. The links between group members indeed exist in our model through somatic spikes. For example, the LGN cell 1’s input connection to dendrite 2 and the LGN cell 2’s input

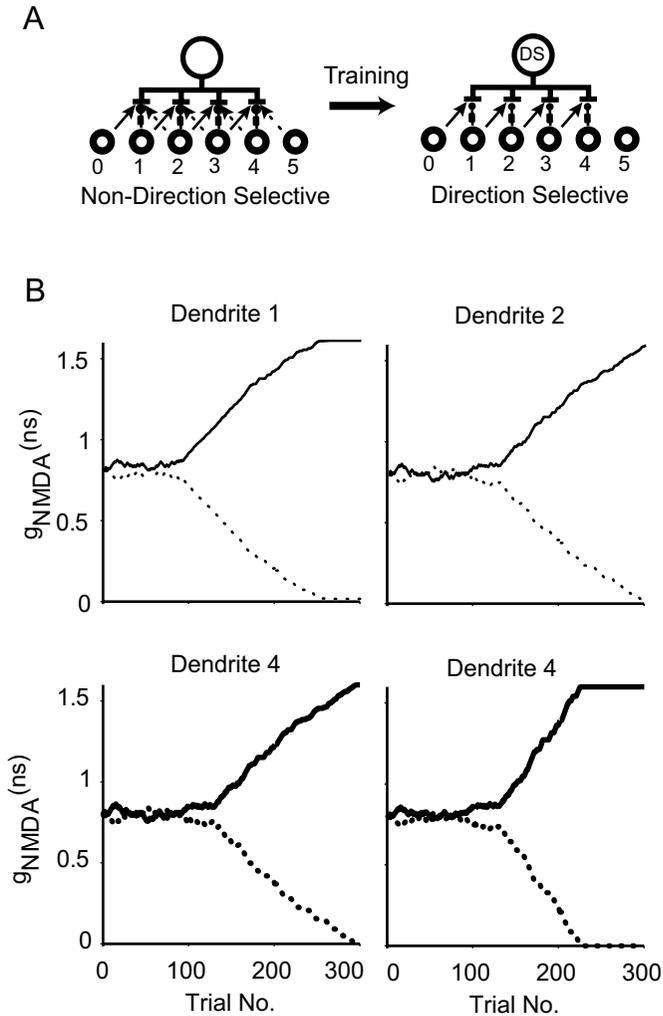


Figure 4: The model with four DS subunits after unsupervised training with a balanced initial setting. (A) All four dendrites initially received balanced excitatory inputs. After training with random bar movements, the input connection from one side is strengthened while the input connection from the other side is weakened. Due to coupling among dendrites and the local nature of shunting inhibition, all subunits are selective for the same direction of motion. (B) Synaptic weight changes for the left (solid curves) and right excitatory inputs (dashed curves) at each dendrite during one simulation run. Learning step size is 0.003 nS. After training, the cell responds only to rightward motion.

connection to dendrite 3 belong to the same left input group. In a rightward movement trial, a late spike caused by LGN cell 1's input will also be counted as a spike caused by LGN cell 2 given the overlap of their input time courses. In case LGN cell 2's connection strength drops below the transition threshold, this would "rescue" it from scenario 2 to scenario 1. The same spike, however, will not help LGN cell 2's connection to dendrite 1, which belongs to the right input group. The inhibitory input to dendrite 1 always opens before LGN cell 2 on dendrite 1 during a rightward movement trial and thus blocks any backpropagating spike from reaching LGN cell 2's connection point. So this "link forward" effect benefits only group friends, not competitors. A similar "link backward" effect also exists.

An early spike caused by LGN cell 2 will also be registered as LGN cell 1's own spike if it happens within 30 ms of LGN cell 1's firing. The "rescue" effort occurs only during the preferred direction movement; there are no linkages among group members during the null direction movement. If a group member is stuck far away from the divergent point, long, consecutive same-direction trials are required to increase its connection strength above the spiking threshold. To speed up convergence, we imposed an additional "majority" rule. We scale the learning step size of each trial with the total number of action potentials generated at the soma during that trial. In such a setting, a group member is increased more in its preferred direction and decreased less in its null direction, once its group responds with more spikes than the other group. This creates direct competition between groups and thus facilitates convergence. We also experimented with a modified version of this majority rule that is more plausible on biophysical grounds: here, the scaling is applied only if the synaptic weight increases. The decrease of weight does not depend on the number of spikes triggered. Both rules yielded the same results.

