

A Model of Corticostriatal Plasticity for Learning Oculomotor Associations and Sequences

Peter Dominey

INSERM Unité 94, France and University of Southern California

Michael Arbib

University of Southern California

Jean-Paul Joseph

INSERM Unité 94, France

Abstract

■ We present models that learn context-dependent oculomotor behavior in (1) conditional visual discrimination and (2) sequence reproduction tasks, based on the following three principles: (1) Visual input and efferent copies of motor output produce patterns of activity in cortex. (2) Cortex influences the saccade system in part via corticostriatal projections. (3) A reinforcement learning mechanism modifies corticostriatal synapses to link patterns of cortical activity to the correct saccade responses during trial-and-error learning. Our *conditional visual discrimination model* learns to associate visual cues with the corresponding saccades to one of two left-right targets. A visual cue produces patterns of neuronal activity in inferotemporal cortex (IT) that projects to the oculomotor region of the striatum. Initially random saccadic "guesses," when directed to the correct target for the current cue, result in increased synaptic strength between the cue-related IT cells and the striatal cells that participate in the correct saccade, increasing the probability that this cue will later elicit the correct saccade.

We show that the model generates "inhibitory gradients" on the striatum as the substrate for spatial generalization. Our *sequence reproduction model* learns, when presented with temporal sequences of spatial targets, to reproduce the corresponding sequence of saccades. At any point in the execution of a saccade sequence, the current pattern of activity in prefrontal cortex (PFC), combined with visual input and the motor efferent copy of the previous saccade, produces a new pattern of activity in PFC to be associated with the next saccade. Like IT, PFC also projects to the oculomotor region of the striatum. Correct guesses for the subsequent saccade in the sequence results in strengthening of corticostriatal synapses between active PFC cells and striatal cells involved in the correct saccade. The sequence is thus reproduced as a concatenation of associations. We compare the results of this model with data previously obtained in the monkey and discuss the nature of cortical representations of spatiotemporal information. ■

INTRODUCTION: TASKS, CONTEXTS, AND BEHAVIOR

While the basal ganglia have traditionally been associated with initiation and control of movement (see Alexander, DeLong, & Strick, 1986), accumulating evidence links basal ganglia additionally to the conditional selection of appropriate state- or context-dependent behavior (e.g., Mishkin, Malamut, & Bachevalier, 1984; Rolls & Williams, 1987; Robbins, Giardini, Jones, Reading, & Sahakian, 1990; Reading, Dunnet, & Robbins, 1991). We present models of conditional behavior (inspired in part by the related ideas of Rolls & Williams, 1987) in which distributed patterns of cortical activity representing visual, spatial, and temporal features selectively interact

with striatum, yielding the execution of context-dependent movements. We consider context in terms of stable patterns of cortical neuronal activity that implement the more abstract notion of *internal state*. Through learning, these patterns of cortical activity become causally linked to selection of the correct one of multiple targets for a saccadic eye movement.

We address two tasks, and two corresponding sources of context, proceeding from data obtained by neurophysiological study of behaving monkeys, to neural network models studied by computer simulation. In the conditional visual discrimination (CVD) task (Fig. 1A1), one of two saccades must be chosen based on the presentation of a fixated visual cue (Dominey, Joseph, & Arbib, 1992). The visual cue provides an *external source*




Figure 1. Association and sequence tasks. (A1) Cue-saccade association (CVD). Central cues are assigned to one of the two targets, and the subject must learn by trial and error to make a saccade to the correct target for each cue. Timeline below. After the cue is presented the subject must fixate it. During this fixation the two targets appear. Removal of the cue is the saccade go-signal. Arrow indicates correct saccade for this cue. (A2) In the generalized cue-target association, the target positions are varied so the subject must learn that a cue specifies the relative, rather than absolute position of the correct cue. Here the correct target is still the rightmost, though its absolute position is now the left of the cue. (B1) Sequence reproduction task (modified from Barone & Joseph, 1989a). The three targets are up (U), right (R), and left (L) of the fixation point (FP). The sequences are of three targets with no repetition, yielding six possible sequences. While FP is fixated, the three targets are presented in a sequence. FP is removed and the targets that have not yet been visited are illuminated three times in succession to form three go-signals. The subject must saccade to the targets in the order they were presented, triggered by the go-signals. This is the tonic-removal format: targets are tonic, and removed after visited. (B2) Modification of sequence target presentation. In this paradigm the targets are presented phasically, and all three targets are presented for each of the three go signals. Unlike the paradigm in B1, for each saccade the subject must choose from all three targets, rather than from those that have not yet been visited. Phasic-nonremoval format: targets are phasic and not removed after visited.

of context, specifying which of the two responses is correct. In the sequence reproduction task (Barone & Joseph 1989a), three targets must be sequentially chosen in the order of their initial presentation (Fig. 1B1 or 1B2). During execution of the sequence, the internal state encoding those targets that have been visited so far provides an *internal source* of context that guides the next movement. We will demonstrate that both tasks can be reduced to forming causal associations between patterns of cortical activity encoding context, and the correct context-dependant saccade choice.

In the CVD task, these patterns of activity arise through perception of the visual cue, and are then linked, by learning, to production of the correct saccades. In the sequence task, the first target provided during sequence presentation creates a pattern of activ-

ity. The second target modifies this pattern, yielding a new pattern that is again modified by the third target. This final pattern should then be able to elicit the first saccade on presentation of the "go" signal. After this saccade, the state is modified again by the resulting new visual input and the saccade motor efferent. This new pattern is set to elicit the second saccade, and so on. In summary (cf. Arbib, 1969), training produces an approximation to finite automaton behavior in which current state and current input determine both the next response and the next internal state. In both tasks, the subject must choose the correct one of multiple targets, and that choice is driven by context—in one case provided by a cue and in the other case provided by visual input and stored state of previous actions. Both tasks are learned in a trial-and-error paradigm, with a correction procedure such that a given trial is repeatedly presented until the correct response is made, then the next trial is randomly selected.

In the simplest versions of the CVD task, the spatial targets are always presented in the same fixed positions. We also consider the problem of spatial generalization, in which this task is learned and performed with the targets presented in different locations. For example, instead of learning that cue R is associated with a saccade to a fixed target 10° right of the cue, one must now learn that cue R is associated with a saccade to the rightmost target, regardless of its absolute location (Fig. 1A2). We present the CVD model¹ that learns the "fixed location" tasks and then show that, when properly trained, it is capable of learning the generalized tasks, including performance with novel target configurations. In the following sections, we will develop first the model for CVD learning, and then build to the sequence reproduction model.

THE ROLE OF BASAL GANGLIA IN THE "BASE" OCULOMOTOR MODEL

The present model extends that of Dominey and Arbib (1992),² which thus provides the "base model" for the current study. There, the control of voluntary saccades to visual and remembered targets is modeled in terms of interactions between posterior parietal cortex, frontal eye fields, the basal ganglia (caudate and substantia nigra), superior colliculus, mediodorsal thalamus, and saccade generator of the brainstem. In the present section, we define the role of the basal ganglia in the oculomotor model, and then show how plasticity may extend and clarify that role.

Briefly (Fig. 2), reflex saccades may be elicited by the projection of the superior colliculus (SC) to the brainstem saccade generator (SG). For voluntary saccades, frontal eye fields (FEF) command (via SC and directly to SG) saccades to visual or remembered targets. Two inhibitory nuclei of the basal ganglia, caudate (CD) and substantia nigra pars reticulata (SNr), are arranged in

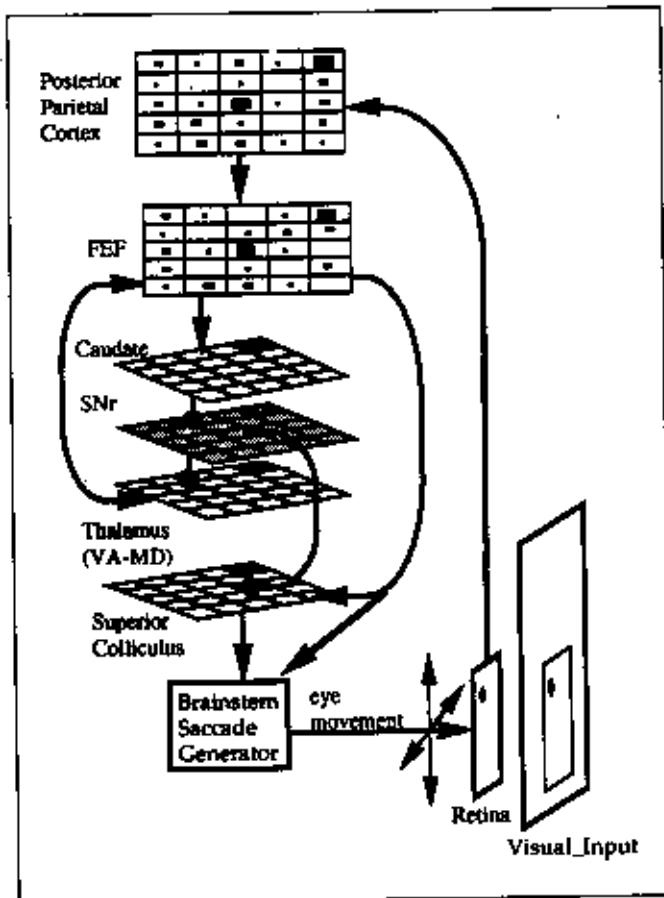


Figure 2. The role of basal ganglia and inhibitory masks in the base model (based on Dominey & Arbib, 1992). Visual input to the retina arrives in PP, influencing FEF effects saccades directly via its projection to SC, and indirectly via the striatonigrocollicular path. SNr tonically inhibits SC, forming an inhibitory mask on SC. This mask is selectively modulated by the influence of CD on SNr, releasing a restricted region of SC from SNr inhibition. Thalamus and FEF interact via reciprocal excitatory projections when SNr's inhibition of thalamus is temporarily removed. Thus, CD manages saccade activity by his influence over the inhibitory SNr that in turn controls the activity of saccade-related activity in SC, thalamus, and FEF. SC and FEF output produce saccades via the brainstem saccade generator, producing an updated retinal image.

series and provide an additional, indirect link between FEF and SC. The FEF has an excitatory topographic projection to caudate that preserves saccade amplitude and direction (Bruce & Goldberg 1984; Segraves & Goldberg 1987; Stanton, Goldberg, & Bruce, 1988a). This link allows FEF to selectively modulate the tonic inhibition of SNr on SC and thalamus (Chevalier, Vacher, Deniau & Desban, 1985; Boussaoud & Joseph, 1985; Hikosaka & Wurtz, 1985) through caudate nucleus (Stanton et al., 1988a). In addition, the basal ganglia pathways provide a mechanism for spatial memory in corticothalamic interactions via the removal of SNr inhibition on the mediodorsal thalamus (MD) (Deniau & Chevalier, 1985; Ilinsky, Jouandet, & Goldman-Rakic, 1985).

Modeling Conventions

The model is implemented in Neural Simulation Language (NSL) (Weitzenfeld, 1991) on a UNIX workstation. Brain regions are modeled as one- or two-dimensional layers of units, each corresponding to a population of neurons. The internal state of each layer is represented by a single variable $m(t)$, which may be interpreted as an array of average membrane potentials of pools of nearby cells of a given type. The time course of m is described by an array of differential equations of the form

$$\tau_m \frac{dm(t)}{dt} = -m(t) + S_m(t)$$

Here τ_m is the time constant for the rate of change of $m(t)$. $S_m(t)$ represents the total input that cells of type m receive from other cells. We use the Euler method to solve the differential equation. The firing rates are then determined by a nonlinear function of the membrane potential

$$M = \text{sigmoid}(m, \text{min_input}, \text{max_input}, \text{min_output}, \text{max_output})$$

where M is the firing rate, m is the membrane potential, and for $\text{min_input} < m < \text{max_input}$, the firing rate is a nonlinear function of m :

$$(\text{max_output} - \text{min_output}) * [S^2 * (3 - 2S)] + \text{min_output}$$

where

$$S = (m - \text{min_input}) / (\text{max_input} - \text{min_input})$$

For $m < \text{min_input}$ or $m > \text{max_input}$, the firing rate is min_output or max_output , respectively.

