Predictive learning of temporal sequences in recurrent neocortical circuits

Rajesh P. N. Rao*[1] and Terrence J. Sejnowski†‡[2]

*Sloan Center for Theoretical Neurobiology, The Salk Institute for Biological Studies, La Jolla, CA 92037, †Howard Hughes Medical Institute, The Salk Institute for Biological Studies, La Jolla, CA 92037 and ‡Department of Biology, University of California at San Diego, La Jolla, CA 92037, USA

Abstract. When a spike is initiated near the soma of a cortical pyramidal neuron, it may back-propagate up dendrites toward distal synapses, where strong depolarization can trigger spike-timing dependent Hebbian plasticity at recently activated synapses. We show that (a) these mechanisms can implement a temporal-difference algorithm for sequence learning, and (b) a population of recurrently connected neurons with this form of synaptic plasticity can learn to predict spatiotemporal input patterns. Using biophysical simulations, we demonstrate that a network of cortical neurons can develop direction selectivity similar to that observed in complex cells in alert monkey visual cortex as a consequence of learning to predict moving stimuli.

2001 Complexity in biological information processing. Wiley, Chichester (Novartis Foundation Symposium 239) p 208–233

Neocortical circuits are dominated by massive excitatory feedback: more than 80% of the synapses made by excitatory cortical neurons are onto other excitatory cortical neurons (Douglas et al 1995, Braitenberg & Schüz 1991). Why is there such massive recurrent excitation in the neocortex and what is its role in cortical computation? Previous modelling studies have suggested a role for excitatory feedback in amplifying feedforward inputs (Douglas et al 1995, Suarez et al 1995, Mineiro & Zipser 1998, Ben-Yishai et al 1995, Somers et al 1995, Chance et al 1999). Recently, it has been shown that recurrent excitatory connections between cortical neurons are modified according to a spike-timing dependent Hebbian learning rule: synapses that are activated slightly before the cell fires are

[1]Present address: Department of Computer Science & Engineering, University of Washington, 114 Sieg Hall, Seattle, WA 98195-2350, USA.

[2]This paper was presented at the Symposium by Terrence J. Sejnowski, to whom correspondence should be addressed.

strengthened whereas those that are activated slightly after are weakened (Markram et al 97) (see also Levy & Steward 1983, Zhang et al 1998, Bi & Poo 1998, Abbott & Blum 1996, Gerstner et al 1996, Senn 1997). Information regarding the postsynaptic activity of the cell is conveyed back to the dendritic locations of synapses by backpropagating action potentials from the soma (Stuart & Sakmann 1994).

Because these recurrent feedback connections can adapt in a temporally specific manner, they may subserve a more general function than amplification, such as the prediction and generation of temporal sequences (Abbott & Blum 1996, Minai & Levy 1993, Montague & Sejnowski 1994, Schultz et al 1997, Softky 1996, Koch 1999, Rao & Ballard 1997). The observation that recurrence can generate sequences has its roots in dynamical systems theory (Scheinerman 1995) and forms the basis of numerous engineering (Kalman 1960) and neural network (Minai & Levy 1993, Rao & Ballard 1997, Jordan 1986, Elman 1990) models for predicting and tracking input sequences. Consider the network of excitatory neurons shown in Fig. 1A. By appropriately learning its recurrent connections, the network can generate sequences of outputs in anticipation of its inputs as depicted in Fig. 1B. The initial activation of a subset of input neurons causes the corresponding set of excitatory neurons to be activated, which in turn activate a different set of excitatory neurons and so on, such that each set of active neurons at a given time step represents the anticipated input at that time step (active neurons are represented as shaded circles in Fig. 1B). The predicted outputs occur just in time to inhibit the input neurons if the external input is excitatory, or excite them if the external input is inhibitory, thereby implementing a stable negative feedback loop and allowing only the unpredicted part of the input to be conveyed to the prediction neurons. Such a model is consistent with some recent ideas regarding cortico-cortical feedback loops (Rao & Ballard 1997, Mumford 1994), predictive coding (Rao & Ballard 1999, Barlow 1998, Daugman & Downing 1995) and visual receptive field development from natural images (Rao & Ballard 1997, Olshausen & Field 1997). In these models, feedback connections from a higher to a lower order cortical area are posited to carry predictions of lower level neural activity, while the feedforward connections are assumed to convey the residual errors in prediction. These errors are used to correct the neural representation at the higher level before generating a subsequent prediction (for example, see Rao & Ballard 1997). Note that for clarity, Fig. 1B shows two different sets of excitatory neurons firing at the two successive time steps, but the model allows arbitrary overlapping subsets of neurons to fire in order to represent temporal sequences with possible overlapping inputs, resulting in sustained firing in some neurons and transient firing in others due to learned recurrent connections.

In this study, we have modelled spike-timing dependent Hebbian synaptic plasticity as a form of 'temporal-difference' learning (Montague & Sejnowski

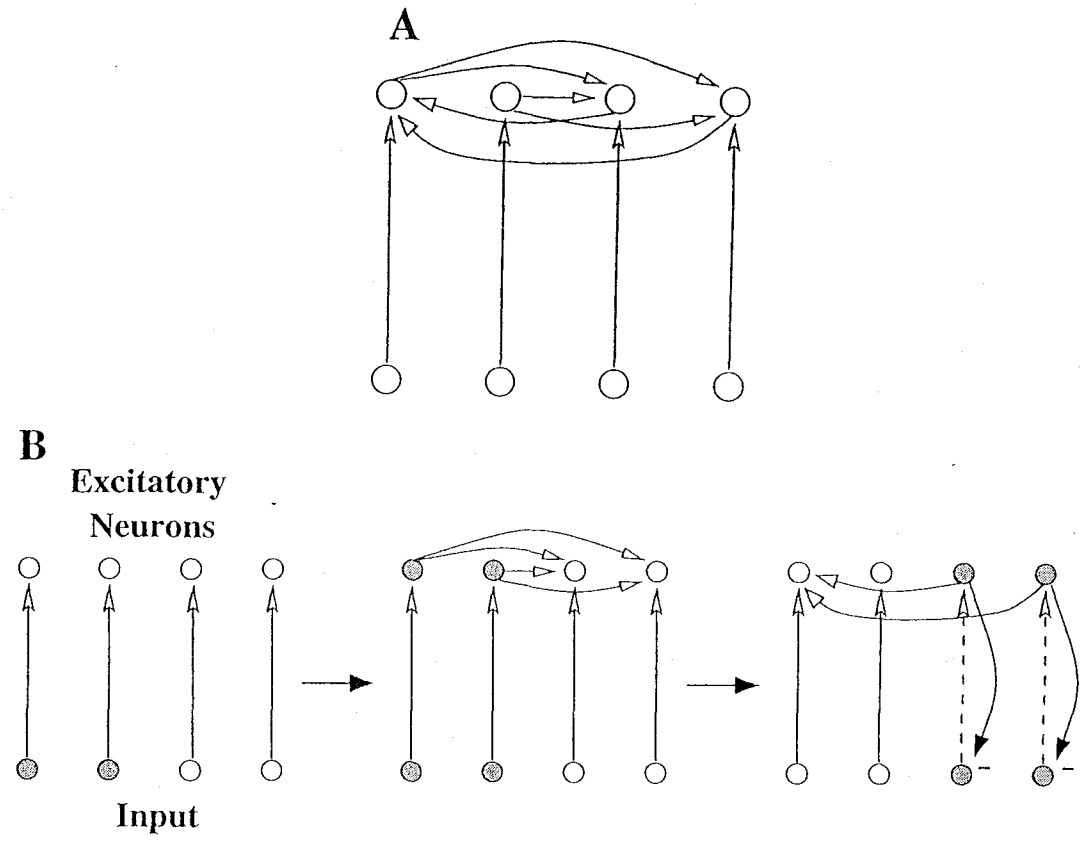

FIG. 1. Prediction using recurrent excitation. (A) An example of a model network of recurrently connected excitatory neurons receiving inputs from a set of input neurons (bottom row). (B) The activation of a subset of input neurons (shaded circles) causes a subset of excitatory neurons to fire which in turn cause a different subset of excitatory neurons to fire due to recurrent excitatory connections. If these recurrent connections are appropriately learned, the second subset of neurons will fire slightly before the expected activation of their corresponding input neurons, allowing inhibition of the inputs and forming a stable negative feedback loop. For clarity, the example shows two different sets of excitatory neurons firing at the two successive time steps, but the learning algorithm allows arbitrary overlapping subsets of neurons to fire in order to represent sequences with possible overlapping inputs, resulting in sustained firing in some neurons and transient firing in others due to the learned recurrent connections.

1994, Schultz et al 1997, Sutton 1988). We have simulated recurrent networks of excitatory and inhibitory cortical neurons possessing this form of synaptic plasticity and have investigated the ability of such networks to learn predictive models of input sequences, focusing in particular on moving stimuli. Detailed compartmental models take into account the temporal dynamics of signal processing in dendrites and the relative timing of spikes in neural populations. Both of these properties were found to be essential in explaining the genesis of complex cell-like direction selectivity in model neocortical neurons.

