Neuron, Vol. 32, 1–20, October 25, 2001, Copyright ©2001 by Cell Press

Cortical Remapping through Spike Timing-Dependent Plasticity

Sen Song and L.F. Abbott[1]
Volen Center for Complex Systems and
Department of Biology
Brandeis University
Waltham, Massachusetts 02454-9110

Summary

Long-term modification of synaptic efficacy can depend on the timing of pre- and postsynaptic action potentials. In model studies, such spike timing-dependent plasticity (STDP) introduces the desirable features of competition among synapses and regulation of postsynaptic firing characteristics. STDP strengthens synapses that receive correlated input, which can lead to the formation of stimulus-selective columns and the development, refinement, and maintenance of selectivity maps in network models. The temporal asymmetry of STDP suppresses strong destabilizing self-excitatory loops and allows a group of neurons that become selective early in development to direct other neurons to become similarly selective. STDP, acting alone without further hypothetical global constraints or additional forms of plasticity, can also reproduce the remapping seen in adult cortex following afferent lesions.

Introduction

The formation of cortical maps during development (Stryker, 1986; Katz and Shatz, 1996; Yuste and Sur, 1999) and their modification during adulthood (Buonomano and Merzenich, 1998; Wall, 1988; Kaas, 1991; O'Leary et al., 1994; Gilbert, 1996; Weinberger, 1995) are fruitful testing grounds for ideas concerning activity-dependent synaptic plasticity. While activity-independent processes undoubtedly play an important role in map formation (Purves and Lichtman, 1985; Ruthazer and Stryker, 1996; Crair et al., 1997, 1998; Crowley and Katz, 2000), manipulations of activity can strongly affect development (Stryker, 1986). Even in adulthood, cortical maps can be remodeled by changes in input patterns, such as those that occur after lesions (Merzenich et al., 1983, 1984; Kaas et al., 1990). A number of models have been proposed to account for activity-dependent aspects of developmental (reviewed in Miller, 1996) and adult (Grajski and Merzenich, 1990; Benuskov et al., 1994; Sutton et al., 1994) plasticity. These share the common feature of requiring, in addition to correlation-based Hebbian synaptic plasticity, some mechanism to ensure competition among synapses. Competition typically arises from constraints imposed on the synaptic strengths (Miller and MacKay, 1994) that are uncorroborated by experimental evidence or additional types of plasticity, such as a sliding threshold (Bienenstock et al., 1982; Abraham, 1997).

[1]Correspondence: abbott@brandeis.edu

Experimental evidence from a number of different preparations indicates that repeated pairing of pre- and postsynaptic action potentials can lead to long-term changes in synaptic efficacy, the sign and amplitude of which depend on relative spike timing (Levy and Steward, 1983; Gustafsson et al., 1987; Debanne et al., 1994; Magee and Johnston, 1997; Markram et al., 1997; Bell et al., 1997; Debanne et al., 1998; Bi and Poo, 1998; Zhang et al., 1998; Egger et al. 1999; Feldman, 2000).

Spike timing-dependent plasticity (STDP) forces synapses to compete with each other for control of the timing of postsynaptic action potentials, and this, by itself, can lead to competitive Hebbian synaptic modification (Kempter et al., 1999; Song et al., 2000). Here, we explore whether STDP alone, without any global constraints on synaptic efficacies or additional forms of plasticity, can account for major features of activity-dependent plasticity of columns and maps during development and adulthood.

An important feature of the form of STDP we discuss is its asymmetry with respect to the timing of pre- and postsynaptic action potentials. Long-term potentiation (LTP) of synapses occurs if postsynaptic action potentials are repeatedly paired with presynaptic spikes that precede them by no more than about 50 ms. Presynaptic spikes that repeatedly follow postsynaptic action potentials, within a time window either similar to the LTP window or considerably longer, depending on the synapse, induce long-term depression (LTD).

The temporal asymmetry of STDP has a number of important consequences. Consider two neurons, A and B, that tend to fire together in the sequence A followed by B. In a time-independent Hebbian model, excitatory synapses from A to B and from B to A would both be strengthened in this situation. STDP, on the other hand, strengthens the synapse from A to B while weakening the synapse from B to A. This allows neuron A to modify the selectivity of neuron B without itself being affected by the changes in B. In more general terms, STDP allows selective groups of neurons with correlated firing patterns to direct the development of nonselective neurons with more random firing patterns.

Many Hebbian models allow only feedforward synapses to be modified by activity because allowing recurrent synapses to strengthen produces strong self-excitatory loops that lead to uncontrollable network activity. This is the network analog of the reciprocal strengthening of synapses between neurons A and B in the two-neuron example just discussed. The temporal asymmetry of STDP suppresses strong recurrent loops leading to stable network models even when all network synapses are subject to activity-dependent modification. In the models we consider, synaptic plasticity affects both feedforward and recurrent network connections, and it does so at all times. Plasticity is not deactivated once the desired structures form. The neuronal selectivities and cortical maps arising from these models are thus stable and persistent.

Results

Each synapse in the models we consider is characterized by a peak conductance g (the peak value of the synaptic conductance following a single presynaptic action potential) that is constrained to lie between 0 and a maximum value g_{max}. Every pair of pre- and postsynaptic spikes can potentially modify the value of g, and the changes due to each spike pair are continually summed to determine how g changes over time. We make the simplifying assumption that the modifications produced by individual spike pairs combine linearly. This is clearly an approximation, as a number of nonlinear affects are observed experimentally (Markram et al., 1997; Sjöström and Nelson, personal communication). However, most of our results depend on only two basic features of STDP: presynaptic spikes arriving slightly before postsynaptic firing produce synaptic potentiation and random pre- and postsynaptic action potentials result in synaptic depression.

A presynaptic spike occurring at time t_{pre} and a postsynaptic spike at time t_{post} modify the corresponding synaptic conductance by $g \rightarrow g + g_{max}F(\Delta t)$, where $\Delta t = t_{pre} - t_{post}$ and

$$F(\Delta t) = \begin{cases} A_+ \exp(\Delta t/\tau_+), & \text{if } \Delta t < 0 \\ -A_- \exp(-\Delta t/\tau_-), & \text{if } \Delta t \geq 0 \end{cases}.$$

If this modification would make g less than 0 or greater than g_{max}, g is set to the appropriate limiting value. The form of the STDP window function, $F(\Delta t)$, (Figure 1A) and the values of the parameters used (Experimental Procedures) are based on experimental data. The time constants τ_+ and τ_- determine the ranges of pre- to postsynaptic spike intervals over which synaptic strengthening and weakening are significant, and A_+ and A_- determine the maximum amount of synaptic modification in each case. The experimental results indicate a value of τ_+ in the range of tens of milliseconds. We have used $\tau_+ = 20$ ms in all our simulations. Values of τ_- fall into two categories depending on the type of synapse being studied. In some cases, $\tau_- \approx \tau_+$ (Markram et al., 1997; Zhang et al., 1998; Bi and Poo, 1998), while in others, $\tau_- >> \tau_+$ (Debanne et al., 1998; Feldman, 2000). We therefore consider two cases: $\tau_- = \tau_+ = 20$ ms and $\tau_- = 5\tau_+ = 100$ ms.

The ratio of the areas under the negative and positive portions of the STDP window function, defined as $B = A_-\tau_-/(A_+\tau_+)$, has a significant impact on our simulations. To avoid uncontrolled synaptic growth, this ratio must be greater than 1, which assures that the total integral of $F(\Delta t)$ over all Δt is negative. In the range above 1, the value of B controls the level of activity in a network model.

STDP and Presynaptic Correlation Times

Before considering network models, we study how STDP affects synapses onto a single postsynaptic neuron. When multiple synapses drive a postsynaptic neuron, STDP tends to segregate them into strong and weak groups, creating a bimodal equilibrium distribution of synaptic strengths (Song et al., 2000). Figure 1B shows the strengths of 1000 excitatory synapses onto a single integrate-and-fire model neuron (Experimental Procedures) after a stable equilibrium distribution has been established by STDP. In this case, the synaptic inputs are uncorrelated, and the division into strong and weak synapses is random with ~50% of the synapses in each group (Figure 1B).

STDP strengthens synapses that are effective at rapidly evoking a postsynaptic action potential, such as groups of presynaptic inputs that fire in a correlated manner (Song et al., 2000). STDP strengthens the synapses made by such inputs provided that the correlation time, which sets the timescale over which the inputs are significantly correlated, is appropriately related to the time constants of the STDP window function (see Figure 1F). For Figure 1C, correlations were introduced among the spike trains for inputs 501 through 1000 (see Experimental Procedures), while the spike trains for inputs 1 through 500 were left uncorrelated. After STDP equilibrated, synapses in the correlated group ended up stronger than those in the uncorrelated group.

