INSTITUTE OF PHYSICS PUBLISHING NETWORK: COMPUTATION IN NEURAL SYSTEMS

Network: Comput. Neural Syst. **15** (2004) 179–198 PII: S0954-898X(04)67913-9

Motion detection and prediction through spike-timing dependent plasticity

A P Shon[1], R P N Rao[1] and T J Sejnowski[2,3]

[1] Department of Computer Science and Engineering, University of Washington, Seattle, WA 98195, USA
[2] Computational Neurobiology Laboratory, The Salk Institute, 10010 N Torrey Pines Road, La Jolla, CA 92037, USA
[3] Department of Biology, University of California at San Diego, La Jolla, CA 92037, USA

E-mail: aaron@cs.washington.edu, rao@cs.washington.edu and terry@salk.edu

Received 17 August 2003, accepted for publication 29 April 2004
Published 26 May 2004
Online at stacks.iop.org/Network/15/179
DOI: 10.1088/0954-898X/15/3/002

Abstract

We describe a possible mechanism for the formation of direction- and velocity-selective cells in visual cortex through spike-timing dependent learning. We contrast the case where only feedforward excitation and inhibition signals are provided to visual neurons with the case where both feedforward and feedback signals are provided. In the feedforward-only case, neurons become selective for a broad range of velocities centered around the training velocity. However, we show that direction selectivity in this case is strongly dependent on delayed feedforward inhibition and in contrast to experimental results, becomes dramatically weaker when inhibition is reduced. When feedback connections are introduced, direction selectivity becomes much more robust due to predictive delays encoded in recurrent activity. Direction selectivity persists in the face of decreasing inhibition in a manner similar to experimental findings. The model predicts that direction-selective cells should exhibit anticipatory activity due to recurrent excitation and suggests a pivotal role for spike-timing dependent plasticity in shaping cortical circuits for visual motion detection and prediction.

1. Introduction

In both the neocortex and in sub-cortical structures such as the hippocampus, researchers have observed a striking dependence of synaptic plasticity on the relative order of pre- and post-synaptic spikes: typically, a synapse is strengthened if an input spike arrives a few milliseconds before an output spike; a reversal in the order of spiking causes a decrease in synaptic strength [21, 4]. This phenomenon has been labeled 'spike-timing dependent plasticity' (STDP).

Modeling studies have demonstrated the importance of STDP in temporal sequence learning [1, 24, 25, 31–33] and coincidence detection [11, 41].

In this paper, we investigate how the interaction between STDP and recurrent connections affects the development of motion detection circuits in the visual cortex. Using networks of integrate-and-fire neurons that adapt their connections using STDP, we show how specific types of neural circuits might develop varying degrees of direction selectivity. Our results show how STDP could enable cortical circuits to develop predictive responses to moving inputs through learned patterns of lateral connections.

1.1. Prior work

A classic model for direction selectivity is the Barlow–Levick model [3] which postulates a spatial discrepancy between excitation and delayed inhibition in the receptive field of a direction selective neuron. The basic idea is that motion in one direction (the 'preferred direction') causes the neuron to fire because excitation arrives before the delayed inhibition. Motion in the opposite direction (the 'null direction') recruits the delayed inhibition first, which arrives just in time to counteract any excitation from the other side of the receptive field (cf figure 4 in this paper).

More recent computational models have postulated a variety of mechanisms for direction selectivity in the visual cortex, including probabilistic feedforward synapses [6], short-term synaptic depression [8, 36], a combination of feedforward spike-timing dependent plasticity and synaptic depression [7, 34], rate-based Hebbian learning [43, 10] and specialized connectivity schemes [2, 15, 35, 40, 20, 19, 27, 26].

1.2. Contributions of this paper

An important aspect of direction selectivity that has so far gone uninvestigated is the interaction between STDP, inhibition and recurrent connectivity in the visual cortex. We investigated this question using a series of simulations based on networks of integrate-and-fire neurons with plastic synapses. We found that the modification of peak conductances of excitatory synapses alone allows single cortical neurons to become direction selective as in the Barlow–Levick model. Starting from simple feedforward-only connectivity schemes, we investigated the development of direction selectivity in networks with increasingly complex (and increasingly realistic) connection schemes as summarized below:

- *Experiment 1*. A single model neuron received a mix of feedforward excitation and inhibition from ON/OFF cells in the lateral geniculate nucleus (LGN). The results demonstrate that a difference in time constants for excitatory and inhibitory currents is sufficient to allow the formation of weak direction selectivity, where cell responses are comparable to background activity rates.
- *Experiment 2*. In addition to receiving feedforward ON/OFF inputs, model neurons inhibited each other recurrently. We found that this leads to competition between the neurons, allowing them to partition the input space and code for different directions of motion.
- *Experiment 3*. We tested the effects of adding feedback excitation and inhibition, in addition to feedforward inputs. The results indicate that such a scheme leads to robust direction selectivity. This selectivity is resilient to changes in inhibition strength, a result consistent with experimental observations.
- *Experiment 4*. We investigated the role of recurrent excitatory connections in mediating prediction and delay compensation in the visual pathway. Our model predicts that

recurrent connections modified using STDP can lead to 'predictive coding' in the neocortex, where cells fire slightly before receiving feedforward inputs from the LGN. As a consequence, cortical cells should continue to propagate a wave of predictive activity for a short time even when the original stimulus is taken away.

2. Methods

Our model consists of two subsystems: a lower level model that represents the retina and LGN and a cortical model. The retina-LGN model takes as input moving one-dimensional bars and generates spike trains. The cortical model receives these spike trains as input and produces voltage traces using integrate-and-fire dynamics [16]:

$$C\frac{\mathrm{d}V_{\mathrm{m}}}{\mathrm{d}t} = \sum_{i=0}^{n} g_{\mathrm{syn}}^{i}(t)\left(E_{\mathrm{syn}}^{i} - V_{\mathrm{m}}\right) + \frac{V_{\mathrm{rest}}}{V_{\mathrm{m}}}R \tag{1}$$

where V_{m} is the membrane potential, $g_{\mathrm{syn}}^{i}(t)$ is the conductance of synapse i at time t, E_{syn}^{i} is the reversal potential of synapse i, V_{rest} is the resting potential of the cell, C is the membrane capacitance and R is the membrane resistance. When the membrane voltage meets or exceeds a threshold voltage, that is, when $V_{\mathrm{m}} \geqslant V_{\mathrm{th}}$, an action potential is generated lasting one simulation step. This is followed by resetting the membrane voltage to $V_{\mathrm{reset}} = -50$ mV for an absolute refractory period of $\tau_{\mathrm{ref}} = 5$ ms.

All simulations used an integration rate of 1 ms per simulation step. To quantify the direction selectivity of a neuron, we denoted the direction of motion that elicited the maximum response as 'preferred' and the opposite direction as 'null', and used the following *direction selectivity index* (DSI):

$$\mathrm{DSI} = 1 - \frac{\text{null spikes}}{\text{preferred spikes}}$$

where the numerator and denominator in the fraction refer to the number of spikes fired by the neuron for a bar moving in the null and preferred directions, respectively.

2.1. Retina-LGN model

The preprocessing step of our system models the ON/OFF center-surround filtering mechanisms in the retina and the lateral geniculate nucleus (LGN) [42]. We modeled these mechanisms using a set of spatiotemporal filters intended to model the combined effects of retino-geniculate processing. Inputs for all experiments in this paper used a one-dimensional light bar 10 retinal pixels (0.1 retinal degrees) in width. The bar moved on every millisecond with a constant velocity set to an integer number of pixels, with the constant velocity varying from experiment to experiment. Figure 1(a) shows an example of one light bar moving at a velocity of 1 pixel per millisecond. Images of these moving bars were preprocessed by convolving them (using a Fourier transform) with a spatial filter $\mathcal{F}_{\mathrm{SPATIAL}}$ given by a difference of Gaussian functions (figure 1(b)). The output of the spatial filter was passed through two temporal filters $\mathcal{F}_{\mathrm{ON}}$ and $\mathcal{F}_{\mathrm{OFF}}$, each formed as a difference of Gaussian functions (figure 1(c) and (d)).