Results

Spike-timing dependent Hebbian plasticity as temporal-difference learning

To accurately predict input sequences, the recurrent excitatory connections between a given set of neurons need to be adjusted such that the appropriate set of neurons are activated at each time step. This can be achieved by using a 'temporal-difference' learning rule (Montague & Sejnowski 1994, Schultz et al 1997, Sutton 1988). In this paradigm of synaptic plasticity, an activated synapse is strengthened or weakened based on whether the difference between two temporally separated predictions is positive or negative. This minimizes the errors in prediction by ensuring that the prediction generated by the neuron after synaptic modification is closer to the desired value than before (see Methods for more details).

In order to ascertain whether spike-timing dependent Hebbian learning in cortical neurons can be interpreted as a form of temporal-difference learning, we used a two-compartment model of a cortical neuron consisting of a dendrite and a soma-axon compartment. The compartmental model was based on a previous study that demonstrated the ability of such a model to reproduce a range of cortical response properties (Mainen & Sejnowski 1996). Figures 2A and 2B illustrate the responses of the model neuron to constant current pulse injection into the soma and random excitatory and inhibitory Poisson-distributed synaptic inputs to the dendrite respectively (see Methods). The presence of voltage-activated sodium channels in the dendrite allowed backpropagation of action potentials from the soma into the dendrite as shown in Fig. 2C.

To study synaptic plasticity in the model, excitatory postsynaptic potentials (EPSPs) were elicited at different time delays with respect to postsynaptic spiking by presynaptic activation of a single excitatory synapse located on the dendrite. Synaptic currents were calculated using a kinetic model of synaptic transmission (Destexhe et al 1997) with model parameters fitted to whole-cell recorded AMPA (α-amino-3-hydroxy-5-methyl-4-isoxazole proprionic acid) currents (see Methods for more details). Other inputs representing background activity were modelled as sub-threshold excitatory and inhibitory Poisson processes with a mean firing rate of 3 Hz. Synaptic plasticity was simulated by incrementing or decrementing the value for maximal synaptic conductance by an amount proportional to the temporal-difference in the postsynaptic membrane potential at time instants $t + \Delta t$ and $t - \Delta t$ for presynaptic activation at time t (see Methods). The delay parameter Δt was set to 5 ms for these simulations; similar results were obtained for other values in the 5–15 ms range. Presynaptic input to the model neuron was paired with postsynaptic spiking by injecting a depolarizing current pulse (10 ms, 200 pA) into the soma. Changes in synaptic efficacy were monitored by applying a test stimulus before and after pairing, and recording the EPSP evoked by the test stimulus.

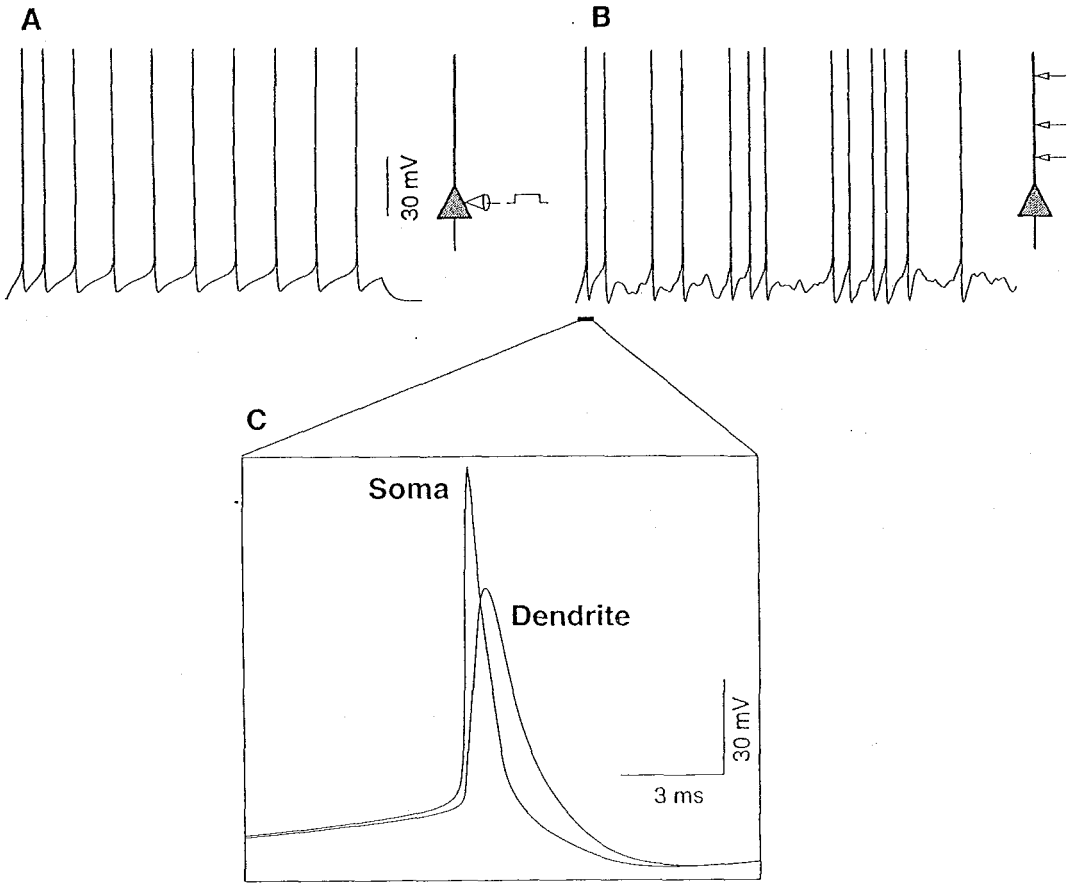

FIG. 2. Model neuron response properties. (A) Response of a model neuron to a 70 pA current pulse injection into the soma for 900 ms. (B) Response of the same model neuron to Poisson distributed excitatory and inhibitory synaptic inputs at random locations on the dendrite. (C) Example of a backpropagating action potential in the dendrite of the model neuron as compared to the corresponding action potential in the soma (enlarged from the initial portion of the trace in [B]).

Figure 3A shows the results of pairings in which the postsynaptic spike was triggered 5 ms after and 5 ms before the onset of the EPSP, respectively. While the peak EPSP amplitude was increased 58.5% in the former case, it was decreased 49.4% in the latter case, qualitatively similar to experimental observations (Markram et al 1997). As mentioned above, such changes in synaptic efficacy in the model are determined by the temporal-difference in the dendritic membrane potential at time instants $t+\Delta t$ and $t-\Delta t$ for presynaptic activation at time t: the difference is positive when presynaptic activation occurs a few milliseconds before a backpropagating action potential invades the dendrite and negative when it occurs slightly after, causing respectively an increase or decrease in synaptic efficacy. The critical window for synaptic modifications in the model depends on the parameter Δt as well as the shape of the backpropagating action potential. This window of plasticity was examined by