We next consider an example inspired by the development of ocular dominance in neurons of the primary visual cortex and lateral geniculate nucleus. Before eye opening, the activity of retinal ganglion cells within each eye is correlated by retinal circuitry due to patterns of activity such as retinal waves (reviewed in Wong, 1999). However, there is little correlation between the activities of the two eyes at this point in development. To simulate this situation, we introduced independent correlations of equal magnitude into inputs 1 through 500 (representing left-eye inputs) and inputs 501 through 1000 (representing right-eye inputs). The two groups were not correlated with each other. This creates a situation in which the two sets of equally correlated inputs compete, and only one ends up dominating the response of the postsynaptic neuron. Which set does this is random. Figure 1D shows the equilibrium distribution of synaptic strengths in a case where inputs 1 through 500 won the competition and formed strong connections to the postsynaptic neuron. Synapses formed by the other correlated group of inputs are weak due to the competitive nature of STDP.

The ultimate distribution of synaptic strengths that arises from STDP depends on the number of pre- and postsynaptic spike pairs that fall into different portions of the STDP window function, $F(\Delta t)$. The average number of presynaptic action potentials occurring at various times before and after a postsynaptic spike is proportional to the correlation function of the pre- and postsynaptic spike trains. An estimate of the overall effect of STDP can be obtained by computing the integral over spike timing differences Δt of the product of $F(\Delta t)$ and this input-output correlation function (Song et al., 2000). The input-output correlation function is flat for the uncorrelated set of inputs in the example of Figure 1C, except for a small excess of presynaptic spikes just before a postsynaptic action potential. The synapses for these inputs are therefore weakened by STDP due to the negative total integral of the window function. On the other hand, the correlation function between the correlated inputs and the postsynaptic action potentials has a prominent peak near time difference 0 (Figure 1E). This peak has a large symmetric component that, by itself, would weaken the synapses when $\tau_- = \tau_+$. How-

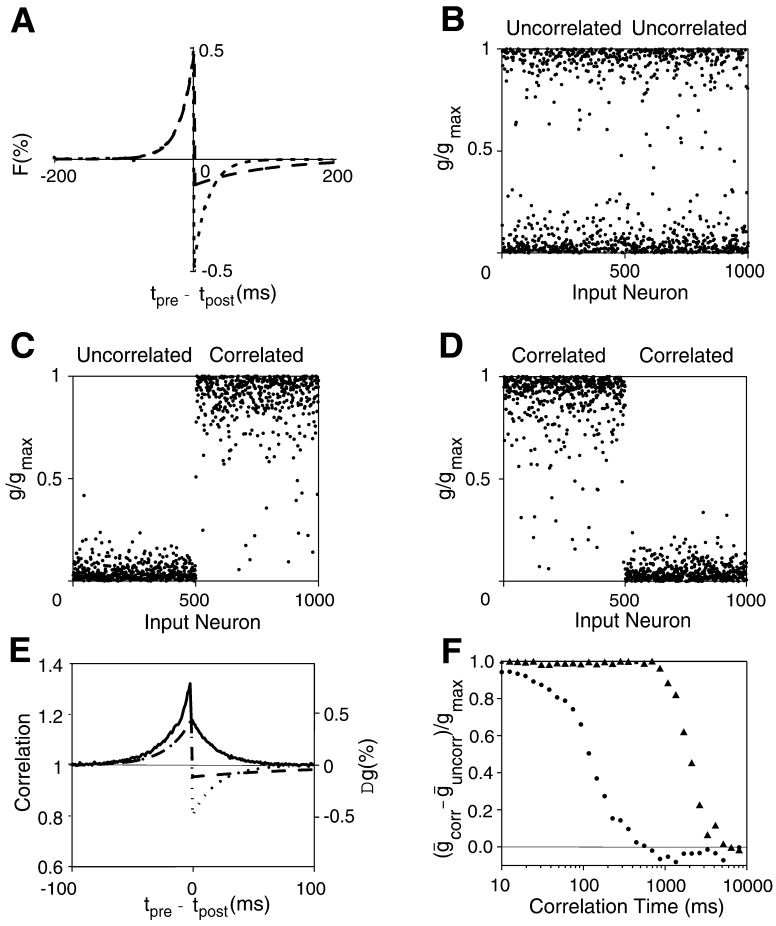

Figure 1. Examples of STDP Involving a Single Postsynaptic Neuron

(A) The STDP window function, which determines the percentage change of peak synaptic conductance (relative to its maximum allowed value) induced by a single pre- and postsynaptic action potential pair at times t_{pre} and t_{post}. The short-dashed curve is for $\tau_- = \tau_+$, and the long-dashed curve is for $\tau_- = 5\tau_+$.

(B) Equilibrium distribution of synaptic strengths for uncorrelated inputs. Each dot represents the strength (relative to the maximum allowed strength) of a synapse from an input neuron to the postsynaptic neuron after STDP has come to equilibrium.

(C) Equilibrium synaptic strengths when the postsynaptic neuron receives both uncorrelated (input neurons 1 through 500) and correlated (input neurons 501 through 1000) inputs. The correlation time was 20 ms and $\tau_- = \tau_+$.

(D) Equilibrium synaptic strengths when the postsynaptic neuron receives input from two groups (input neurons 1 through 500 and 501 through 1000) that were both equally correlated but uncorrelated with each other.

(E) Average correlation between presynaptic action potentials of the correlated group of inputs in (C) and the postsynaptic spike train. The solid curve indicates the relative probability of a presynaptic spike occurring at time t_{pre} when a postsynaptic spike occurs at time t_{post}. The curve is normalized so that a value of 1 arises from chance occurrences of such pairs. The dotted and the dashed curves show the STDP window functions for $\tau_- = \tau_+$ and $\tau_- = 5\tau_+$.

(F) The result of a sequence of runs similar to that shown in (C) but with different correlation times. The difference between the average value of the synaptic strengths (divided by the maximum synaptic strength) for the correlated group of input neurons (inputs 501 through 1000) and the uncorrelated input neurons (inputs 1 through 500) is plotted against the correlation time. Circles are for $\tau_- = \tau_+$ and triangles for $\tau_- = 5\tau_+$.

ever, there is also an excess of presynaptic spikes before the postsynaptic response due to the input integration performed by the postsynaptic neuron. The portion of the peak to the left of time differences of 0 in Figure 1E is larger than the portion to the right. This excess causes the synaptic strengths of the correlated group of inputs to grow.

In this example, the decay time constant of the input-output correlations (i.e., the correlation time) is close to the membrane time constant of the neuron and to the decay times $\tau_- = \tau_+ = 20$ ms of the STDP window function. For STDP with $\tau_- = 5\tau_+$, the basic result of strengthening of correlated groups of synapses is preserved, but τ_- no longer matches the decay time of the input-output correlation for positive time differences (Figure 1E). As a result, the symmetric component of the peak in the input-output correlation function around zero time difference causes synaptic strengthening by STDP, and correlated inputs are strengthened even more effectively when $\tau_- = 5\tau_+$ than when $\tau_- = \tau_+$.

Another important difference between these two cases is that STDP with $\tau_- = 5\tau_+$ is sensitive to correlations over much longer timescales than STDP with $\tau_- = \tau_+$. To study this, we performed a series of simulations similar to the one shown in Figure 1C but with different correlation times for the correlated group of inputs (a correlation time of 20 ms was used in Figure 1C). Figure 1F shows the difference between the average synaptic conductance of the correlated and uncorrelated groups after STDP has come to equilibrium. For the case $\tau_- = \tau_+$, large differences between the two groups, as seen in Figure 1C, begin to shrink for correlation times greater than 100 ms and vanish for a correlation time of about 500 ms. For even larger correlation times, the correlated group ends up with synapses that are slightly weaker than those of the uncorrelated group. When $\tau_- = 5\tau_+$, the correlated group develops much stronger synapses than the uncorrelated group for correlation times up to 1 s, and the difference between the two groups does not vanish until the correlation time becomes greater than about 5 s.