The spatiotemporal outputs of the filters described above model the relative firing rates over time for LGN 'ON' and 'OFF' units. These time-varying firing rates were fed as input to a Poisson-rate spike generator, which generated two spike trains (one for the 'ON' units, another for the 'OFF' units). Rectification is performed so that the minimum spike rate is clamped at 0. In addition to these input spikes, uniformly distributed spikes were added to

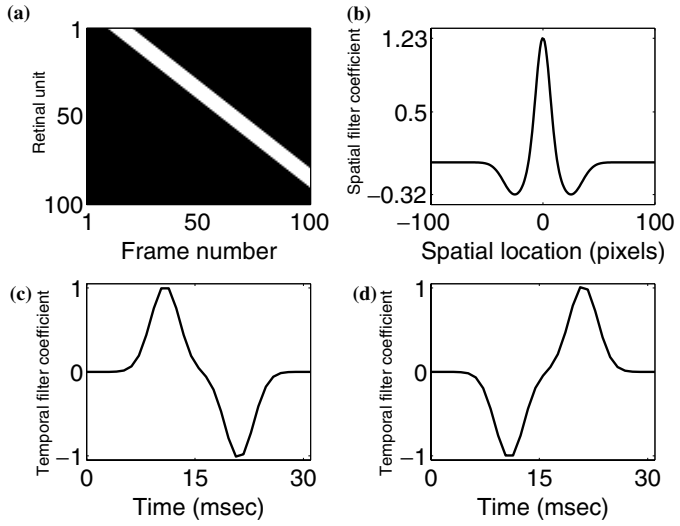

Figure 1. Raw input and LGN filters: (a) example of raw inputs for a light bar moving at a velocity of 1 pixel per ms. Note that the bar shown here is plotted vertically at any given frame number. The bar moves in slowly from the top of the figure. (b) Spatial filter used in the LGN module. The filter is formed as the difference of Gaussian curves. (c) Temporal filter for ON cells in the LGN, formed as the difference of Gaussians. (d) Temporal filter for LGN OFF cells, formed as the difference of Gaussians.

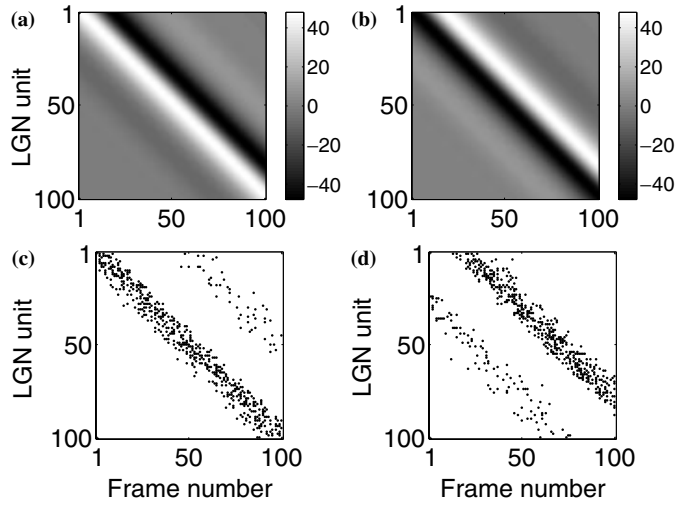

Figure 2. Convolved input traces and spike rasters. (a) Convolved input traces in time domain following LGN FFT processing for ON units. Greyscale values represent (unrectified) Poisson spike generator rates. (b) Traces following LGN FFT processing for OFF units. (c) Spike raster corresponding to the activity of LGN ON cells, produced using a Poisson process. (d) Spike raster corresponding to the activity of LGN OFF cells.

each LGN spike train to model noisy spontaneous activity in the LGN and retina. Figures 2(a) and (b) show the spatiotemporal outputs of the 'ON' and 'OFF' filters to the moving bar

Table 1. Synapse parameters.

Parameter	Value	Units
Excitatory reversal potential (E_{syn})	0	mV
Excitatory peak time $\left(\tau_{\text{peak}}^{\text{exc}}\right)$	10	ms
Excitatory maximum peak conductance $\left(g_{\text{syn}}^{\max}\right)$	0.02 (feedforward, experiments 1, 3, 4); 0.05 (recurrent, experiments 3, 4); 0.006 (experiment 2)	μS
Inhibitory reversal potential $\left(E_{\text{syn}}^{\text{inh}}\right)$	−80	mV
Feedforward inhibitory peak time $\left(\tau_{\text{peak}}^{\text{FFinh}}\right)$	40	ms
Feedforward inhibitory peak conductance $\left(g_{\text{syn}}^{\text{inhFF}}\right)$	0.0018	μS
Feedback inhibitory peak time $\left(\tau_{\text{peak}}^{\text{FBinh}}\right)$	5	ms
Feedback inhibitory peak conductance $\left(g_{\text{syn}}^{\text{inhFB}}\right)$	0.025 (experiment 3); 0.0055 (experiment 4)	μS

stimulus depicted in figure 1(a). Figures 2(c) and (d) show the corresponding 'ON' and 'OFF' spike trains that the cortical model receives as input.

The LGN spatial receptive field size in figure 1 is about 100 pixels. If we assume this corresponds to the size of a typical LGN receptive field (i.e., about $1 \times 1^{\circ}$), we get the relation that one pixel corresponds to approximately 0.01°. The velocities used during training ranged from 1 to 5 retinal pixels per millisecond, which corresponds to velocities in the range $10–50^{\circ}\ \text{s}^{-1}$. This can be considered a reasonable range of velocities to expect in the visual field. Overlap in LGN receptive fields varies depending on the part of the retina to which the receptive fields correspond. In our model, LGN inputs have a high degree of overlap with their neighbors, approximately 98 retinal pixels (0.98° of the visual field). We assume a cortical neuron receptive field size of 2°, giving a retinal eccentricity of approximately $2–5^{\circ}$ from the fovea.

2.2. Synapse model

We model synaptic impulse responses using the alpha function:

$$g(t) = \frac{t \exp(-t/\tau_{\text{peak}})}{\tau_{\text{peak}} \exp(-1)} \tag{2}$$

where τ_{peak} defines the peak time of the alpha function, and both t and τ_{peak} are relative to a spike input to the synapse at $t = 0$. Table 1 lists the values of synaptic parameter settings that were held constant during each experiment.

2.3. Cortical neuron model

We modeled cortical neurons as leaky integrate-and-fire neurons [16]. A second-order Runge–Kutta solver was used to perform integration of the neural membrane voltage. In addition to the input spike trains from the retina-LGN model, each cortical neuron also received exponentially distributed current to model noisy background inputs to the neurons. We chose exponential noise as an approximation to the positive half of a zero-mean Gaussian. Thus, we assume only excitatory noise, with inhibition being more tightly controlled. This is in keeping with simulation [37] and *in vitro* [39] studies that demonstrate a need for precise inhibitory control to maintain acceptable activity levels in recurrent neural networks. Note that our results do not depend on any particular noise model; we have also run several of our experiments using zero-mean and non-zero mean Gaussian noise, with no qualitative effect on the results.

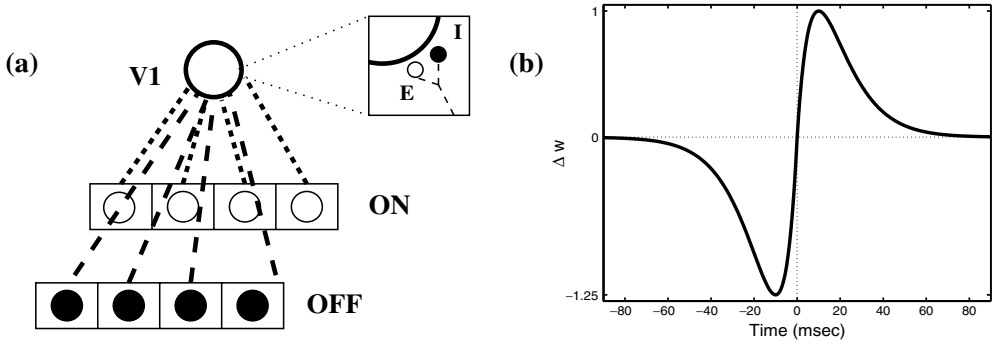

Figure 3. Experiment 1: feedforward architecture and STDP learning window. (a) In the first group of experiments, a single cortical cell receives inputs from ON- and OFF-selective cells from the LGN. Each connection consists of a plastic excitatory connection and a fixed-strength inhibitory connection. (b) STDP learning window used in our experiments. The negative lobe of the window is larger than the positive lobe to facilitate competition between synaptic weights of individual neurons.

