

Motion Detection and Prediction through Spike-Timing Dependent Plasticity

A. P. Shon[†], R. P. N. Rao[†] and T. J. Sejnowski^{‡§}

[†]Department of Computer Science and Engineering
University of Washington
Seattle, WA 98195

[‡]Computational Neurobiology Laboratory
The Salk Institute
10010 N. Torrey Pines Road

La Jolla, CA 92037

[§]Department of Biology
University of California at San Diego
La Jolla, CA 92037

E-mail:

`aaron@cs.washington.edu, rao@cs.washington.edu, terry@salk.edu`

Submitted to: *NET*

PACS numbers: 87.18b, 87.18Hf, 87.18Sn

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 [17, 4]. This phenomenon has been labeled “spike-timing dependent plasticity” (STDP). Modeling studies have demonstrated the importance of STDP in temporal sequence learning [19, 1, 24, 18, 25, 26] and coincidence detection [10, 32].

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. Specifically, we show the following:

- (i) Feedforward-only excitation and inhibition can cause emergence of direction selectivity. Modifying excitatory connections using an STDP learning rule leads to an asymmetry in the weights which, coupled with slow inhibition, can cause neurons to spike preferentially in one direction of motion. However, only a few spikes are emitted by each neuron, meaning that appreciable amounts of noise may disrupt direction selectivity in realistic neural circuits.
- (ii) Addition of fixed-strength recurrent inhibitory connections to the above model allows competition between neurons, enabling multiple directions to be encoded.
- (iii) Addition of recurrent excitatory connections modified through STDP results in much stronger and more robust direction selective responses, compared to cases where only feedforward excitation is present.

- (iv) 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 lateral geniculate nucleus (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.

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 neurons. 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 (the cf. Fig. 5 in this paper).

More recent computational models have postulated a variety of mechanisms for direction selectivity in the visual cortex, including probabilistic feedforward synapses [5], short-term synaptic depression [7, 29], a combination of feedforward spike-timing dependent plasticity and synaptic depression [6, 27], rate-based Hebbian learning [34, 9], and specialized connectivity schemes [2, 14, 28, 31, 16, 15, 21, 20].

1.2. Contributions of this Paper

An important aspect of direction selectivity that has thus 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 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 LGN. The results demonstrate that a difference in time constants for excitatory and inhibitory currents is sufficient to allow to 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 velocities.
- **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

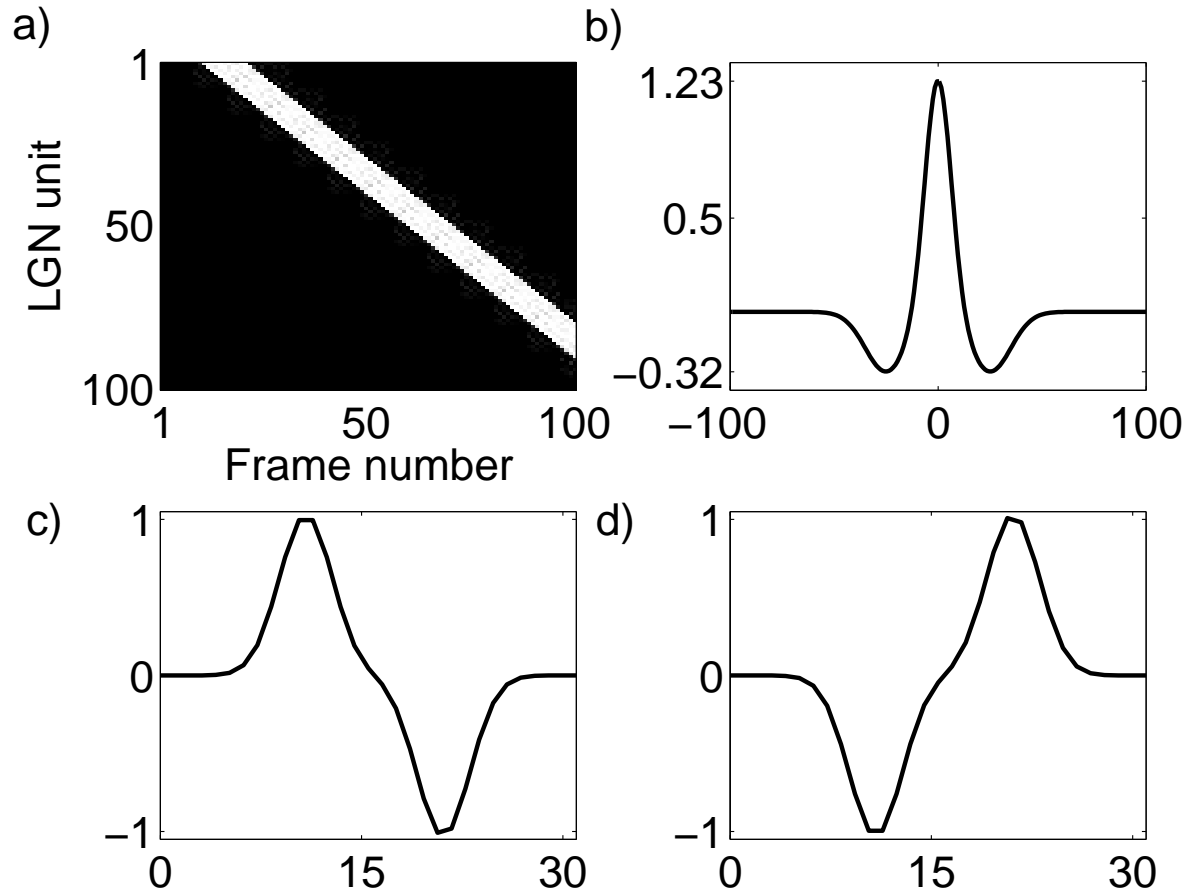


Figure 1. Raw input and LGN filters: (a) Example of raw inputs for a light bar moving at a velocity of 1 LGN unit per time step. (b) Spatial filter used in the LGN module. The filter is formed as the difference of Gaussian curves. (c) Temporal filter for ON cells in the LGN, formed as the difference of Gaussians. (d) Temporal filter for LGN OFF cells, formed as the difference of Gaussians.

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. We found that STDP allows recurrently connected model cortical neurons to establish “predictive waves” of activity, allowing them to fire in anticipation of inputs from the LGN.