Development of Selectivity and Columns

Previous models have demonstrated that Hebbian synaptic plasticity can cause neurons to become selective to specific aspects of their input provided that appropriate global constraints or additional plasticity mechanisms are included (reviewed in Miller, 1996). To test whether STDP can generate such selectivity by itself, we applied it to a network model intended to simulate

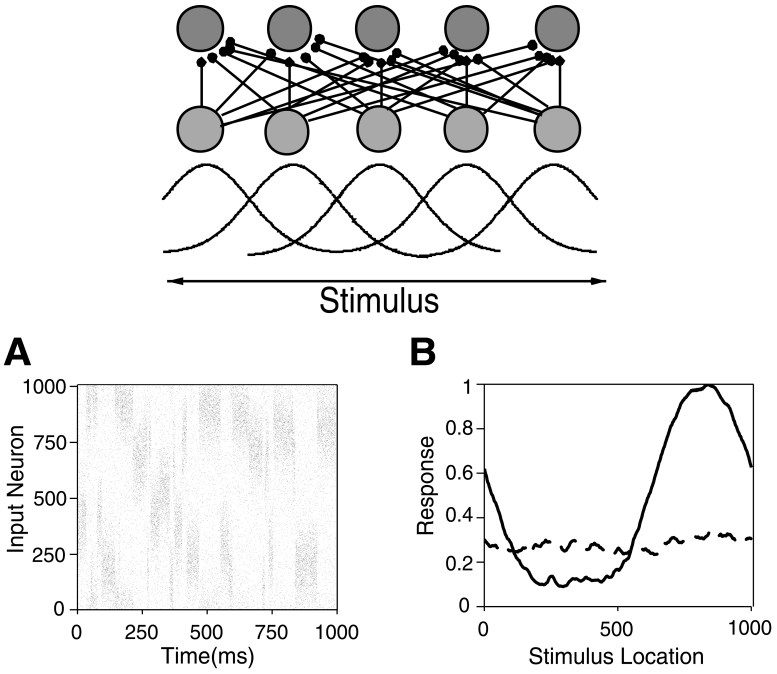

Figure 2. STDP Leads to the Development of Selectivity

The upper panel shows the network used in this example with Gaussian firing rate curves for the input neurons (lower row of circles) and feedforward connections to the network neurons (upper row of circles) but recurrent connections absent. (A) A typical pattern of input neuron activity. Each dot is a spike, and all 1000 input neurons are shown for 1 s. (B) Firing rate curves for network neuron 40 before (dashed) and after (solid) STDP. The initial values of the feedforward synapses were chosen randomly and uniformly between 0 and g_{max}, resulting in little initial selectivity. In this and similar figures, responses are plotted as firing rates normalized to a maximum value of 1, and periodic boundary conditions were imposed on the network and firing rate curves.

thalamic relay neurons providing feedforward input to a recurrently interconnected cortical circuit. The model has 1000 input neurons with responses generated by a hypothetical stimulus. The stimulus is parameterized by a single variable that might, for example, represent the location of a touch stimulus on the skin or the location or orientation of a visual image on the retina. The firing rate of each input neuron in response to the stimulus is determined by a Gaussian firing rate curve (Experimental Procedures and upper panel of Figure 2) that reaches its maximum value for a stimulus location that we call the preferred location of that cell. The preferred stimulus locations progress smoothly across the array of input neurons, which induces correlations in the firing of neighboring input neurons. To simplify the discussion, we identify a specific stimulus location by giving the label of the input neuron that has that location as its preferred stimulus location. For example, stimulus location 20 is the value that elicits the maximum response from input neuron 20.

The 1000 input neurons drive 200 network neurons (except for Figure 5, where 250 neurons were used) through sparse random (20% connection probability) excitatory feedforward connections. The network neurons are interconnected in an all-to-all manner by excitatory synapses, and each network neuron also receives background input that makes it spontaneously active (Experimental Procedures).

During simulated development, input to the network is generated by a series of brief presentations of the stimulus at random locations. Each presentation lasts for a short period of time chosen from an exponential distribution with a mean of either 20 or 100 ms. Figure 2A shows a raster of typical input neuron activities generated by such stimuli. Each dot in the plot is a spike, and the high-density areas reveal the stimulus locations. Activity corresponding to a given stimulus location lasts for a short amount of time before switching to activity generated by another random stimulus location.

In a set of initial simulations, we disabled the recurrent connections and allowed STDP to modify feedforward synapses while random stimuli were presented. We show the case $\tau_- = \tau_+$, but similar results were obtained for $\tau_- = 5\tau_+$. At the start of the simulation, the strengths of the feedforward connections were assigned random values. Because of the random connectivity, the neurons in the network were initially nonselective, responding almost equally for all stimulus locations (dashed curve in Figure 2B). However, when STDP came to equilibrium, each neuron in the network had developed input selectivity (solid curve in Figure 2B). Due to the competitive nature of the rule, the strengthening of synapses associated with one group of correlated inputs suppresses other synapses, eventually leading to strong feedforward connections exclusively from a contiguous set of input neurons.

Without recurrent connections, different network neurons develop different selectivities with random preferred locations. Recurrent connections lead to the formation of a single column of neurons with similar selectivities for both $\tau_- = \tau_+$ and $\tau_- = 5\tau_+$, but again we show the former case. In these simulations, the feedforward connections were set to random initial values, and the recurrent connection strengths were initialized to 0. Figures 3A and 3B show equilibrium synaptic strengths for the feedforward and recurrent synapses, respectively, in grayscale plots. The shading at each point represents the strength of the corresponding synapse. The horizontal stripe in Figure 3A indicates that all the network neurons receive strong synapses from input neurons in the neighborhood of input 800, while synaptic connections from other input neurons are weak. This pattern of connectivity confers similar selectivity to all the neurons in the network, as seen in the firing rate

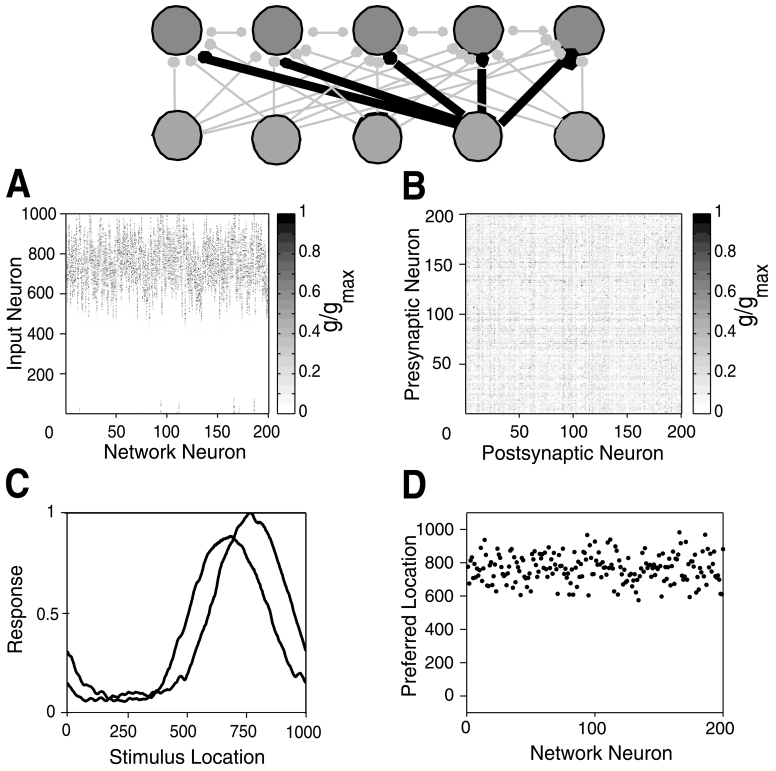

Figure 3. Formation of a Selective Column

In the network figure at the top, the thick black connections represent strong synapses. Other synapses are weak after STDP comes to equilibrium. The resulting configuration makes all of the network neurons sensitive to stimuli that excite the same subset of input neurons. (A) Grayscale plot of the strengths of feedforward synapses between input and network neurons after STDP has come to equilibrium. The *x* value of each point corresponds to the label of the network neuron and the *y* value to the label of the input neuron. The horizontal band reflects the strong synapses illustrated in the network diagram above. (B) Grayscale plot of the strengths of recurrent synapses between network neurons after STDP has come to equilibrium. The *x* value of each point corresponds to the label of the postsynaptic network neuron and the *y* value to the label of the presynaptic network neuron. Recurrent synapses are all weak in this example. (C) Firing rate curves for network neurons 40 and 130, which have similar selectivities as do all network neurons. (D) The preferred stimulus locations for the network neurons all take values near 800.

curves of two representative network neurons (Figure 3C) and in the similar preferred stimulus locations near input location 800 for all the network neurons (Figure 3D). At the final stage, the recurrent synapses between network neurons are all quite weak (Figure 3B). However, as we will see, recurrent synapses play an important role in the formation of the column.

To study the developmental sequence leading to selectivity and column structure in the network neurons, it proved useful to "seed" the selectivity of the network. For Figure 4, network neurons 81 through 120 were given initial feedforward weights (seeded) that made them selective for stimulus locations in the range 401 through 600. The patch in the center of Figure 4A reveals the initial feedforward synaptic strengths that provided this seeding. The strengths of the recurrent synapses were initially set to 0 (Figure 4B), and the unseeded network neurons displayed little initial selectivity. On the other hand, the seeded network neurons were selective for nearby input locations and therefore fired in a correlated manner from the beginning of the simulation.