Initially, the model cell received balanced inputs at each of its dendrites (see Figure 4B) and is not direction selective. After 300 training trials, the entire left input group wins over the right input group, and the cell develops four DS subunits sensitive to rightward motion. Each subunit reached its divergent point at different trials and went through different weight change trajectories. Note that initially, LGN cell 2 provided excitatory input to both dendrites 1 and 3. After training, only the connection to dendrite 3 remains. Therefore, the learning process is indeed branch specific. We carried out 100 simulations with a 0.032 nS learning step size, which is increased linearly with the number of action potentials generated according to our majority rule. In all cases, the model cell achieved uniform DS subunit structures within 200 trials with  $DI = 1$ . The model converged to a right-direction-selective unit during 47 simulation runs and a left-direction-selective unit during the remaining runs.

In the above simulations, all dendrites receive balanced input from each side. We further tested our learning model in a "random start" configuration. The total input connection strength to a dendrite is still fixed, but the rel-

ative contribution from the left input cell and the right input cell is randomly assigned (see Figure 5A). After 200 training trials, the right input group won at dendrite 1, and the left input group was leading at dendrites 3 and 4. Because of the majority rule and the link-forward and link-backward effects, the left input group finally increased its “friends” connection strength at dendrite 1 and 2 and destabilized its “competitors.” After 1000 training trials, the cell developed four rightward-motion-selective subunits. To test the stability of our learning model in the random start condition, we ran 10 simulations, each with four different learning step sizes: 0.1 nS, 0.032 nS, 0.01 nS, and 0.003 nS. Training periods are 100 trials, 500 trials, 1000 trials, and 2000 trials, respectively. DI converged to 1 in all conditions.

Another possible scenario during development is that initially all the excitatory geniculate inputs to the V1 cell are very weak—insufficient to generate a spike. As the input connections are strengthened, the cell starts to spike, and competition among input synapses begins. We tested our learning model in such a developmental configuration. Initially, all connection weights are zero. They are then gradually increased because of the “competition” rule. The total excitatory connection strength to a dendrite is low, so at each trial, both inputs are increased by the maximum allowed learning step size (see Figure 5B). During the first 100 trials, there are no spikes, and all the input connections increase at each trial. Once the input connection strength reaches about 0.3 nS, the model cell starts to spike, and the connection strength of different input groups starts to diverge. After 600 training trials, the model cell develops four rightward-motion-selective subunits. To test the stability of our learning model in the development condition, we carried out 10 simulations each at four different learning step sizes. The model cell converged to  $DI = 1$  under all conditions.

**3.4 Differential Excitation-Inhibition Learning.** Up to now, the inhibitory connection strength was kept fixed throughout the simulations. Here we investigate what happens when we relax this constraint and let inhibition converge to its optimal connection strength. Given the fact that all inhibitory synaptic connection strengths are the same on all dendrites, we do not need a local, dendrite-specific, learning rule. We do not need a direction-selective one either, since the excitatory learning rule accomplished this purpose.

We used a simple inhibition learning rule that links the amplitude of inhibition to the model cell’s average response, as described by Soto-Trevino, Thoroughman, Marder, and Abbott (2001). We set a target spike number of 3 spikes per trial and keep a 30-trial spiking history. At each trial, the inhibitory synaptic weight change is linearly related to the difference between the target value and the average spike number:  $g_{inhibition\ new} = g_{inhibition\ old} + (average - target) * 0.1\ nS/spike$ . We cap the inhibitory synaptic weight at 20 nS. This learning rule is not a direction-selective one: a DS cell that spikes six times in its preferred direction and zero in its null direction or a non-DS cell that spikes three times in both directions both fulfill the re-