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 1-dimensional bars and generates spike trains. The cortical model receives these spike trains as input and produces voltage traces using integrate-and-fire dynamics [13]. All simulations used an integration rate of 1 msec 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):

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

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 direction 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) [33]. 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 1-dimensional light bar 10 retinal/LGN units in width. The bar moved on every time step with a constant velocity set to an integer number of retinal/LGN units, with the constant velocity varying from experiment to experiment. Fig. 1 (a) shows an example of one light bar moving at a velocity of 1 unit per time step. Images of these moving bars were preprocessed by convolving them (using a Fourier transform) with a spatial filter $\mathcal{F}_{SPATIAL}$ given by a difference of Gaussian functions (Fig. 1 (b)). The output of the spatial filter was passed through two temporal filters \mathcal{F}_{ON} and \mathcal{F}_{OFF} , each formed as a difference of Gaussian functions (Fig. 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). In addition to these input spikes, uniformly-distributed spikes were added to each LGN spike train to model noisy spontaneous activity in the LGN and retina. Fig. 2 (a) and (b) show the spatiotemporal outputs of the “ON” and “OFF” filters to the moving bar stimulus depicted in Fig. 1 (a). Fig. 2 (c) and (d) show the corresponding “ON” and “OFF” spike trains that the cortical model receives as input.

2.2. Synapse model

We model synaptic impulse responses using the alpha function:

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

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 $t = 0$.

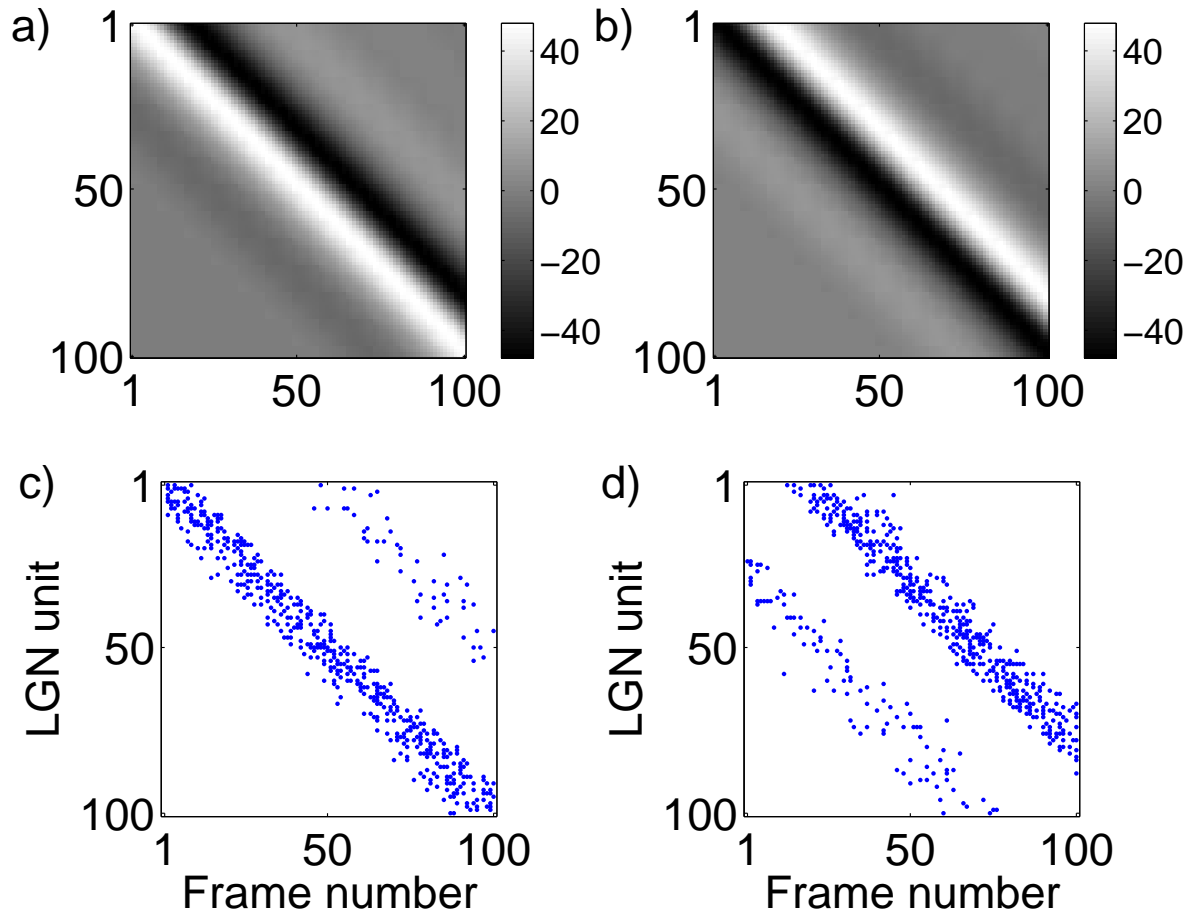


Figure 2. FFT traces and spike rasters: (a) Traces following LGN FFT processing for ON units. (b) Traces following LGN FFT processing for OFF units. (c) Spike raster corresponding to activity of LGN ON cells, produced using a Poisson process. (d) Spike raster corresponding to activity of LGN OFF cells.

Parameter	Value	Units
Excitatory reversal potential (E_{syn})	0	mV
Excitatory peak time (τ_{peak}^{exc})	10	msec
Excitatory maximum peak conductance (g_{syn}^{max})	0.02	μS
Inhibitory reversal potential (E_{syn}^{inh})	-80	mV
Feedforward inhibitory peak time (τ_{peak}^{FFinh})	40	msec
Feedforward inhibitory peak conductance (g_{syn}^{inhFF})	0.0018	μS
Feedback inhibitory peak time (τ_{peak}^{FBinh})	5	msec
Feedback inhibitory peak conductance (g_{syn}^{inhFB})	0.025 (Exp. 3) 0.0055 (Exp. 4)	μS