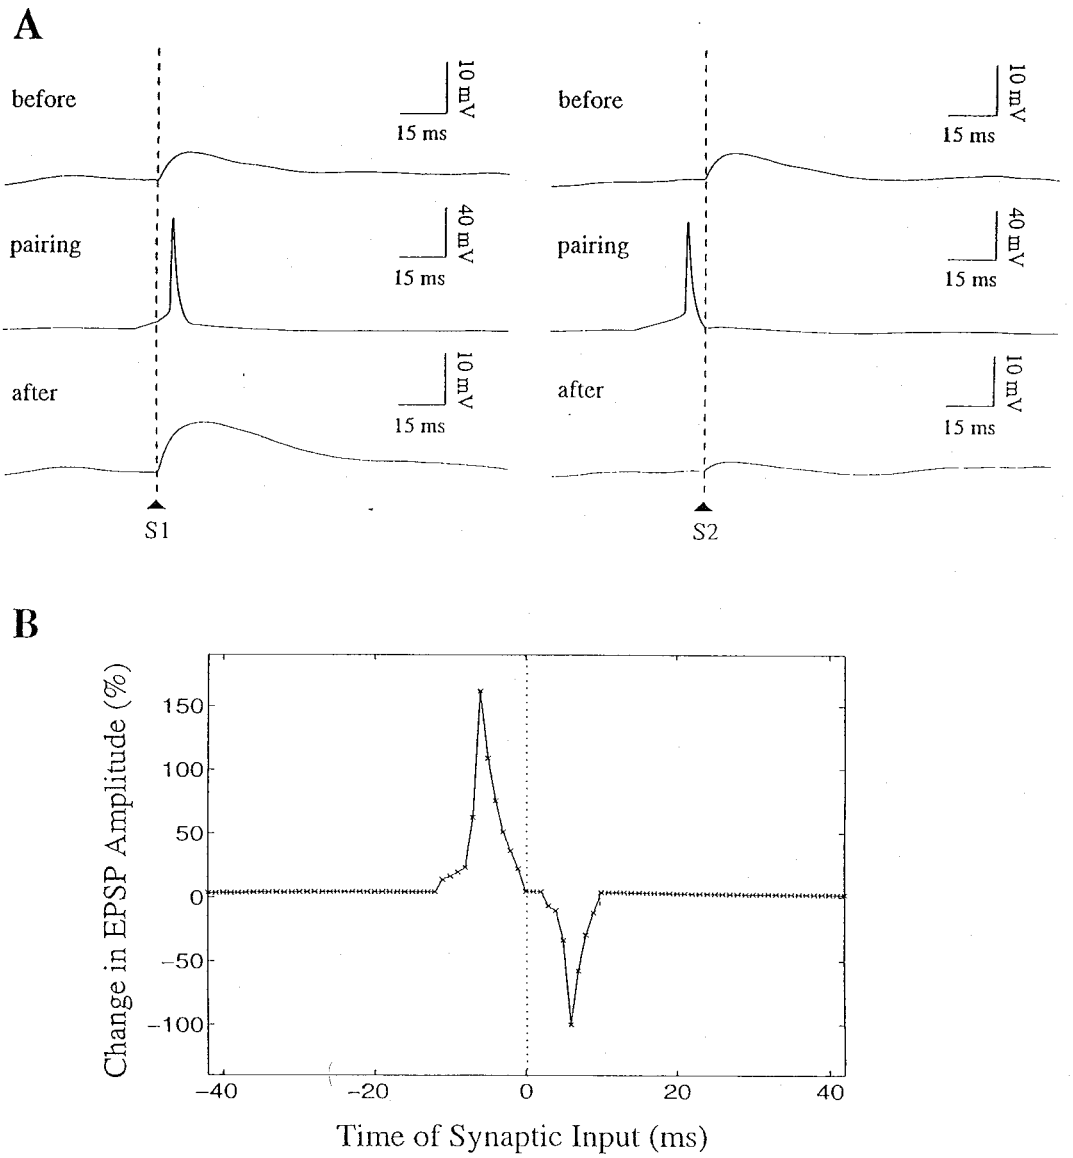

FIG. 3. Synaptic plasticity in a model neocortical neuron. (A) (*Left panel*) The response at the top ('before') is the EPSP evoked in the model neuron due to a presynaptic spike (S1) at an excitatory synapse. Pairing this presynaptic spike with postsynaptic spiking after a 5 ms delay ('pairing') induces long-term potentiation as revealed by an enhancement in the peak of the EPSP evoked by presynaptic simulation alone ('after'). (*Right panel*) If presynaptic stimulation (S2) occurs 5 ms after postsynaptic firing, the synapse is weakened resulting in a decrease in peak EPSP amplitude. (B) Critical window for synaptic plasticity obtained by varying the delay between presynaptic and postsynaptic spiking (negative delays refer to cases where the presynaptic spike occurred before the postsynaptic spike).

varying the time interval between presynaptic stimulation and postsynaptic spiking (with $\Delta t = 5$ ms). As shown in Fig. 3B, changes in synaptic efficacy exhibited a highly asymmetric dependence on spike timing similar to physiological data (Bi & Poo 1998). Potentiation was observed for EPSPs that

occurred between 1 and 12 ms before the postsynaptic spike, with maximal potentiation at 6 ms. Maximal depression was observed for EPSPs occurring 6 ms after the peak of the postsynaptic spike and this depression gradually decreased, approaching zero for delays greater than 10 ms. As in rat neocortical neurons (Markram et al 1997), *Xenopus* tectal neurons (Zhang et al 1998), and cultured hippocampal neurons (Bi & Poo 1998), a narrow transition zone (roughly 3 ms in the model) separated the potentiation and depression windows. Note that the exact duration of the potentiation and depression windows in the model can be adapted to match physiological data by appropriately choosing the temporal-difference parameter Δt and/or varying the distribution of active channels in the dendrite the synapse is located on.

Learning to predict using temporal-difference learning

To see how a network of model neurons can learn to predict sequences using the learning mechanism described above, consider the simplest case of two excitatory neurons N1 and N2 connected to each other, receiving inputs from two separate input neurons I1 and I2 (Fig. 4A). Suppose input neuron I1 fires before input neuron I2, causing neuron N1 to fire (Fig. 4B). The spike from N1 results in a sub-threshold EPSP in N2 due to the synapse S2. If input arrives from I2 any time between 1 and 12 ms after this EPSP and the temporal summation of these two EPSPs causes N2 to fire, the synapse S2 will be strengthened. The synapse S1, on the other hand, will be weakened because the EPSP due to N2 arrives a few milliseconds after N1 has fired. Thus, on a subsequent trial, when input I1 causes neuron N1 to fire, it in turn causes N2 to fire several milliseconds *before* input I2 occurs due to the potentiation of the recurrent synapse S2 in previous trial(s) (Fig. 4C). Input neuron I2 can thus be inhibited by the predictive feedback from N2 just before the occurrence of imminent input activity (marked by an asterisk in Fig. 4C). This inhibition prevents input I2 from further exciting N2. Similarly, a positive feedback loop between neurons N1 and N2 is avoided because the synapse S1 was weakened in previous trial(s) (see arrows in Figs 4B and 4C). Figure 4D depicts the process of potentiation and depression of the two synapses as a function of the number of exposures to the I1–I2 input sequence. The decrease in latency of the predictive spike elicited in N2 with respect to the timing of input I2 is shown in Fig. 4E. Notice that before learning, the spike occurs 3.2 ms after the occurrence of the input whereas after learning, it occurs 7.7 ms before the input. This simple example helps to illustrate how subsets of neurons may learn to selectively trigger other subsets of neurons in anticipation of future inputs while maintaining stability in the recurrent network.

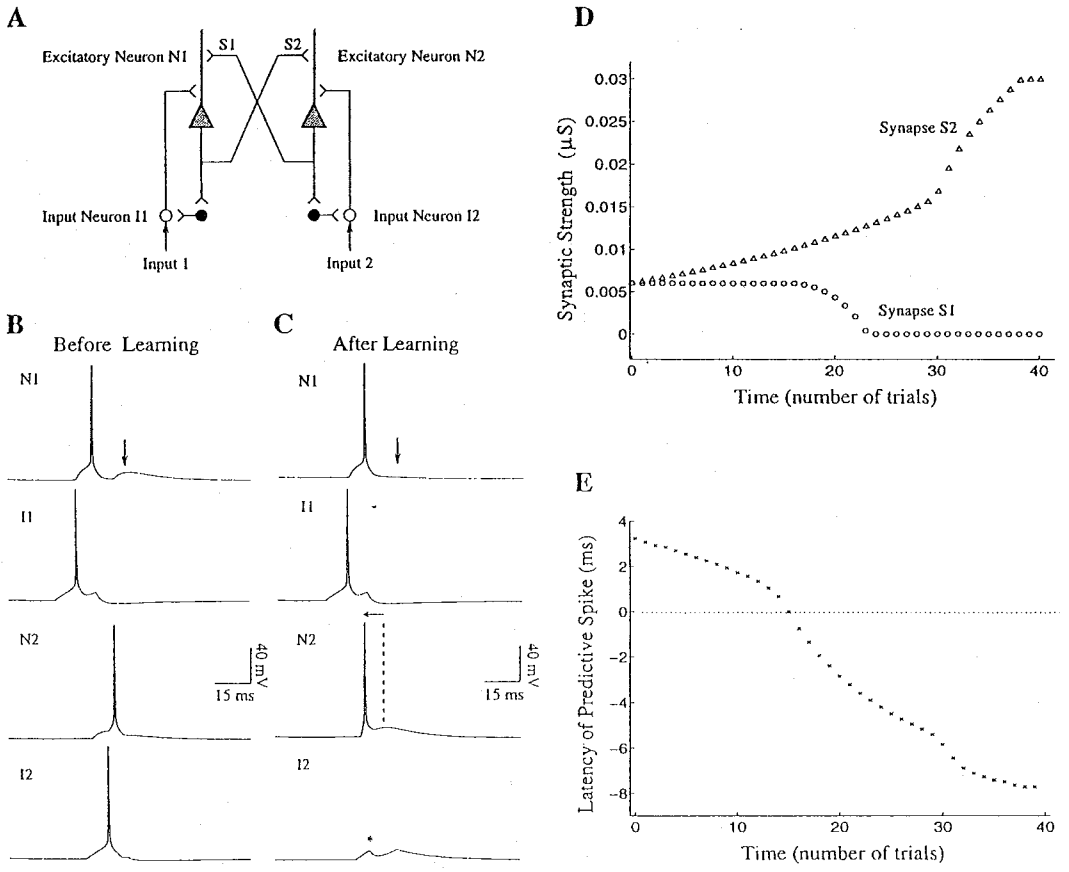

FIG. 4. Learning to predict using spike-timing dependent Hebbian plasticity. (A) A simple network of two model neurons N1 and N2 recurrently connected via excitatory synapses S1 and S2. Sensory inputs are relayed to the two model neurons by input neurons I1 and I2. Feedback from N1 and N2 inhibit the input neurons via inhibitory interneurons (darkened circles). (B) Activity in the network elicited by the input sequence I1 followed by I2. Notice that N2 fires after its input neuron I2 has fired. (C) Activity in the network elicited by the same input sequence after 40 trials of learning. Notice that due to the strengthening of synapse S2, neuron N2 now fires several milliseconds before the time of expected input from I2 (dashed line), allowing it to inhibit I2 (asterisk). On the other hand, synapse S1 has been weakened, thereby preventing re-excitation of N1 (downward arrows show the corresponding decrease in EPSP). (D) Potentiation and depression of synapses S1 and S2 respectively during the course of learning. Synaptic strength was defined as maximal synaptic conductance in the kinetic model of synaptic transmission (see Methods). (E) Latency of the predictive spike in neuron N2 during the course of learning measured with respect to the time of input spike in I2 (dotted line). Note that the latency is initially positive (N2 fires after I2) but later becomes negative, reaching a value of up to 7.7 ms before input I2 as a consequence of learning.