Table 2. Cortical neuron parameters.

Parameter	Value	Units
Capacitance (C)	0.5	nF
Resistance (R)	40	MΩ
Resting potential (E_{leak})	−60	mV
Threshold voltage (V_{th})	−40	mV
Reset voltage (V_{reset})	−50	mV
Refractory period (τ_{ref})	5	ms
Exponentially distributed noise magnitude (η)	0.35	nA

Each model cortical neuron received a separate set of feedforward excitatory connections from the LGN 'ON' cells and 'OFF' cells. Each excitatory connection was paired with a fixed-strength feedforward inhibitory connection (see figure 3(a)). In experiments 2, 3 and 4, fixed-strength feedback inhibitory connections were present between all cortical neurons (with no self-connections). In experiments 3 and 4, excitatory feedback connections were also present. Table 2 summarizes neural parameters used in our experiments.

2.4. Learning rule

Learning rules for STDP are typically based on a temporally asymmetric window that determines the sign and amount of synaptic modification as a function of the time difference between pre- and post-synaptic spiking (e.g. [4]). The learning window we used captures the shape and temporal extent of the window observed in physiological experiments and is shown in figure 3(b). Note that in keeping with the observations for firing rate stability noted in [14] (see also [38]), which state that the negative lobe of the synaptic learning kernel should be larger than the positive lobe, our learning rule has a negative lobe 1.25 times the size of the positive lobe. The learning window is multiplied by a learning rate parameter Δ_g to determine the magnitude of synaptic modification for a given time step. In all simulations shown here, Δ_g for all synapses was set to the constant value 10^{-4}.

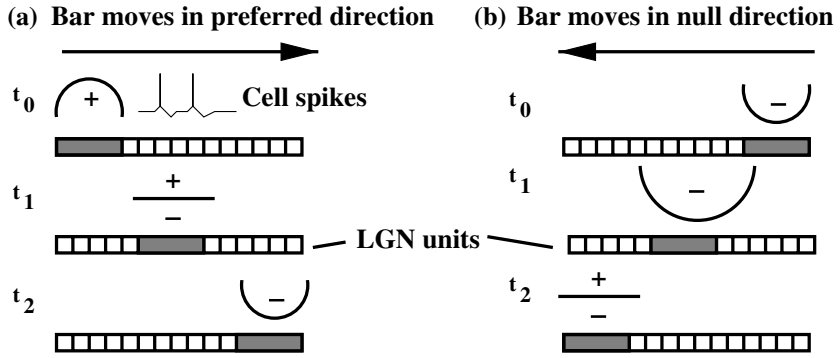

Figure 4. Direction selectivity with feedforward connections relies on interaction between inhibition and weight asymmetry. (a) Spike-timing dependent plasticity induces an asymmetry in synaptic weight values. When a light bar moves in the preferred direction, the bar activates LGN units with high-strength connections at time t_0, *before* the slower inhibition can compensate at time t_2, causing the cell to spike. (b) When the bar moves in the null direction, weak synaptic weights are encountered at time t_0. By time t_2, when the bar reaches the region of strong synaptic weights, slow-acting inhibition has already risen to compensate, preventing any firing from occurring.

3. Results

3.1. Experiment 1: feedforward connections only

A mismatch in the time constants for inhibitory and excitatory synapses, in conjunction with a spike-timing dependent rule, can lead to the development of direction selectivity. Because of the asymmetry in the synaptic learning kernel, repeated exposure to bars moving in the same direction will cause an asymmetry in the excitatory LGN synaptic weights. When a bar moves across the retina in the learned preferred direction, it will first encounter a group of high-valued synaptic weights, causing the cortical neuron to fire. In contrast, when the bar moves in the opposite (null) direction, it will first encounter a group of synapses with low peak conductance. By the time the bar reaches the high-valued weights, feedforward inhibition will have risen sufficiently to prevent the cortical neuron from firing. Figure 4 demonstrates this idea, similar to the direction-selective neural detector first proposed by Barlow and Levick [3]. More recently, Mo and Koch have proposed a model in the context of reverse-phi effects (illusory motion in the presence of rapid reversal in image contrast) that implements a similar idea for feedforward direction selectivity [28]. Their model demonstrates that inhibitory interaction between 'ON' and 'OFF' synapses is necessary to explain the reverse-phi effect in V1 complex cells. As an initial proof of concept, experiment 1 shows how a single neuron receiving only feedforward excitation and inhibition and exponentially distributed noise can learn direction selectivity.

3.1.1. Training paradigm. We trained a single cortical neuron using a light bar that moved at a velocity of five retinal pixels (0.05°) per simulation step. We performed ten training iterations, where each iteration consisted of moving the bar from left to right across the simulated retina and applying spike-timing dependent learning to modify the excitatory connections from the retinal/ LGN system. Each iteration used a different Poisson-generated raster of LGN input spikes, ensuring that the feedforward excitatory weights are biased in general toward forward motion without overtraining for one particular sequence of input spikes. Each iteration ran for 350 ms.

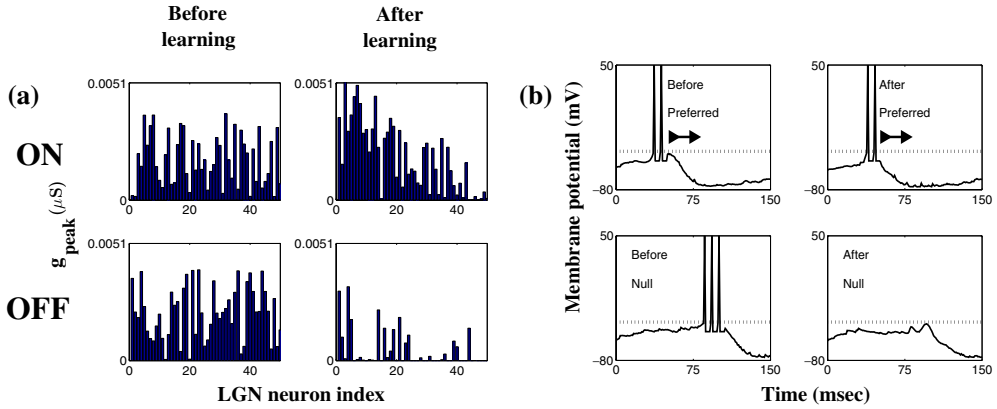

Figure 5. Learned weight asymmetry leads to direction selectivity. (a) These bar graphs show feedforward weights from 50 ON cells and 50 OFF cells onto a single cortical neuron. Before learning (two left graphs), weights are distributed uniformly on the interval [0, 0.004] μS. After learning (two right graphs), the spike-timing dependent learning rule generates weight asymmetry. Most OFF cell inputs to this cortical neuron are depressed, while the shape of the ON cell inputs becomes skewed between a region of strong weights (to the left) and weaker weights (to the right). (b) The weight asymmetry after learning leads to direction selectivity. Before learning, the cortical cell responds vigorously (with three spikes) to movement in the null direction. After learning, the cell only spikes in response to the preferred direction of motion.

3.1.2. Weak direction selectivity. Figure 5 shows how, despite a uniform distribution before training, a marked asymmetry forms in the excitatory feedforward connections as a result of applying STDP over multiple iterations. Figure 7 demonstrates how the neuron responds with two spikes in the preferred direction and three spikes in the null direction before training, and responds only in the preferred direction after training, with no spikes in the null direction. The weight asymmetry prevents response in the null direction and allows response in the preferred direction.

Unfortunately, as figure 7 shows, the response in the preferred direction remains weak; only two spikes are generated after training. Although the asymmetry between feedforward excitation and inhibition can generate direction selectivity, it is clearly inadequate for creating enough spikes to overcome large-scale noise fluctuations. We defer a discussion of robust direction selectivity using recurrent connections to experiment 3 in section 3.3.

Figure 6 shows that the shape of the spike-timing dependent modification window is critical—not all Hebbian learning rules will lead to direction selectivity, even when the training inputs are biased toward one direction of motion. In figure 6(c), we show the same learning window as elsewhere, but with the negative lobe of the window reflected above the x axis. After training, the neuron displays almost no direction selectivity, with 15 spikes in the preferred direction and 13 spikes in the null direction. This result demonstrates that arbitrary causal learning rules do not necessarily lead to direction selectivity. Figure 6(a) shows that weights are still uniformly distributed after training; unlike figure 5(a), no spatial asymmetry in synaptic weights occurs. Note that, even if a weight decay term were included for normalization, the model neuron would still respond equally vigorously to both directions of motion.