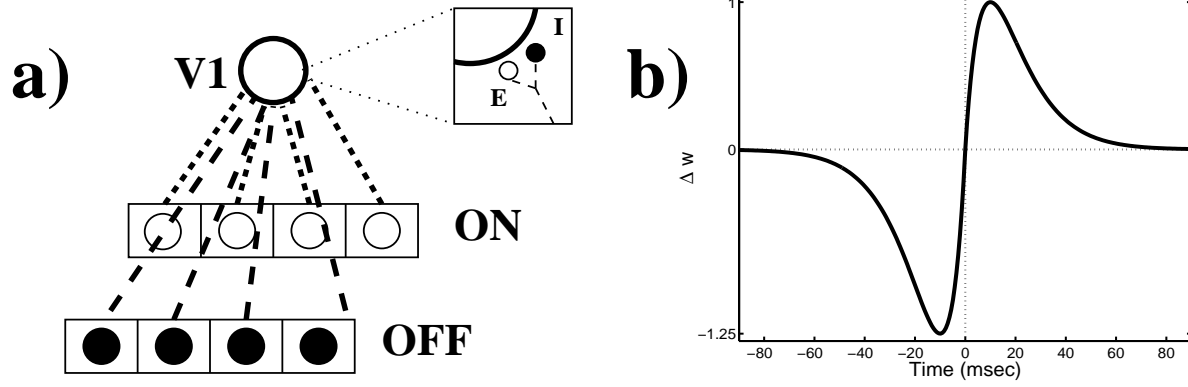


Figure 3. Experiment 1: Feedforward architecture and STDP learning window: (a) In the first group of experiments, a single cortical cell receives inputs from ON and OFF-selective cells from the LGN. Each connection consists of a plastic excitatory connection and a fixed-strength inhibitory connection. (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 2.2: Synapse parameters

2.3. Cortical neuron model

We modeled cortical neurons as leaky integrate-and-fire neurons [13]. 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 Poisson-distributed current to model noisy background inputs to the neurons. 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 Fig. 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 2.3 summarizes neural parameters used in our experiments.

Parameter	Value	Units
Capacitance (C)	0.5	nF
Resistance (R)	40	M Ω
Resting potential (E_{leak})	-60	mV
Threshold voltage (V_{th})	-40	mV
Refractory period (τ_{ref})	5	msec
Poisson-distributed noise magnitude (η)	0.35	nA

Table 2.3: Cortical neuron parameters

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 postsynaptic 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 Fig. 3 (b). Note that in keeping with the observations for firing rate stability noted in [12] (see also [30]), 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} .

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 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. Fig. 5 demonstrates this idea, similar to the direction-selective neural detector first proposed by Barlow and Levick [3]. As an initial proof of concept, experiment 1 shows how a single neuron receiving only feedforward excitation and inhibition and Poisson-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 5 LGN units per simulation step. We performed 10 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 msec.

3.1.2. Weak direction selectivity Fig. 4 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. Fig. 6 demonstrates how the neuron responds with 2 spikes in the preferred direction and 3 spikes in the null direction before training, and

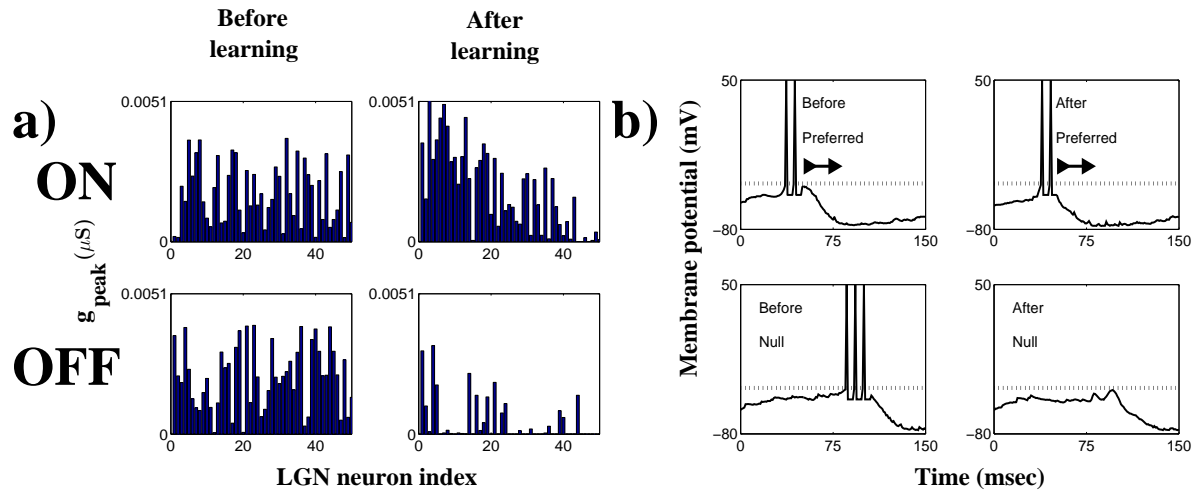


Figure 4. Learned weight asymmetry leads to direction selectivity: (a) These bar graphs show feedforward weights from 50 ON cells and 50 OFF cells onto a single cortical neuron. Before learning (two left graphs), weights are distributed uniformly. After learning (two right graphs), the spike-timing dependent 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 4 spikes) to movement in the null direction. After learning, the cell only spikes in response to the preferred direction of motion.




Figure 5. Direction selectivity with feedforward connections relies on interaction between inhibition and weight asymmetry: (a) When a light bar moves in the preferred direction, the bar activates LGN units with high-strength connections *before* the slower inhibition can compensate, causing the cell to spike. (b) When the bar moves in the null direction, weak synaptic weights are encountered early on. By the time the bar reaches the region of strong synaptic weights, inhibition has already risen to compensate, preventing any firing from occurring.