For each layer, we specify the time constant, the inputs to the layer, and the threshold values on the sigmoid function used to compute the firing rate from the membrane potential. Connections between these two dimensional arrays of neurons are defined in terms of interconnection masks that describe the synaptic weights. Consider the following equation where A, B, and C are layers of neurons and $M1$ is a 3×3 connection mask:

- $\tau_A = 10$ msec
- $S_A = C + B * M1$

This states (a) that the membrane time constant for A, τ_A , is 10 msec, and (b) that for each cell ij in layer A the input S_A is the sum of the output of the ij cell in C, plus the sum of the outputs of the 9 cells in B centered at ij times their corresponding weights in $M1$. That is, the $*$ operator in " $B * M1$ " indicates that mask $M1$ is spatially convolved with B. In biological terms, this mask $M1$ defines the receptive field of A in terms of the projection from B. The $*$ operator is "overloaded," that is, when applied to a layer and a scalar it represents simple multiplication; when applied to a layer and a mask it represents convolution.

The Base Model

Here we define the base model, starting with retinal input, through the processing that leads to a saccade. After each saccade, the external visual world in space (Visual_Input) is shifted on the retina an amount equal and opposite to the saccade:

$$\text{retina}(k, l) = \text{Visual_Input}(k + [l - 2], l + [j - 2]) \quad (1)$$

where the pair (i, j) is the eye position in the visual field, calculated in terms of the previous eye position and the saccade that updates this position. Retina, and all other arrays of cells (except Visual_Input which is larger, and V4, which is 1×6) will be a 5×5 array (with the coordinates of each 5×5 array running from 0 to 4) in the simulations reported here (though Dominey & Arbib, 1992 use 9×9 arrays). This coarse grain model is adequate for the present challenge of showing the adequacy of corticostriatal plasticity to mediate conditional visual discrimination and sequence reproduction, but will be refined in subsequent studies that address finer details of the interaction between the various brain regions.

Posterior parietal (PP) cortex gets this retinal input, and random activity from the Noise layer. (Noise was not considered by Dominey & Arbib, 1992.) Values from retina are 0 for no target and 70 for a target. Individual values in the Noise layer vary on a given trial between 0 and 15. "Fovea on" (FOn) cells receive input from the central element of PP, and are used to provide a signal that indicates whether a target is being fixated or not.

$$\begin{aligned} \tau_{pp} &= 0.01 \\ S_{pp} &= \text{Retina} + \text{Noise} \\ PP &= \text{sigmoid}(pp, 0, 85, 0, 110) \\ \text{FOn} &= PP(2, 2) \end{aligned} \quad (2)$$

FEF is influenced by PP and thalamus (Thal), and gets an inhibitory influence during fixation (i.e., when the central element of PP is active), preventing distractibility during fixation.

$$\begin{aligned} \tau_{fef} &= 0.01 \\ S_{fef} &= 0.4 * PP + 0.8 * Thal - 0.6 * \text{FOn} \\ \text{FEF} &= \text{sigmoid}(fef, 0, 100, 0, 100) \end{aligned} \quad (3)$$

Caudate receives input from FEF. Competition between input signals is implemented by a 5×5 convolution mask that produces lateral inhibition, Inhibitory-Collaterals, with 0.5 in the central element, and -0.1 in all other locations. This simulates the lateral inhibitory effect of inhibitory GABAergic axon collaterals of the medium spiny neurons in the striatum (Wilson & Groves, 1980).

$$\begin{aligned} \tau_{cd} &= 0.01 \\ S_{cd} &= \text{FEF} + \text{CD} * \text{InhibitoryCollaterals} \\ \text{CD} &= \text{sigmoid}(cd, 0, 100, 0, 75) \end{aligned} \quad (4)$$

The tonic activity of substantia nigra is then inhibited by caudate (Chevalier et al., 1985):

$$\begin{aligned} \tau_{snr} &= 0.01 \\ S_{snr} &= 75 - \text{CD} \\ \text{SNr} &= \text{sigmoid}(snr, 0, 75, 100, 0) \end{aligned} \quad (5)$$

The superior colliculus is driven by FEF to the extent that it is disinhibited by caudate's inhibition of SNr. SC is also inhibited by cortical fixation related activity [FOn from Eq. (2)]. A winner-take-all function in SC (Didday, 1976) is used to select the maximal activity for saccade generation. The spatial coordinates of this saccade activity are combined with the current eye position to produce the new eye position that allows the corresponding new retinal inputs to be computed in Eq. (1).

$$\begin{aligned} \tau_{sc} &= 0.01 \\ S_{sc} &= \text{winner_take_all}(\text{FEF} - \text{SNr} - 0.6 * \text{FOn}) \\ \text{SC} &= \text{sigmoid}(sc, 30, 110, 0, 100) \end{aligned} \quad (6)$$

The thalamus has a reciprocal excitatory interaction with FEF and also with prefrontal cortex (PFC), described below in the sequencing model. We will see that this reciprocal interaction provides the sustained thalamocortical activity responsible for a form of spatial memory in sequence and association performance.

$$\begin{aligned} \tau_{thal} &= 0.01 \\ S_{thal} &= \text{PFC} + \text{FEF} - \text{SNr} \\ \text{THAL} &= \text{sigmoid}(thal, 0, 75, 0, 100) \end{aligned} \quad (7)$$

In the model's production of a visually guided saccade, the fixation point goes off just as the peripheral target goes on. The visual target information from retina, via posterior parietal cortex (PP), excites the FEF element corresponding to the location of the new target, which projects topographically to the caudate. This excitation, combined with the loss of the FOn inhibition, activates the corresponding CD element, which projects an inhibitory signal to SNr, resulting in the release of SNr's inhibition of the topographically corresponding cells of SC. The topographic excitatory FEF projection to SC, combined with the release of the SNr inhibition, activates the SC element that generates a motor burst signal of saccade direction and amplitude corresponding to the target location coded in FEF.

Although we will not go into the details here, the full Dominey-Arbib model shows how working memory (which stores visual targets when they are no longer visible but will later be responded to) may be supplied by reciprocal connections between FEF and mediodorsal thalamus (MD), and how posterior parietal cortex (PP) may provide the dynamic remapping of retinotopic representations when multiple saccades are made to remembered targets. The simulations reproduce simple, memory, and double saccades, without adaptation, to demonstrate a functionally correct model in which:

1. Managing the inhibitory projection from SNr to SC allows selective cortical control of target locations for voluntary saccades to immediate and remembered targets. This control is directed by FEF via CD.

2. Saccades can be driven by representational memory that is hosted in reciprocal connections between MD and FEF that are governed by connections to MD from SNr. A state change is thus driven by a representation of a stimulus in the absence of the stimulus itself, indicating a primitive symbolic processing capability.

The key point for the present paper is the observation that *managing the inhibitory projection from SNr to SC allows selective cortical control of target locations*. In the previous model, the role of CD and SNr seems functionally gratuitous (although it is necessary for a biologically plausible model that addresses the data of, e.g., Hikosaka, Sakamoto, & Usui, 1989a,b,c), since the CD disinhibition simply mimics the combined effect of the removal of FOn inhibition and the issuance of the topographic command from FEF to SC. The thesis of this paper is shown in Figure 3. Cortical and subcortical commands for a saccade are combined in a topographic map within the deep layers of SC, where a winner-take-all (WTA) network selects the largest peak of activity to form the saccade command to be executed by the brainstem saccade generator SG. However, this command will be executed only if FOn and SNr allow the activity in deep SC to reach a critical level. In the models considered above, FEF relayed the location of only one target to SC and so the role of CD and SNr did not seem crucial. However, in the present paper, we consider cases in which FEF relays the location of multiple targets. Our hypothesis is that context-dependent learning can modify the pattern of disinhibition applied at the caudate, thus favoring one target over the other(s).

CONDITIONAL VISUAL DISCRIMINATION

To address the CVD task, in which cue-saccade associations must be learned, we augment the base model with a cortical structure involved in representation of visual

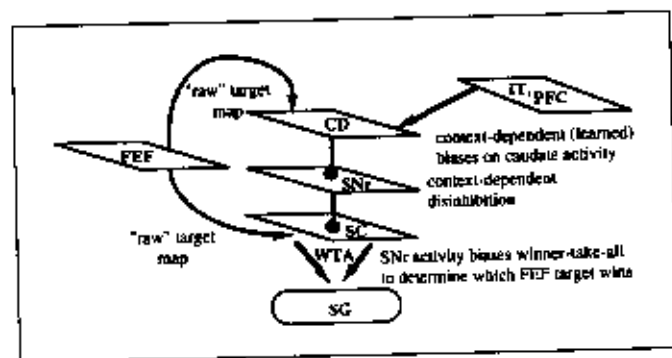


Figure 3. Corticostriatal plasticity for selective disinhibition. In this simplified illustration of the model, we demonstrate that when multiple targets are represented in the FEF command to caudate and SC, the choice between these targets can be made through the influence of additional corticostriatal influences that will favor one of the targets. IT and PFC corticostriatal projections are modified by learning to bias caudate in favor of the correct one of multiple saccade targets dependent on the association or sequence context.

cues, corresponding to the inferior temporal cortex (IT), and a learning mechanism that can modify the corticostriatal projections (Figs. 3 and 4). IT corticostriatal projections form longitudinal bands that distribute information along the anterior-posterior extent of the striatum, including target areas adjacent to and overlapping with FEF projections to the head of the caudate nucleus (Selemon & Goldman-Rakic, 1985), the striatal component of the oculomotor loop. Via this pathway, we suggest that IT can influence saccade generation as cue-related activity produces a bias in caudate that favors one over the other saccade targets.

The Conditional Visual Discrimination Model

Visual input to the model is provided to the 5x5 retina in which each element can code either a visual target or a visual cue. Such a cue is defined by its color (Red, Green, or Blue value), and shape (Circle, Square, or Diamond). These features are represented as the 6-tuple "FeatureVector" (R,G,B; C,S,D). These are graded values, allowing us to blend colors and shapes to create a variety of cues. The model will produce outputs that code eye movements by activation of an element in a 5x5 spatial map that we will consider to be an eye movement map in SC [Eqs. (1) and (6)]. For each trial the simulation will compare the model's response to the desired response and provide a reward or punishment signal as appropriate.

Stimulus Representation

The inferior temporal lobe, of the "what" visual system, provides a cortical area in which the properties of visual stimuli including color and shape are encoded in populations of cells with feature preferences (Desimone, Albright, Gross, & Bruce, 1984; Fuster, 1990; Tanaka, Saito, Fukada, & Moriya, 1991). In our model (Fig. 4), visual input is decomposed into *foveal* form (color and shape), which feeds IT, and position information on the *overall* visual field, which feeds PP. We do not model the pathways to IT and PP in any detail. Instead, a map of target positions, including that of the cue, is transferred directly to PP, while "form" is coded as a 6-element vector in an array labeled "V4."