3.2. Experiment 2: competitive feedback inhibition

Our previous experiment demonstrated the ability of a single neuron to learn direction selectivity when the neuron is exposed to light bars moving in a single direction. In the

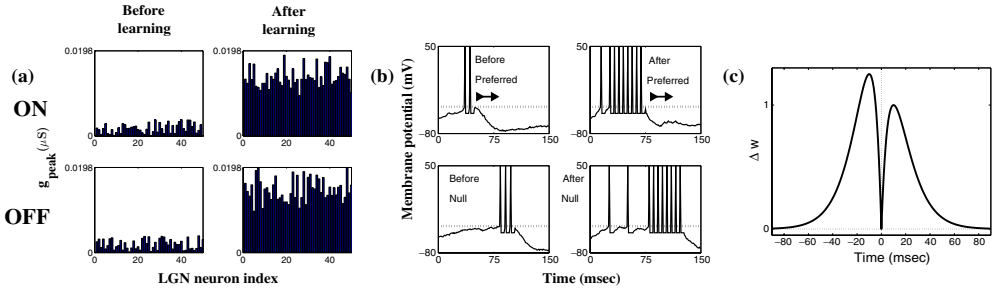

Figure 6. Other Hebbian learning windows do not lead to direction selectivity. (a) These bar graphs show feedforward weights from 50 ON cells and 50 OFF cells onto a single cortical neuron. Before learning (two left graphs), weights are distributed uniformly on the interval [0, 0.004] μS. After learning (two right graphs), a Hebbian learning rule maintains weight symmetry. (b) The weight symmetry after learning prevents direction selectivity. Before learning, the cortical cell responds vigorously (with three spikes) to movement in the null direction. After learning with the Hebbian kernel, the cell spikes even more vigorously in response to both the preferred and the null directions of motion. (c) Non-STDP Hebbian learning window. Here the negative lobe of the STDP window shown in figure 3(b) has been reflected back above the x axis.

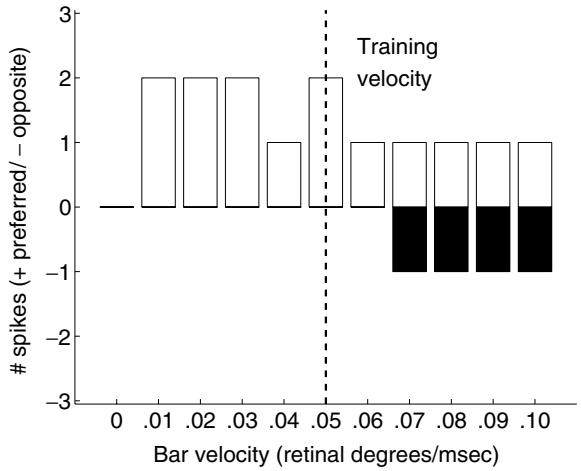

Figure 7. Learning causes direction selectivity across a range of velocities. A single cortical neuron trained using a velocity of five retinal pixels (0.05°) per ms demonstrates direction selectivity across a range of velocities, with a notable decline in selectivity for higher velocities.

next set of simulations, we investigated whether groups of neurons recurrently connected by inhibitory synapses can learn to code for multiple directions of motion simultaneously.

3.2.1. Training paradigm. We trained a group of three cortical neurons on two different directions of motion. The first direction involved a bar moving left to right at a velocity of 1 retinal pixel per millisecond; the second direction involved a bar moving right to left at a velocity of 1 retinal pixel per millisecond. Each pass of the light bar over the retina constituted one iteration of the simulation. We applied the spike-timing dependent learning rule over 20 total iterations, 10 for the left-to-right bar and 10 for the right-to-left bar. Again, each iteration ran for 350 ms. The neurons were all-to-all connected (with no self-connections) using inhibitory synapses with a peak synaptic conductance $g_{\text{peak}}^{\text{inhFB}} = 0.025\,\mu\text{S}$.

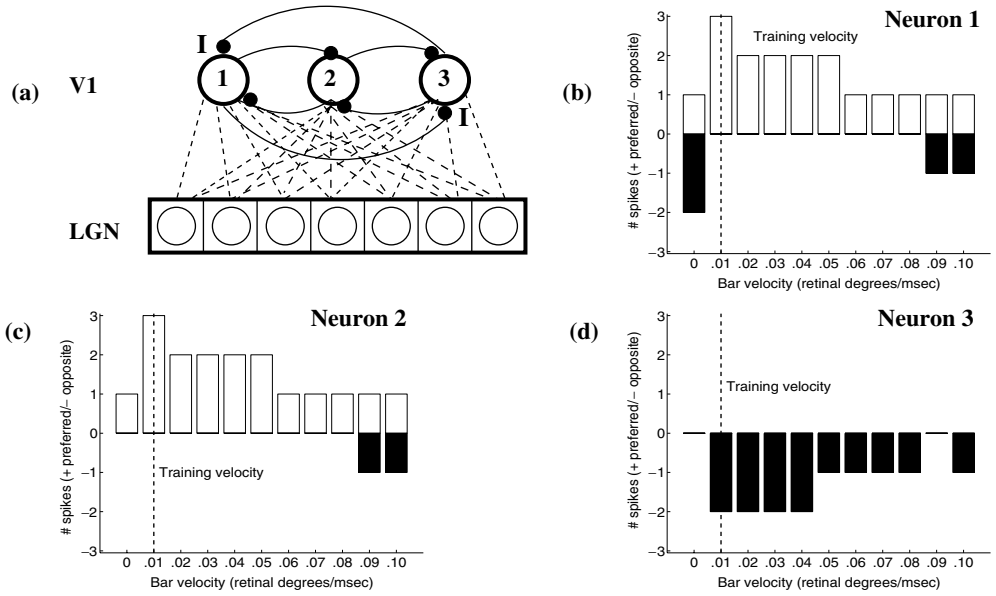

Figure 8. Experiment 2: inhibitory feedback connections permit learning of multiple velocities. (a) In the second group of experiments, a collection of three cortical cells are interconnected using fixed-strength inhibitory synapses. The three neurons are exposed to two different velocities (in this case, +1 and −1, representing two different directions of motion). (b), (c), (d) Neurons 1 and 2 code for the opposite direction of motion as neuron 3.

3.2.2. Partitioning the set of input sequences. Given slight initial biases in the feedforward excitatory weights, some neurons may be expected to respond more vigorously than others to bars moving in a particular direction. These vigorously responding neurons will inhibit the less-responsive neurons. In turn, only the vigorously responding neurons will modify their feedforward excitatory synapses sufficiently to create an asymmetry that codes for motion in one particular direction. Figure 8(a) shows a schematic diagram of this arrangement. Figures 8(b), (c) and (d) demonstrate responses of the three cortical neurons to motion in each direction. Two of the neurons code for motion in the left-to-right direction; the other codes for motion in the right-to-left direction.

Figure 9 shows the resulting asymmetry in feedforward excitatory weights. Note that neurons 1 and 2 display weights that code for the opposite direction of motion as neuron 3.

Figure 10 shows that learning direction selectivity through competitive inhibition scales to larger numbers of neurons. Figure 10(a) shows how a system of 100 neurons learns direction selectivity. Fortyfive of the 100 neurons have a direction selectivity greater than or equal to 0.5 in either the left-to-right or right-to-left direction when tested at the training velocity of 1 retinal pixel per millisecond. Thirty of these direction-selective neurons prefer left-to-right motion; the other 15 prefer the right-to-left direction. Figure 10(b) shows the number of spikes generated in the preferred and opposite directions by each neuron. Most neurons that do become directionally selective exhibit weak responses of only one or two spikes. However, a relatively small number of neurons do fire vigorously in response to bars moving in the preferred direction (with five or more spikes), with no spikes in response to bars moving in the opposite direction. Peak conductance $g_{\text{peak}}^{\text{inhFB}}$ of feedback connections was set to 0.000 81.