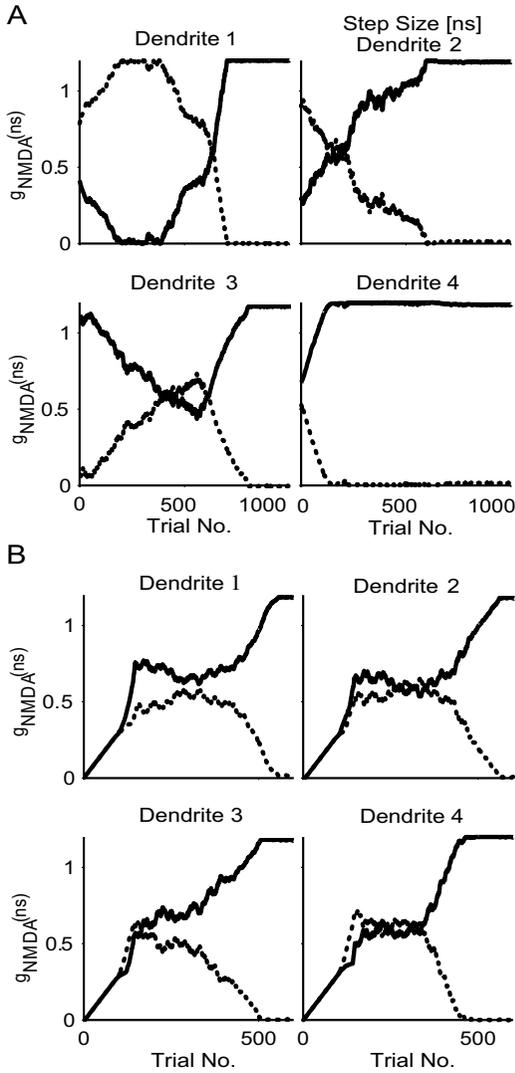


Figure 5: The model with four DS subunits after unsupervised training under the “random start” setting and “development” setting. (A) All four dendrites initially received random excitatory inputs. The total excitatory connection strength at each branch was set to 1.2 nS, and the learning step size was 0.003 nS. After training, the cell responded only to rightward motion. (B) The excitatory input connection strengths to the model were initially zero. Synaptic weight changes for the left and right excitatory inputs. Learning step size was 0.003 nS. After training, the cell responded only to rightward motion.

quirement. Synaptic weight changes for one such simulation are shown in Figure 6. Initially, all excitatory (see Figure 6C) and inhibitory (see Figure 6B) synapses' weights are zero, and so is the average spike number. Because of the "competition" rule, the weights of the excitatory synapses increase gradually, until they exceed the spiking threshold. The learning model is not direction selective because there is no inhibition. The average number of spikes per trial quickly increases over the target value: three spikes per trial. As the inhibitory synapse connection strengths start to increase, so does the competition among excitatory synapses. The inhibitory synapse weights overshoot until it gradually settles down (see Figure 6B) and the model cell stabilizes to become direction selective. We carried out 20 simulations: the model cell converged to  $DI = 0.97$  (on average) after 200 trials in each case.

**3.5 Two-cell Network Learning.** Finally, we investigated if the same learning principle, together with mutual inhibition, could develop tuning for opposite directions of motion between two neurons. We assume that both cortical cells share the same connections from the LGN input array (see Figure 7A). In addition, each cell inhibits the other cell at the soma. This inhibition has a fixed delay of 10 ms, a time course given by the difference of two exponentials with time constants of 0.5 ms and 4 ms, a reversal potential  $E_{GABA} = -70$  mV, and a constant conductance  $g_{GABA} = 5.0$  nS. The associated postsynaptic conductance change rises to its maximum value within the first 2 ms and decays in the following 20 ms. This corresponds to a  $GABA_A$  inhibitory synapse, generally located close to the soma. In order to model inhibitory feedback from nearby cells, an additional inhibitory synapse, identical to the previous one that connects the two-cell network, is added to each soma (see Figure 7A). The time of inhibitory input from nearby regions is randomly selected within the first 20 ms of every trial for each cell in this two-cell network. Due to this stochastic fluctuation, the activities of the two cells are not identical. Cell 1's somatic firing induces inhibition on cell 2. This inhibition prevents cell 2 from firing and thus pushes its direction selectivity more away from cell 1's. Indeed, from a theoretical point of view, the evolution of the network reaches equilibrium when both cells have  $DI = 1$  and are tuned toward opposite directions; mutual inhibition has no effect when the bar moves in cell 1's preferred direction (since cell 2 will not fire, so it will not inhibit cell 1's firing) and guarantees that cell 1 will not fire when the bar moves in cell 1's null direction (since cell 2 will fire in its preferred direction and thus inhibit cell 1). In reality, however, a limited number of simulation trials as well as the model's stochastic effect determines that not all simulation scenarios end in the global sink.

We carried out 60 simulations of this two-cell network learning with a 0.032 nS learning step size for 2000 trials. The weights of all excitatory synapses for both cells were initially set to 0.6 nS. Similar to the definition of the index of direction selectivity,  $DI$ , we define  $RE_{cell\ pair} = |(\text{REF}_{cell1} -$