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

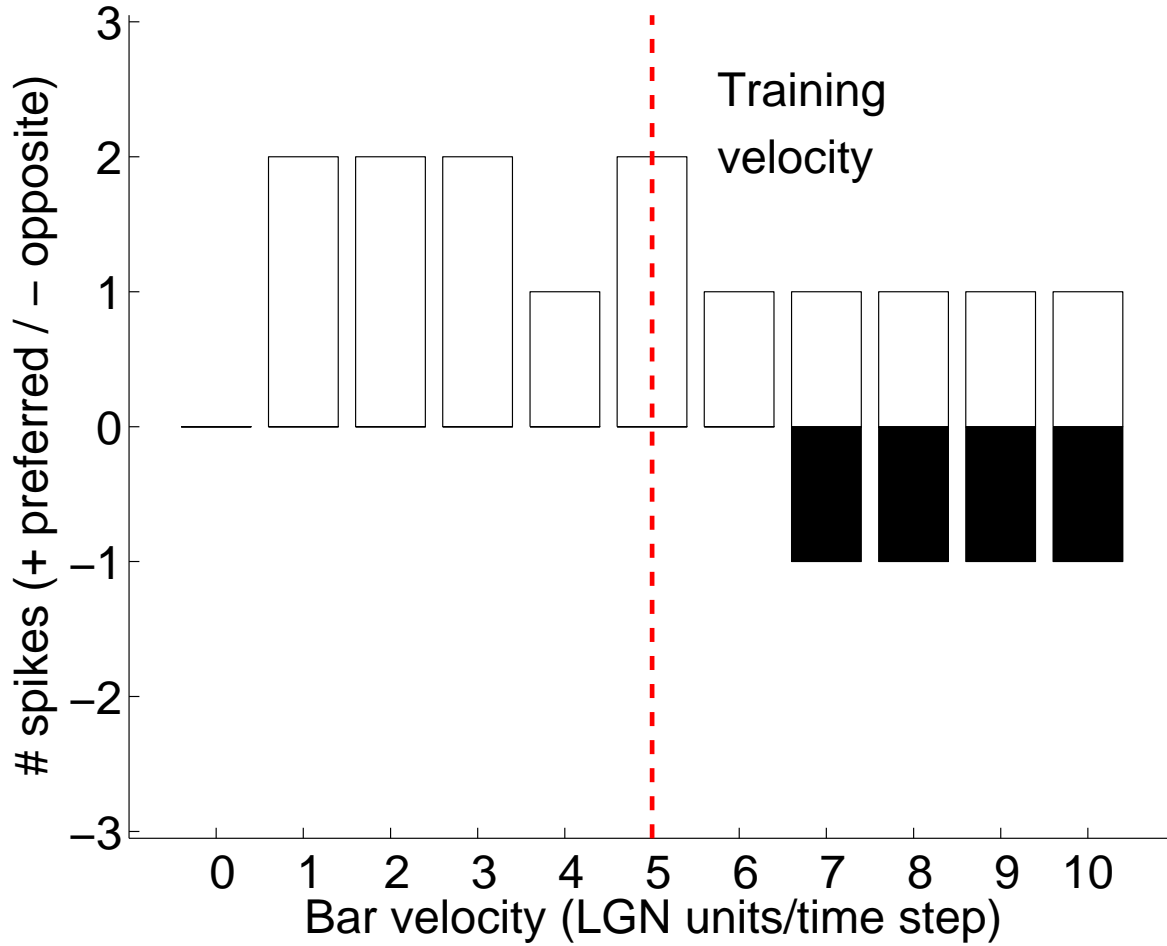


Figure 6. Learning causes direction selectivity across a range of velocities: A single cortical neuron trained using a velocity of 5 LGN units per msec demonstrates direction selectivity across a range of velocities, with a notable decline in selectivity for higher velocities.

direction.

Unfortunately, as Fig. 6 shows, the response in the preferred direction remains weak; only 2 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 below.

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 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.

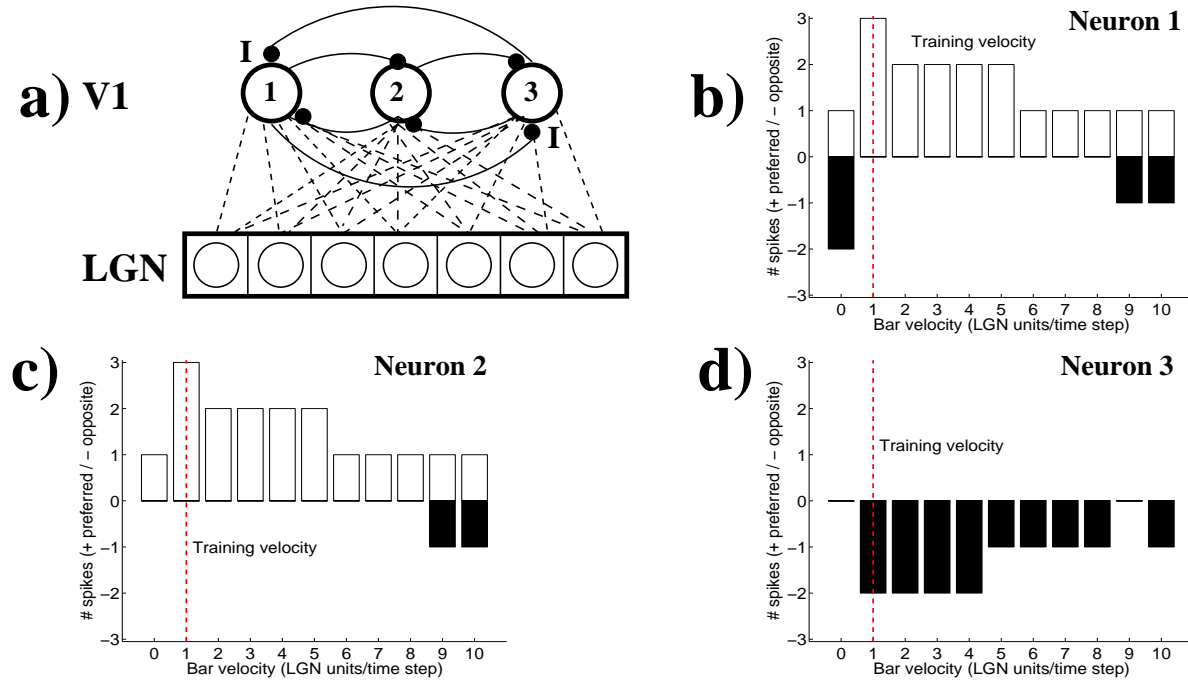


Figure 7. Experiment 2: Inhibitory feedback connections permit learning of multiple velocities: (a) In the second group of experiments, a collection of 3 cortical cells are interconnected using fixed-strength inhibitory synapses. The 3 neurons are exposed to 2 different velocities (in this case, +1 and -1, representing two different directions of motion). (b,c,d) Neurons 1 and 2 code for the opposite direction of motion as neuron 3.