Figures 4C–4F show snapshots of the strengths of the feedforward (4C and 4E) and recurrent (4D and 4F) synapses during the simulation. Initially, STDP strengthens synapses from the seeded group of network neurons, which fire in a correlated manner, to other network neurons. The horizontal band of increased synaptic strength seen in the middle of Figure 4D represents the synapses made from network neurons 81 through 120 to other neurons in the network. This strengthening precedes the modification of the feedforward synapses seen faintly in Figure 4C. Recurrent synapses are strengthened before feedforward synapses because there is only one strongly correlated group of network neurons, while there are multiple competing correlated groups of input neurons.

As the recurrent synapses grow stronger, the seeded network neurons, which respond directly to input neurons 401 through 600, begin to drive unseeded network neurons. As a result, the unseeded network neurons start to fire slightly after input neurons 401 through 600 are excited by the stimulus. This pattern of firing, input neurons 401 through 600 followed by unseeded network neurons, is exactly the pre- before postsynaptic sequence that causes STDP to strengthen synapses. As a result, synapses from input neurons 401 through 600 to all of the unseeded network neurons become strong (Figure 4E).

The final step in the development of a selective column of network neurons is the weakening of the recurrent synapses. Once feedforward synapses from input neurons 401 through 600 to all the network neurons have strengthened sufficiently, they compete with recurrent synapses for further strengthening by STDP. STDP favors short latency inputs over longer latency inputs (Song et al., 2000), so the feedforward synapses win this competition. Figures 4E and 4F show the final synaptic strengths after the system has reached equilibrium. The horizontal band of strong feedforward synapses seen in Figure 4E indicates that all the neurons now have similar input selectivities. Figure 4F shows that the recurrent synapses end up quite weak.

When the same stimulation is run with $\tau_- = 5\tau_+$, column development occurs in a similar manner, but the equilibrium values of the recurrent synapses are larger than for $\tau_- = \tau_+$, although the pattern of recurrent synapses is fairly random. The stronger final recurrent connection strengths arise because the symmetric compo-

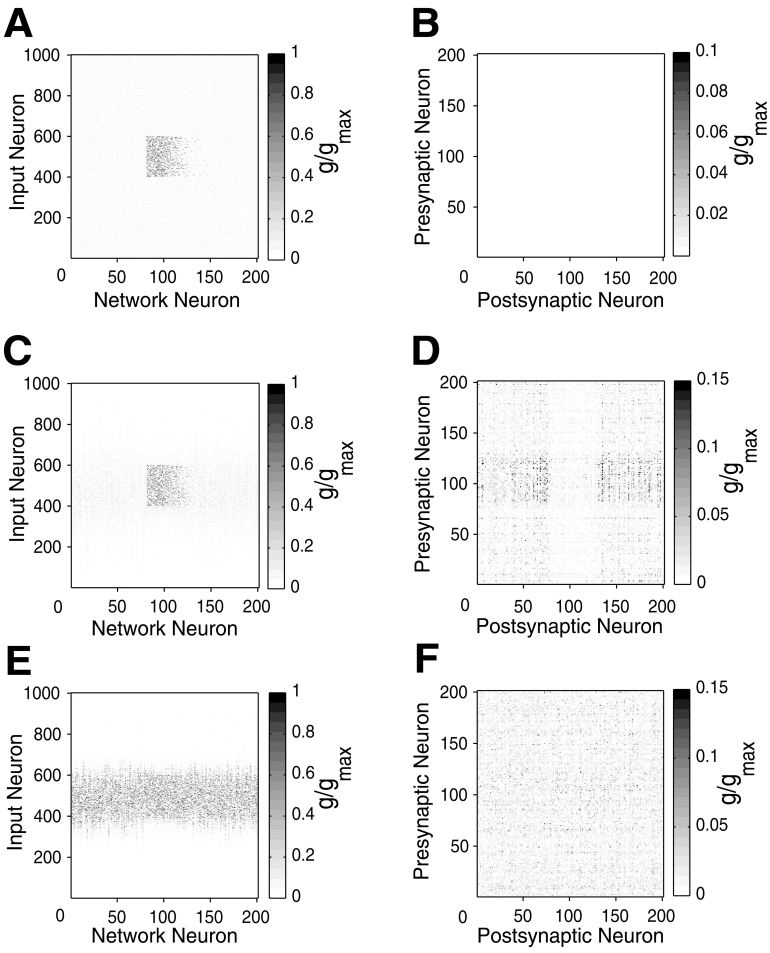

Figure 4. Synaptic Strengths at Various Stages of Column Development

Graphs on the left are for feedforward synapses and on the right are for recurrent synapses. Data are presented in grayscale plots as in Figure 3. (A) Initial feedforward synaptic strengths. The block in the center corresponds to the seed, which makes network neurons 81 through 120 selective for inputs 401 through 600. (B) Initial recurrent synaptic strengths were initially set to 0. (C and D) Feedforward and recurrent synaptic strengths after some time but before equilibrium has been reached. The horizontal band in the center of (D) corresponds to strong synapses formed by the seeded cluster of network neurons onto other network neurons. (E and F) Equilibrium synaptic strengths. A well-formed column selective to inputs 401 through 600 can be seen in (E), and the recurrent synapses shown in (F) have become quite weak.

nent of the input-output correlation function leads to synaptic strengthening in this case. The recurrent connections tend to stabilize clusters of network neurons with their own selectivities, and this makes the columnar structure less tight when $\tau_- = 5\tau_+$ than when $\tau_- = \tau_+$.

If all the network neurons receive direct feedforward input, STDP constructs a column driven either exclusively (for $\tau_- = \tau_+$) or primarily (for $\tau_- = 5\tau_+$) by feedforward input. However, not all cortical neurons receive direct input from the thalamus. To simulate this situation, we removed some of the connections from the input neurons to the network neurons in the model (we also increased the number of network neurons from 200 to 250). Figure 5 shows the results of a simulation with $\tau_- = \tau_+$ in which network neurons 101 through 200 were not connected to the input neurons (Figure 5A). The network neurons that are directly connected to the input neurons develop similar selectivities through the tuning of feedforward synapses, as in the previous simulation. However, the recurrent connections from the feedforward-driven network neurons to network neurons that are not directly connected to the inputs remain strong (Figure 5B), unlike the connections between network neurons receiving feedforward input. The network neurons disconnected from the input neurons can be viewed as an additional layer within the network, so a hierarchical architecture has formed. Network neurons end up with the same selectivity whether or not they receive direct feedforward input (Figure 5C), and the network neurons still form a column with tightly bunched preferred stimulus locations (Figure 5D).

Column formation through STDP can occur over a range of parameter values. The process is most sensitive to B, the ratio of areas under the negative and positive parts of the STDP window function. The value of this ratio must be appropriately adjusted for both the recurrent (B_{recur}) and feedforward (B_{ff}) connections. The simulations produce well-formed columns if $1 < B_{recur} < B_{ff} < 1.07$. Columns with a larger degree of dispersion in their preferred locations arise if $B_{ff} < B_{recur} < B_{ff} + 0.02$. Larger values of B_{recur} lead to the formation of preferences at random locations, as if there were no recurrent connections at all. If $B_{ff} > 1.07$, the development of selectivity is disrupted. The simulations are not too sensitive to the overall size of the conductance changes arising from STDP, set by the parameter A_+. However if this is too large, especially in the case $\tau_- >> \tau_+$, the recurrent synapses can exhibit oscillatory behavior, and the model never stabilizes.

Refinement of Cortical Maps

In our simulations to this point, the competitive nature of STDP leads to a winner-take-all situation that favors the formation of a single column. To create a continuous map of selectivities rather than a single column, we need to restrict the spread of selectivity from one neuron to