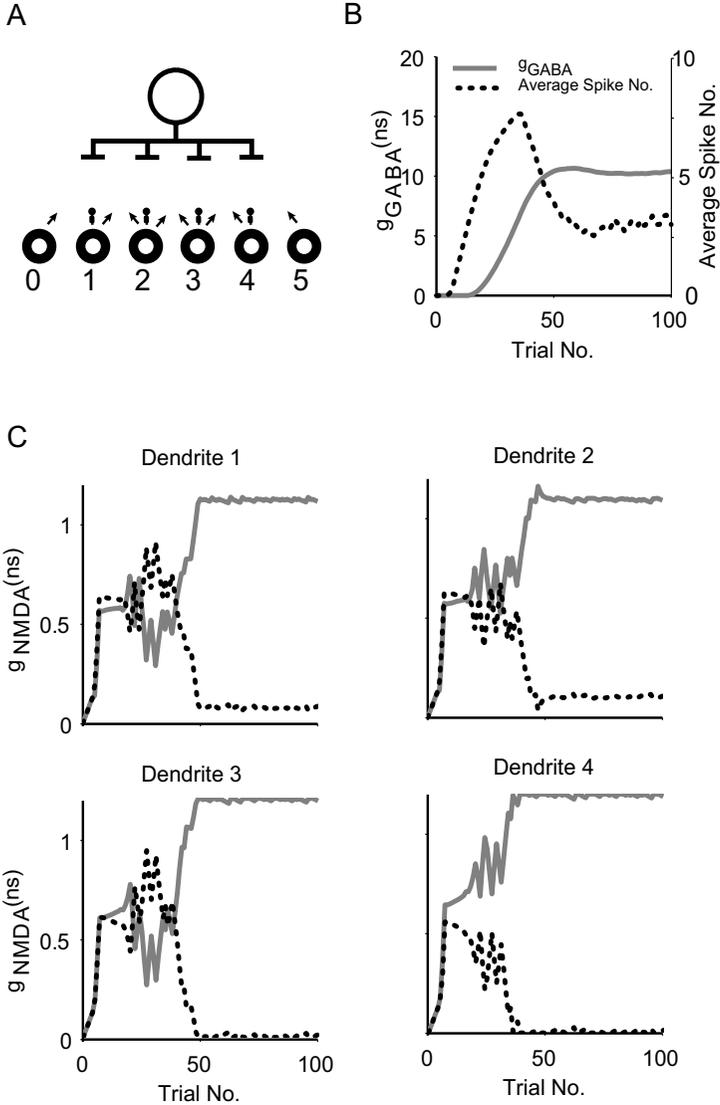


Figure 6: Differential learning of both excitation and inhibition. (A) Schematic drawing of the starting condition. All initial excitatory and inhibitory connection strengths are zero. (B) Synaptic weight changes for an inhibitory input cell (solid curve) and the average spike number of the model cell (dotted curve, averaged over 30 trials) during a simulation run. (C) Synaptic weight changes for the left input cells (solid curves) and the right input cells (dotted curves) at each dendrite during one simulation run. The model cell is direction selective after learning.