3.2.1. Training paradigm We trained a group of 3 cortical neurons on two different directions of motion. The first direction involved a bar moving left to right at a velocity of 1 retinal/LGN unit per time step; the second direction involved a bar moving right to left at a velocity of 1 unit per time step. Each pass of the light bar over the retinal/LGN unit 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 msec. The neurons were all-to-all connected (with no self-connections) using inhibitory synapses with a peak synaptic conductance $g_{peak}^{inhFB} = 0.025\mu S$.

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. Fig. 7 (a) shows a schematic diagram of this arrangement. Fig. 7 (b), (c), and (d) demonstrate responses of the 3 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.

Fig. 8 shows the resulting asymmetry in feedforward excitatory weights. Note that

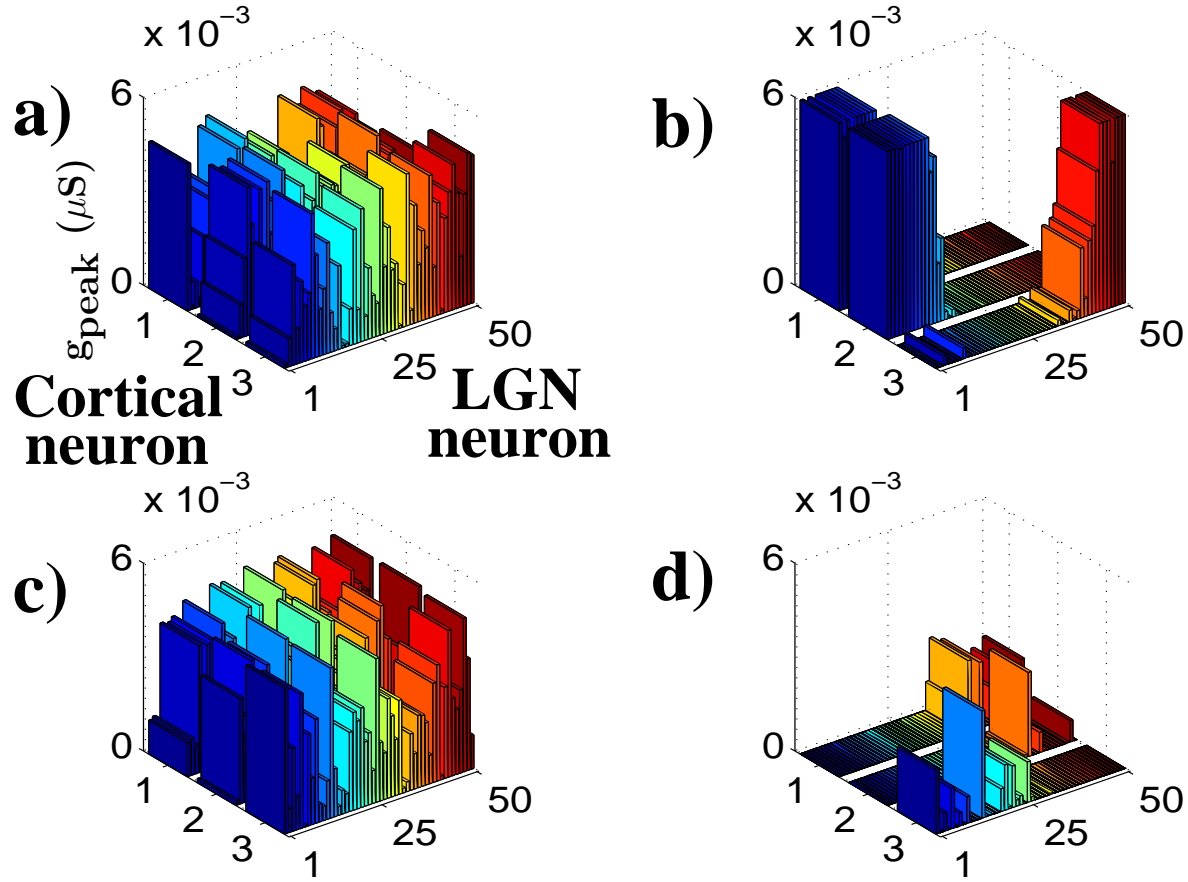


Figure 8. Feedforward weights reflect competition between cortical cells: (a,c) Before learning, weights are uniformly distributed. (b,d) After learning, competition between cortical neurons has caused an asymmetry in the weights of neurons 1 and 2 that causes spiking in one direction of motion, while neuron 3's weights respond to motion in the opposite direction.

neurons 1 and 2 display weights that code for the opposite direction of motion as neuron 3.

Fig. 9 demonstrates how the weak direction selectivity developed as a result of STDP in feedforward-only connections drops off as feedforward and feedback inhibitory strengths 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 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 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