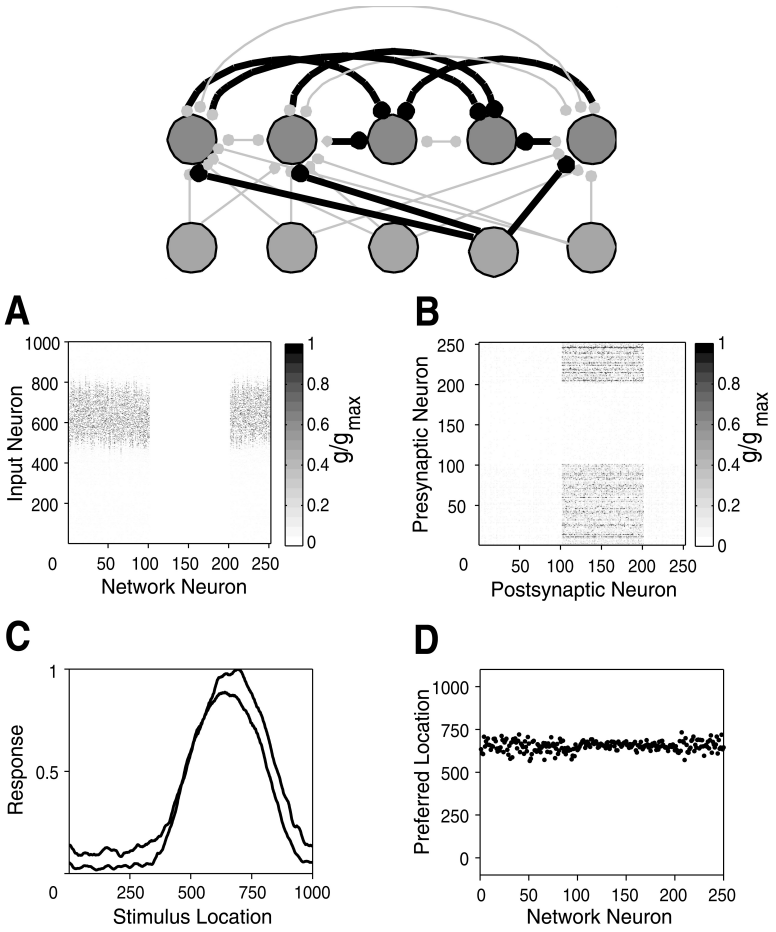

Figure 5. Formation of Multiple Layers

The upper panel shows the equilibrium configuration of synaptic strengths with strong synapses from input neurons 601 through 800 to the network neurons they enervate and strong synapses from these network neurons to other network neurons that do not receive direct feedforward input. (A) Feedforward synaptic strengths at equilibrium in a grayscale plot. The horizontal band indicates the formation of a single column selective to inputs in the neighborhood of 700. The hole in the middle of this band corresponds to the neurons that received no feedforward input. (B) Equilibrium recurrent synaptic strengths in a grayscale plot. The vertical band indicates that strong synapses have formed from network neurons that are driven by feedforward input (network neurons 1 through 100 and 201 through 250) to network neurons that do not receive direct feedforward input (network neurons 101 through 200). (C) Firing rate curves for a feedforward-driven network neuron (neuron 40) and a recurrently driven network neuron (neuron 150) show similar selectivities. (D) Preferred stimulus locations are similar across the network. There is less dispersion in the preferred stimulus locations of the recurrently driven network neurons than the feedforward-driven neurons because the recurrent connections are all-to-all.

another. We do this, in part, by limiting the range of the recurrent connections between network neurons to local neighborhoods, rather than allowing them to be all-to-all (Experimental Procedures). In addition, we choose the parameters B_{recur} and B_{ff} so that feedforward connections tend to dominate recurrent connections. In particular, the values of B_{recur} and B_{ff} must allow for the formation of a column, as described in the previous section. In addition, B_{recur} must be large enough to prevent the recurrents from transferring selectivity from one group of neurons to other neurons with well-established but different selectivities of their own.

Recent experiments have stressed the importance of activity-independent processes in map formation (Crowley and Katz, 1999; Hubener and Bonhoeffer, 1999; Crowley and Katz, 2000). Therefore, we begin by studying the effects of STDP on map development when a coarse map is set up initially in an activity-independent manner (i.e., the map is seeded). Later in this section, we study map formation without any structure imposed initially. In the seeded simulations, a map-like structure was imposed on the network model by setting the initial strengths of the feedforward synapses from the input neurons to the network neurons in the manner indicated by the diagonal stripe of strong feedforward connections in Figure 6A. The crudeness of the map is evident from the width of this stripe and from the weak initial selectivities of the network neurons shown in Figure 6C.

STDP tightens and refines this map. After STDP has equilibrated, the band of strong feedforward synapses is narrower (Figure 6B), and the firing rate curves of the network neurons are sharper, as shown in Figure 6D. The tightness of the final map is primarily determined by the width of the input correlations. The ordered progression of preferred locations of the network neurons seen in Figure 6E reveals the well-defined map and is much tighter than the initial distribution of preferred locations (Figure 6F). The pattern of synaptic strengths seen in Figure 6B is stable. We have simulated many hours of activity and observed no significant changes in the network map.

Simulations performed with $\tau_- = 5\tau_+$ give similar results. The major difference is that some recurrent synapses are retained as in the case of the single column discussed previously (see Figure 8D).

Unseeded Development of Cortical Maps

Although activity-independent processes appear to act early in development to initialize map formation, we have explored whether STDP by itself, without any seeding, can lead to map development. If the network model used to this point is run without initial seeding, a single column structure forms. However, an orderly map can arise solely through STDP from random initial conditions if inhibitory connections are introduced into the network. In a set of simulations, we introduced all-to-all uniform inhibitory interactions of fixed strength between network units, in addition to their plastic local excitatory connec-

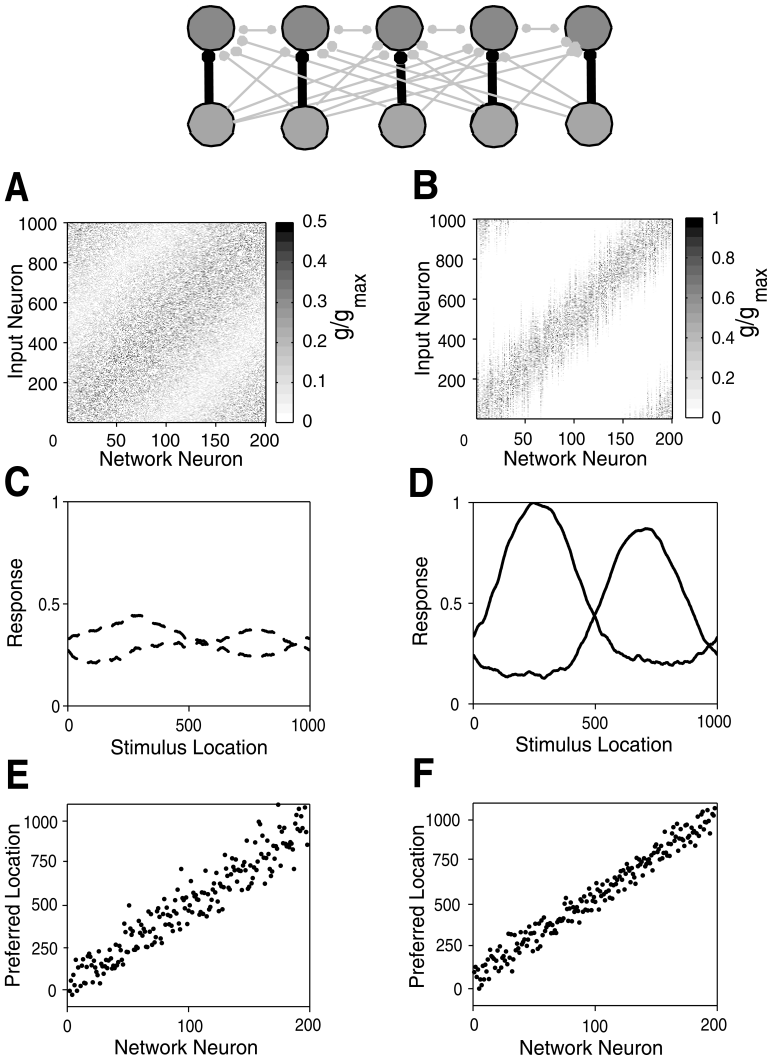

Figure 6. Synaptic Strengths, Firing Rate Curves, and Preferred Stimulus Locations Before and After Map Refinement

The upper panel shows the final stage when strong feedforward synapses form a topographic map from the input neurons to the network neurons. In the remaining panels, graphs on the left are before STDP has been applied, and graphs on the right are after STDP has come to equilibrium. (A) Grayscale plot of the initial feedforward synaptic strengths in an example with initial seeding of the map. The diagonal band of strengthened synapses indicates that they form a rough map. (B) Feedfoward synaptic strengths after STDP has come to equilibrium and refined the map in (A). The refined map structure is visible as a tightened diagonal band of strengthened synapses. (C) Firing rate curves of neurons 40 and 130 before the application of the STDP. The curves are quite wide and shallow, showing that the selectivity is quite weak. (D) Firing rate curves of the same two neurons after STDP have equilibrated showing the increased selectivity. (E and F) Preferred stimulus locations for the network neurons before (E) and after (F) the application of STDP.

tions (see Experimental Procedures). These inhibitory connections tend to make different neurons in the network develop different location preferences, whereas the excitatory recurrent interactions favor similar preferences. When the excitatory connections are restricted to local neighborhoods, these opposing forces can lead to the formation of a smoothly changing cortical map. Short-range excitation and long-range inhibition is important for map formation in previous models as well (reviewed in Miller, 1996). Our model differs from previous work in that the excitatory connections are allowed to be plastic.