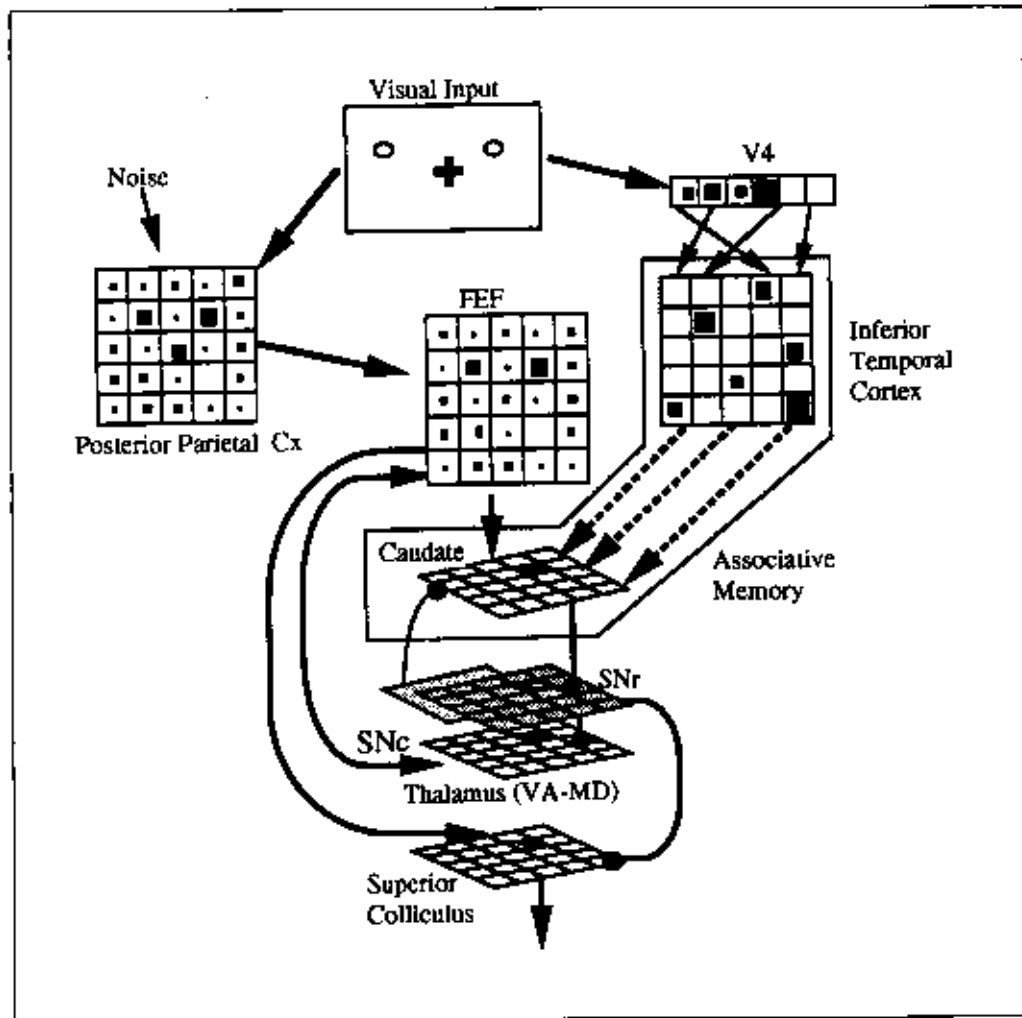
$$\tau_{v4} = 0.01$$

$$S_{v4} = 4 * \text{FeatureVector}$$

$$V4 = \text{sigmoid}(v4, 0, 40, 0, 40) \quad (8)$$

This 6 element vector that codes the cue's shape and color is transformed by a random, unmodifiable set of synapses connecting V4 and IT. V4_IT_Synapses is the 6x25 matrix of connection weights varying between -0.5 and +0.5 that, multiplied by V4, yields the input to IT. The resulting IT cells encode features, and conjunctions and disjunctions of features due to the mixed excitatory and inhibitory connections.

Figure 4. Schematic of association model. The model in Figure 2 is augmented with "V4," IT, and modifiable IT-caudate synapses, and the SNc dopamine system. The shaded region indicates the modifiable corticostriatal system. See text for explanation.



$$\begin{aligned} \tau_{it} &= 0.01 \\ J_{it} &= V4 * V4_IT_Synapses \\ IT &= \text{sigmoid}(it, 0, 40, 0, 60) \end{aligned} \quad (9)$$

$$\begin{aligned} \tau_{cd} &= 0.01 \\ S_{cd} &= 0.1(IT * IT_CD_Synapses) + 0.4 * FEF \\ &\quad + CD * InhibitoryCollaterals \\ CD &= \text{sigmoid}(cd, 0, 100, 0, 75) \end{aligned} \quad (4.1)$$


A Site for Sensory Motor Association

The caudate nucleus of the striatum participates in the preparation and execution of voluntary, context dependent saccades (Hikosaka et al., 1989a,b,c), and receives projections containing external context information from the occipitotemporal pathway—including IT—and spatial movement information from posterior parietal cortex and frontal eye fields (Selemon & Goldman-Rakic, 1985; Stanton et al., 1988a) [Eq. (4.1)].

In the association model, we update the base model Eq. (4), so that caudate receives an additional input from IT, via the modifiable IT_CD_Synapses. The IT and FEF connections are scaled by 0.1 and 0.4 based on the relative projection strengths of FEF and IT to the anterior caudate (Selemon & Goldman-Rakic, 1985). This results in initially small IT influences in CD that increase with training as the IT_CD_Synapses are modified.

In this view, caudate provides an anatomical site for the formation of context-behavior associations (Mishkin et al., 1984; Rolls & Williams 1987; Yeterian & Pandya, 1991). The intact function of the striatum is necessary for learning CVD tasks, as their learning and performance are impaired by selective striatal lesion (Reading et al., 1991). These investigators (Robbins et al., 1990) also demonstrated that the intact nigrostriatal dopamine system is necessary for learning CVD tasks. In the model, IT units encode features of the cue, and the IT corticostriatal path (IT_CD_Synapses) provides the caudate with access to this information. In the naive state, the connections in this nontopographic projection are randomized values between 0 and 1. By learning, described below, these synapses are modified so that IT cells active for a given cue excite caudate cells that participate in the associated saccade, inhibiting SNr [Eq. (5)], allowing SC to select the correct saccade. A key role for the caudate,

Figure 5. Dopamine regulation of corticostriatal activity. (A) Feedback circuit: excitatory inputs represented with arrows, inhibitory with blunt ends. Corticostriatal activation inhibits the SNr. Excessive inhibition of SNr releases inhibition on SNc, allowing an increase in SNc nigrostriatal dopamine production, attenuating corticostriatal activity. (B) Regulation of activity for two targets (L and R) in striatum. Without DA modulation both activities are above caudate saturation threshold, so their relative intensity information is lost. The effect of dopamine release in striatum is to attenuate the corticostriatal activity, so that the relative intensity information is again available.



then, is to participate in the selection of the correct one of multiple saccade options.

Regulation of Corticostriatal Input

For a given pattern of activity in IT, and multiple targets represented in FEF, multiple caudate cells with different saccade preferences may be activated. Ideally, the caudate will use the cue-driven IT bias, established by learning, to select the correct target. However, if the combined activity of IT bias and target exceeds the threshold of caudate sensitivity at more than one topographic location in CD, then the bias information from IT that would normally favor one site over another will be lost. To bring the caudate activity back into a range in which the IT bias can be detected, we introduce a negative feedback circuit that decreases corticostriatal excitability in this situation.

The circuit is based on the knowledge that the release of dopamine in striatum decreases the excitability of striatal cells by corticostriatal afferents (Mercuri et al., 1985; Garcia-Munoz, Young, & Groves, 1991; Nisenbaum & Berger 1992; Calabresi, Mercuri, Sancsario, & Bernardi, 1993), and thus attenuates corticostriatal activity, allowing striatum once again to detect the effect of IT's bias. In our model (Fig. 5A), when strong cortical inputs overstimulate striatum, SNr is strongly inhibited. This inhibition of SNr leads to the disinhibition of SNc (Carlsson & Carlsson, 1990). The resulting increase in activity of SNc DA cells leads to the increased release of DA in striatum, which attenuates the corticostriatal signal (Calabresi et al., 1993), bringing the most active cells back to the maximum input capacity of striatum, and reducing the activity of the less active surrounding striatal cells, thus increasing the discriminatory capacity of the network (Fig. 5B). We currently model this circuit and SNc DA function in the simplified form of Eqs. (10)

and (4.2). This allows caudate to make better use of cue-related information from IT, and assists in the learning and execution of cue-saccade associations by providing a kind of contrast enhancement between strong, competing striatal inputs, resulting in a 30–40% increase in performance (percent correct) during a training session.

DA_Modulation represents our implementation of the dopamine normalization of cortical inputs. In the base model of Domincy and Arbib (1992), this modulation is not present, i.e., it is effectively set to 1. Normally the term DA_Modulation is 1, but when the maximum inputs to caudate, $cd_{max}()$, is above a threshold of caudate sensitivity, DA_Modulation is reduced. This simulates the reduction in corticostriatal excitation due to increased striatal dopamine levels, evoked by disinhibition of dopamine-producing cells in substantia nigra pars compacta. This reduction of the effect of cortical activity on striatum brings it back to the normal range of operation for caudate.

$$DA_Modulation = \begin{cases} CD_Upper_Threshold/cd_max & \text{if } cd_max > CD_Upper_Threshold \\ 1 & \text{otherwise} \end{cases} \quad (10)$$

$$\tau_{cd} = 0.01$$

$$S_{cd} = DA_Modulation * (0.1(IT * IT_CD_Synapses) + 0.4 * FEF + CD * InhibitoryCollaterals)$$

$$CD = \text{sigmoid}(cd, 0, 100, 0, 75) \quad (4.2)$$

An Adaptive Mechanism

In addition to its short-term modulatory role, the observations of reward-related DA release during learning, and the impairment of learning with DA depletion suggest that dopamine plays an additional important role in long-term synaptic changes required for learning stimu-

lus-response habits. First, Ljungberg, Apicella, and Schultz, (1991) found that during learning a delayed alternation task, nigrostriatal dopamine-producing cells are activated by the reward, and cues associated with the reward, and that these responses are qualitatively preserved during conditioning, postconditioning, and overtraining. On error trials, a depression of activity was seen at the time when a reward would have been given (Ljungberg et al., 1991). Second, depletion of dopamine from the dorsal striatum in rodents impairs their ability to learn a cued-choice task (Robbins et al., 1990) similar to our cued-saccade task. Finally, the relatively high level of NMDA receptors in striatum (Monaghan & Cotman, 1985) and the observation that striatal NMDA receptor density varies inversely with striatal DA levels in patients with Parkinson's disease (Weihmuller, Ulas, Nguyen, Cotman, & Marshall, 1992) suggest a relation between DA and learning-related corticostriatal plasticity.

It thus appears that phasic, reward-related DA activity participates in long-term modification of corticostriatal synapses that link contextual sensory inputs with striatal cells involved in producing the correct behavior. Two candidate mechanisms for this synaptic plasticity, long-term potentiation (LTP) and depression (LTD), have been observed in the striatum, and at least LTD has been observed to specifically require the presence of dopamine (Calabresi, Maj, Pisani, Mercuri, & Bernardi, 1992). Repetitive activation of cortical inputs produces LTD in striatal cells whose NMDA channels are inactive, and LTP in those whose NMDA channels are active, effects that can be readily observed by manipulating the Mg^{2+} NMDA block (Calabresi, Pisani, Mercuri, & Bernardi, 1992; Walsh & Dunia, 1993). We can simplify this by saying that sufficient depolarization (which removes the Mg^{2+} block of the NMDA receptors) in striatal cells makes these cells candidates for LTP in the presence of additional cortical input, while insufficiently depolarized cells will be candidates for LTD. Recall that by reducing all striatal activity so that only the maximally active cells fire at their maximum rate, our modeled striatal regulation by DA actually pushes the less active cells into the LTD region. We now take a first step in relating these observations to a formal learning rule for our model, with the acknowledgment that the problem of assigning physiological mechanisms to the formal elements of a learning model is neither trivial nor complete.

Learning Rule

We employ a reinforcement-based learning rule (see Barto, 1990 for a review) to modify corticostriatal synaptic strength based on correct/incorrect behavior. The learning rule produces incremental changes in synaptic connections between IT (source) and caudate (target) neurons. Correct responses strengthen synapses between IT and caudate cells that participated in the response, while errors weaken these synapses. These

changes will tend to produce IT-driven caudate activity that increasingly favors correct behavior, and disfavors incorrect behavior, since IT cells specific for a given cue will become more strongly connected to caudate cells active in the associated saccade, and more weakly connected to other caudate cells. We also implement a form of competition so that for a given source cell, increases in its synaptic influence on some target cells will be compensated for by small decreases in its synaptic influence on other target cells. These increases and decreases correspond roughly to the LTP and LTD described above. The most active synapses are strengthened (LTP) and, via normalization, the less active synapses are weakened (LTD).