Direction selectivity from predictive sequence learning

To facilitate comparison with published neurophysiological data, we have focused specifically on the problem of predicting moving visual stimuli. Previous modelling studies have suggested that recurrent excitation may play a crucial role in generating direction selectivity in cortical neurons by amplifying their weak

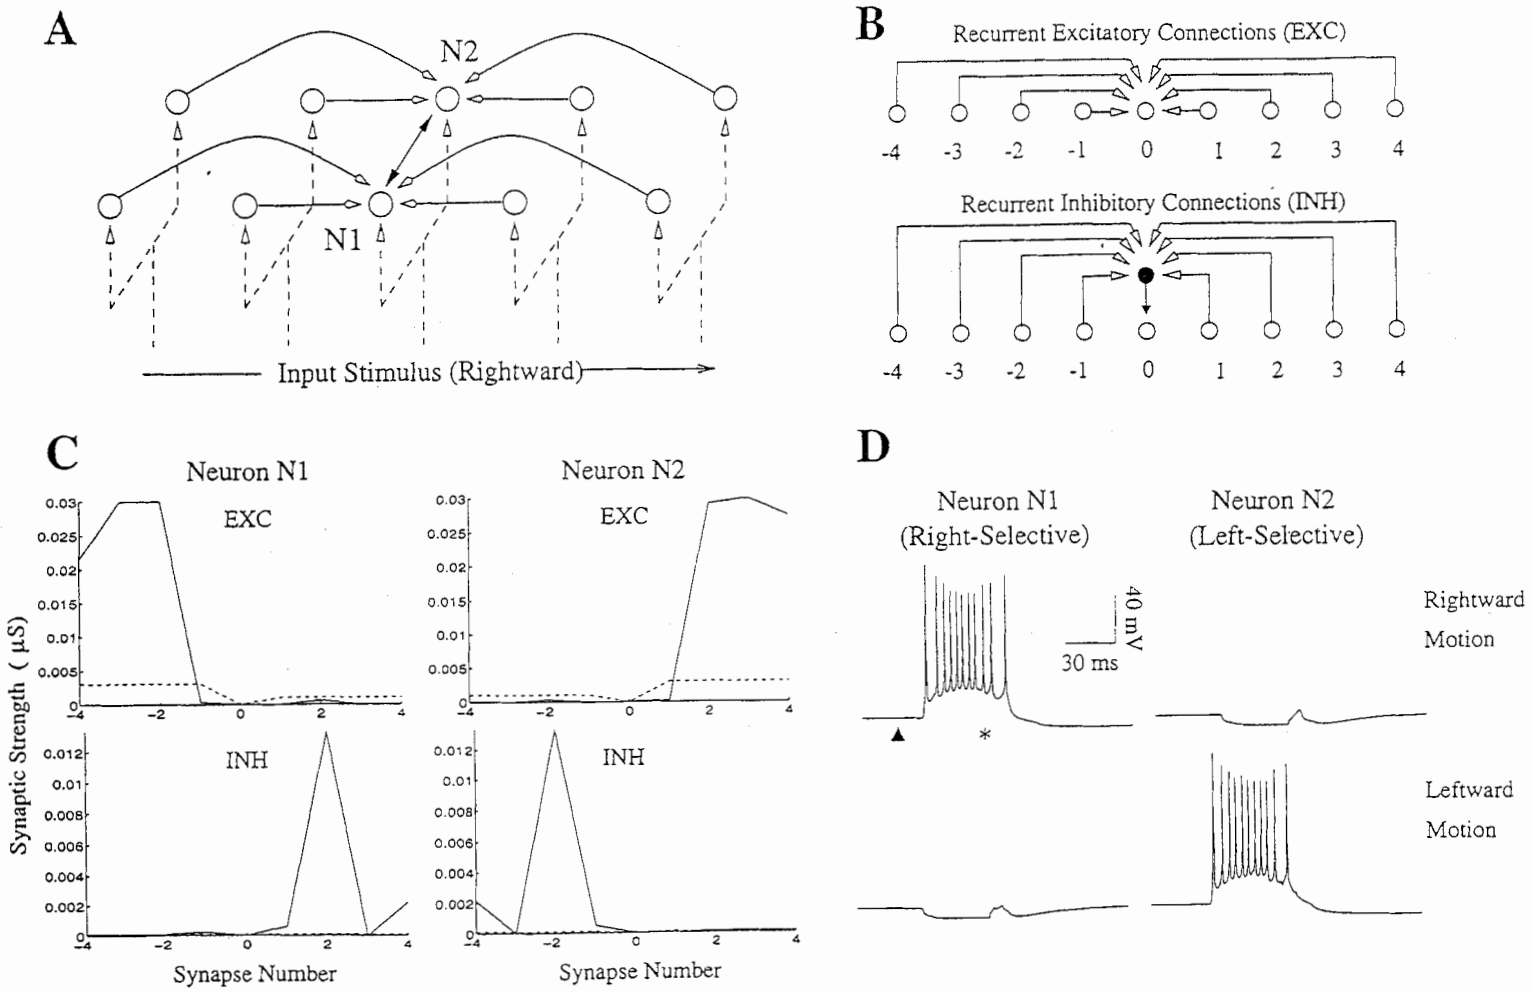
A
N2
N1
Input Stimulus (Rightward)
B
Recurrent Excitatory Connections (EXC)
-4 -3 -2 -1 0 1 2 3 4
Recurrent Inhibitory Connections (INH)
-4 -3 -2 -1 0 1 2 3 4
C
Neuron N1
Neuron N2
EXC
INH
Synaptic Strength (μS)
Synapse Number
D
Neuron N1
(Right-Selective)
Neuron N2
(Left-Selective)
40 mV
30 ms
Rightward
Motion
Leftward
Motion

feedforward inputs (Douglas et al 1995, Suarez et al 1995, Mineiro & Zipser 1998). Our simulations suggest that a network of recurrently connected neurons can develop direction selectivity as a consequence of learning to predict moving stimuli. We used a network of recurrently connected excitatory neurons as shown in Fig. 5A receiving retinotopic sensory input consisting of moving pulses of excitation (8 ms pulse of excitation at each neuron) in the rightward and leftward directions. The task of the network was to predict the sensory input by learning appropriate recurrent connections such that a given neuron in the network can fire a few milliseconds before the arrival of its input pulse of excitation. The network was comprised of two parallel chains of neurons with mutual inhibition (dark arrows) between corresponding pairs of neurons along the two chains. The network was initialized such that within a chain, a given excitatory neuron received both excitation and inhibition from its predecessors and successors. This is shown in Fig. 5B for a neuron labelled '0'. Inhibition at a given neuron was mediated by an inhibitory interneuron (dark circle) which received excitatory connections from neighbouring excitatory neurons (Fig. 5B, lower panel). The interneuron received the same input pulse of excitation as the nearest excitatory neuron. Excitatory and inhibitory synaptic currents were calculated using kinetic models of synaptic transmission based on properties of AMPA and $GABA_A$ (γ-aminobutyric acid A) receptors as determined from whole-cell recordings (see Methods). Maximum conductances for all synapses were initialized to small positive values (dotted lines in Fig. 5C) with a slight asymmetry in the recurrent excitatory connections for breaking symmetry between the two chains. The initial asymmetry elicited a single spike slightly earlier for neurons in one chain than neurons in the other chain for a given motion direction, allowing activity in the other chain to be inhibited.

FIG. 5. Emergence of direction selectivity in the model. (A) A model network consisting of two chains of recurrently connected neurons receiving retinotopic inputs. A given neuron receives recurrent excitation and recurrent inhibition (white-headed arrows) as well as inhibition (dark-headed arrows) from its counterpart in the other chain. (B) Recurrent connections to a given neuron (labelled '0') arise from 4 preceding and 4 succeeding neurons in its chain. Inhibition at a given neuron is mediated via a GABAergic interneuron (darkened circle). (C) Synaptic strength of recurrent excitatory (EXC) and inhibitory (INH) connections to neurons N1 and N2 before (dotted lines) and after learning (solid lines). Synapses were adapted during 100 trials of exposure to alternating leftward and rightward moving stimuli. (D) Responses of neurons N1 and N2 to rightward and leftward moving stimuli. As a result of learning, neuron N1 has become selective for rightward motion (as have other neurons in the same chain) while neuron N2 has become selective for leftward motion. In the preferred direction, each neuron starts firing several milliseconds before the actual input arrives at its soma (marked by an asterisk) due to recurrent excitation from preceding neurons. The dark triangle represents the start of input stimulation in the network.