The form of the map arising from STDP in the unseeded case depends on the range of the local excitatory connections. For a range of 40 (20 to the left and 20 to the right) to around 100, a single smooth map usually forms in the simulations initialized with random weights. An example of such a map is shown in Figure 7B. This simulation started with the random initial feedforward connection strengths seen in Figure 7A, but nevertheless the final feedforward connection strengths (Figure 7B) exhibit a map-like structure similar to that observed in Figure 6B.

Because the absolute location of the map and its orientation with respect to the network are not determined by the initial conditions, the map that forms in these simulations can be arranged in either direction and can be located at any point across the network. Occasionally, a "double" map can arise, with the stripe pattern wrapping twice around the network (remember that periodic boundary conditions have been imposed). This frequently happens if the range of recurrent connections is less than 40. For a recurrent range of more than 100, a partial map generally forms, with a continuous variation of preferred locations that does not cover the full range of stimulus locations. More complicated patterns can also arise, but the variation of preferred locations is always smooth within these structures. Finally, to illustrate the nature of the activity generated by the network after the map has formed, we show in Figure 7C a raster of the action potentials fired by the network neurons in response to a stimulus that sweeps steadily across the full range of values.

Adult Plasticity of Cortical Maps

Cortical maps can be reorganized in the adult brain as a result of injury or behavioral training. This typically involves normal or highly active regions of a cortical

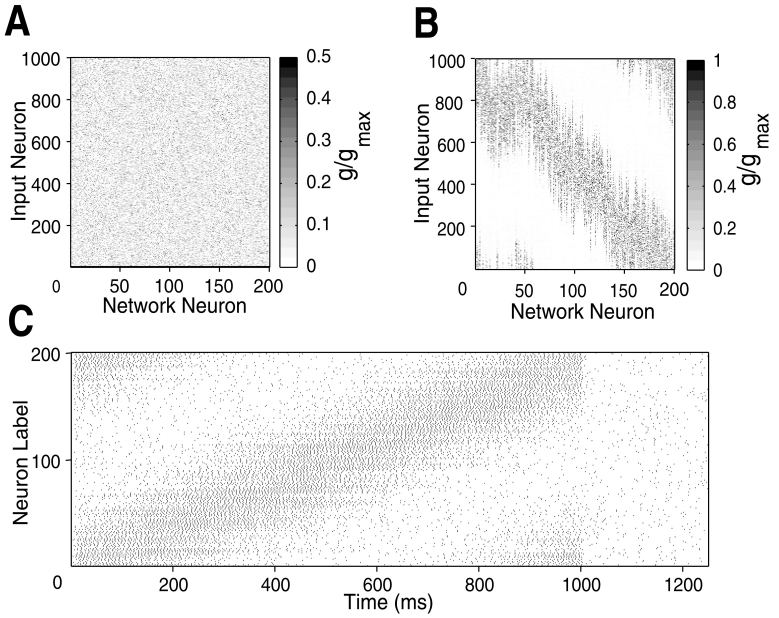

Figure 7. Synaptic Strengths Before and After Unseeded Map Development in a Network with Additional All-to-All Recurrent Inhibitory Connections

(A) Grayscale plot of the initial feedforward synaptic strengths showing the lack of seeding. Synaptic strengths are random.

(B) Feedforward synaptic strengths after STDP have come to equilibrium. The diagonal band of strong synapses reveals the formation of a continuous map from the random initial condition in (A).

(C) A raster showing the activity of all the network neurons in response to a stimulus that sweeps across the map at a steady rate and then terminates at 1000 s. Each dot represents an action potential from the corresponding neuron. The effect of the periodic boundary conditions is apparent in the network activity.

map expanding into regions that have lesioned inputs or that are less highly activated during training (reviewed in Buonomano and Merzenich, 1998). Lesion-induced plasticity can arise from STDP (Feldman, 2000). We have simulated this situation by removing all the feedforward connections from a subset of the input neurons to network neurons that have already formed a map as in the previous section. Figure 8A shows the strengths of the feedforward synapses immediately after the lesion. The hole in the middle of the band of strong synapses reflects the removal of all feedforward connections made from inputs 301 through 700. The synaptic modifications induced by this lesion follow a progression similar to that shown in Figure 4 for the formation of a single column. As noted previously, feedforward inputs are favored by STDP over recurrent synapses. However, when the feedforward inputs are lesioned, recurrent connections become the primary source of selectivity. Recurrent connections to the network neurons with lesioned inputs grow in strength when their normal competitors, the feedforward inputs, are removed. In this way, the neurons with lesioned inputs adopt the selectivity of neighboring network neurons with intact feedforward inputs.

The strengthening of recurrent synapses is only the first of the changes induced by an input lesion to the network. As in the case of a single column, the strengthened recurrent connections drive the network in a way that favors the strengthening of feedforward synapses to the network neurons with lesioned inputs from input neurons that did not initially form strong connections to them. Figure 8B shows the strengths of the feedforward synapses after the system has reached equilibrium for the case $\tau_- = \tau_+$. Figure 8E gives the corresponding preferred stimulus locations for the network neurons. Neurons 51 through 150, which were previously responsive to input locations 301 through 700, have now acquired input connections similar to those of neurons 151 through 200 (Figure 8B). Their preferred stimulus locations have also shifted to around 800 (Figure 8E), so these network neurons have developed selectivities similar to their neighboring neurons.

Figures 8C and 8D show the final synaptic strengths for a simulation done with $\tau_- = 5\tau_+$. In this case, network neurons 101 through 150 took on the same selectivities as neurons 151 through 200, and network neurons 51 through 100 acquired the selectivities of neurons 1 through 50 (Figure 8F). In simulations done with $\tau_- = \tau_+$, connection and selectivity patterns similar to those in Figures 8C and 8F were observed for a while, but over time network neurons 51 through 100 switched to acquire a selectivity around 800, similar to the tuning of network neurons 101 through 150 (Figure 8E). This process is slow because of competition between two correlated groups, neurons 1 through 50 and neurons 151 through 200. The stronger recurrent synapses allowed by STDP with $\tau_- = 5\tau_+$ (seen in Figure 8D) stabilize the split remapping observed in Figures 8C and 8F.

The parameter range for which remapping following a lesion occurs is similar to the range where maps form, except that if B_{recur} is too large, the lesioned area fails to acquire the selectivities of neighboring areas and remains unresponsive.

Discussion

Correlation-based synaptic modification has proved valuable for the study of developmental and adult plasticity in cortical maps, but it suffers from a number of problems. Synapses are strengthened whenever coincident pre- and postsynaptic activity occurs, which could happen by chance rather than reflecting a causal relationship. This can lead to the nonselective strengthening of all synapses, which is clearly an undesirable outcome. Furthermore, correlation-based synaptic plasticity is not by itself competitive, so additional constraints or plasticity mechanisms must be imposed. STDP can solve both of these problems.

The temporal asymmetry of STDP with respect to spike timing provides a mechanism for transferring selectivity across a network. Any group of neurons that becomes selective to a particular set of inputs and be-

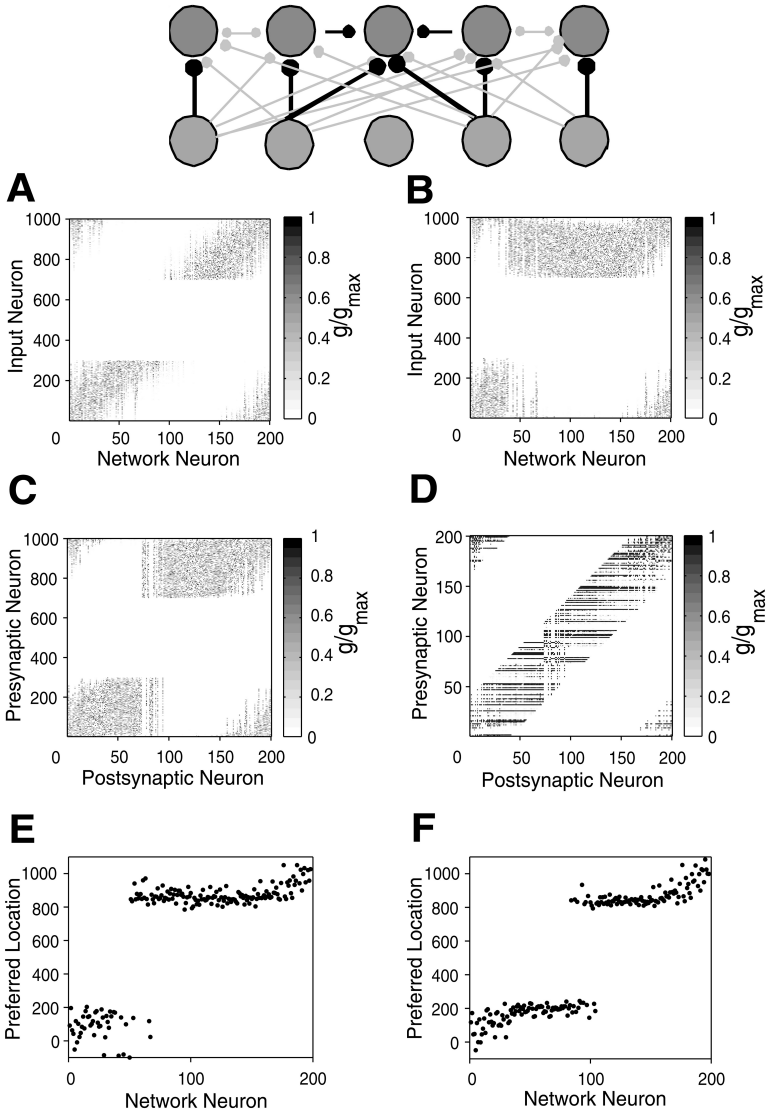

Figure 8. Synaptic Strengths and Preferred Stimulus Locations After Lesioning by Disconnecting Feedforward Synapses from Inputs 301 through 700