The learning-related updating is performed by

$$w_{ij}(t+1) := w_{ij}(t) + DA_Modulation * (RewardContingency - 1) * C1 * F_i * F_j \quad (11)$$

$$w_{ij}(t+1) := w_{ij}(t+1) * \frac{\sum_j w_{ij}(t)}{\sum_j w_{ij}(t+1)} \quad (12)$$

where the ":= " denotes assignment rather than equality. Here, w_{ij} is the strength of the synapse connecting IT cell i to caudate cell j . These w s make up the elements of IT_CD_Synapses in Eq. (4.2), and are modified after each simulated saccade. Equation (10) sets DA_Modulation as a function of its role in the SNc feedback loop. In Eq. (11), based on the arrival or denial of expected reward, we simulate SNc reward-related modulation by the term RewardContingency, which is 1.5 for correct trials, 0.5 for incorrect trials, and 1 when no reward or punishment is applied, corresponding to the increases and decreases in SNc activity for reward and error trials, respectively (Schultz, Apicella, & Ljungberg, 1993). F_i and F_j are the firing rates of the IT cells coding the cue, and the caudate cells coding the saccade, respectively. The term "DA_Modulation * (RewardContingency - 1)" will be positive on rewarded trials and negative on error trials. C1 is a constant that specifies the learning rate, and is set to $2.5e^{-5}$.

A form of competition is provided by a weight normalization procedure that conserves the total synaptic weight that each IT cell can distribute to its striatal synapses [Eq. (12)]. If, after learning has occurred, one synapse from cell i to cell j was increased, then since the total synaptic weight from cell i is conserved, the result of this increase is a small decrease in all other synapses from i . Similarly, when a weight is decreased due to an incorrect response, the other synapses from i are increased. Via this normalization process, the postsynaptic cells compete for influence from presynaptic cells, producing cue discrimination. After each saccade, the simulator applies Eqs. (11) and (12) with the appropriate value for RewardContingency. Particularly in the case of rewarded trials, Eqs. (11) and (12) approximate how LTP may occur in the most active postsynaptic cells

[via Eq. (11)], and LTD in the others [via the normalization of Eq. (12)]. While captured here in the same equations, it may be that learning on error trials also involves other structures including the nucleus accumbens (Robbins et al., 1990; Reading et al., 1991) and we consider this in the discussion.

In summary, the interaction of DA's feedback and plasticity roles is captured in Eqs. (10), (11), and (12). During initial learning, Eq. (10) will set DA_Modulation to 1, since the corticostriatal inputs are initially weak. In Eq. (11), based on RewardContingency, synapses active in saccade generation will be strengthened or weakened for correct and incorrect saccades, respectively. When synapses are strengthened for saccade related cells on correct trials, the surrounding synapses are weakened by the normalization of Eq. (12). After significant learning, strong IT corticostriatal inputs may be above striatal threshold for more than one target, even though this input is biased toward the correct target. Equation (10) reduces DA_Modulation so that the maximum IT-striatal input is just equal to the striatal threshold, increasing the signal-to-noise ratio so that the correct bias is now perceived in striatum, leading to the correct choice. In addition, by reducing DA_Modulation, the synaptic change for this well learned cue is reduced (but still positive) in Eq. (11), helping to prevent an excessive allocation of synaptic weight to an already well-learned association.

Simulation

In this section we describe the evolution of behavior from an initial naive state to a final trained state for our model of CVD learning, and compare the behavior and single unit activity of the model with those of primates performing related tasks. Recall from Figure 4 that the site of adaptation is in the IT corticostriatal synapses, so that representation of cues in IT can influence saccade generation. These modifiable connections are initialized to random values on the interval (0,1), so a given cue has no large a priori preference for any saccade.

A trial is initiated by the onset of the central fixation point. Once the model initiates fixation, the fixation point is replaced by a colored shape (cue). The "form vector" (color and shape) of the cue is projected by "V4" to IT, and the IT activity projects to caudate via the modifiable synapses, IT - CD Synapses [Eq. (4.2)].

After a delay, the two peripheral targets appear, leading to the activation of PP, which in turn signals to FEF and then to caudate. Each element of PP receives its topographic spatial input, to that is added a "noise pattern" that contributes up to 21% of the total PP activity, and provides a "symmetry breaking" function—a form of guessing between the two targets during initial learning.³ During the period of cue and target overlap, the combined activity of FEF and IT produces "preparatory" activity in the caudate, leading to its inhibition of SNr,

which in turn disinhibits thalamic nuclei that project back to FEF [Eq. (7)]. Reciprocal excitation between disinhibited thalamic and FEF cells, combined with the lateral inhibition in caudate, tends to amplify slight differences in corticostriatal inputs, favoring the stronger and reducing the weaker. This is similar to the reciprocal FEF-thalamic interactions that maintain spatial memory in Dominey and Arbib (1992), but in the present case, the reciprocal activity contributes to both selection and memory functions, as we will see below.

At the go signal (removal of the cue), inhibitory fixation-related cells in FEF (FOn) are deactivated, allowing SC to respond to inputs from FEF and SNr. In the naive state, the cue-driven IT activity projects with randomized strength to caudate, while the effects of noise in PP are seen in caudate and downstream, and drive the winner-take-all (WTA) mechanism in SC resulting in a saccadic "guess." After a correct "guess" the IT-caudate synapses between cue and saccade-related cells are strengthened, increasing the probability that the cue, when repeated, will drive the correct caudate cells [Eqs. (11) and (12)]. Incorrect trials result in a reduction of the active corticostriatal synapses. After significant learning, the IT influence on caudate dominates that of noise, and the performance approaches 100%.

We trained the model on four sets of cue-target associations using two fixed targets, left and right, one row above the fixation point (Table 1). Two cues were associated with the left, and two with the right target. As the training proceeds, learning overpowers the effect of noise and performance improves. We observed that the rate of learning varies inversely with the level of noise in PP [Eq. (2)], and that for a primate trained in this task, the learning rate improves with the number of new associations learned (also see Mitz, Godshalk, & Wise, 1991). This suggests a progressive reduction of noise in the animal as familiarity with the CVD task itself is acquired.

We examined the task-related activity of caudate cells during initial training, intermediate, and overtrained

Table 1. Learning Progression for four Cue-Target Associations.^a

Assoc. No.	Cue Feature Vector	Direction	E1	E2	E3
1	(10, 0,0;10, 0,0)	Left	11/15	15/17	16/16
2	(0, 10, 0;0, 10, 0)	Right	11/14	15/15	16/16
3	(0, 0, 10;0, 0, 10)	Left	11/15	14/14	16/16
4	(8, 2,0; 0,2 .8)	Right	11/17	14/17	16/16
Correct by epoch (%)			72	92	100

^a Each cue is represented by a six-element feature vector (Cue Feature Vector), and is arbitrarily paired with either the left or right target. Results for successive training epochs display the performance improvement as the associations are learned, and overpower the effects of noise.






Figure 6. Changes in caudate response due to learning. Upper: Cue-related activity in IT, caudate, and SNr before learning the cue-saccade association. IT influence on caudate, for a given cue, is non-specific and does not significantly contribute to selection of a saccade. Middle: After learning the IT_Caudate_Synapses are modified, and the same cue now produces activation of caudate and inhibition of SNr appropriate for the associated saccade. Lower: After generalization learning. The IT influence on caudate now forms a spatial gradient of activity that will favor the rightmost of any two targets presented on the horizontal row above the fixation point.

epochs. Recall that in the model, CD receives input from FEF [Eq. (4.2)], so that even in the naive state, caudate cells have saccade-related activity with a spatial preference. As training proceeds, cue-related activity in IT becomes increasingly associated with activation of the correct saccade-related cells in CD. In Figure 6 we display the changes in IT-striatal synapses and the corresponding changes in cue-driven striatal activity that

result from learning. The upper panel of Figure 6 shows the activity produced by a "go right" cue in IT; the initial (naive) synaptic weights between IT and caudate, and the corresponding influence produced in caudate and SNr by that cue. Before training, the cue-driven caudate activity is low and homogeneous. Figure 6 (middle panel) shows the change in synaptic weights and the corresponding change in caudate and SNr activity elic-

ited by the same cue after training to 95% performance for two (left and right) cues. Note that via training, the weights are now redistributed into the two columns corresponding to learned association between two cues and the left and right targets. The "go right" cue produces an increased activation of cells related to rightward saccades, and a small decrease, due to our normalization model of LTD, in other cells. This cue also produces some activity for the left saccade target because its representation in IT overlaps partially with that of the "go left" cue. Thus we see how, through learning, cortical activity driven by a cue produces an inhibitory mask in SNr that guides the correct saccade choice.

Comparison of Model and Data

During the intermediate stages of training, before this transition between noise-driven vs. cue-driven saccades occurs, we see that activation of caudate cells during the cued period accurately predicts the direction of the upcoming saccade, independent of the cue. Figure 7 illustrates a cell with leftward saccade preference in different phases of training. In the first row, the cell is shown with well learned cues for left and right saccades. In the second and third rows during the intermediate phases of training when mistakes were still being made, we see that activation of this cell during the cued period predicts saccade direction independent of the cue, or the correctness of the saccade. Note on the second row, the differences in activity of the cell during the cue-only

period on these two trials. When noise favors the leftward saccade, a stronger response is seen both during the cue and cue with targets periods. When noise favors a rightward saccade, the influence of the cue is much less apparent, indicating how at this stage of learning, noise is more influential than cue-related activity. On the incorrect saccade (middle row, right column) the cell shows an intermediate response for its preferred direction, even though the opposite saccade was made. Once the IT-caudate learning dominates the influence of noise, errors are greatly or completely reduced, and activation of these cells predicts correct cue-guided saccades, dependent both on the cue and the correct saccade. A noteworthy property of these cells is that equivalent cues (i.e., having different color and shape, but associated with the same saccade) produce similar activation in their associated saccade-related caudate cells. This can be seen for the four correct (arrow and shaded target match) cue-guided saccades in Figure 7.

Our preliminary data from a cue-guided saccade experiment in rhesus macaque monkeys (Dominey, Joseph, & Arbib, 1993) indicate that task-related caudate cells encode combinations of cue and saccade properties, as suggested by the model. While few cells showed purely cue-related responses, almost one-third of the task related cells respond to cue onset with an increase in activity that reliably predicted subsequent saccade direction (Dominey et al., 1993). A typical cell of this type is shown in Figure 8. Like the simulated cell in Figure 7, for intermediate stage learning this cell responds during

Figure 7. Simulated saccade preparation cell. Arrow indicates direction of saccade, dark square indicates correct direction. All traces are from a cell that has preference for leftward saccades. Left column: leftward saccades. Right column: rightward saccades. See main text.