To evaluate the consequences of synaptic plasticity, the network of neurons was exposed alternately to leftward and rightward moving stimuli for a total of 100 trials. The excitatory connections (labelled 'EXC' in Fig. 5B) were modified according to the asymmetric Hebbian learning rule in Fig. 3B while the excitatory connections onto the inhibitory interneuron (labelled 'INH') were modified according to an asymmetric anti-Hebbian learning rule that reversed the polarity of the rule in Fig. 3B. In other words, if presynaptic activity occurred before (after) the postsynaptic spike in the interneuron, the excitatory connection to the inhibitory interneuron was weakened (strengthened). Although not yet reported in the neocortex, such a rule for inhibitory interneurons is consistent with the spike-timing dependent anti-Hebbian plasticity observed in inhibitory interneurons in a cerebellum-like structure in weakly electric fish (Bell et al 1997).

The synaptic conductances learned by two neurons (marked N1 and N2 in Fig. 5A) located at corresponding positions in the two chains after 100 trails of exposure to the moving stimuli are shown in Fig. 5C (solid line). Initially, for rightward motion, the slight asymmetry in the initial excitatory connections of neuron N1 allows it to fire slightly earlier than neuron N2 thereby inhibiting neuron N2. Additionally, since the EPSPs from neurons lying on the left of N1 occur before N1 fires, the excitatory synapses from these neurons are strengthened while the excitatory synapses from these same neurons to the inhibitory interneuron are weakened according to the two learning rules mentioned above. On the other hand, the excitatory synapses from neurons lying on the right side of N1 are weakened while inhibitory connections are strengthened since the EPSPs due to these connections occur after N1 has fired. The synapses on neuron N2 and its associated interneuron remain unaltered since there is no postsynaptic firing (due to inhibition by N1) and hence no backpropagating action potentials in the dendrite. Similarly, for leftward motion, neuron N2 inhibits neuron N1 and the synapses associated with N2 are adapted according to the two learning rules. As shown in Fig. 5C, after 100 trials, the excitatory and inhibitory connections to neuron N1 exhibit a marked asymmetry, with excitation originating from neurons on the left and inhibition from neurons on the right. Neuron N2 exhibits the opposite pattern of connectivity.

As expected from the learned pattern of connectivity, neuron N1 was found to be selective for rightward motion while neuron N2 was selective for leftward motion (Fig. 5D). Moreover, when stimulus motion is in the preferred direction, each neuron starts firing a few milliseconds before the time of arrival of the input stimulus at its soma (marked by an asterisk) due to recurrent excitation from preceding neurons. Conversely, motion in the non-preferred direction triggers recurrent inhibition from preceding neurons as well as inhibition from the active neuron in the corresponding position in the other chain. Thus, the learned pattern of connectivity allows the direction-selective neurons comprising the two chains in

the network to conjointly code for and predict the moving input stimulus in each direction.

The role of recurrent excitation and inhibition

To ascertain the role of recurrent excitation in the model, we gradually decreased the value of the maximum synaptic conductance between excitatory neurons in the network, starting from 100% of the learned values. For a stimulus moving in the preferred direction, decreasing the amount of recurrent excitation increased the latency of the first spike in a model neuron and decreased the spike count until, with less than 10% of the learned recurrent excitation, the latency equalled the arrival time of the input stimulus and the spike count dropped to 1 (Figs 6A and 6B). These results demonstrate that recurrent excitation plays a crucial role in generating predictive activity in model neurons and enhances direction-selective responses by increasing the spike count in the preferred direction.

To evaluate the role of inhibition in maintaining direction selectivity in the model, we quantified the degree of direction selectivity using the direction index: 1−(number of spikes in non-preferred direction)/(number of spikes in preferred direction). Figures 6C and 6D show the distribution of direction indices with and without inhibition in a network of two chains containing 35 excitatory and 35 inhibitory neurons. In the control case, most of the excitatory neurons and inhibitory interneurons receiving recurrent excitation are highly direction selective. Blocking inhibition significantly reduces direction selectivity in the model neurons but does not completely eliminate it, consistent with some previous physiological observations (Sillito 1975, Nelson et al 1994). The source of this residual direction selectivity in the model in the absence of inhibition can be traced to the asymmetric recurrent excitatory connections in the network which remain unaffected by the blockage of inhibition.

Comparison with awake monkey complex cell responses

Similar to complex cells in primary visual cortex, model neurons are direction selective throughout their receptive field because at each retinotopic location, the corresponding neuron in the chain receives the same pattern of asymmetric excitation and inhibition from its neighbours as any other neuron in the chain. Thus, for a given neuron, motion in any local region of the chain will elicit direction-selective responses due to recurrent connections from that part of the chain. This is consistent with previous modelling studies (Chance et al 1999) suggesting that recurrent connections may be responsible for the spatial-phase invariance of complex cell responses. Assuming a 200 μm separation between excitatory model neurons in each chain and utilizing known values for the

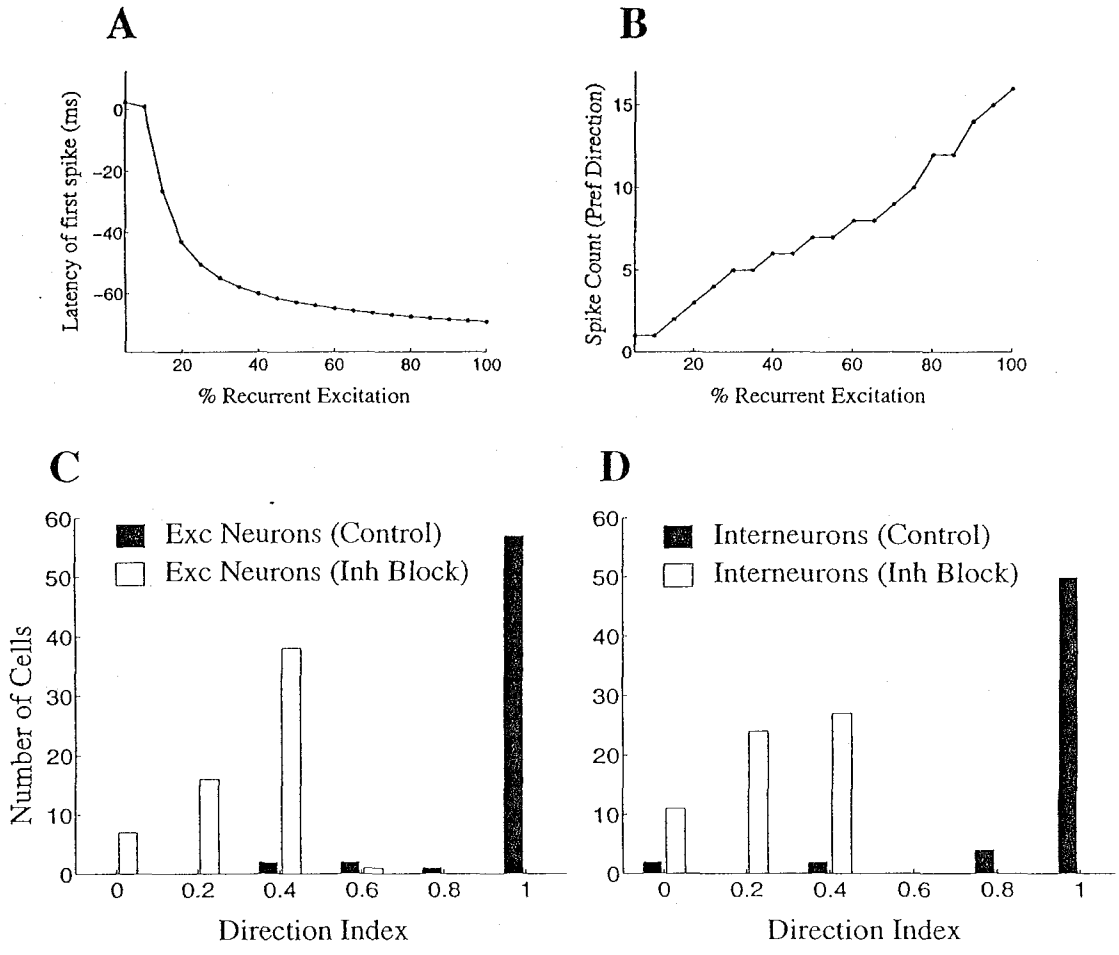

FIG. 6. The role of recurrent excitation and inhibition. (A) & (B) Latency of the first spike and number of spikes elicited in an excitatory neuron in the preferred direction as a function of the strength of recurrent excitation in a model network (100% corresponds to the learned values of recurrent connection strength). The network comprised of two chains, each containing 35 excitatory neurons and 35 inhibitory interneurons (mutual inhibition between corresponding neurons in the two chains was mediated by a separate set of inhibitory neurons that were not plastic). (C,D) Distribution of direction selectivity in the network for excitatory and inhibitory interneurons respectively with (Control) and without GABAergic inhibition (Inh Block) as measured by the direction index: 1−(Non-Preferred Direction Response)/(Preferred Direction Response).

cortical magnification factor in monkey striate cortex (Tootell et al 1988), one can estimate the preferred stimulus velocity of model neurons to be 3.1°/s in the fovea and 27.9°/s in the periphery (at an eccentricity of 8°). Both of these values fall within the range of monkey striate cortical velocity preferences (1°/s to 32°/s) (van Essen 1985, Livingstone 1998).