The upper panel shows the network after STDP has come to equilibrium. The network neurons represented by the center circle in the upper row, which have lost their feedforward input, are partially driven by neighboring network neurons and also receive strengthened synapses from input neurons surrounding the lesion. (A) Feedforward synaptic strengths immediately after the lesion. The diagonal band shows the preexisting map, and the hole in the middle reflects the lesioning of the feedforward synapses from input neurons 301 through 700. (B) Feedforward synaptic strengths after STDP with $\tau_- = \tau_+$ has reequilibrated following the lesion. The network neurons with lesioned inputs now receive strong feedforward synapses from input neurons 701 through 1000. (C and D) Feedforward and recurrent synaptic strengths after STDP with $\tau_- = 5\tau_+$ has reequilibrated following the lesion. The network neurons with lesioned inputs receive strong feedforward synapses from input neurons 1 through 300 and 701 through 1000 (C), and the recurrents retain a nonzero strength (D). (E and F) Preferred stimulus locations of the network neurons following recovery from the lesion for $\tau_- = \tau_+$ (E) and $\tau_- = 5\tau_+$ (F).

gins to fire in a correlated manner can direct the development of selectivity in other network neurons. The selective "teacher" group drives nonselective "student" neurons through strengthened recurrent connections, causing them to become similarly selective. If no feedforward input is available to the student neurons, the process stops there, but if feedforward input is available, the teacher neurons provide an instructive signal that induces selective strengthening of feedforward synapses. This process terminates with the weakening of the instructive recurrent synapses, leading to a stable column or map of selective neurons. In agreement with this scheme, it has been suggested that structure within primary visual cortex precedes and guides the development of thalamocortical inputs (Ruthazer and Stryker, 1996; Crair et al., 1997, 1998; Trachtenberg et al., 2000).

In the case of an input lesion, the process outlined in the previous paragraph repeats. First, recurrent synapses from nearby network neurons strengthen to provide selectivity to the network neurons with lesioned inputs. The strengthened recurrents provide an instructive signal that guides the strengthening of new feedforward inputs to the network neurons with lesioned inputs. This sequence appears to match that seen in animal studies where evidence suggests that thalamocortical organization is guided by earlier intracortical changes. For example, the initial reorganization in rat barrel cortex following whisker clipping appears to involve the potentiation of intracortical synapses, while later modifications affect thalamic afferents (Diamond et al., 1993; Armstrong-James et al., 1994; Glazewski and Fox, 1996).

The most distinctive predictions of a model that uses STDP for developmental and adult plasticity concern the time course of plastic changes. Loss of input correlation (as opposed to loss of activity) leads to a rapid decrease in synaptic strength in these models. In the absence of competition from a feedforward source of input, STDP will potentiate the most correlated set of intracortical inputs to a given neuron. However, once these inputs are strengthened, they can act as a training signal, allowing feedforward synapses to strengthen. The shorter latency of feedforward over recurrent inputs

leads to their ultimate dominance. Thus, a signature of STDP-based plasticity is a transient increase in the strength of intracortical pathways that are then weakened once strong feedforward pathways become established.

The timescale of synaptic weakening under STDP varies for different preparations. STDP with $\tau_- = \tau_+$ favors feedforward architectures when direct feedforward input is available but promotes recurrent network connections when feedforward input is absent while discouraging the formation of strong recurrent loops. STDP with $\tau_- >> \tau_+$ allows stronger recurrent connections to persist, which could be advantageous in some situations. A large value of τ_- also ensures sensitivity of STDP to groups of inputs with long correlation times. These differences suggest that regional and developmental variabilities in STDP timing properties may reflect different functional roles for synaptic plasticity.

Experimental Procedures

The integrate-and-fire neuron used in the single neuron and network models follows a standard implementation (Troyer and Miller, 1997). The membrane potential obeys

$$\tau_m \frac{dV}{dt} = V_{rest} - V + g_{ex}(t)(E_{ex} - V).$$

with $\tau_m = 20$ ms, $V_{rest} = -74$mV, and $E_{ex} = 0$mV. In addition, the neuron fires an action potential when the membrane potential reaches a threshold value of -54mV, and the membrane potential is then reset to -60mV. Synaptic inputs are modeled as conductance changes with instantaneous rise times and exponential decays, so that a single presynaptic action potential at time 0 generates a synaptic conductance $g_{ex}(t) = g\ exp(-t/\tau_{ex})$, with $\tau_{ex} = 5$ ms. The synaptic conductances, including the peak conductance parameter g, are measured in units of the leakage conductance of the neuron and are thus dimensionless.

For all the simulations shown in Figure 1, we used $A_+ = 0.005$, and the presynaptic firing rates were generated randomly from Gaussian distributions. We divide time into intervals chosen from an exponential distribution with mean interval equal to the correlation time being investigated. At the start of each interval, we generate $N + 1$ random numbers y and x_a for $a = 1, 2, ..., N$ from Gaussian distributions with mean 0 and standard deviation 1. The firing rate of input neuron a is then set to $10(1 + 0.3x_a + 0.3y)$ Hz, if the neuron belongs to a correlated cluster, and $10(1 + 0.3\sqrt{2}x_a)$ Hz, if it belongs to an uncorrelated cluster. These rates are held constant until the start of next interval. When two correlated clusters are used (Figure 1D), we choose independent interval start times and independent x_a and y values for the two groups, so they were uncorrelated with each other. We used $g_{max} = 0.015$ and $B = 1.05$ for the simulations shown in Figure 1.

The input neurons of the network model generate spikes through a Poisson process with a time-varying firing rate. Periodic boundary conditions are imposed to avoid edge effects. The firing rate of input neuron a in response to a stimulus at location s is

$$r_a = R_0 + R_1(e^{-(s-a)^2/2\sigma^2} + e^{-(s+1000-a)^2/2\sigma^2} + e^{-(s-1000-a)^2/2\sigma^2}),$$

with $R_0 = 10$ Hz, $R_1 = 80$ Hz, and $\sigma = 100$. The latter two Gaussian functions are used to give approximate periodicity. The firing rates of the input neurons are generated over variable time intervals chosen from an exponential distribution with a mean of 20 ms (or, for some simulations, 100 ms). At the start of each interval, a stimulus location s is chosen randomly between 1 and 1000 from a uniform distribution, and the input firing rates computed from the above equation are then held constant until the start of the next interval.

The network neurons are described by the integrate-and-fire model (see above) and driven by excitatory conductances evoked by the spikes of the input neurons. Each network neuron receives an additional background Poisson input at 500 Hz through a synapse of strength 0.096, except in Figure 5, where the strength was set to 0.116 to give the network neurons a higher initial firing rate. This single input represents the effects of a large number of afferents. The amount of background input is not critical as long as the network neurons keep firing.

In the network model, we used $g_{max} = 0.02$. The initial feedforward connection strengths for the maps we discussed were set to

$$0.5g_{max}\ \exp\left(-\frac{1}{2}\left(\frac{d}{100}\right)^2\right),$$

where, for the synapse from input neuron i to network neuron j, $d = i/5 - j$. If $d > 100$, it was set to $200 - d$, and if $d < -100$, it was set to $200 + d$. This was done to satisfy the periodicity condition. For the examples involving maps, recurrent connections between network neurons were limited in range so that network neuron j only connected to other network neurons in the range $j - 40$ to $j + 40$, modulo the periodic boundary conditions imposed on the network (except where this range was varied to study the consequences). For the simulations with all-to-all inhibitory connections, the strength of the inhibitory connections was 0.3 g_{max}, and the initial strengths of the excitatory feedforward connections were chosen randomly and uniformly between 0 and 0.5 g_{max}. The strengths of the excitatory recurrent connections were set initially to 0.