Figure 11 demonstrates how the weak direction selectivity developed as a result of STDP in feedforward-only connections drops off as feedforward and feedback inhibitory strengths

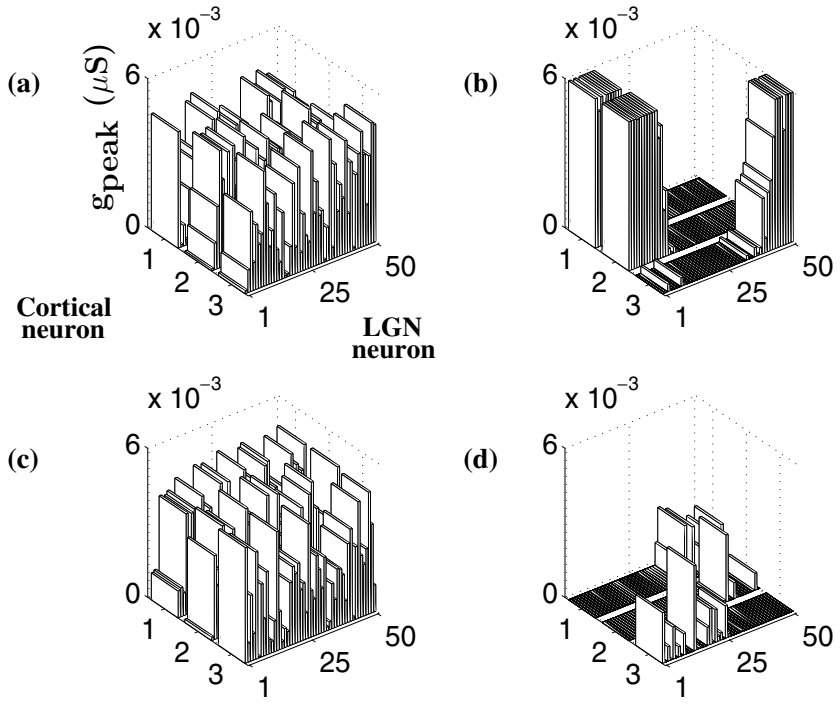

Figure 9. Feedforward weights reflect competition between cortical cells. (a), (c) Before learning, weights are uniformly distributed on the interval $[0, 0.005]$ μS. Cortical neuron indices in this graph match the neuron index numbers given in figure 8, i.e. cortical neuron 1 here matches the neuron whose responses are given in figure 8(b), etc. (b), (d) After learning, competition between cortical neurons has caused an asymmetry in the weights of neurons 1 and 2 that causes spiking in one direction of motion, while neuron 3's weights respond to motion in the opposite direction.

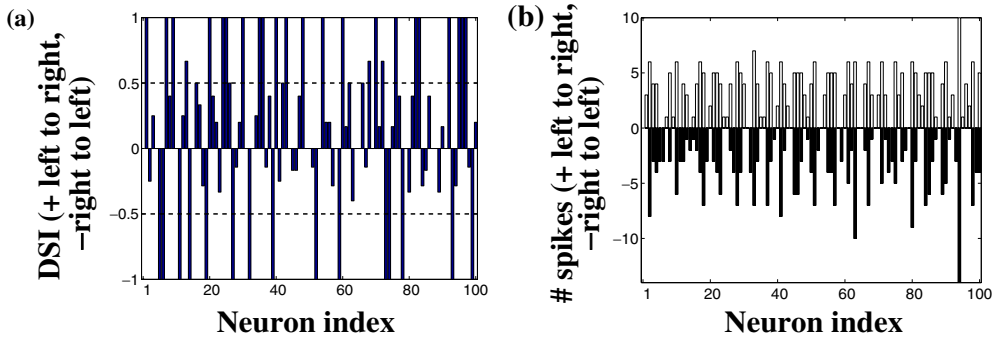

Figure 10. Neural competition scales to larger networks. (a) Direction selectivity index plotted over a network of 100 model cortical neurons, each interacting through inhibitory feedback connections. Horizontal bars denote direction selectivities of 0.5 in the left-to-right (+0.5) and right-to-left (-0.5) directions. (b) Spike counts taken from a single testing simulation run of the 100-neuron network. Most neurons fired on the order of 2–5 spikes, with a few outliers firing up to 14 spikes.

are reduced from 100% of training inhibition down to 0% in 20% decrements. This result contrasts with biological findings that complex cells maintain direction selectivity even when inhibition is greatly reduced. This leads us to conclude that although competitive inhibition

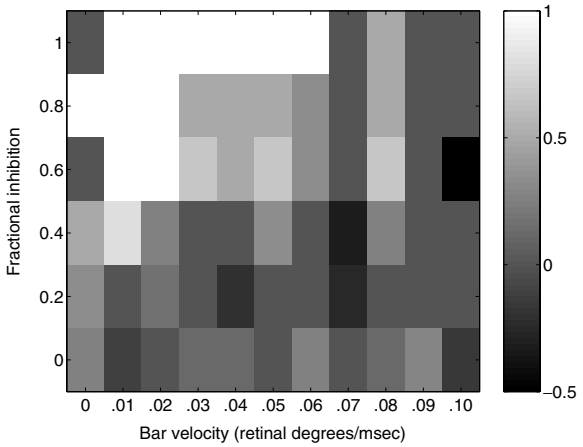

Figure 11. Feedforward-only learning causes dropoff in direction selectivity as inhibition is decreased. A single neuron is trained without feedback inhibition or feedback excitation on sequences with velocities of five retinal pixels (0.05°) per ms. After training, direction selectivity is tested over a range of velocities on the range 1–10. For each velocity, direction selectivity is measured when feedforward inhibition is set at 100% of training inhibition, 80% ,..., 0%. Selectivity drops off markedly for higher velocities than the training velocity and for reduced inhibition. Colours (see bar at right) represent direction-selectivity index values of the model cortical cell (cf section 2), from light (strong preference for left-to-right motion) to dark (weak preference for right-to-left motion).

can allow cortical neurons to code for direction selectivity, additional, excitatory recurrent synapses (cf section 3.3) are necessary to cause robust learning of direction selectivity.

3.3. Experiment 3: feedback excitatory and inhibitory connections

Our previous experiments showed that the mismatch between feedforward excitation and inhibition time constants is sufficient to create an asymmetry in feedforward excitatory weights, causing direction selectivity, and that mutually inhibiting groups of cortical neurons can compete to code for stimuli moving in different directions. However, since experiments 1 and 2 only employed feedforward connections, the direction selectivity developed by the cortical neurons was relatively weak. Intuitively, having recurrent excitatory connections biased in the learned direction of motion should help the cortical neurons code much more strongly for stimuli moving in that direction. In experiment 3, we investigated the effects of adding recurrent excitatory synapses that are modified by STDP, along with weak non-plastic recurrent inhibitory synapses.

3.3.1. Training paradigm. Our simulated cortical network for this experiment consists of a chain of 11 integrate-and-fire neurons connected all-to-all (no self-connections) with excitatory and inhibitory synapses (figure 12). The recurrent inhibitory synapses are initialized to constant fixed values of $g_{\text{peak}}^{\text{inhFB}} = 0.0055\,\mu\text{S}$, and the recurrent excitatory synapses are initialized to random, uniformly distributed values from 0 to 0.005 μS. We assume that each neuron receives input from a patch of the retinal/LGN system that does not overlap with the receptive field of any other cortical cell in the chain. Further, the feedforward excitatory weights were set to the values learned in the previous experiments, so an asymmetry biasing the network toward weak responses in the preferred direction already exists. The recurrent excitatory weights

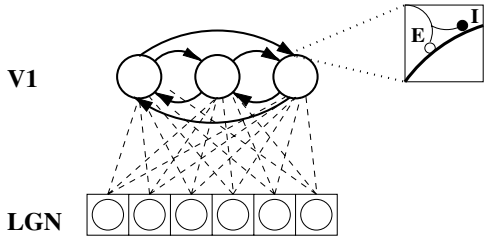

Figure 12. Experiment 3: Network with recurrent excitation and inhibition. The third set of experiments covers the case where feedback excitation and inhibition and feedforward excitation and inhibition are present in the network.