The model predicts that the neuroanatomical connections for a direction-selective neuron should exhibit a pattern of asymmetrical excitation and inhibition similar to Fig. 5C. A recent study of direction-selective cells in awake

monkey V1 found excitation on the preferred side of the receptive field and inhibition on the null side consistent with the pattern of connections learned by the model (Livingstone 1998). For comparison with this experimental data, spontaneous background activity in the model was generated by incorporating Poisson-distributed random excitatory and inhibitory alpha synapses on the dendrite of each model neuron. Post-stimulus time histograms (PSTHs) and space-time response plots were obtained by flashing optimally oriented bar stimuli at random positions in the cell's activating region. As shown in Fig. 7, there is good qualitative agreement between the response plot for a direction-selective complex cell and that for the model. Both space-time plots show a progressive shortening of response onset time and an increase in response transiency going in the preferred direction; in the model, this is due to recurrent excitation from progressively closer cells on the preferred side. Firing is reduced to below background rates 40–60 ms after stimulus onset in the upper part of the plots; in the model, this is due to recurrent inhibition from cells on the null side. The response transiency and shortening of response time course appears as a slant in the space-time maps, which can be related to the neuron's velocity sensitivity (see Livingstone 1998 for more details).

Discussion

Our results show that a network of recurrently connected neurons endowed with a temporal-difference based asymmetric Hebbian learning mechanism can learn a predictive model of its spatiotemporal inputs. Using a biophysical model of neocortical neurons, we showed that a temporal-difference learning rule for prediction when applied to backpropagating action potentials in dendrites produces asymmetric learning windows similar to those observed in physiological experiments (see Senn 1997, Egelman & Montague 1998) for possible biophysical mechanisms based on *N*-methyl-D-aspartate (NMDA) receptor activation and voltage-dependent Ca^{2+} channels). When exposed to moving stimuli, neurons in a simulated network with recurrent excitatory and inhibitory connections learned to fire a few milliseconds before the expected arrival of an input stimulus and developed direction selectivity as a consequence of learning. The model predicts that a direction-selective neuron should start responding a few milliseconds before the preferred stimulus arrives at the retinotopic location of the neuron in primary visual cortex. Such predictive neural activity has recently been reported in ganglion cells in the rabbit and salamander retina (Berry et al 1999).

The development of direction selectivity in our model requires a slight initial bias in cortical connectivity (Fig. 5C) which is then sharpened by visual experience of moving stimuli. This is consistent with experimental evidence

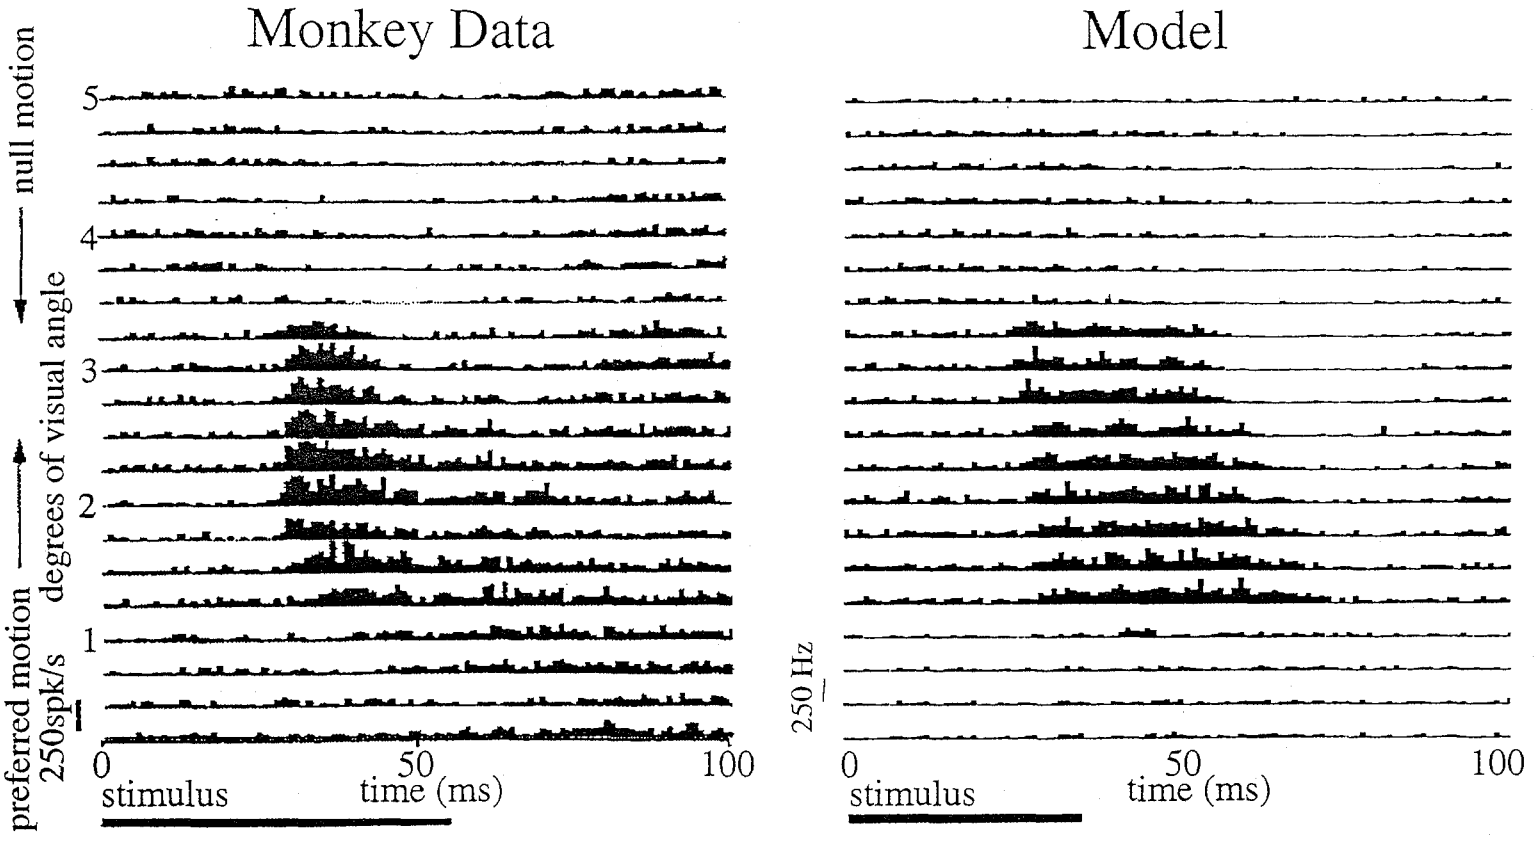

FIG. 7. Comparison of monkey and model space–time response plots. (*Left*) Sequence of PSTHs obtained by flashing optimally oriented bars at 20 positions across the 5°-wide receptive field (RF) of a complex cell in alert monkey V1 (from Livingstone 1998). The cell's preferred direction is from the part of the RF represented at the bottom towards the top. Flash duration = 56 ms; inter-stimulus delay = 100 ms; 75 stimulus presentations. (*Right*) PSTHs obtained from a model neuron after stimulating the chain of neurons at 20 positions to the left and right side of the given neuron. Lower PSTHs represent stimulations on the preferred side while upper PSTHs represent stimulations on the null side.

indicating that (a) some cells in cat visual cortex show some amount of direction selectivity before eye opening (Movshon & van Sluyters 1981) and (b) visual experience during a critical period can profoundly affect the development of direction selectivity (for example, direction selectivity can be abolished by strobe rearing; Humphrey & Saul 1998). Although several models for the development of direction selectivity have been proposed (Feidler et al 1997, Wimbauer et al 1997), the roles of spike timing and asymmetric Hebbian plasticity have not been previously explored. An interesting question currently being investigated is whether the explicit dependence of visual development on spike timing in our model can account for the fact that only low frequencies of stroboscopic illumination (approximately 8 Hz or below) lead to a loss of direction selectivity.