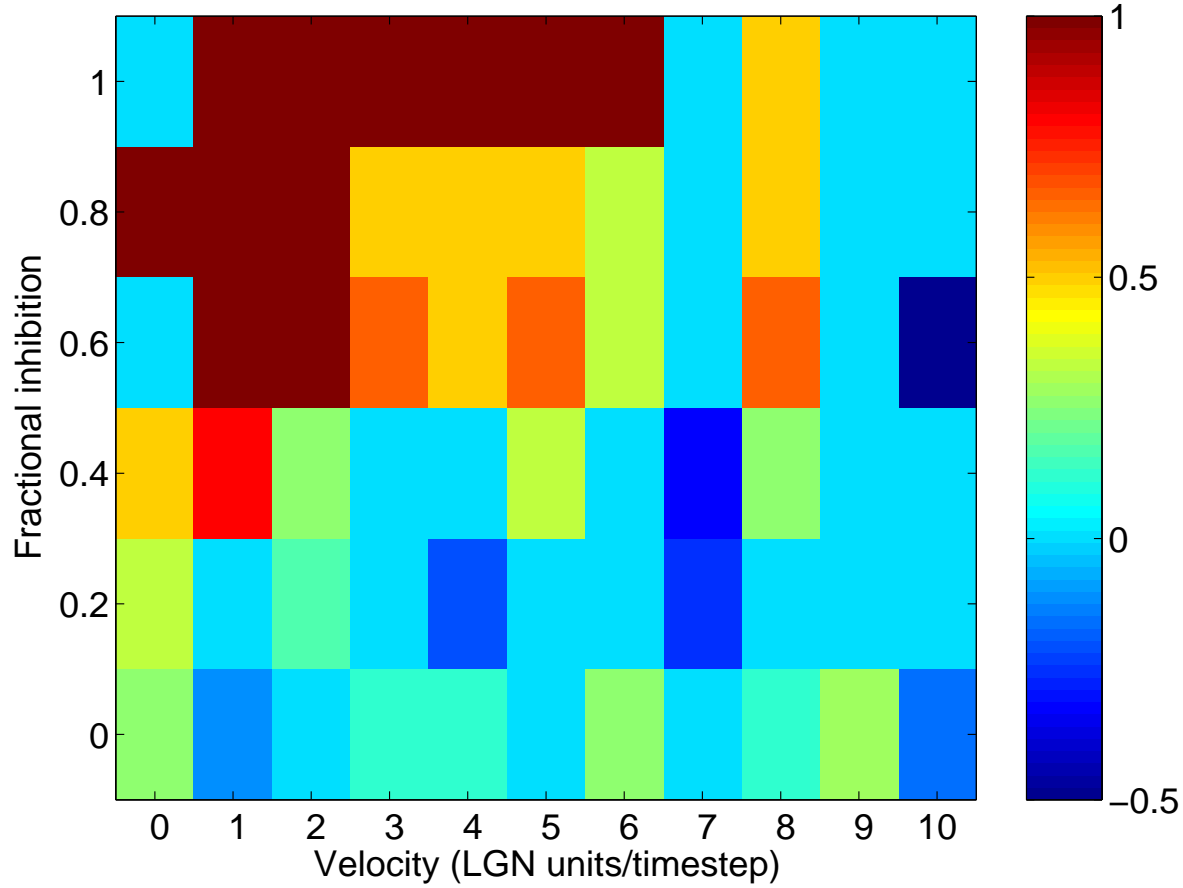


Figure 9. Feedforward-only learning causes dropoff in direction selectivity as inhibition is decreased: A single neuron is trained without feedback inhibition or feedback excitation on sequences with velocities of 5 LGN units per time step. After training, direction selectivity is tested over a range 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.

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 (Fig. 10). The recurrent inhibitory synapses are initialized to constant fixed values of $g_{peak}^{inhFB} = 0.0055\mu S$, and the recurrent excitatory synapses are initialized to random, uniformly-distributed values from 0 to $0.005\mu S$. We assume that each neuron receives input from a patch of the retinal/LGN system that does not overlap with the




Figure 10. 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.

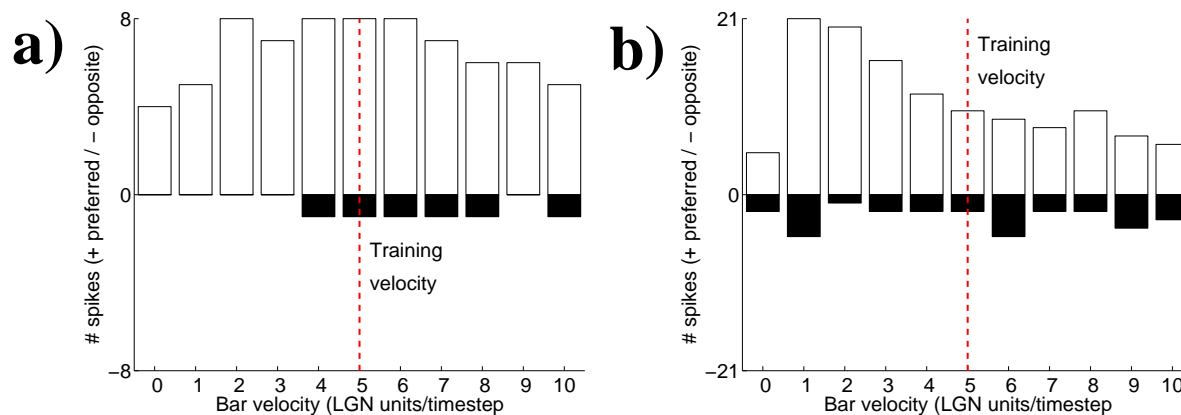


Figure 11. Recurrent connections allow robust direction selectivity: In contrast to the single neuron shown in Fig. 6, responses in the preferred direction are represented by numerous spikes. Responses in the null direction remain weak. (a) shows responses with inhibition at 100% of initial; (b) shows responses when inhibition is set at 60% of initial. Note the difference in scale of the y axes for the two subgraphs.

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 were modified according to the STDP rule as the network was exposed to a moving bar 10 retinal/LGN units wide moving at a velocity of 5 retinal/LGN units per time step. Again, each pass of the light bar over the retinal/LGN unit constituted one iteration of the simulation, and each iteration ran for 725 msec. 10 iterations comprised the total training set for the network.

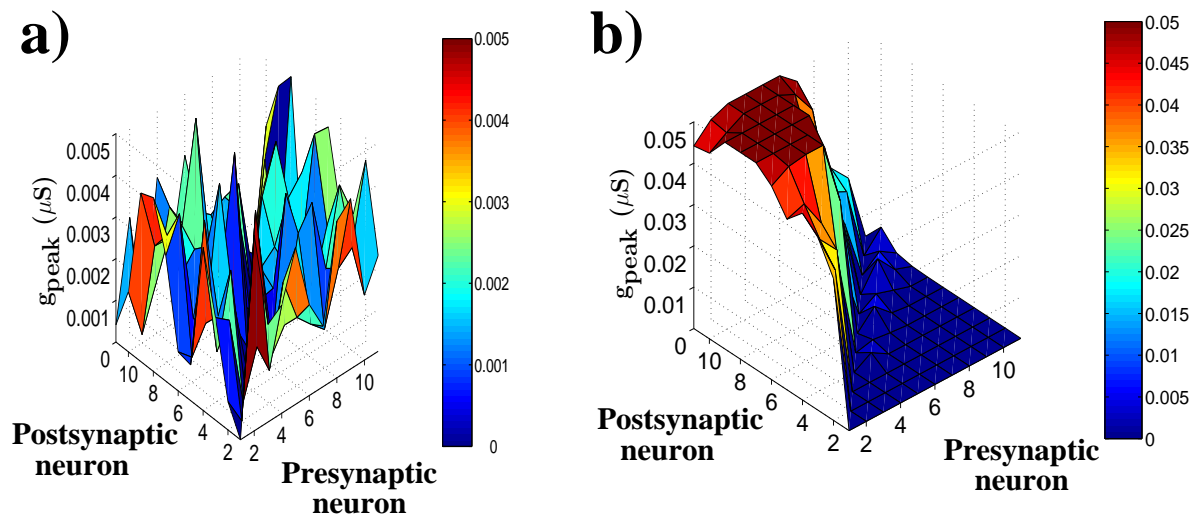


Figure 12. Recurrent weights show asymmetry after training: (a) Shows recurrent synaptic weights before training. (b) Shows weights from the same network after training. A clear asymmetry results from being trained on bars moving in the preferred direction. Note the difference in scale from (a).