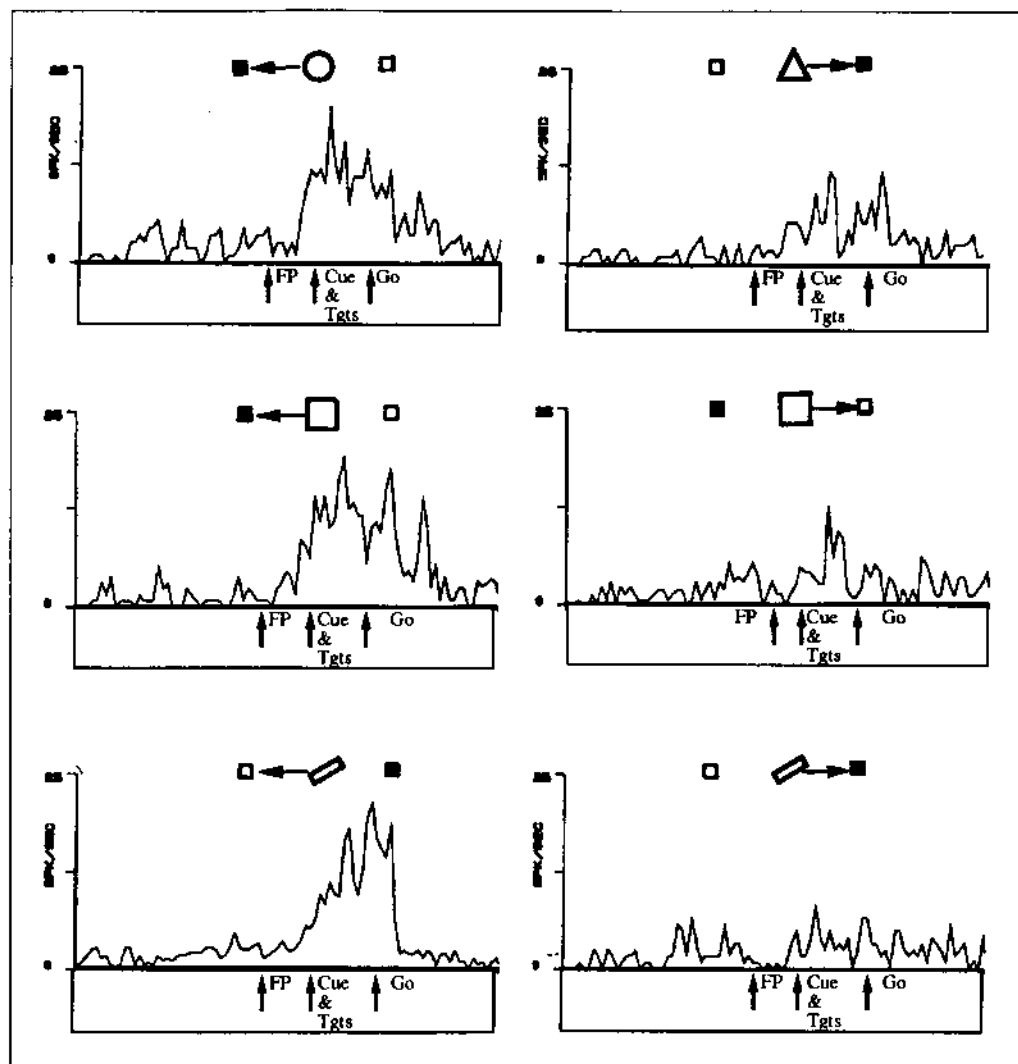


Figure 8. Saccade preparation cell. Same notation and arrangement as in Figure 7. Vertical axis firing rate 0-25 spikes/second. Horizontal axis, trial time: 9000 ms, aligned on cue onset. See text.



the cue-target overlap, and predicts saccade direction, independent of the cue, whereas for well learned cues it predicts cue and saccade direction. The two cues in the top row were well-learned and the only errors were fixation related. On the second and third rows, we see correct and error trials for cues that were being learned. As in the model, for this intermediate learning, the cell predicts saccade direction, independent of cue. Interestingly, in both model and animal, the activity increases near the onset of the cue, well before the saccade itself. In the model, this early response is produced by combined cue-related activity from IT, and the noise/target-related activity in FEF in PP. Competition in striatum via inhibitory axon collaterals (Wilson & Groves, 1980), along with the FEF thalamic interactions (Illinsky et al., 1985), tends to increase the activity in the area favored by noise or learning, decreasing the activity related to the disfavored target, thus producing the effect of pre-saccade direction prediction.

Hikosaka et al. (1983c) note that a common characteristic of caudate cells is their predictive or anticipatory activity. We see this in Figure 8, where, before the cue

onset, the cell begins to fire as if anticipating the cue (or preparing for its preferred saccade), and continues to fire if it is a preferred cue, while stopping if not. This anticipatory activity appears to demonstrate a temporal (as opposed to visuospatial) association capability that may represent preparation for, or expectation of a future event, and poses a challenge for future modeling.

It has been noted that many responses in basal ganglia neurons to external stimuli are not primarily sensory, but instead are related to the meaning of the sensory inputs (Schultz, 1989). For example, in a CVD go-nogo task roughly one-third of caudate cells responded to the trigger stimulus, and one-quarter of these responding cells were activated only in the go trials. When the go-nogo meaning of the stimuli was reversed, the majority of these differentially responding neurons also reversed their stimulus preference, indicating that the neuronal response was related to the behavior itself, and not to the stimulus that triggered that behavior (Rolls, Thorpe, & Maddison, 1983). In our own modeled reversal experiments the same is true. In reversal, a behavior that was previously rewarded is now punished. As the initially

strong IT-caudate synapses are weakened during the reversal, there is a perseverance until the original behavior is sufficiently weakened so that noise can again contribute to guessing. The new behavioral association is gradually formed, first via guess and then via increased effects of learning the new association, yielding finally a cue-guided response for the opposite saccade direction.

Prediction #1

In a related set of experiments on remembered saccades, Hikosaka et al. (1989a) demonstrated sustained activity during the delay period of a memory saccade task in caudate cells that had spatial preferences for the remembered saccade locations. This activity was elicited both in cases where the target was presented transiently before the memory period, as well as in cases where the saccades alternated between two left-right targets so that the next saccade direction could be predicted from the previous saccade. In this second case, cells with spatially selective movement/memory fields were activated without a visual stimulus falling in that field. In




Figure 9. Cue-guided memory saccade. During fixation, a well-learned cue is briefly presented, then replaced by a neutral fixation point. At removal of the fixation point, the two targets appear, and the subject must saccade to the correct target. The brief cue presentation activates the topographic region of caudate corresponding to the learned saccade. This in turn inhibits SNr, releasing SNr inhibition on thalamus. Background cortical activity now excites thalamus, and the reciprocal interaction between FEF and thalamus amplifies this activity, seen as a rise in FEF and thalamus traces. After removal of the cue this activity remains, in a form of spatial memory that is activated by central cue. Following the go signal (fixation point removal), this memory effect of sustained reduction of SNr activity at the topographic location for the correct saccade leads to collicular activation for that saccade, now that the inhibitory influence of FEF on SC and CD is removed.

Figure 9, we show the results of a modified version of this memory saccade task in which the cue itself (rather than a target flash or the previous saccade) informs the correct saccade. The cue is presented phasically during fixation, activating the reciprocal thalamocortical circuit via striatonigral disinhibition. A "spatial memory," evoked by a central cue rather than the target, is stored in this circuit until the go-signal occurs and the saccade is executed without the target being presented. In this case, we predict that for well-learned cues a form of conditioning will be observed: saccade-related caudate cells will be driven by the central cue, and continue to fire after the cue is removed, through the go signal until the correct saccade, without a target stimulus falling in the cells' receptive field.

Learning of Spatial Generalizations

The above model demonstrates how a cue-driven pattern of neural activity in inferotemporal cortex (IT) can become associated with a saccade to the correct one of two fixed spatial targets. An example is illustrated in Figure 6 (middle panel) where cue-driven IT activity modulates SC's inhibitory mask, SNr, to select for activation in SC appropriate for a saccade to the fixed right target. We are now concerned with learning the more general sensorimotor association that will allow the same cue to guide the choice of the rightmost of any two targets, as illustrated in Figure 1A2, without ad hoc modification of the model to accommodate this requirement. Thus, the same cue-related IT activity should modulate SNr so as to select the rightmost of any two targets, even if these two targets are arranged in new positions.

We investigated the learning of spatial generalizations with the model, using multiple target configurations with two well-known cues (from associations 1 and 2 in Table 1), one each for left and right. Three experiments were performed to compare generalization under different conditions. We observe that for a given horizontal row of five positions, there are 10 different left-right configurations of two targets with which we can test spatial generalization.

In Experiment 1, with two of the well-learned cues, we first trained the model using 6 of the 10 target configurations as a training set (generalization learning, GL), and then tested its ability to perform on the remaining four target configurations (generalization test, GT). As a control (Experiment 2), we then trained the model on a *single* configuration of targets for the same number of trials as in the training set above (nongeneralization learning, NGL), and then tested it on the generalization configurations (GT). Finally, in Experiment 3, we tested the untrained model only on the test set (GT). In Table 2, we see that performance on the test set (GT) is best if some generalization training precedes the testing (Exp. 1), and that the model in the naive state performs better

Table 2. Summary of Three Experiments on Generalization as a Function of Pretraining.^a

Experiment 1				Experiment 2				Experiment 3
GL1	GL2	GL3	GT1	NGL1	NGL2	NGL3	GT1	GT1
48/64	60/66	60/66	52/66	56/65	67/67	67/67	3/63	29/64
75%	91%	91%	79%	86%	100%	100%	5%	45%

^a In Experiment 1, we expose the model to a generalization learning set (GL), and then a generalization test set (GT). In control Experiment 2, we train an equal time on a fixed, nongeneralization learning set (NGL), and then test the generalization (GT). Note the poor generalization in this case. In Experiment 3 we expose the naive model to the GT.

on the test set (Exp. 3) than the model that was trained on only a single configuration of the targets (Exp. 2). This is because biases toward the two fixed targets can now cause them to be chosen over the new targets in GT, even if incorrect (Exp. 2), and these biases must now be "unlearned."

In Figure 6 (lower panel), we display the activity in caudate and SNr during presentation of a rightward cue after generalization learning on the entire set at performance above 95% correct. Note that the activity in caudate forms a "spatial gradient" that is monotonically increasing as we move from left to right. When combined with target activities for left and right targets, this gradient ensures that the caudate will favor the rightmost target. The cue-related SNr gradient is more localized, but is functionally equivalent when target and cue-related activity are combined in CD and then SNr. Thus, instead of forming an inhibitory mask that allows activity at only a restricted site (as in Fig. 6, middle) the generalized learning yields a mask that selects the leftmost of two targets in any location on the trained row. A more conceptual depiction of the gradient for spatial generalization is shown in Figure 10. What is significant is that the theoretical solution to spatial generalization provided by the "disinhibition gradient" of Figure 10 was "discovered" by our learning model with the same circuitry used to relate cues to targets in *fixed* positions.

We note that extensive training with the targets distributed above the fixation point resulted in an over-

representation of this area in the synaptic weights, and a corresponding underrepresentation of the lower region, leading to poor performance when targets were presented there. A similar problem was seen in the initial training of one monkey, who initially refused to respond to targets in the lower visual field after extensive training with targets only in the upper field (unpublished observations).

The continuous shape of this gradient results from the simple statistical properties of the distribution of rewards for each target location. Consider a cue to be associated with saccades to the rightmost target. In our set of 10 target pairs, as a target is farther to the absolute right, the number of target pairs in which it is rightmost increases. Thus, over many trials, more rewards will be received at more rightward positions so that these positions will receive more influence from the cue, thus forming a cue-driven gradient of activity in caudate that is strongest at the right and diminishes to the left. It is known that randomized presentation of training data leads to a greater generalization capability than does blocking of trials (e.g., Lee, Magill, & Weeks, 1985). We have observed this in simulation runs where blocked trials tend to produce humps in the middle of the gradient, leading to subsequent errors, whereas a random distribution of target configurations tends to preserve the monotonic nature of the gradient. While we will not present the details here, we have also observed that presenting the trials so that the gradient is built up at the rightward extreme and gradually extends to the left produces a smooth gradient in a minimal number of training examples and with a minimal number of errors.

Prediction #2

In a related study of target selection based on relative position, Niki (1974) used a left-right delayed alternation task, with multiple target configurations. Recording single prefrontal cortex units, he found one population of cells that was selective for relative target position during movement preparation, and another population selective for absolute target position that was active during movement execution. Our model predicts that, once the gradient is learned, cells near the favored side of the gradient will predict *relative direction* of the forthcom-




Figure 10. Spatial gradient implements spatial generalization. In solving the spatial-generalization problem, by applying a gradient of disinhibition prior to a winner-take-all process: for two targets of equal strength, the given gradient ensures that the rightmost target wins. The gradient itself, then, implements the concept of "rightmost."

ing movement, as soon as the cue is present, whereas cells at the locations specific for the actual saccades will predict the *actual saccade direction* following the target presentation, prior to and during the saccade.