Temporally asymmetric Hebbian learning has previously been suggested as a possible mechanism for sequence learning in the hippocampus (Levy & Steward 1983, Abbott & Blum 1996) and as an explanation for the asymmetric expansion of hippocampal place fields during route learning (Mehta et al 1997). Some of these theories require relatively long temporal windows of synaptic plasticity (on the order of several hundreds of milliseconds) (Abbott & Blum 1996) while others have utilized temporal windows in the sub-millisecond range for coincidence detection (Gerstner et al 1996). Prediction and sequence learning in our model is based on a window of plasticity in the tens of milliseconds range which is roughly consistent with recent physiological observations (Markram et al 1997, Zhang et al 1998, Bi & Poo 1998). Although a fixed learning window (roughly 15 ms of potentiation/depression) was used in the simulations, the temporal extent of this window can be modified by changing the parameter Δt. The temporal-difference model predicts that the shape and width of the asymmetric learning window should be a function of the backpropagating action potentials in the dendrite that the synapse is located on. In the case of hippocampal neurons and cortical neurons, the width of backpropagating action potentials in apical dendrites has been reported to be in the range of 10–25 ms, which is comparable to the size of potentiation/depression windows for synapses located on these dendrites (Bi & Poo 1998, Stuart & Sakmann 1994). Additionally, in order to account for the off regions observed in the receptive fields of cortical direction-selective cells (Livingstone 1998), we included synaptic plasticity of excitatory synapses on inhibitory interneurons by assuming that the sign of the spike-timing dependent Hebbian learning window was inverted from that found on pyramidal neurons. This inversion has been found in excitatory synapses on inhibitory interneurons in a cerebellum-like brain structure in weakly electric fish (Bell et al 1997), but remains a prediction of our model for the cortex.

In vitro experiments involving cortical and hippocampal slices suggest the possibility of short-term plasticity in synaptic connections onto pyramidal neurons (Thomson & Deuchars 1994, Tsodyks & Markram 1997, Abbott et al

1997). The kinetic model of synaptic transmission used in the present study can be extended to include short-term plasticity with the addition of a parameter governing the level of depression caused by each presynaptic action potential (Chance et al 1999, Tsodyks & Markram 1997, Abbott et al 1997). The adaptation of this parameter may allow finer control of postsynaptic firing in the model in addition to the coarse-grained control offered by modifications of maximal synaptic conductance. As suggested by previous studies (Chance et al 1999, Abbott 1997), we expect the addition of synaptic depression in our model to enhance the transient response of model neurons to stimuli such as flashed bars (see Fig. 7) and to broaden the response to drifting stimuli, due to the reduced sensitivity of postsynaptic neurons to mean presynaptic firing rates. In preliminary simulations, the inclusion of short-term plasticity did not significantly alter the development of direction selectivity in recurrent network models as reported here.

The idea that prediction and sequence learning may constitute an important goal of the neocortex has previously been suggested in the context of statistical and information theoretic models of cortical processing (Minai & Levy 1993, Montague & Sejnowski 1994, Mumford 1994, Daugman & Downing 1995, Abbott & Blum 1996, Schultz et al 1997, Rao & Ballard 1997, Barlow 1998, Rao 1999). Our biophysical simulations suggest a possible implementation of such models in cortical circuitry. Several authors have observed the general uniformity in the basic structure of the neocortex across different cortical areas (Hubel & Wiesel 1974, Creutzfeldt 1977, Sejnowski 1986, Douglas et al 1989). Given the universality of the problem of encoding and generating temporal sequences in both sensory and motor domains, the hypothesis of predictive sequence learning in recurrent neocortical circuits may help provide a unifying principle for understanding the general nature of cortical information processing (Creutzfeldt 1977, Sejnowski 1986).

Methods

Temporal-difference learning. The simplest example of a temporal-difference learning rule arises in the problem of predicting a scalar quantity z using a neuron with synaptic weights $w(1), \ldots w(k)$ (represented as a vector $\mathbf{w}$). The neuron receives as presynaptic input the sequence of vectors $\mathbf{x}_1, \ldots \mathbf{x}_m$. The output of the neuron at time t is assumed to be given by: $P_t = \sum_i w(i)x_t(i)$. The goal is to learn a set of synaptic weights such that the prediction P_t is as close as possible to the target z. One method for achieving this goal is to use a temporal-difference (TD[0]) learning rule (Sutton 1988). The weights are changed at time t by an amount given by:

$$\Delta \mathbf{w}_t = \alpha(P_{t+1} - P_t)\mathbf{x}_t \quad (1)$$

where α is a learning rate or gain parameter and the final prediction P_{m+1} is defined to be z. Note that in such a learning paradigm, synaptic plasticity is governed by the temporal difference in postsynaptic activity at time instants $t+1$ and t in conjunction with presynaptic activity $\mathbf{x}_t$ at time t.

Neocortical neuron model. Two-compartment model neocortical neurons consisting of a dendritic compartment and a soma-axon compartment (Mainen & Sejnowski 1996) were implemented using the simulation software *Neuron* (Hines 1993). Four voltage-dependent currents and one Ca^{2+}-dependent current were simulated: fast Na^+, I_{Na}; fast K^+, I_{Kv}; slow non-inactivating K^+, I_{Km}; high voltage-activated Ca^{2+}, I_{Ca}; and Ca^{2+}-dependent K^+ current, I_{KCa} (see Mainen & Sejnowski 1996 for references). Conventional Hodgkin–Huxley-type kinetics were used for all currents (integration time step = 25 μs, temperature = 37 °C). Ionic currents I were calculated using the ohmic equation: $I = \bar{g}A^xB(V-E)$ where $\bar{g}$ is the maximal ionic conductance density, A and B are activation and inactivation variables, respectively (x denotes the order of kinetics; see Mainen & Sejnowski 1996 for further details), and E is the reversal potential for the given ion species ($E_K = -90$ mV, $E_{Na} = 60$ mV, $E_{Ca} = 140$ mV, $E_{leak} = -70$ mV). The following active conductance densities were used in the dendritic compartment (in pS/μm^2): $\bar{g}_{Na} = 20$, $\bar{g}_{Ca} = 0.2$, $\bar{g}_{Km} = 0.1$, and $\bar{g}_{KCa} = 3$, with leak conductance 33.3 μS/cm^2 and specific membrane resistance 30 kΩ/cm^2. The soma–axon compartment contained $\bar{g}_{Na} = 40\,000$ and $\bar{g}_{Kv} = 1400$. For all compartments, the specific membrane capacitance was 0.75 μF/cm^2. Two key parameters governing the response properties of the model neuron are (Mainen & Sejnowski 1996): the ratio of axo-somatic area to dendritic membrane area (ρ) and the coupling resistance between the two compartments (κ). For the present simulations, we used the values $\rho = 150$ (with an area of 100 μm^2 for the soma–axon compartment) and a coupling resistance of $\kappa = 8$ MΩ. Poisson-distributed synaptic inputs to the dendrite were simulated using alpha function (Koch 1999) shaped current pulse injections (time constant = 5 ms) at Poisson intervals with a mean presynaptic firing frequency of 3 Hz.

Model of synaptic transmission and plasticity. Synaptic transmission at excitatory (AMPA) and inhibitory ($GABA_A$) synapses was simulated using first order kinetics of the form:

$$\frac{dr}{dt} = \alpha[T](1 - r) - \beta r \quad (2)$$

where $r(t)$ denotes the fraction of postsynaptic receptors bound to the neurotransmitter at time t, $[T]$ is the neurotransmitter concentration, and α and β are the forward and backward rates for transmitter binding. Assuming receptor binding directly gates the opening of an associated ion channel, the resulting synaptic current can be described as (Destexhe et al 1998):

$$I_{syn} = \bar{g}_{syn} r(t)(V_{syn}(t) - E_{syn}) \tag{3}$$

where $\bar{g}_{syn}$ is the maximal synaptic conductance, $V_{syn}(t)$ is the postsynaptic potential and E_{syn} is the synaptic reversal potential. For the simulations, all synaptic parameters were set to values that gave the best fit to whole-cell recorded synaptic currents (see Destexhe et al 1998). Parameters for AMPA synapses: $\alpha = 1.1 \times 10^{-6}\,\mathrm{M^{-1}s^{-1}}$, $\beta = 190\,\mathrm{s^{-1}}$, and $E_{AMPA} = 0\,\mathrm{mV}$. Parameters for $GABA_A$ receptors: $\alpha = 5 \times 10^{-6}\,\mathrm{M^{-1}s^{-1}}$, $\beta = 180\,\mathrm{s^{-1}}$, and $E_{GABAA} = -80\,\mathrm{mV}$. Synaptic plasticity was simulated by adapting the maximal synaptic conductance $\bar{g}_{AMPA}$ for recurrent excitatory synapses onto excitatory neurons and GABAergic interneurons according to the learning mechanism described in the text. Inhibitory synapses were not adapted since evidence is currently lacking for their plasticity. We therefore used the following fixed values for $\bar{g}_{GABAA}$ (in μS): 0.04 for Fig. 4, 0.05 for mutual inhibition between the two chains and 0.016 for recurrent inhibitory connections within a chain for the simulations in Fig. 5.