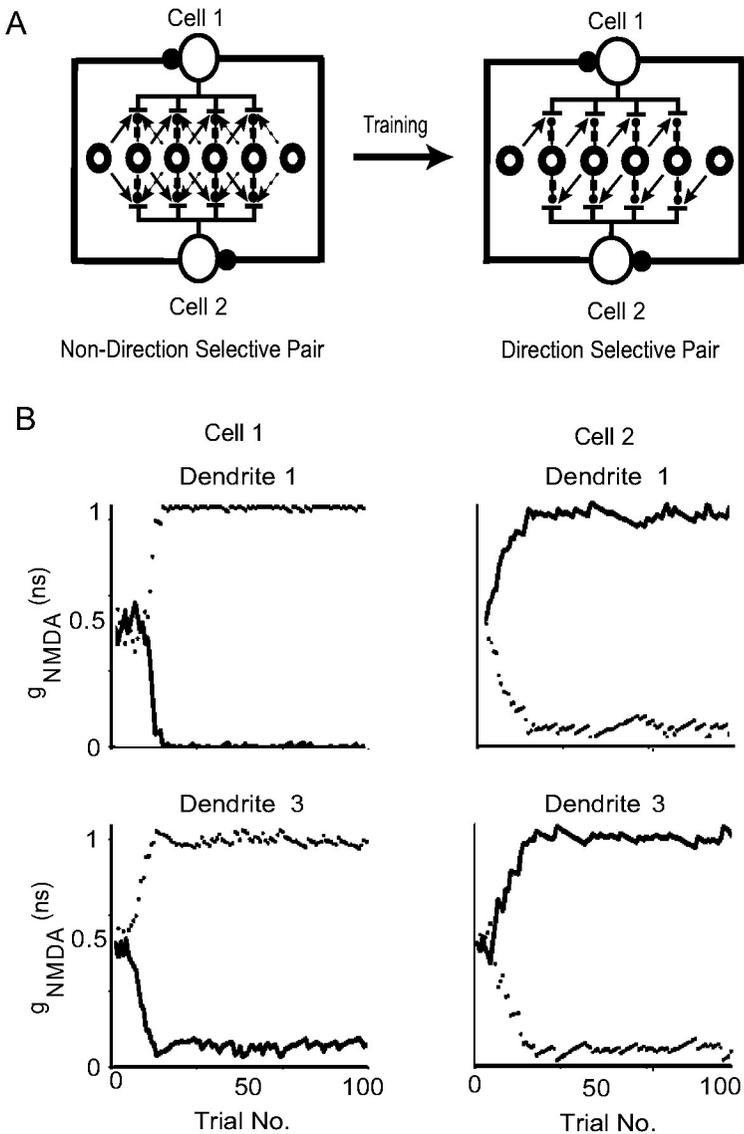


Figure 7: Illustration of the training of a two-cell network. (A) Each neuron receives the same geniculate input. Each cell inhibits the other one and is also the target for common inhibition from other sources of inhibition (not shown). After training (on the right), both cells become selective for motion in the two opposing directions. (B) Synaptic weight changes of two dendrites (out of four) in each neuron. Synaptic weight changes for the other two dendrites have similar trajectories. After training, cell 1 develops tuning for rightward direction of motion; cell 2 develops tuning for leftward direction of motion, with  $DI = 1$ .

$REF_{cell2})/2|$  for each two-cell network ( $REF = DI$  when a cell is tuned more toward the left direction of motion and  $REF = -DI$  when the cell is tuned more toward the right direction of motion) to assess how far apart the two cells' direction selectivities are. In an ideal case, when the two cells develop the maximal extent of selectivity toward opposite directions of motion,  $RE$  for this pair will be 1. For all 60 runs,  $DI$  averages at 0.94, and  $RE$  averages at 0.89. Figure 7B shows the synaptic weight changes for one simulation of such mutually connected two-cell network learning. The two cells individually developed tuning for opposite directions of motion.

#### 4 Discussion

---

Our scheme depends on several testable assumptions. First, learning of DS is broken down into four different scenarios, depending on the absence or presence of postsynaptic inhibition and whether the cell spikes (see Figure 1C). These could be tested in calcium imaging experiments. Shunting inhibition can be mimicked using the dynamic clamp (Chance, Abbott, & Reyes, 2002). This, together with two-photon calcium imaging, can be used to determine if shunting inhibition indeed can direct local synaptic modifications via localized changes in the calcium concentration. Second, our "competition" rule suggests that neurons with direction-selective subunit structures on its dendrites have the ability to control, independently, the total excitatory input connection strengths to each of their major dendrites. Evidence for such a mechanism, albeit operating at the whole cell level, has been provided by Turrigiano, Leslie, Desai, Rutherford, and Nelson (1998). Our rule is different from other rules that depend on initial bias (Rao & Sejnowski, 2000) in that our rule is experience driven. This suggests that cats reared in an environment with little motion in one particular direction—for instance, by having the animals wear LCD goggles—will show a deficit in direction-selective cells tuned for that direction relative to the opposite direction of motion. In our learning model, random motion plays a key role in breaking the balance between the left and right input cells. Symmetry could also be broken by a bias in the initial connection strength between the left and right input.

Our learning scheme relies on a few additional assumptions. Delayed inhibition is of the shunting type, which is crucial to our multiple subunits learning model in achieving the branch-specific veto of excitation and the branch-specific blocking of backpropagating spikes. We compared the single unit and the multiple subunits model with four, six, and eight subunits and shunting inhibition (at  $-60$  mV) to hyperpolarizing inhibition with  $E_{GABA}$  between  $-60$  mV and  $-90$  mV (the cell's resting potential is  $-60$  mV). While single unit learning is not dependent on shunting inhibition, the more subunits a cell has, the more it relies on shunting inhibition (data not shown; Mo, 2003).

The model is quite robust to perturbations in the initial conditions and in the amplitude of the delay between excitation and feedforward inhibition. We reduced this delay from its standard value of 10 ms to 5 and 2.5 ms (while doubling the inhibition rise time from 1 to 2 ms) without any effect on DI. We also varied the velocity of the speeding bar from its fixed value at 10 degrees per second to one of four values (5, 10, 15, and 20 degrees per second) randomly selected for each trial during learning. When a bar moves very fast in the null direction, delayed inhibition acts too slowly to reduce excitation significantly. This reduces direction selectivity. However, in combination, slow velocities served to stabilize the response to higher velocities, with the final DI converging to an average of 0.9 for all velocities.

This model assumes a one-dimensional input space. It will be a challenge to extend it to direction selectivity in two dimensions. In that case, both orientation and direction selectivity will have to be simultaneously considered.