SEQUENCE LEARNING

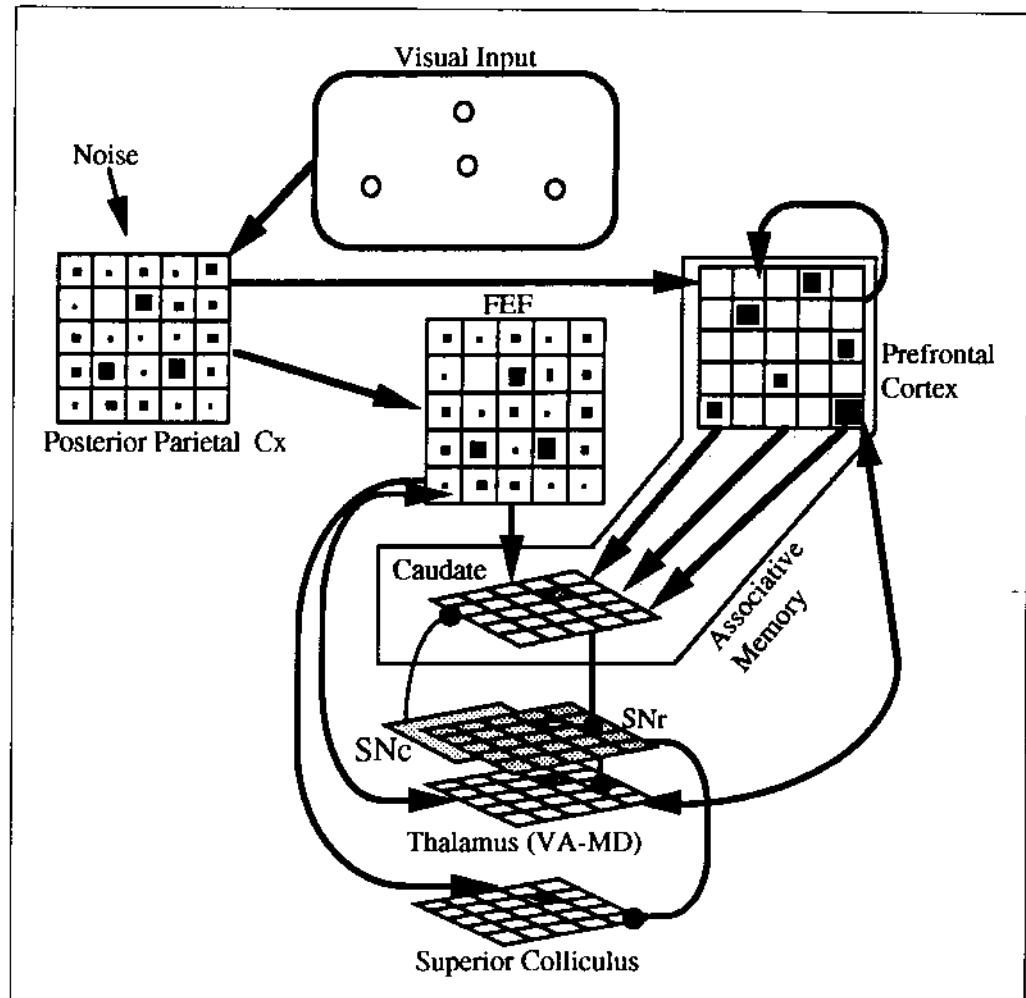
The previous model learned static input:output mappings. In the sequence task the mappings learned are time varying. To address this temporal processing, we augment the model with a cortical structure involved in representation of spatial and temporal features, corresponding to prefrontal cortex (PFC ≈ superior arcuate area and caudal part of the principal sulcus), along with its corticostriatal projections, and a learning mechanism that can modify these connections (Fig. 11). Like those of IT, the PFC corticostriatal projections form longitudinal bands that include target areas adjacent to and overlapping with the FEF projections to caudate nucleus (Selemon & Goldman-Rakic, 1985), the striatal component of the oculomotor loop. In the model this is reflected by overlapping the PFC and FEF projections in the caudate. We will identify information processing re-

quirements for the sequence task, and then present our implementation based on known biological constraints.

Neural Coding for Sequence Reproduction

Barone and Joseph (1989a) recorded prefrontal cortex neurons in monkeys trained to observe and reproduce sequences of three lighted push buttons arranged, respectively, above, to the left, and to the right of a central lighted fixation point. In the first phase of each trial, the three targets were turned on in random order; in the second phase, the animal had to saccade to and press each target in the order of their illumination, cued by full illumination of the targets that remained to be pressed three successive times, as illustrated in our Figure 1B1. Three hundred and two cells were recorded from the superior arcuate area and caudal part of the principal sulcus, and were classified as visual tonic (VT), fixation, context, saccade-related, and visual-phasic cells. VT cells (35/302 = 11.5%) showed sustained activation during fixation of the central fixation point (FP) following onset of one of the three targets, and were spatially selective, similar to the sustained activity of real (Barone & Joseph, 1989b; Funahashi, Bruce, & Goldman-Rakic,

Figure 11. Schematic of sequencing model. Caudate saccade-related cells are influenced by topographic projections from FEF and also by modifiable, nontopographic projections from PFC. Shaded region indicates that the modifiable PFC-to-caudate projection is the adaptive component of the model.



1989) and simulated (Dominey & Arbib, 1992) FEF cells during memory saccades. In addition, the VT cells also had a preference for the temporal order in which their spatially preferred target appeared. Fourteen VT cells were selective for the first target in the sequence, 10 cells were active for the first or second targets in the sequence, 8 were active just for the second target, and 3 cells were active for the second or third target. In one group of VT cells, the saccade to the target reset the cell to its pretrial firing rate, whereas a second group of VT cells was not reset with the targeting saccade.

Context cells made up the largest class of task-related units ($116/302 = 38.5\%$). They were activated during the fixation of a given target in the sequence execution phase, contingent on which other targets had been, or were going to be pressed. Of these cells 96 were spatially selective, in that they had a preference for spatial location of the target. These cells seem to reflect the state of execution of the sequence and could be considered to provide a state-transition function for the system.

For all six of the possible sequences of three elements, the first trigger is fixation point removal, and the lighting at full strength of all targets. This ambiguous trigger must be augmented with some internal information to resolve the ambiguity. Recall that many of the VT cells showed sustained activity only when the target in their visual field was first in the sequence. The pairing of these "rank 1" VT cells, activity with the ambiguous trigger is now a unique trigger. Consider first the sequence UR. On presentation of U, a subpopulation of the VT cells with the appropriate spatial receptive field becomes active. We will refer to this subpopulation as VT1. On presentation of R, the activity in VT1 combined with stimulus 2 will activate subpopulation VT2. Now, when the trigger is presented, the animal will have an ensemble of activity in the VT cells, which, when combined with the ambiguous trigger, will produce a unique combination of internal and external activity. The point is that cortex is organized so that presenting the different sequences produces different patterns of cortical activity.

During the learning required to execute this new two-item sequence, following the presentation of the two targets and then the first "trigger," the animal will at some point randomly saccade to the correct first target, U, and get rewarded,⁴ strengthening the association between the response to target R and the cortical pattern of activity or "motor set" produced by observing sequence UR. The new internal state, including activity of context cells that respond after the saccade to U, resulting from the new visual input and saccade efferent copy, forms a new "motor set" that will be associated, through learning, with target R. Thus, it appears that temporal structure is encoded in terms of transitions between internal representations of context. In the training of the animals, all six two-item sequences were learned first,

before moving to three-item sequences. This suggests that a three-item sequence might be learned as a combination of a two-item sequence and a third element. Indeed, in Barone and Joseph's (1989a) report, cells were found that were activated for a target only when it came first, either in a two-item or three-item sequence.

The Sequence Learning Model

As in the association model, the visual input is provided by a 5×5 spatial array. As illustrated in Figure 1B, the input consists of a fixation point and a sequence of three visual targets. After a delay, the model must reproduce the target sequence as a sequence of saccades to the targets in the order they were presented. Each saccade moves the "eye" so that the retina is centered on the saccade target [Eq. (1)]. For each trial the simulation will compare the model's response to the desired response and provide a reward or punishment signal as appropriate, triggering the application of Eqs. (11) and (12) for the modification of PFC to caudate synapses.

Processing Requirements

We consider sequence reproduction in terms of transitions between successive internal states, such that each state evokes the correct next saccade, and its execution leads to a transition to the next state. The ability to learn to recognize and then reproduce a sensorimotor sequence requires the following: (1) After the ordered presentation of the targets, the model must be in a state to produce a saccade to the first target. (2) After the n th saccade, the model must enter a state in which it will produce the $(n+1)$ th saccade, and so on. As illustrated in Figure 1, in this sequence task, the central point is fixated while the three peripheral targets are illuminated in some order. This activity pattern can then be associated with the first saccade of the sequence. Execution of this saccade results in a new visual input, since the eyes move to the first target. Modification of the previous internal state by the combination of the new visual input and a "copy" of the previous saccade produces a new pattern of activity that should be sufficient to specify the next move, and so on.

The sequencing problem requires that the brain contains a site that receives (1) visual input from the "where" system, (2) a form of motor efferent copy of saccade outputs, and finally (3) some form of self-input for maintaining state information through time.

Internal Representation of Temporal Events

PFC provides a likely anatomical site for representation of context from internal and external sources in this sequencing task. Visuospatial activity from PP provides the spatial "where" input to PFC (Goldman-Rakic, 1987),

satisfying (1). Thalamic inputs to PFC provide a form of simple spatial memory, as seen in the FEFmem cells of our original model (Dominey & Arbib, 1992), and the real FEF (Bruce & Goldberg, 1984) during memory saccades. Postsaccadic activity corresponding to the previous saccade(s), as seen in the FEF (Bruce & Goldberg, 1984) is provided to PFC via a thalamic relay of SC activity, SC_POST [Eq. (15)], satisfying (2). The final input to the PFC [Eq. (13)] is a damped self input [Eq. (14)] with randomized projections varying between -0.6 and 0.4 , satisfying (3). This allows for complex recurrent memory loops, and also for the inhibitory interactions that will produce order selectivity. The negative bias (i.e., weights ranging between -0.6 and 0.4) prevents a runaway self-excitatory saturation of PFC. SC_PFC, PP_PFC are connection matrices with synaptic strengths that vary between -0.5 and 0.5 , respectively, providing a balanced inhibitory and excitatory signal of SC and PP activity, respectively. The constants for each term represent their relative contribution to the state maintenance function of PFC.

$$\begin{aligned}\tau_{pfc} &= 0.01 \\ S_{pfc} &= 2.0(PP * PP_PFC) + \\ &\quad 1.5(\text{Damped_PFC} * PFC_PFC) + \\ &\quad 5(SC_POST * SC_PFC) + 4 * THAL \\ PFC &= \text{sigmoid}(pfc, 0, 100, 0, 100)\end{aligned}\quad (13)$$

Damped_PFC gets topographic input from PFC. Its 25 cells are split into groups of five that have five different time constants [represented as (0.1, 0.6, 1.1, 1.6, 2.1) in Eq. (14)]. This diversity of time constants was determined to accommodate different durations of target presentations (the phasic and tonic target presentation paradigms that produce brief and prolonged input, respectively, described below), and thus place different requirements on the temporal sensitivity of PFC.

$$\begin{aligned}\tau_{\text{damped_pfc}} &= (0.1, 0.6, 1.1, 1.6, 2.1) \\ S_{\text{damped_pfc}} &= PP_PFC \\ \text{Damped_PFC} &= \text{sigmoid}(pfc, 0, 100, 0, 100)\end{aligned}\quad (14)$$

The SC_POST layer provides a record of postsaccadic activity to PFC.

$$\begin{aligned}\tau_{sc_post} &= 0.2 \\ S_{sc_post} &= 5 * SC \\ SC_POST &= \text{sigmoid}(sc_post, 0, 100, 0, 100)\end{aligned}\quad (15)$$

Via the learning and normalization rules [Eqs. (11) and (12)] that apply reward or punishment after each saccade, the patterns of PFC activity that precede each saccade influence caudate [Eq. (4.3)] so that each pattern activates the correct saccade related cells. This leads to inhibition of the SNr [Eq. (5)], and disinhibition of SC [Eq. (6)] for production of the correct saccade. Inputs from PFC to CD are via the modifiable synapses in PFC_CD_Synapses.

$$\begin{aligned}\tau_{cd} &= 0.01 \\ S_{cd} &= DA_Modulation * [0.1(PFC \\ &\quad * PFC_CD_Synapses) + 0.4 * FEF \\ &\quad + CD * InhibitoryCollaterals] \\ CD &= \text{sigmoid}(cd, 0, 100, 0, 75)\end{aligned}\quad (4.3)$$

Simulation

This PFC architecture was developed so that at each point in the execution of each of the sequences, there will be a unique pattern of activity in PFC. Consider the two sequences right-left-up (RLU) and right-up-left (RUL). Though the first saccade is the same in both sequences, the resulting patterns of cortical activity must be different if these patterns are to be reliably associated with correct *second* saccades appropriate for these two different sequences. Due to the order preservation provided by the mixed excitatory and inhibitory PFC inputs, and the recurrent connections in PFC, perception of these two sequences produces different patterns in PFC that remain different even after identical saccades to R, and identical visual input resulting from this saccade. Figure 12 illustrates the PFC and caudate activity after the saccades to R in these two sequences. Note the different patterns of PFC activity, and the resulting preference in CD for the U and L targets, respectively, produced by these divergent patterns. The CVD learning model presented above has demonstrated that a unique pattern of activity (generated by a visual cue) can be associated with the correct one of two targets via corticostriatal trial-and-error learning. In the current simulation, the model learns to associate PFC activity patterns (generated by the context of the sequence) with the correct one of three targets, again via corticostriatal reinforcement learning.