Synaptic plasticity was simulated by changing maximal synaptic conductance $\bar{g}_{AMPA}$ by an amount equal to $\Delta\bar{g}_{AMPA} = \alpha(P_{t+\Delta t} - P_{t-\Delta t})$ for each presynaptic spike at time t, where P_t denotes the postsynaptic membrane potential at time t. The conductance was adapted whenever the absolute value of $\bar{g}_{AMPA}$ exceeded 10 mV with a gain α in the range 0.02–0.03 μS/V. The maximum value attainable by a synaptic conductance was set equal to 0.03 μS. Note that the learning rule above differs from the pure TD(0) learning rule in that it depends on postsynaptic activity Δt ms in the future as well as Δt ms in the past whereas the TD(0) rule depends on future and current postsynaptic activity (see Equation 1). This phenomenological model of synaptic plasticity is consistent with known biophysical mechanisms such as calcium-dependent and NMDA receptor-dependent induction of long-term potentiation (LTP) and depression (LTD) (see Senn 1997, Egelman & Montague 1998, for possible biophysical implementations).

Acknowledgments

This work was supported by the Alfred P. Sloan Foundation and Howard Hughes Medical Institute. We thank Margaret Livingstone, Dmitri Chklovskii, David Eagleman, and Christian Wehrhahn for comments and suggestions.

References

Abbott LF, Blum KI 1996 Functional significance of long-term potentiation for sequence learning and prediction. Cereb Cortex 6:406–416

Abbott LF, Varela JA, Sen K, Nelson SB 1997 Synaptic depression and cortical gain control. Science 275:220–224

Barlow H 1998 Cerebral predictions. Perception 27:885–888

Bell CC, Han VZ, Sugawara Y, Grant K 1997 Synaptic plasticity in a cerebellum-like structure depends on temporal order. Nature 387:278–281

Ben-Yishai R, Bar-Or RL, Sompolinsky H 1995 Theory of orientation tuning in visual cortex. Proc Natl Acad Sci USA 92:3844–3848

Berry MJ 2nd, Brivanlou IH, Jordan TA, Meister M 1999 Anticipation of moving stimuli by the retina. Nature 398:334–338

Bi GQ, Poo MM 1998 Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. J Neurosci 18:10464–10472

Braitenberg V, Schüz AA 1991 Anatomy of the cortex. Springer-Verlag, Berlin

Chance FS, Nelson SB, Abbott LF 1999 Complex cells as cortically amplified simple cells. Nat Neurosci 2:277–282

Creutzfeldt OD 1977 Generality of the functional structure of the neocortex. Naturwissenschaften 64:507–517

Daugman JG, Downing CJ 1995 Demodulation, predictive coding, and spatial vision. J Opt Soc Am A 12:641–660 (erratum: 1995 J Opt Soc Am A 12:2077)

Destexhe A, Mainen ZF, Sejnowski TJ 1998 Kinetic models of synaptic transmission. In: Koch C, Segev I (eds) Methods in neuronal modeling, 2nd edn: From ions to networks. MIT Press, Cambridge, MA, p 1–25

Douglas RJ, Martin KAC, Whitteridge D 1989 A canonical microcircuit for neocortex. Neural Computation 1:480–488

Douglas RJ, Koch C, Mahowald M, Martin KA, Suarez HH 1995 Recurrent excitation in neocortical circuits. Science 269:981–985

Egelman DM, Montague PR 1998 Computational properties of peri-dendritic calcium fluctuations. J Neurosci 18:8580–8589

Elman JL 1990 Finding structure in time. Cognit Sci 14:179–211

Feidler JC, Saul AB, Murthy A, Humphrey AL 1997 Hebbian learning and the development of direction selectivity: the role of geniculate response timings. Network: Comput Neural Syst 8:195–214

Gerstner W, Kempter R, van Hemmen JL, Wagner H 1996 A neuronal learning rule for sub-millisecond temporal coding. Nature 383:76–81

Hines M 1993 In: Eeckman FH, (ed) Neural systems: analysis and modeling. Kluwer, Boston, MA, p 127–136

Hubel DH, Wiesel TN 1974 Uniformity of monkey striate cortex: a parallel relationship between field size, scatter, and magnification factor. J Comp Neurol 158:295–305

Humphrey AL, Saul AB 1998 Strobe rearing reduces direction selectivity in area 17 by altering spatiotemporal receptive-field structure. J Neurophysiol 80:2991–3004

Jordan MI 1986 Attractor dynamics and parallelism in a connectionist sequential machine. In: Proceedings of the Eighth Annual Conference of the Cognitive Science Society, Amherst, MA, August 1986. Lawrence Erlbaum, Hillsdale, NJ, p 531–546

Kalman RE 1960 A new approach to linear filtering and prediction theory. Trans ASME J Basic Eng 82:35–45

Koch C 1999 Biophysics of computation: information processing in single neurons. Oxford University Press, New York

Levy WB, Steward O 1983 Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. Neuroscience 8:791–797

Livingstone MS 1998 Mechanisms of direction selectivity in macaque V1. Neuron 20:509–526

Mainen ZF, Sejnowski TJ 1996 Influence of dendritic structure on firing pattern in model neocortical neurons. Nature 382:363–386

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

Mehta MR, Barnes CA, McNaughton BL 1997 Experience-dependent, asymmetric expansion of hippocampal place fields. Proc Natl Acad Sci USA 94:8918–8921

Minai AA, Levy WB 1993 Sequence learning in a single trial. In: Proceedings of the 1993 world congress on neural networks, July 1993, Portland OR, II, Lawrence Erlbaum, Hillsdale, NJ, p 505–508

Mineiro P, Zipser D 1998 Analysis of direction selectivity arising from recurrent cortical interactions. Neural Comput 10:353–371

Montague PR, Sejnowski TJ 1994 The predictive brain: temporal coincidence and temporal order in synaptic learning mechanisms. Learn Mem (Cold Spring Harb)1:1

Movshon JA, Van Sluyters RC 1981 Visual neural development. Annu Rev Psychol 32:477–522

Mumford D 1994 Neuronal architectures for pattern-theoretic problems. In: Koch C, Davis JL (eds) Large-scale neuronal theories of the brain. MIT Press, Cambridge, MA p 125–152

Nelson S, Toth L, Sheth B, Sur M 1994 Orientation selectivity of cortical neurons during intracellular blockade of inhibition. Science 265:774–777

Olshausen BA, Field DJ 1997 Sparse coding with an overcomplete basis set: a strategy employed by V1? Vision Res 37:3311–3325

Rao RPN 1999 An optimal estimation approach to visual perception and learning. Vision Res 39:1963–1989

Rao RPN, Ballard DH 1997 Dynamic model of visual recognition predicts neural response properties in the visual cortex. Neural Comp 9:721–763

Rao RPN, Ballard DH 1999 Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nat Neurosci 2:79–87

Scheinerman ER 1995 Invitation to dynamical systems. Prentice Hall, New Jersey

Schultz W, Dayan P, Montague PR 1997 A neural substrate of prediction and reward. Science 275:1593–1599

Sejnowski TJ 1986 Open questions about computation in cerebral cortex. In: McClelland JL, Rumelhart DE (eds) Parallel distributed processing, vol 2: Explorations in the microstructure of cognition. MIT Press, Cambridge, MA, p 372–389

Senn W, Tsodyks M, Markram, H 1997 An algorithm for synaptic modification based on exact timing of pre- and post-synaptic action potentials. In: Gerstner W, Germond A, Hasler M, Nicoud JD (eds) Artificial neural networks—ICANN'97, proceedings of the 7th international conference, Lausanne, Switzerland. Lect Notes Comput Sci 1327:121–126

Sillito AM 1975 The contribution of inhibitory mechanisms to the receptive field properties of neurones in the striate cortex of the cat. J Physiol (Lond) 250:305–329

Softky WR 1996 Unsupervised pixel-prediction. In: Touretzky DS, Mozer MC, Hasselmo ME (eds) Advances in neural information processing systems 8: proceedings of the 1995 Conference, MIT Press, Cambridge, MA p 809–815

Somers DC, Nelson SB, Sur M 1995 An emergent model of orientation selectivity in cat visual cortical simple cells. J Neurosci 15:5448–5465

Stuart GJ, Sakmann B 1994 Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. Nature 367:69–72

Suarez H, Koch C, Douglas R 1995 Modeling direction selectivity of simple cells in striate visual cortex within the framework of the canonical microcircuit. J Neurosci 15:6700–6719

Sutton RS 1988 Learning to predict by the method of temporal differences. Machine Learning 3:9–44

Thomson AM, Deuchars J 1994 Temporal and spatial properties of local circuits in neocortex. Trends Neurosci 17:119–126

Tootell RB, Switkes E, Silverman MS, Hamilton SL 1988 Functional anatomy of macaque striate cortex. II. Retinotopic organization. J Neurosci 8:1531–1568

Tsodyks MV, Markram H 1997 The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. Proc Natl Acad Sci USA 94:719–723 (erratum: 1997 Proc Natl Acad Sci USA 94:5495)

van Essen DC 1985 Functional organization of primate visual cortex. In: Peters A, Jones EG, (eds) Cerebral cortex, vol 3: Visual cortex. Plenum Press, New York p 259–329

Wimbauer S, Wenisch OG, Miller KD, van Hemmen JL 1997 Development of spatiotemporal receptive fields of simple cells: I. Model formulation. Biol Cybern 77:453–461

Zhang LI, Tao HW, Holt CE, Harris WA, Poo M 1998 A critical window for cooperation and competition among developing retinotectal synapses. Nature 395:37–44