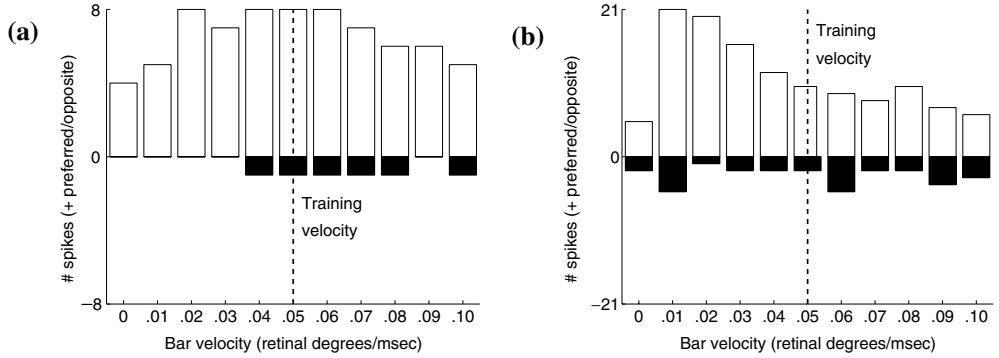

Figure 13. Recurrent connections allow robust direction selectivity. In contrast to the single neuron shown in figure 7, responses in the preferred direction are represented by numerous spikes. Responses in the null direction remain weak. (a) Responses with inhibition at 100% of initial; (b) responses when inhibition is set at 60% of initial. Note the difference in scale of the y axes for the two subgraphs.

were modified according to the STDP rule as the network was exposed to a moving bar 10 pixels (0.1°) wide moving at a velocity of five pixels (0.05°) per millisecond. Again, each pass of the light bar over the retina constituted one iteration of the simulation, and each iteration ran for 725 ms. Ten iterations comprised the total training set for the network.

3.3.2. Strong direction selectivity. Figures 13 and 14 show that STDP causes an asymmetry in the excitatory recurrent connections that leads to robust direction selectivity. Figure 13(a) shows the number of spikes fired by the neuron in the middle of the chain when presented with bars moving at velocities from $0, \ldots, 10$. Compared to the results for feedforward-only excitation in figure 7, the neuron displays much more vigorous activity in the preferred direction, while firing either 0 or 1 spikes in the null direction for all velocities except 0. Strong direction selectivity persists until inhibition is lowered to 20% to 0% of normal (see figure 15). Note the higher velocities in particular display much more robust direction selectivity as compared to the feedforward-only case. Figure 14 shows how STDP causes an initially uniform distribution of recurrent excitatory weights to develop a marked asymmetry after learning. Our findings are consistent with other modeling studies, for example, the work of Suarez *et al* [40], who found that asymmetric, excitatory recurrent connections are necessary to replicate biological data and ensure robust direction selectivity.

We also ran our simulations on a larger recurrently connected network of 100 model cortical neurons (figure 16). Here, each cortical neuron received 50 LGN inputs, and LGN

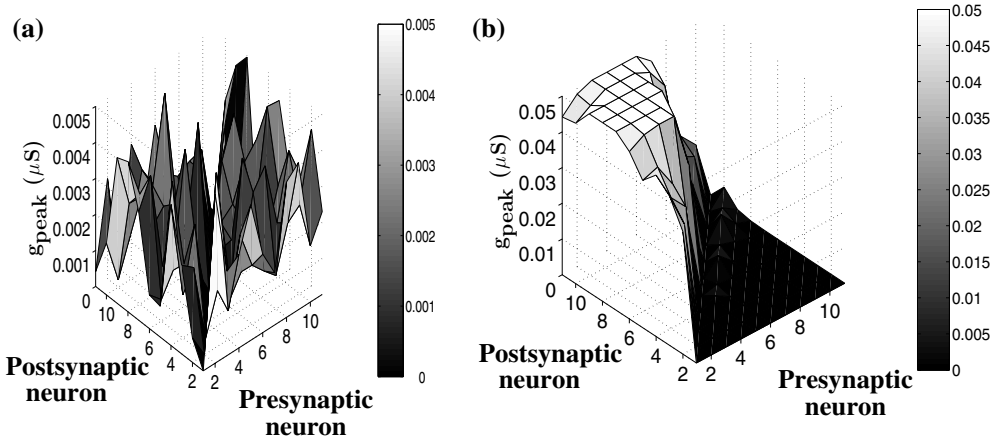

Figure 14. Recurrent weights show asymmetry after training. (a) Recurrent synaptic weights before training. Colour bar shows relative values of recurrent synaptic weights, from light (high peak conductance) to dark (low peak conductance). (b) Weights from the same network after training. A clear asymmetry results from being trained on bars moving in the preferred direction. Note the difference in scale from (a).

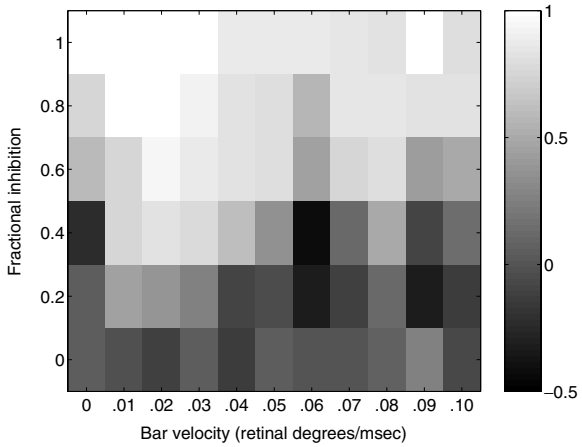

Figure 15. Learning with recurrent excitatory feedback causes robust direction selectivity even when inhibition is dropped. A mix of inhibitory and excitatory connections on both the feedforward and the feedback synapses is trained on a moving bar. The resulting network displays robust direction selectivity across a range of velocities, even as feedforward inhibition is decreased. As in figure 11, greyscale values represent direction-selectivity index values (cf section 2).

receptive fields for adjacent cortical cells overlapped (each cortical neuron's receptive field shared 45 of 50 LGN neurons with each neighbor). We adjusted the maximum peak synaptic conductance $g_{\mathrm{syn}}^{\mathrm{max}}$ to 0.1; all other parameter settings were identical to the 11-neuron network as described in section 2. Figure 16(a) demonstrates the robustness of direction selectivity as feedforward inhibition is lowered from 100% of initial to 0%. The selectivity is measured from cortical neuron 45 of 100. Many neurons in the middle of the array responded weakly compared to the middle neuron in the 11-neuron network; this is likely due to the interaction of the overlapping receptive fields and recurrent inhibition, causing activity to 'jump' between distant recurrent neurons along the chain while suppressing responses from intervening neurons. This

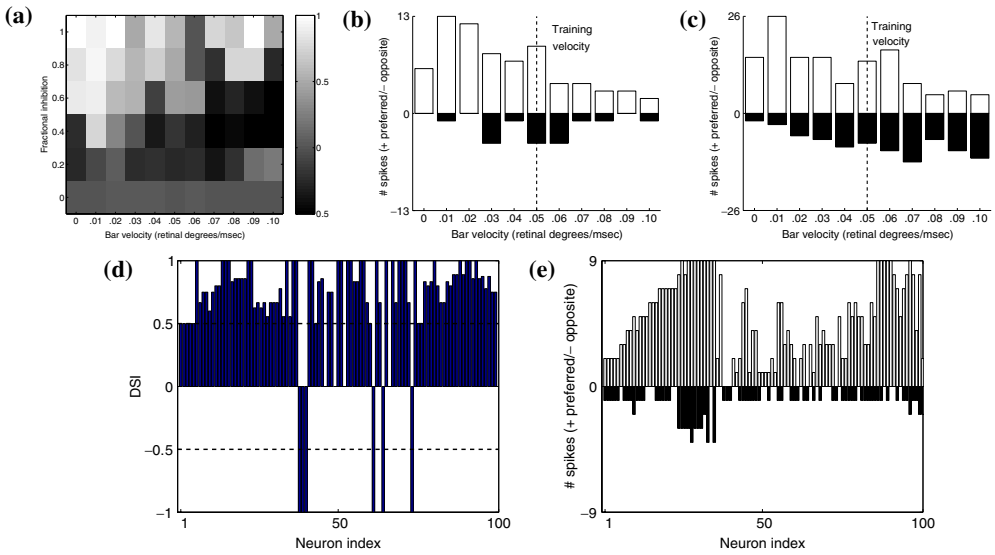

Figure 16. Direction selectivity scales to large recurrent networks with overlapping fields. (a) A network of 100 neurons, with overlapping receptive fields (see text), was trained on a moving bar. The resulting network displays robust direction selectivity across a range of velocities, even as feedforward inhibition is decreased. As in figure 11, greyscale values represent direction-selectivity index values (cf section 2). (b) Number of spikes for preferred and null directions, plotted for neuron number 45 from the network, plotted for 100% initial inhibition. (c) Number of spikes for neuron 45 at 60% initial inhibition. (d) Direction selectivity of all 100 neurons for the training velocity (5 retinal pixels ms^{-1}, that is, 0.05° ms^{-1}), plotted for 100% initial inhibition. (e) Spike counts in preferred and null directions for all 100 neurons at the training velocity (5 retinal pixels ms^{-1}, that is, 0.05° ms^{-1}), plotted for 100% initial inhibition.

effect is notable in figure 16(e), where many neurons at each end of the chain fire a large number of spikes, while only a few in the middle spike. Different sizes of receptive field overlap and settings for inhibition would lead to different levels of activity for neurons in the middle of the chain.