3.3.2. Strong direction selectivity Fig. 11 shows that STDP causes an asymmetry in the excitatory recurrent connections that leads to robust direction selectivity. Fig. 11(a) shows the number of spikes fired by the neuron in the middle of the chain when presented with bars moving at velocities from $0 \dots 10$. Compared to the results for feedforward-only excitation in Fig. 6, 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 Fig. 13). Note the higher velocities in particular display much more robust direction selectivity as compared to the feedforward-only case. Our findings are consistent with other modeling studies, for example, the work of Suarez, Koch, and Douglas [31], who found that asymmetric, excitatory recurrent connections are necessary to replicate biological data and ensure robust direction selectivity.

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 [1, 24, 25, 26]. We investigated the implications of these findings within the context of motion detection and direction selectivity.

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




Figure 13. 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.

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 5 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.




Figure 14. Recurrent connections cause predictive waves of activity: A moving stimulus (green 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 (red line), the model predicts a peak in activity only after the flash has occurred.

4.2. Predictive firing

Fig. 14 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 msec in the model) before the arrival of the moving stimulus on the receptive 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.

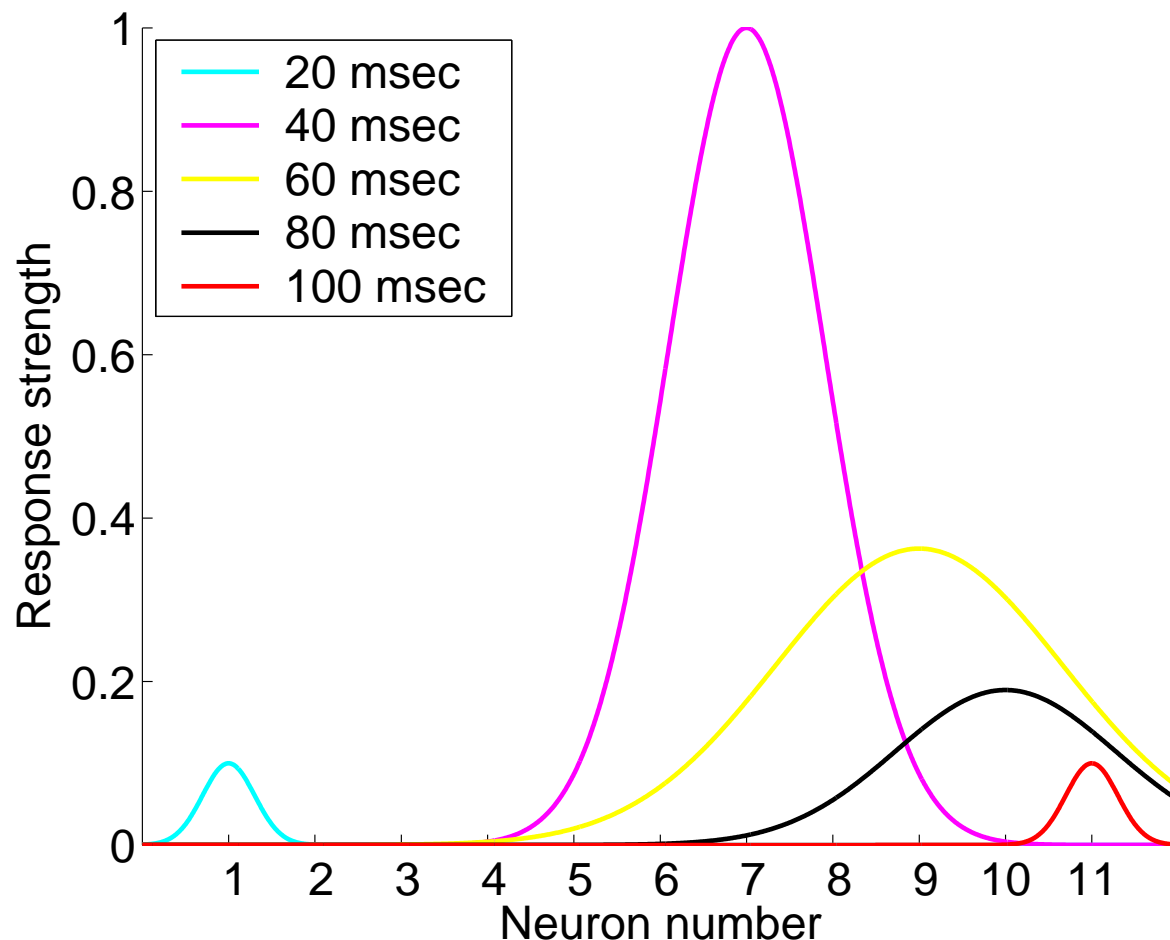


Figure 15. Predictive activity continues to propagate when moving bar is stopped: After 20 msec of exposing the network to a moving bar, we turn off all inputs from the LGN. Recurrent connections continue to propagate the activity even in the absence of external input. In a larger cortical model, this 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.

4.3. Continuing predictive activity

Fig. 15 shows the results from the second experimental setup. We begin with a bar moving at a velocity of 5 LGN units per simulation step. After 20 milliseconds, input from the LGN stops completely. 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. Fig. 15 shows this effect in the model network; the figure plots the mean location of activity within the network of 11 cortical neurons over time. Cortical activities were measured every 20 msec, 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 cortical activity. In long chains of recurrently connected neurons, this 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.




Figure 16. 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.

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 4 main results:

- (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 as well as 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 [22, 23].

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 Fig. 16). 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 [8, 11] from statistics and machine learning, 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.