For the network simulations, we used $A_+ = 0.001$. For $\tau_- = \tau_+$, we used $B_{ff} = 1.06$ and $B_{recur} = 1.04$. We used $B_{ff} = 1.15$ and $B_{recur} = 1.13$ for simulations with $\tau_- = 5\tau_+$.

To generate the response firing rate curves for the network neurons, we ramped the stimulus from 0 to 1000. The response curve was calculated by counting spikes over 1000 repetitions, and the resulting curve was smoothed by averaging over 20 neighboring points. Smoothing did not affect the general shape of the curve, but it removed sampling noise. Periodic boundary conditions were enforced when smoothing. The preferred stimulus location is defined as the peak of the resulting firing rate curve.

Acknowledgments

Research supported by the National Institute of Mental Health (MH58754), the Sloan Center for Theoretical Neurobiology at Brandeis University, the W.M. Keck Foundation, and a Howard Hughes Predoctoral Fellowship to S.S. We thank Sacha Nelson and Jesper Sjöström for useful discussions.

Received March 9, 2001; revised August 17, 2001.

References

Abraham, W.C. (1997). Metaplasticity, a new vista across the field of synaptic plasticity. Prog. Neurobiol. *52*, 303–323.

Armstrong-James, M., Diamond, M.E., and Ebner, F.F. (1994). An innocuous bias in whisker use in adult rats modifies receptive fields of barrel cortex neurons. J. Neurosci. *14*, 6978–6991.

Bell, C., Han, V., Sugawara, Y., and Grant, K. (1997). Synaptic plasticity in a cerebellum-like structure depends on temporal order. Nature *387*, 278–281.

Benuskov, L., Diamond, M., and Ebner, F. (1994). Dynamic synaptic modification threshold, computational model of experience-dependent plasticity in adult rat barrel cortex. Proc. Natl. Acad. Sci. USA *91*, 4791–4799.

Bi, G.-Q., and Poo, M.-M. (1998). Activity-induced synaptic modifications in hippocampal culture, dependence on spike timing, synaptic strength and cell type. J. Neurosci. *18*, 10464–10472.

Bienenstock, E.L., Cooper, L.N., and Munro, P.W. (1982). Theory for the development of neuron selectivity, orientation specificity and binocular interaction in visual cortex. J. Neurosci. *2*, 32–48.

Buonomano, D., and Merzenich, M. (1998). Cortical plasticity: from synapses to maps. Ann. Rev. Neurosci. *21*, 149–186.

Crair, M., Ruthazer, E., Gillespie, D., and Stryker, M. (1997). Relationship between the ocular dominance and orientation maps in visual cortex of monocularly deprived cats. Neuron *19*, 307–318.

Crair, M., Gillespie, D., and Stryker, M. (1998). The role of visual experience in the development of columns in cat visual cortex. Science *279*, 566–570.

Crowley, J.C., and Katz, L.C. (1999). Development of ocular dominance columns in the absence of retinal input. Nat. Neurosci. *2*, 1125–1130.

Crowley, J.C., and Katz, L.C. (2000). Early development of ocular dominance columns. Science *290*, 1321–1324.

Debanne, D., Gahwiler, B.H., and Thompson, S.M. (1994). Asynchronous pre- and postsynaptic activity induces associative long-term depression in area CA1 of the rat hippocampus in vitro. Proc. Natl. Acad. Sci. USA *91*, 1148–1152.

Debanne, D., Gahwiler, B., and Thompson, S. (1998). Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. J. Physiol. *507*, 237–247.

Diamond, M.E., Armstrong-James, M., and Ebner, F.F. (1993). Experience-dependent plasticity in adult rate barrel cortex. Proc. Natl. Acad. Sci. USA *90*, 2082–2086.

Egger, V., Feldemeyer, D., and Sakmann, B. (1999). Coincidence detection and efficacy changes in synaptic connections between spiny stellate neurons of the rat barrel cortex. Nat. Neurosci. *2*, 1098–1105.

Feldman, D. (2000). Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. Neuron *27*, 45–46.

Gilbert, C. (1996). Plasticity in visual perception and physiology. Curr. Opin. Neurobiol. *6*, 269–274.

Glazewski, S., and Fox, K. (1996). Time course of experience-dependent synaptic potentiation and depression in barrel cortex of adult rats. J. Neurophysiol. *75*, 1714–1729.

Grajski, K., and Merzenich, M. (1990). Hebb-type dynamics is sufficient to account for the inverse magnification rule in cortical somatotopy. Neural Comp. *2*, 71–84.

Gustafsson, B., Wigstrom, H., Abraham, W.C., and Huang, Y.-Y. (1987). Long-term potentiation in the hippocampus using depolarizing current pulses as the conditioning stimulus to single volley synaptic potentials. J. Neurosci. *7*, 774–780.

Hubener, M., and Bonhoeffer, T. (1999). Eyes wide shut. Nat. Neurosci. *2*, 1043–1045.

Kaas, J. (1991). Plasticity of sensory and motor maps in adult mammals. Annu. Rev. Neurosci. *14*, 137–167.

Kaas, J., Krubitzer, L., Chino, Y., Langston, A., Polley, E., and Blair, N. (1990). Reorganization of retinotopic cortical maps in adult mammals after lesions of the retina. Science *248*, 229–231.

Katz, L., and Shatz, C. (1996). Synaptic activity and the construction of cortical circuits. Science *274*, 1133–1138.

Kempter, R., Gerstner, W., and van Hemmen, J. (1999). Hebbian learning and spiking neurons. Phys. Rev. E *59*, 4498–4514.

Levy, W.B., and Steward, D. (1983). Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. Neuroscience *8*, 791–797.

Magee, J.C., and Johnston, D. (1997). A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. Science *275*, 209–213.

Markram, H., Lubke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science *275*, 213–215.

Merzenich, M., Kaas, J., Wall, J., Nelson, R., Sur, M., and Felleman, D. (1983). Topographic reorganization of somatosensory cortical areas 3b and 1 in adult monkeys following restricted deafferentation. Neuroscience *8*, 33–55.

Merzenich, M., Nelson, R., Stryker, M., Cynader, M., Schoppmann, A., and Zook, J. (1984). Somatosensory map changes following digit amputation in adult monkeys. J. Comp. Neurol. *224*, 591–605.

Miller, K. (1996). Receptive fields and maps in the visual cortex: models of ocular dominance and orientation columns. In Models of Neural Networks, III, E. Domany, J. van Hemmen, and K. Schulten, eds. (New York: Springer-Verlag), pp. 55–78.

Miller, K.D., and MacKay, D.J.C. (1994). The role of constraints in Hebbian learning. Neural Comp. *6*, 100–126.

O'Leary, D., Ruff, N., and Dyck, R. (1994). Developmental, critical period plasticity, and adult reorganization of mammalian somatosensory systems. Curr. Opin. Neurobiol. *4*, 535–544.

Purves, D., and Lichtman, J. (1985). Principles of Neural Development (Sunderland, MA: Sinauer Associates).

Ruthazer, E., and Stryker, M. (1996). The role of activity in the development of long-range horizontal connections in area 17 of the ferret. J. Neurosci. *16*, 7253–7269.

Song, S., Miller, K., and Abbott, L. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. Nat. Neurosci. *3*, 919–926.

Stryker, M. (1986). The role of neural activity in rearranging connections in the central visual system. In The Biology of Change in Otolaryngology, R.J. Ruben, T.R. Van De Water, and E.W. Rubel, eds. (Amsterdam: Elsevier), pp. 211–224.

Sutton, G., Reggia, J., Armentrout, S., and D'Autrechy, C. (1994). Cortical map reorganization as a competitive process. Neural Comp. *6*, 1–13.

Trachtenberg, J.T., Trepel, C., and Stryker, M.P. (2000). Rapid extragranular plasticity in the absence of thalamocortical plasticity in the developing visual cortex. Science *287*, 2029–2032.

Troyer, T.W., and Miller, K.D. (1997). Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. Neural Comp. *9*, 971–983.

Wall, J. (1988). Variable organization in cortical maps of the skin as an indication of the lifelong adaptive capacities of circuits in the mammalian brain. Trends Neurosci. *11*, 549–557.

Weinberger, N. (1995). Dynamic regulation of receptive fields and maps in the adult sensory cortex. Annu. Rev. Neurosci. *19*, 129–158.

Wong, R.O. (1999). Retinal waves and visual system development. Ann. Rev. Neurosci. *22*, 29–47.

Yuste, R., and Sur, M. (1999). Development and plasticity of the cerebral cortex: from molecules to maps. J. Neurobiol. *41*, 1–6.

Zhang, L., Tao, H., Holt, C., Harris, W., and Poo, M.-M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. Nature *95*, 37–44.