Figures 16(b) and (c) demonstrate the number of spikes fired in each direction of motion for neuron 45, with feedforward inhibition respectively set at 100% and 60% of initial. The neuron spikes vigorously for all velocities, and continues to display direction selectivity except for very high velocities with lowered inhibition. Figure 16(d) displays the DSI values (positive values indicating a preference for the training direction of left to right, negative values indicating a preference for right to left), with inhibition at 100% of initial, for the training velocity of five pixels (0.05 retinal degrees) per simulation step. Figure 16(e) displays the corresponding spike counts in each direction.

4. Model predictions

Our model of STDP-driven selectivity for direction and motion makes several experimentally testable predictions. In particular, it is known that STDP allows a network of neurons to predictively encode sequences [5, 1, 31–33, 22–24, 14]. We investigated the implications of these findings within the context of motion detection and direction selectivity.

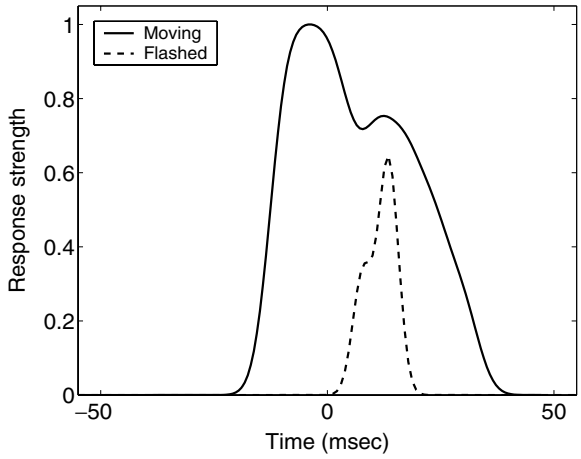

Figure 17. Recurrent connections cause predictive waves of activity. A moving stimulus (solid line) causes recurrent connections to fire predictively, allowing cells to spike before the stimulus reaches their receptive fields. In contrast, when a flashed stimulus is provided (dashed line), the model predicts a peak in activity only after the flash has occurred.

4.1. Experimental paradigm

We used the 11-neuron cortical network described in section 3.3 (after training the recurrent excitatory weights) to determine whether moving bars could generate predictive activity. In this experiment, we assume that synaptic weights have stabilized to represent a preferred direction of motion, and therefore turn off synaptic plasticity while running our simulations. We present two experimental setups:

- *Predictive firing.* In this setup, we move a bar across the retina-LGN system, then examine the activity of the five middle neurons to determine whether the onset of activity precedes the appearance of the bar in the receptive field of the neurons. The bar is considered to impinge on the receptive field of the neurons as soon as the rightmost edge of the bar encounters the leftmost retina-LGN unit corresponding to the leftmost cortical neuron.
- *Continuing predictive activity.* In this setup, we examined the dynamics of model neuron responses in the network when a moving bar was abruptly stopped after an initial period of motion.

Preliminary physiological results that are suggestive of our simulation results have been reported by Jancke *et al* [13].

4.2. Predictive firing

Figure 17 shows the results from the first experimental setup. We contrast the results when the input is a moving bar with the results when the input is a single flashed bar. To generate the curves shown here, we convolve the mean activity of the middle 3 neurons in the case of moving stimuli with a Gaussian with mean 0 and standard deviation 0.1, and in the case of flashed stimuli with mean 0 and standard deviation 0.05.

Our model predicts a sharp onset of activity for the flashed stimulus, and that activity for the moving stimulus should not be as sharply peaked and begin a few milliseconds (approximately 20 ms in the model) before the arrival of the moving stimulus on the receptive

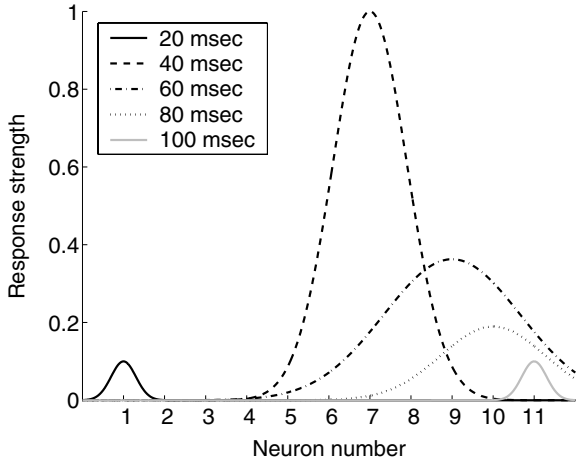

Figure 18. Predictive activity continues to propagate when the moving bar is stopped. After 20 ms of exposing the network to a moving bar, we turn off all inputs from the LGN. Recurrent connections continue to propagate the activity even in the absence of external input. In a larger cortical model, this activity would gradually be reduced due to recurrent inhibition. Here, the small number of simulated recurrent neurons and asymmetry in the weights acts to reduce activity as time passes.

field of the leftmost neuron. Further, our model predicts that no such predictive firing will occur in direction-selective cells when exposed to a bar moving in the null direction.

4.3. Continuing predictive activity

Figure 18 shows the results from the second experimental setup. We begin with a bar moving at a velocity of five retinal pixels (0.05 retinal degrees) per millisecond. After 20 ms, input from the LGN stops completely. At the point where input from the LGN stops, the trailing edge of the moving bar has only reached the end of the receptive field for model cortical neuron 2 (with the bar itself also impinging on the receptive field for model cortical neuron 3). Our model predicts that, even in the absence of continued LGN input, a propagating wave of activity should continue for some time as a result of recurrent excitation. Figure 18 shows this effect in the model network; the figure plots the mean location of activity within the network of 11 cortical neurons over time. Model cortical activities were measured every 20 ms, and plotted as Gaussians whose means are located at the mean locus of cortical activity and whose standard deviations are given by the standard deviation of model cortical activity. Successive bumps in activity continue from left to right along the neural chain, propagating because of a learned left-to-right asymmetry in the recurrent weights. Feedback inhibition prevents cells 3 and 4 from responding strongly. In long chains of recurrently connected neurons, continued predictive firing will eventually cease as recurrent inhibition overcomes excitation; in our small simulated network, activity stops as the wave of firing reaches the end of the chain.

5. Conclusions

We have shown how STDP causes different configurations of model cortical neurons to learn to detect motion in a particular direction. We demonstrated four main results:

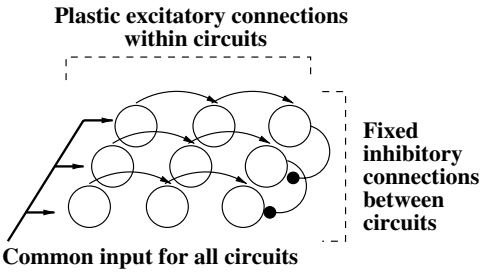

Figure 19. Proposal for cortical architecture. We propose that cortical columns act to partition the set of input sequences, with nearby columns competitively inhibiting one another and excitatory intracolumn connections coding for temporal correlations. This setup is reminiscent of the expectation maximization (EM) approach to finding a model for probabilistic input data, and of Kohonen's SOM model.

(i) STDP allows single neurons receiving feedforward excitatory and inhibitory connections to develop weak direction selectivity.
(ii) A network of mutually competitive, inhibitory neurons can learn to code for multiple different directions.
(iii) A network with both feedforward and STDP-driven recurrent excitatory connections develops robust direction selectivity.
(iv) Recurrently connected networks of direction-selective neurons can predictively encode the direction of stimulus motion and fire in anticipation of feedforward inputs.

Our model predicts that recordings from visual cortical columns should reveal moving waves of activity when the retina is exposed to moving bars, and that the activity should begin slightly before the bar reaches the receptive field of the cortical neuron (i.e. it should be 'predictive'). Waves of activity should continue for a short while in cortex even after the stimulus is removed. These trends may represent a general coding strategy used throughout the neocortex: chains of excitatory neurons provide a top–down prediction of how a stimulus will evolve with time, biasing the activity of lower level sensory neurons, an idea consistent with recent predictive coding models [29, 30].