Rapidly rising calcium concentration changes in dendritic spines mediated by action potentials and long, sustained rising of calcium concentration induced by synaptic inputs have been observed in calcium imaging experiments (Sabatini, Oertner, & Svoboda, 2002). The measured calcium decay constant is 12 ms at spines and 15 ms at small dendrites. We used a single decay constant of 15 ms for both calcium sources and assumed instantaneous dendrite-to-spine diffusion. We have no evidence to suggest that a more sophisticated treatment of calcium dynamics will change our conclusion appreciably. Experimental evidence suggests voltage-gated calcium channels exist in spines, while little calcium diffuses between the spine and the dendritic shaft in either direction (Sabatini et al., 2002). Such a scheme is computationally equivalent to our model setting given that we used instantaneous dendrite-to-spine unidirectional calcium diffusion and a single electrical compartment for the spine and the dendritic shaft. We choose the N-type voltage-gated calcium channel to have a voltage-sensitive calcium dynamics different from calcium flowing through NMDA channels. The high-threshold L-type voltage-gated calcium channels should also serve our purpose.

We exploit a calcium gain control mechanism that dynamically shifts the learning curve according to the average, local activity level. The key to the stability of the original BCM learning rule (Bienenstock et al., 1982) is a nonlinear threshold that decreases and increases faster than the average response. Such a sliding threshold control requires that the tuning curve be narrowed for small responses and broadened for large responses. For simplicity, we simulated a linear sliding of the learning curve without changing its shape. This can be implemented in many ways, such as by an increased calcium pump within the spine with respect to time-averaged calcium exposure or adaptation of calcium-dependent enzyme activities. Both the duration and the amplitude of postsynaptic calcium concentration affect long-term potentiation (LTP) and long-term depression (LTD) (Yang, Tong,

& Zucker, 1999). Brief and large postsynaptic calcium concentration changes lead to LTP, while sustained and moderate changes cause LTD. However, brief and moderate calcium concentration change can lead to either LTP or LTD. A calcium gain-control mechanism or sliding learning threshold can explain such phenomena. Sustained calcium concentration elevation may shift the learning curve as well as the LTD-LTP transition threshold toward high calcium concentration and thus increase the probability of LTD formation.

Learning of multiple direction-selective subunits requires a rule whereby the amplitude of the learning step scales with the total number of action potentials generated at the soma during that trial (which, in our model, is roughly proportional to the peak calcium concentration at the synapse; data not shown): a group of synapses on a dendrite have their weights increased more in their preferred direction and decreased less in the null direction, once the soma responds to these synapses with more spikes than to the other group of synapses associated with opposite direction. This creates direct competition between the two groups and thus facilitates convergence. We also experimented with a modified version of this majority rule that is more plausible: here, the scaling is applied only if the synaptic weight increases. The weight reduction does not depend on the number of spikes triggered. Both rules yielded the same results.

Spike-time-dependent plasticity (STDP) is a temporal asymmetry Hebbian learning rule. The synaptic weight change depends on the relative timing of the presynaptic input and the backpropagating spike. Our learning rule is different from the "prediction and sequence learning" mechanism (Montague & Sejnowski, 1994; Montague et al., 1995; Markram et al., 1997). In our learning scenario 1, the excitatory connection is increased if there are backpropagating spikes within a certain period following synaptic activation; in learning scenario 3, the excitatory connection is decreased if there is a spike a few milliseconds before its opening and the backpropagating spike is clamped by inhibition at opening. These fit into the general framework of the STDP with the exception that our learning rule takes not only the temporal sequence between the input and the output into account but also local inhibition (via its effect on calcium). This difference is critical for the learning of branch-specific DS subunits.

Our model depends on action potentials generated at the soma that carry the direction-selective signal propagating back to the excitatory synapses whose weights need to be adjusted. If these synapses were located in remote parts of the distal tree where the action potential may fail to propagate (Golding, Staff, & Spruston, 2002; Mehta, 2004), the learning process would fail to converge. Thus, we would predict that these synapses, which carry the visual signals, are close enough to the soma for the backpropagating action potential to create a reliable signal for guiding the learning process.