The model training progresses gradually, based on the shaping used in primate training, moving progressively to sequences of one, two, and then three items. In a typical run of our simulation of the sequence learning task, the model attained over 95% correct performance after exposure to approximately 800 trials of the 6 three item sequences. After the learning had taken place, we plotted the responses of the 25 PFC cells during all six sequences, and analyzed them according to the criteria defined by Barone and Joseph (1989a).

Table 3 shows the progression of performance while the model learns the 6 sequences during 6 training epochs. The first epoch (E1) is a "no choice" condition, i.e., for each of the 6 sequences the three targets are presented while the model fixates, then the fixation point is removed, and the single targets are re-presented in the same order (go-signals). The model simply saccades to the targets as they are presented, and thus performs at 100% correct. In the subsequent epochs (E2-E6), the "choice" paradigm illustrated in Figure 1B1 is used, i.e., after the fixation point is removed, all the

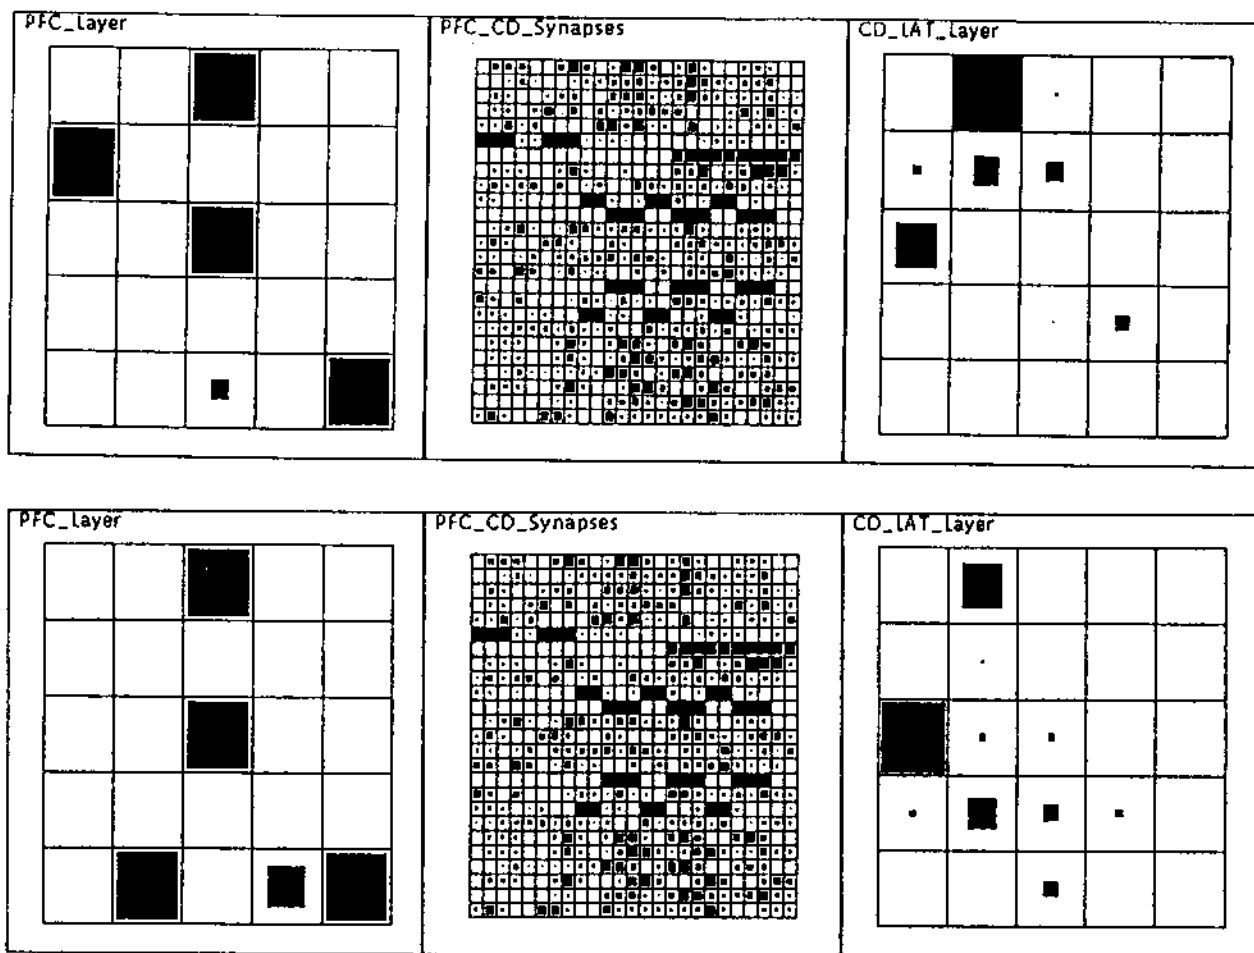


Figure 12. Snapshots of PFC, PFC_CD synapses, and caudate activity after the first saccade to R in the sequences RUL (above) and RLU (below). In both cases the model is now fixating R, so targets U and L are shifted 2 units left and 1 unit up on the updated retinal image, as seen in caudate. In both RUL and RLU the model must select from these two targets. Though the first saccades and the subsequent retinal inputs were identical for RUL and RLU, the resulting patterns of cortical activity in PFC are different. These patterns are transformed to caudate activity via the PFC_CD_Synapses. In caudate, we see that for sequence RUL (above), the upper target is favored, while in RLU, (below), the left target is favored.

Table 3. Training Evolution for Six Sequences.^a

Seq. No.	Targets	E1	E2	E3	E4	E5	E6
1	LUR	14/14	13/23	11/13	10/16	12/19	1/1
2	RUL	14/14	13/26	10/18	10/10	11/11	1/1
3	URL	13/13	12/15	10/10	10/10	11/11	1/1
4	LRU	13/13	12/44	10/17	10/26	11/15	1/1
5	RLU	13/13	12/49	10/17	10/13	11/16	1/1
6	URL	13/13	12/26	10/10	10/10	11/12	1/1
Correct by epoch (%)		100	40	71	70	80	100

^a E1 "no-choice." E2-E6 require "choice" and demonstrate progressive improvement from 40 to 100%.

targets that have not yet been visited are presented three times in succession (go-signals) and the model must learn by trial and error the correct sequences. The training in E1 provides a basis for subsequent performance, but alone is not sufficient, since the "no-choice" and "choice" go-signal visual inputs from PP to PFC are different, yielding different patterns of activity in PFC. In E2-E6 we see a progressive improvement from 40% (chance is 0.17%, i.e., $1/3 \times 1/2$) to 100%.

We observed that a high learning-rate [C1 in Eq. (11)] is useful in early epochs, but can produce a form of instability as learning proceeds, such that synaptic changes made during the correct performance of one sequence can interfere with the correct performance of a previously learned sequence. While we did not systematically study the optimal trajectory of learning-rate change, we did determine that by reducing the learning-rate in later epochs, this instability can be elimi-

nated. Indeed, for the learning that precedes 100% correct performance, the learning-rate for correct trials (learning-rate) can be reduced to zero, so that synaptic change occurs only on error trials (forgetting-rate). In E1-E2 the learning- and forgetting-rates are $2.5e-5$. In E3 the learning-rate is reduced to $5e-6$, and in E4-E6 the learning- and forgetting-rates are 0 and $2.5e-6$, respectively. Thus the final improvement from 70 to 100% occurs via "forgetting" rather than "learning." We note that leaving both learning- and forgetting-rates fixed at $2.5e-5$ produces a final performance between 85 and 90% correct.

Comparison of Model and Data

In Figure 13A we see a typical VT cell from our simulation, and in Figure 13B a typical VT cell from Barone and Joseph (1989a). The simulated cell responds to the left target (L) when it appears first, as seen in the first two traces, which are for sequences LUR and LRU. In the remaining sequences, the cell is also activated before

saccades to the left. This same combination of activity is seen in a primate VT cell, which responds to the upper target (U) when it appears first, and before saccades to that target. Figure 14 compares typical context cells (A) from the model and (B) Barone and Joseph (1989a). Both the animal and model cells start firing after the saccade to U and stop firing after the saccade to R or L in the sequences URL and ULR. Figure 15 presents the population of 25 PFC cells during the execution of all six sequences.

Table 4 compares the distribution of simulated cell types, from Figure 15, with those found in primates (Barone & Joseph, 1989a). Our percentage of context cells compares well with those in the monkey, and while we see a larger percentage of VT cells than in the monkey, in both the model and monkey there are less VT than context cells. These relative proportions suggest that more effort is required to maintain the dynamic internal state during execution (with context cells), than to construct the initial state at the start of sequence execution (with VT cells). Recall from Eq. (13) that the

Figure 13. Comparison of simulated and real visual-tonic cells. (A) Simulated VT cell activity in each of the six sequences. Target presentation marked with long arrows, and cue presentation with short arrows in "Trial progress." (B) A typical VT cell from Barone and Joseph (1989a). Vertical arrows indicate target presentations; the vertical bar indicates removal of fixation point, and small triangles indicate saccades.