- [1] L. F. Abbott and K. I. Blum. Functional significance of long-term potentiation for sequence learning and prediction. *Cereb. Cortex*, 6:406–416, 1996.
- [2] E. H. Adelson and J. Bergen. Spatiotemporal energy models for the perception of motion. *J. Opt. Soc. Am. A*, 2(2):284–299, 1985.
- [3] H. Barlow and W. Levick. The mechanism of directionally selective units in rabbit’s retina. *J. Physiol. (Lond.)*, 178:477–504, 1965.
- [4] G. Bi and M. Poo. Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.*, 18(24):10464–10472, 1998.
- [5] N. Buchs and W. Senn. Learning direction selectivity through spike-timing dependent modification of neurotransmitter release probability. *Neurocomputing*, 38–40:121–127, 2001.
- [6] N. Buchs and W. Senn. Spike-based synaptic plasticity and the emergence of direction selective simple cells: simulation results. *J. Computational Neurosci.*, 13:167–186, 2002.
- [7] F. Chance, S. Nelson, and L. Abbott. Temporal characteristics of v1 cells arising from synaptic depression. In J. Bower, editor, *Computational Neuroscience, Trends in Research 1998*. Amsterdam: Elsevier Press, 1998.
- [8] A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Statistical Soc., Series B*, 39(1):1–38, 1977.
- [9] J. C. Feidler, A. B. Saul, A. Murthy, and A. L. Humphrey. Hebbian learning and the development of direction selectivity: the role of geniculate response timings. *Network: Computation in Neural Systems*, 8:195–214, 1997.
- [10] W. Gerstner, R. Kempter, J. L. van Hemmen, and H. Wagner. A neuronal learning rule for sub-millisecond temporal coding. *Nature*, 383:76–81, 1996.
- [11] H. Hartley. Maximum likelihood estimation from incomplete data. *Biometrics*, 14:174–194, 1958.
- [12] R. Kempter, W. Gerstner, and J. L. van Hemmen. Intrinsic stabilization of output rates by spike-time dependent Hebbian learning. *Neural Computation*, 13:2709–2742, 2001.
- [13] C. Koch. *Biophysics of Computation: Information Processing in Single Neurons*. New York: Oxford University Press, 1999.
- [14] C. Koch and T. Poggio. The synaptic veto mechanism: does it underlie direction and orientation selectivity in the visual cortex? In D. Rose and V. Dobson, editors, *Models of the Visual Cortex*, pages 408–419. New York, NY: Wiley, 1985.
- [15] M. Livingstone. Mechanisms of direction selectivity in macaque V1. *Neuron*, 20:509–526, 1998.

- [16] R. Maex and G. A. Orban. Model circuit of spiking neurons generating directional selectivity in simple cells. *J. Neurophysiol.*, 75(4):1515–1545, 1996.
- [17] H. Markram, J. Lübke, M. Frotscher, and B. Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic apss and epsps. *Science*, 275:213–215, 1997.
- [18] M. R. Mehta and M. Wilson. From hippocampus to V1: Effect of LTP on spatiotemporal dynamics of receptive fields. In J. Bower, editor, *Computational Neuroscience, Trends in Research 1999*. Amsterdam: Elsevier Press, 2000.
- [19] A. A. Minai and W. Levy. Sequence learning in a single trial. In *Proceedings of the 1993 INNS World Congress on Neural Networks II*, pages 505–508. New Jersey: Erlbaum, 1993.
- [20] P. Mineiro and D. Zipser. Analysis of direction selectivity arising from recurrent cortical interactions. *Neural Computation*, 10(2):353–371, 1998.
- [21] M. Miyashita, D.-S. Kim, and S. Tanaka. Cortical direction selectivity without directional experience. *NeuroReport*, 8(5):1187–1191, 1997.
- [22] R. P. N. Rao and D. H. Ballard. Dynamic model of visual recognition predicts neural response properties in the visual cortex. *Neural Computation*, 9(4):721–763, 1997.
- [23] R. P. N. Rao and D. H. Ballard. Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive field effects. *Nature Neuroscience*, 2(1):79–87, 1999.
- [24] R. P. N. Rao and T. J. Sejnowski. Predictive sequence learning in recurrent neocortical circuits. In *Advances in Neural Information Processing Systems 12*, pages 164–170. Cambridge, MA: MIT Press, 2000.
- [25] R. P. N. Rao and T. J. Sejnowski. Spike-timing dependent Hebbian plasticity as temporal difference learning. *Neural Computation*, in press, 2001.
- [26] R. P. N. Rao and T. J. Sejnowski. Self-organizing neural systems based on predictive learning. *Phil. Trans. R. Soc. Lond. A*, 361, 2003.
- [27] W. Senn and N. Buchs. Spike-based synaptic plasticity and the emergence of direction selective simple cells: mathematical analysis. *J. Computational Neurosci.*, 14:119–138, 2003.
- [28] R. M. Shapley, R. C. Reid, and R. Soodak. Spatiotemporal receptive fields and direction selectivity. In M. S. Landy and J. A. Movshon, editors, *Computational Models of Visual Processing*, pages 109–118. Cambridge, MA: MIT Press, 1992.
- [29] A. P. Shon and R. P. N. Rao. Learning temporal patterns by redistribution of synaptic efficacy. In J. Bower, editor, *Computational Neuroscience, Trends in Research 2003*. Amsterdam: Elsevier Press, (to be published).
- [30] S. Song, K. D. Miller, and L. F. Abbott. Competitive Hebbian learning through spike-timing dependent synaptic plasticity. *Nature Neuroscience*, 3:919–926, 2000.
- [31] H. Suarez, C. Koch, and R. Douglas. Modeling direction selectivity of simple cells in striate visual cortex within the framework of the canonical microcircuit. *J. Neurosci.*, 15(10):6700–6719, 1995.
- [32] R. E. Suri and T. J. Sejnowski. Spike propagation synchronized by temporally asymmetric hebbian learning. *Biological Cybernetics*, 87(5–6):440–445, 2002.
- [33] B. Wandell. *Foundations of Vision*. Sunderland, MA: Sinauer, 1995.
- [34] S. Wimbauer, O. Wensch, K. Miller, and J. van Hemmen. Development of spatiotemporal receptive fields of simple cells: I. Model Formulation. *Biological Cybernetics*, 77:453–461, 1997.