In addition to feedforward connections, there are extensive feedback interactions among V1 cells, and these feedback currents are likely to be impor-

tant for sharpening directional tuning (Douglas, Koch, Mahowald, Martin, & Suarez, 1995; Maex & Orban, 1996). Nonlinear shunting inhibition is likely to act during the initial stage of visual cortical processing to set the balance between opponent “on” and “off” responses in different locations of the visual receptive field (Borg-Graham et al., 1988). Our simulation demonstrates that using mutual inhibition, it is possible to generate two direction-selective cells (or, presumably, two tightly coupled groups of such cells) tuned toward opposite directions of motion. It will be interesting to test if mutual excitation among cells selective to the same direction of motion helps to achieve this architecture. In addition, in V1, many columns of neurons that are selective for a given orientation are subdivided into patches preferring opposite directions of motion (Weliky, Bosking, & Fitzpatrick, 1996; Shmuel & Grinvald, 1996). We believe mutual inhibition can also serve as a possible mechanism to organize the distribution of direction-selective cells within an orientation-tuned column. However, the actual mechanisms that the visual system uses at its developmental stage to achieve an overall balanced directional tuning as well as an organized distribution of such tuning within single orientation columns still remain to be tackled experimentally. It will be exciting to see more experimental evidence to support or refute our model.

### Acknowledgments

---

We thank Patrick Wilken for an insightful discussion of the BCM learning rule and Bartlett Mel for providing us with cell models, help, and criticism. Thanks to Ying Gong for a critical reading of the manuscript. This research was supported by grants from the NSF-sponsored Engineering Research Center at Caltech, the National Institutes of Health, and the National Institute of Mental Health.

### References

---

- Abbott, L. F., & Nelson, S. B. (2000). Synaptic plasticity: Taming the beast. *Nature Neuroscience*, 3 (Suppl.), 1178–1183.
- Anderson, J. S., Carandini, M., & Ferster, D. (2000). Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *J. Neurophysiology*, 84, 909–926.
- Archie, K. A., & Mel, B. W. (2000). A model for intradendritic computation of binocular disparity. *Nature Neuroscience*, 3, 54–63.
- Barlow, H. B., & Levick, R. W. (1965). The mechanism of directional selectivity in the rabbit's retina. *J. Physiology*, 173, 477–504.
- Benison, G., Keizer, J., Chalupa, L. M., & Robinson, D. W. (2001). Modeling temporal behavior of postnatal cat retinal ganglion cells. *J. Theor. Biol.*, 210, 187–199.

- Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *J. Neuroscience*, *2*, 32–48.
- Blais, B., Cooper, L. N., & Shouval, H. (2000). Formation of direction selectivity in natural scene environments. *Neural Computation*, *12*, 1057–1066.
- Borg-Graham, L. J., Monier, C., & Frégnac, Y. (1998). Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature*, *393*, 369–373.
- Buchs, N. J., & Senn, W. (2002). Spike-based synaptic plasticity and the emergence of direction selective simple cells: Simulation results. *J. Comput. Neurosci.*, *13*, 167–186.
- Chance, F. S., Abbott, L. F., & Reyes, A. D. (2002). Gain modulation from background synaptic input. *Neuron*, *35*, 773–782.
- Cynader, M., & Chernenko, G. (1976). Abolition of direction selectivity in the visual cortex of the cat. *Science*, *193*, 504–505.
- Destexhe, A., Mainen, Z. F., & Sejnowski, T. J. (1994). Synthesis of models for excitable membranes, synaptic transmission and neuromodulation using a common kinetic formalism. *J. Comput. Neurosci.*, *1*, 195–230.
- Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A. C., & Suarez, H. H. (1995). Recurrent excitation in neocortical circuits. *Science*, *269*, 981–985.
- Emerson, R. C., Citron, M. C., Vaughn, W. J., & Klein, S. A. (1987). Nonlinear directionally selective subunits in complex cells of cat striate cortex. *J. Neurophysiology*, *58*, 33–65.
- Feidler, J. C., Saul, A. B., Murthy, A., & Humphrey, A. L. (1997). Hebbian learning and the development of direction selectivity: The role of geniculate response timings. *Network: Comput. Neural Syst.*, *8*, 195–214.
- Golding, N. L., Staff, N. P., & Spruston, N. (2002). Dendritic spikes as a mechanism for cooperative long-term potentiation. *Nature*, *418*, 326–331.
- Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. New York: Wiley.
- Hines, M. L., & Carnevale, N. T. (1997). The NEURON simulation environment. *Neural Computation*, *9*, 1179–1209.
- Humphrey, A. L., & Saul, A. B. (1998). Strobe rearing reduces direction selectivity in area 17 by altering spatiotemporal receptive-field structure. *J. Neurophysiol.*, *22*, 2945–2955.
- Koch, C. (1999). *Biophysics of computation: Information processing in single neurons*. New York: Oxford University Press.
- Koch, C., & Poggio, T. (1985). The synaptic veto mechanism: Does it underlie direction and orientation selectivity in the visual cortex? In D. Rose & V. Dobson (Eds.), *Models of the visual cortex* (pp. 408–419). New York: Wiley.
- Koch, C., & Poggio, T. (1987). Biophysics of computation: Neurons, synapses, and membranes. In G. M. Edelman, W. E. Gall, & W. M. Cowan (Eds.), *Synaptic function* (pp. 637–697). New York: Wiley.
- Koch, C., Poggio, T., & Torre, V. (1982). Retinal ganglion cells: A functional interpretation of dendritic morphology. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.*, *298*, 227–263.