Our results suggest the following model for the development of cortical direction selectivity and sequence learning in general: chains of recurrently connected neurons learning to code for a particular direction of motion interact competitively with other chains through recurrent inhibition (see figure 19). For any given sequence of inputs, one recurrent chain may 'win out' over its neighbors to code for a particular temporal sequence, spiking sufficiently to prevent neurons in other chains from spiking significantly. The 'winning' chain of neurons would modify their recurrent excitatory synapses according to the STDP learning rule, so that the winning chain is more likely to respond vigorously to the temporal input sequence in the future. The mutually competitive interaction of chains of recurrently connected neurons, interspersed with STDP-driven learning, is strongly reminiscent of the well-known expectation maximization (EM) algorithm [9, 12] from statistics and machine learning (and also of Kohonen's self-organizing maps (SOMs) [17, 18]), raising the intriguing possibility that the neocortex utilizes statistically motivated principles for learning temporal sequences.

Acknowledgments

This work is being supported by the Sloan Foundation, an NSF award from the BITS program and a National Defense Science and Engineering Graduate Fellowship to APS. We thank the

anonymous reviewers for their helpful comments, and Dirk Jancke and colleagues for alerting us to their unpublished results that inspired figures 17 and 18.

References

[1] Abbott L F and Blum K I 1996 Functional significance of long-term potentiation for sequence learning and prediction *Cereb. Cortex* **6** 406–16
[2] Adelson E H and Bergen J 1985 Spatiotemporal energy models for the perception of motion *J. Opt. Soc. Am.* A **2** 284–99
[3] Barlow H and Levick W 1965 The mechanism of directionally selective units in rabbit's retina *J. Phsyiol.* **178** 477–504
[4] Bi G and Poo M 1998 Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type *J. Neurosci.* **18** 10464–72
[5] Blum K I and Abbott L F 1996 A model of spatial map formation in the hippocampus of the rat *Neural Comput.* **8** 85–93
[6] Buchs N and Senn W 2001 Learning direction selectivity through spike-timing dependent modification of neurotransmitter release probability *Neurocomputing* **38–40** 121–7
[7] Buchs N and Senn W 2002 Spike-based synaptic plasticity and the emergence of direction selective simple cells: simulation results *J. Comput. Neurosci.* **13** 167–86
[8] Chance F, Nelson S and Abbott L 1998 Synaptic depression and the temporal response characteristics of V1 simple cells *J. Neurosci.* **18** 4785–99
[9] Dempster A, Laird N and Rubin D 1977 Maximum likelihood from incomplete data via the EM algorithm *J. R. Stat. Soc. Ser.* B **39** 1–38
[10] Feidler J C, Saul A B, Murthy A and 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
[11] Gerstner W, Kempter R, van Hemmen J L and Wagner H 1996 A neuronal learning rule for sub-millisecond temporal coding *Nature* **383** 76–81
[12] Hartley H 1958 Maximum likelihood estimation from incomplete data *Biometrics* **14** 174–94
[13] Jancke D, Erlhagen W, Schöner G and Dinse H 2004 Shorter latencies for motion trajectories than for flashes in population responses of cat primary visual cortex *J. Physiol.* **556** 971–82
[14] Kempter R, Gerstner W and van Hemmen J L 2001 Intrinsic stabilization of output rates by spike-time dependent Hebbian learning *Neural Comput.* **13** 2709–42
[15] Koch C and Poggio T 1985 The synaptic veto mechanism: does it underlie direction and orientation selectivity in the visual cortex? *Models of the Visual Cortex* ed D Rose and V Dobson (New York: Wiley) pp 408–19
[16] Koch C 1999 *Biophysics of Computation: Information Processing in Single Neurons* (New York: Oxford University Press)
[17] Kohonen T 1982 Self-organized formation of topologically correct feature maps *Biol. Cybern.* **43** 59–69
[18] Kohonen T 2000 *Self-Organizing Maps* 3rd edn (Berlin: Springer)
[19] Livingstone M 1998 Mechanisms of direction selectivity in macaque V1 *Neuron* **20** 509–26
[20] Maex R and Orban G A 1996 Model circuit of spiking neurons generating directional selectivity in simple cells *J. Neurophysiol.* **75** 1515–45
[21] Markram H, Lübke J, Frotscher M and Sakmann B 1997 Regulation of synaptic efficacy by coindence of postsynaptic aps and epsps *Science* **275** 213–5
[22] Mehta M R, Barnes C A and McNaughton B L 1997 Experience-dependent, asymmetric expansion of hippocampal place fields *Proc. Natl Acad. Sci. USA* **94** 8918–21
[23] Mehta M R, Quirk M C and Wilson M A 2000 Experience-dependent, asymmetric shape of hippocampal receptive fields *Neuron* **25** 707–15
[24] Mehta M R and Wilson M 2000 From hippocampus to V1: effect of LTP on spatiotemporal dynamics of receptive fields *Computational Neuroscience, Trends in Research 1999* ed J Bower (Amsterdam: Elsevier)
[25] Minai A A and Levy W 1993 Sequence learning in a single trial *Proc. 1993 INNS World Congress on Neural Networks II* (New Jersey: Erlbaum) pp 505–8
[26] Mineiro P and Zipser D 1998 Analysis of direction selectivity arising from recurrent cortical interactions *Neural Comput.* **10** 353–71
[27] Miyashita M, Kim D-S and Tanaka S 1997 Cortical direction selectivity without directional experience *Neuroreport* **8** 1187–91
[28] Mo C-H and Koch C 2003 Modeling reverse-phi motion-selective neurons in cortex: double synaptic-veto mechanism *Neural Comput.* **15** 735–59

[29] Rao R P N and Ballard D H 1997 Dynamic model of visual recognition predicts neural response properties in the visual cortex *Neural Comput*. **9** 721–63
[30] Rao R P N and Ballard D H 1999 Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive field effects *Nat. Neurosci*. **2** 79–87
[31] Rao R P N and Sejnowski T J 2000 Predictive sequence learning in recurrent neocortical circuits *Advances in Neural Information Processing Systems 12* (Cambridge, MA: MIT Press) pp 164–70
[32] Rao R P N and Sejnowski T J 2001 Spike-timing dependent Hebbian plasticity as temporal difference learning *Neural Comput*. **13** 2221–37
[33] Rao R P N and Sejnowski T J 2003 Self-organizing neural systems based on predictive learning *Phil. Trans. R. Soc. Lond*. A 361
[34] Senn W and Buchs N 2003 Spike-based synaptic plasticity and the emergence of direction selective simple cells: mathematical analysis *J. Comput. Neurosci*. **14** 119–38
[35] Shapley R M, Reid R C and Soodak R 1992 Spatiotemporal receptive fields and direction selectivity *Computational Models of Visual Processing* ed M S Landy and J A Movshon (Cambridge, MA: MIT Press) pp 109–18
[36] Shon A P and Rao R P N 2003 Learning temporal patterns by redistribution of synaptic efficacy *Computational Neuroscience, Trends in Research 2002* vol 52-54C, ed J Bower (Amsterdam: Elsevier) pp 13–8
[37] Smith A C, Wu X and Levy W B 2000 Controlling activity fluctuations in large, sparsely connected random networks *Neuron: Comput. Neural Syst*. **11** 63–81
[38] Song S, Miller K D and Abbott L F 2000 Competitive Hebbian learning through spike-timing dependent synaptic plasticity *Nat. Neurosci*. **3** 919–26
[39] Steele P M and Mauk M D 1999 Inhibitory control of LTP and LTD: stability of synapse strength *J. Neurophysiol*. **81** 1559–66
[40] Suarez H, Koch C and Douglas R 1995 Modeling direction selectivity of simple cells in striate visual cortex within the framework of the canonical microcircuit *J. Neurosci*. **15** 6700–19
[41] Suri R E and Sejnowski T J 2002 Spike propagation synchronized by temporally asymmetric Hebbian learning *Biol. Cybern*. **87** 440–5
[42] Wandell B 1995 *Foundations of Vision* (Sunderland, MA: Sinauer)
[43] Wimbauer S, Wenisch O, Miller K and van Hemmen J 1997 Development of spatiotemporal receptive fields of simple cells: I. Model formulation *Biol. Cybern*. **77** 453–61