- Livingstone, M. S. (1998). Mechanisms of direction selectivity in macaque V1. *Neuron*, *20*, 509–526.
- Livingstone, M. S., Pack, C. C., & Born, R. T. (2001). 2-D substructure of MT receptive-fields. *Neuron*, *30*, 781–793
- Maex, R., & Orban, G. A. (1996). Model circuit of spiking neurons generating directional selectivity in simple cells. *J. Neurophysiol.*, *75*, 1515–1545.
- Markram, H., Lubke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, *275*, 213–215.
- Mehta, M. R. (2004). Cooperative LTP can map memory sequences on dendritic branches. *Trends in Neuroscience*, *27*, 69–72.
- Mel, B. W. (2002). Have we been Hebbing down the wrong path? *Neuron*, *34*, 275–288.
- Mo, C. H. (2003). *Synaptic learning rules for local synaptic interaction: Theory and application to direction-selectivity*. Unpublished doctoral dissertation, California Institute of Technology.
- Mo, C. H., & Koch, C. (2003). Modeling reverse-phi motion selective neurons in cortex: Double synaptic veto mechanism. *Neural Computation*, *15*, 735–759.
- Montague, P. R., Dayan, P., Person, C., & Sejnowski, T. J. (1995). Bee foraging in uncertain environments using predictive Hebbian learning. *Nature*, *377*, 725–728.
- Montague, P. R., & Sejnowski, T. J. (1994). The predictive brain: Temporal coincidence and temporal order in synaptic learning mechanisms. *Learn. Mem.*, *1*, 1–33.
- Poggio, T. (1982). *Visual algorithms* (A.I. Memo No. 683). Cambridge, MA: Artificial Intelligence Laboratory, Massachusetts Institute of Technology.
- Poggio, T., & Torre, V. (1978). A new approach to synaptic interactions. In R. Heim & G. Palm (Eds.), *Lecture notes in biomathematics: Theoretical approaches to complex systems*, *21* (pp. 89–115). Berlin: Springer-Verlag.
- Poirazi, P., Brannon, T., & Mel, B. W. (2003). Pyramidal neuron as 2-layer neural network. *Neuron*, *37*, 989–999.
- Rao, R. P. N., & Sejnowski, T. J. (2000). Predictive sequence learning in recurrent neocortical circuits. In S. A. Solla, T. K. Lee, & K. R. Muller (Eds.), *Advances in neural information processing systems*, *12* (pp. 164–170). Cambridge, MA: MIT Press.
- Rao, R. P. N., & Sejnowski, T. J. (2001). Spike-timing-dependent Hebbian plasticity as temporal difference learning. *Neural Computation*, *13*, 2221–2237.
- Sabatini, B. L., Oertner, T. G., & Svoboda, K. (2002). The life cycle of Ca(2+) ions in dendritic spines. *Neuron*, *33*, 439–452.
- Saul, A. B., & Feidler, J. C. (2002). Development of response timing and direction selectivity in cat visual thalamus and cortex. *J. Neurosci.*, *22*, 2945–2955.
- Shmuel, A., & Grinvald, A. (1996). Functional organization for direction of motion and its relationship to orientation maps in cat area 18. *J. Neurosci.*, *16*, 6945–6964.
- Song, S., Miller, K. D., & Abbott, L. F. (2002). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*, *3*, 919–926.

- Soto-Trevino, C., Thoroughman, K. A., Marder, E., & Abbott, L. F. (2001). Activity-dependent modification of inhibitory synapses in models of rhythmic neural networks. *Nature Neuroscience*, *4*, 297–303.
- Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., & Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, *391*, 892–896.
- Weliky, M., Bosking, W. H., & Fitzpatrick, D. (1996). A systematic map of direction preference in primary visual cortex. *Nature*, *379*, 725–728.
- Yang, S. N., Tang, Y. G., & Zucker, R. S. (1999). Selective induction of LTP and LTD by postsynaptic  $[Ca^{2+}]_i$  elevation. *J. Neurophysiology*, *81*, 781–787.
- Zador, A. M., Koch, C., & Brown, T. (1990). Biophysical model of a Hebbian synapse. *PNAS*, *87*, 6718–6722.

---

Received July 22, 2003; accepted May 4, 2004.