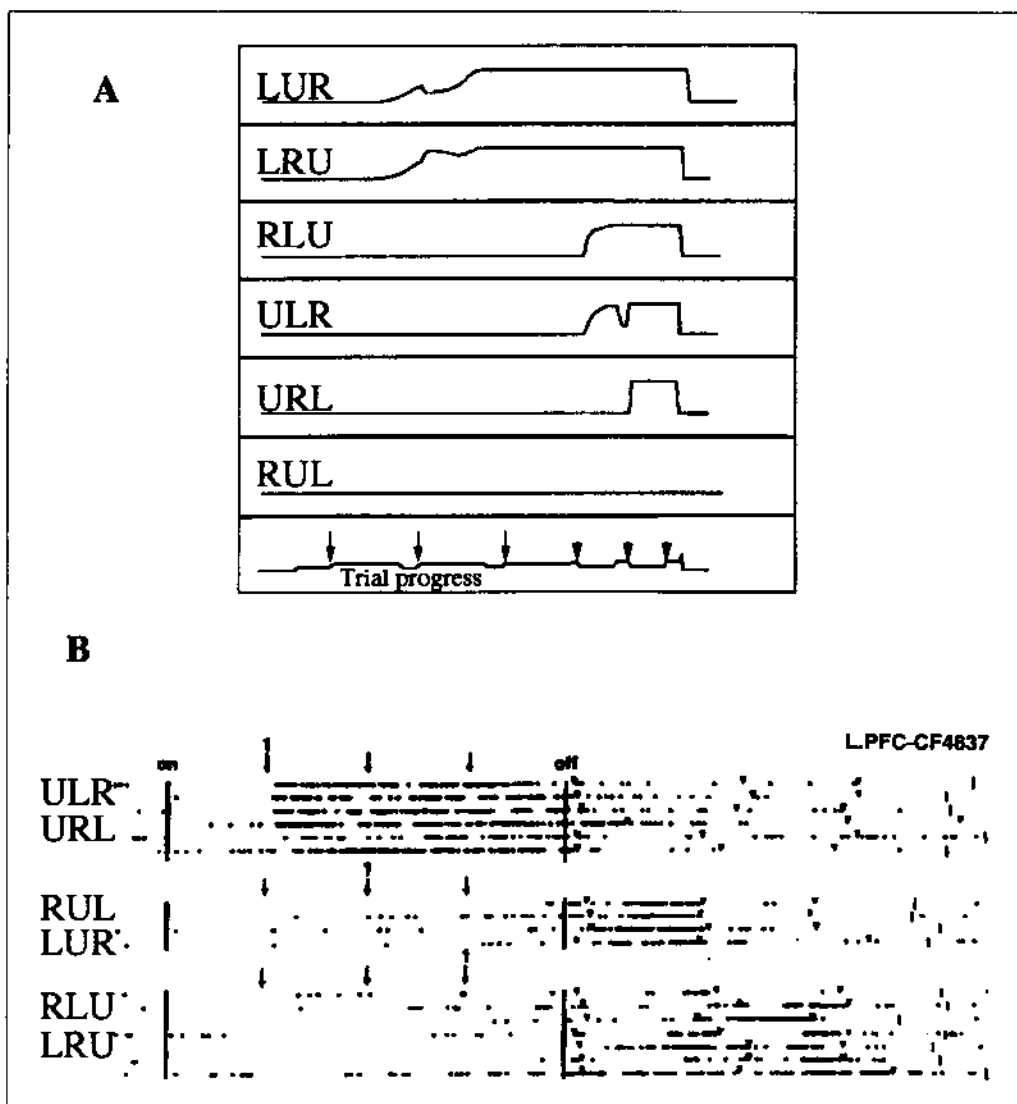
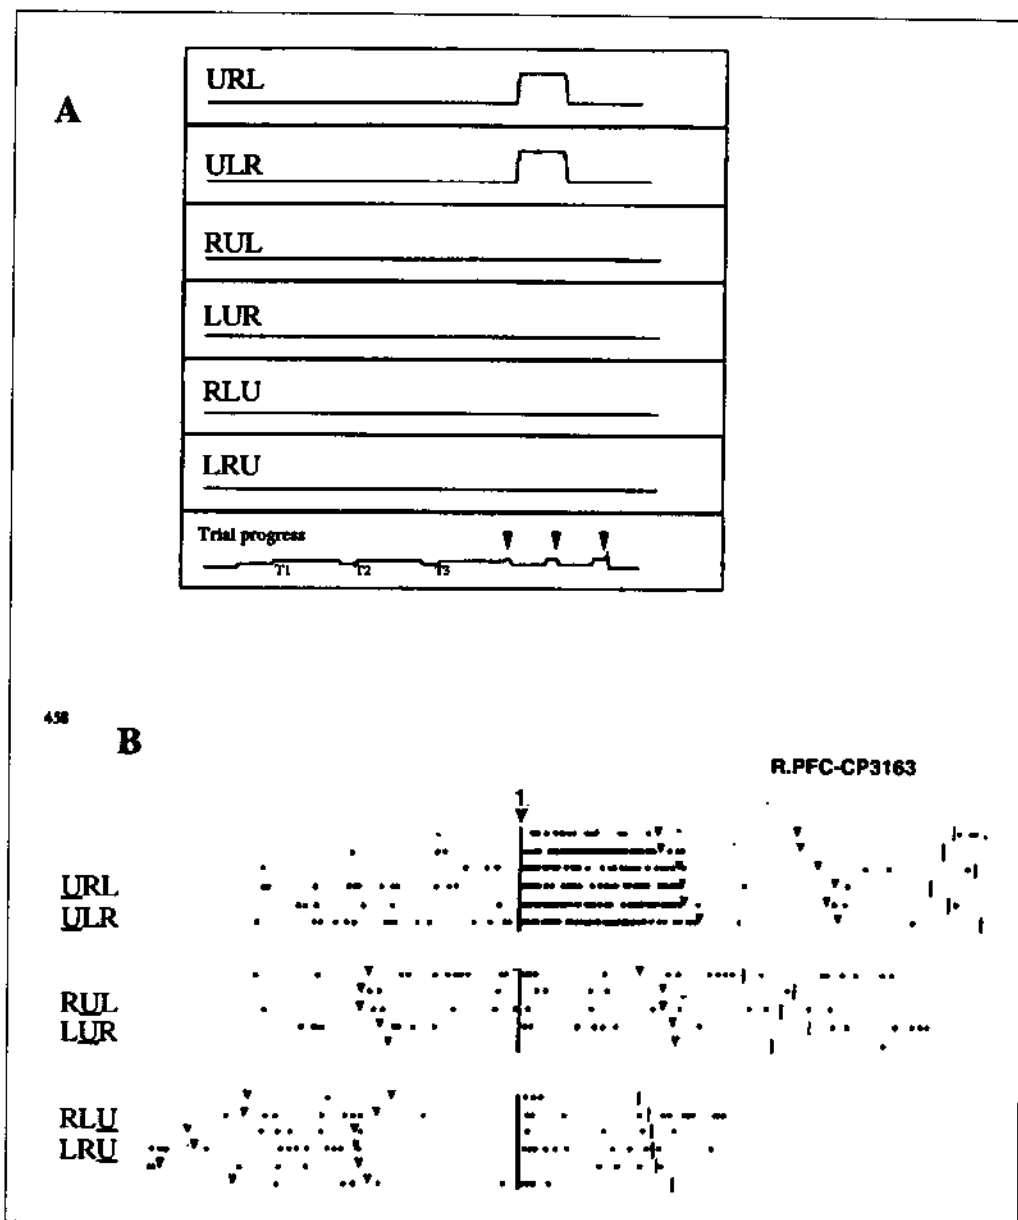


Figure 14. Comparison of simulated and real context cells. Context cells respond after saccades to a particular target, dependent on which targets have already been, or remain to be, chosen. (A) Response of a simulated PFC cell during execution of the six sequences. Like the real PFC cell, this simulated cell is active after saccades to U only when it is the first saccade in the sequence. Target presentation indicated in "Trial progress" as in Figure 13. (B) Context cell from Barone and Joseph (1989). This cell is activated after the saccade to U in the sequences ULR and URL, and in no other sequences.



modeled PFC is constructed to receive spatial input from PP, a damped self input, postsaccadic input, and a form of spatial memory indirectly from thalamus, satisfying the minimum information requirements we defined previously, and providing the architectural basis that yields VT and context cell types. Regarding our lack of saccade-, fixation-, and arm movement-related cells, it is likely that the real PFC receives other inputs in addition to the minimal set we defined for this task. Indeed, in the animal task, there is a reaching movement associated with each target while our simulated task involves only saccades. Barone and Joseph conclude that the activation of spatially selective fixation neurons intervenes only during preparation for visually guided reaching, based on their own results and those of Suzuki and Azuma (1977), indicating that visually guided reaching was required for activation of the fixation cells. The absence of a visually

guided reaching component to our sequential saccade task may remove the need for fixation and movement-related cells and justify their absence in our population.

Figure 16 shows snapshots of the patterns of activity in PFC at the time of each of the 18 saccades made during the execution of 6 sequences. We can see that while there is some overlap in these patterns, each one is sufficiently distinct that it can reliably be associated, via PFC-IT plastic synapses, with its corresponding saccade.

Effects of Target Presentation Format

In the 1989 paper, Barone and Joseph presented the stimuli in a "tonic-removal" format, in which once a target came on in the sequence presentation phase it remained on until after the first saccade was completed.

Figure 15. PFC activity during sequence execution. Time traces for 25 PFC cells during execution of the six sequences. The lowest trace indicates presentation of targets and saccades as in Figures 13 and 14. Classification based on criteria in Barone and Joseph (1989): VT, visual tonic; CX, context; SI, signal relate; NA, not active; O, other.




Table 4. Number and Classification of Prefrontal Cortex Cells

	<i>Barone and Joseph (1989)</i>		<i>Simulation-Tonic</i>		<i>Cells from Figure 15</i>
	N	%	N	%	
Visual tonic	35	12	8	35	2,4,6,7,8,14,21,23
Fixation	65	21	—	—	
Context	116	38	9	39	5,10,12,15,16,20,22,24,25
Saccade-related	33	11	—	—	
Movement-related	9	3	—	—	
Signal-related	24	8	3	14	9,11,18
Others	20	7	3	14	3,13,19

In the sequence reproduction phase, once a target had been visited, it was removed from subsequent go signals. This format raises two questions concerning unit activity and performance. First, is the visual-tonic effect directly related to the tonic presentation of the targets, or would it still be seen if they were phasically presented? Second, the tonic-removal format is easier since for the third saccade, there is only one target to choose from. What

would happen to performance if all three targets were lighted for the three go signals?

We trained the model using the "phasic-nonremoval" paradigm illustrated in Figure 1B2. In this case each target is presented phasically, and the three go signals use all three targets, rather than just the targets that have not yet been visited. The model performed nearly as well, consistently above 85% for the 6 sequences. While

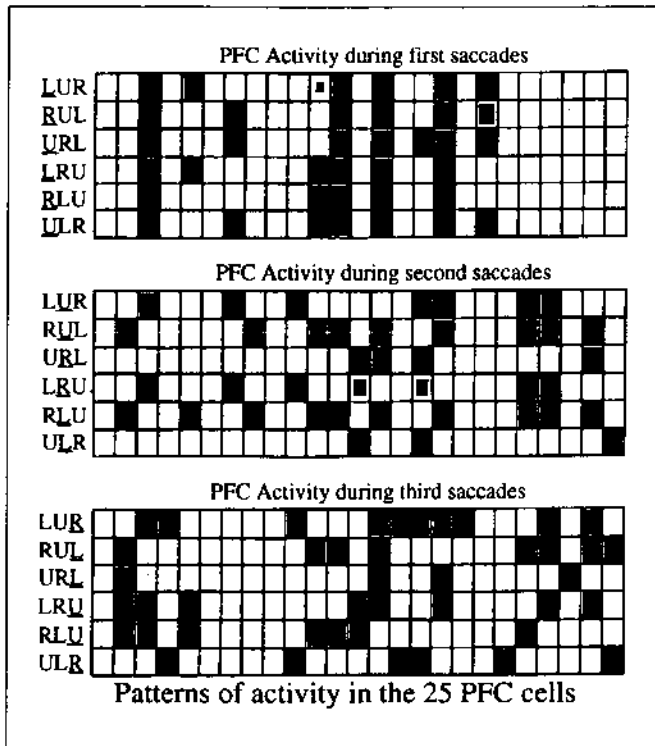


Figure 16. Pattern of activity in PFC at the time each of the 18 saccades was made (6 sequences \times 3 saccades per sequence = 18 saccades). The three arrays represent the PFC activity during each of the first, second, and third saccades in the six sequences. The 25 columns represent the 25 PFC cells, and the size of the solid filling indicates the degree of firing for each cell. Annotation to the left indicates the sequence and the current correct saccade (underlined) that is associated with the pattern of activity in PFC. The associative memory must correctly "recognize" these different patterns and associate them with the correct saccades. These patterns are distinct enough so that the model can perform above the level of 95% correct as seen in well-trained primates.

the performance of the model in both cases is quite good, the unit response characteristics in the "phasic-nonremoval" are less obvious than in the "tonic-removal." These differences in activation are likely due primarily to the change to phasic presentation of the targets, which removes a source of tonic input. It is noteworthy that although the temporal structure of these PFC cells is less obvious and the task is more difficult, the model performs well. This suggests that the state coding properties are in the patterns of activity in the entire PFC population, rather than in stereotypic behavior of a few cells.

DISCUSSION

We have developed minimal extensions of our previous model of saccade generation (Dominey & Arbib, 1992) that solve the conditional visual discrimination (CVD) and sequence learning tasks. In doing so, we have dem-

onstrated that for both tasks, correct performance can be achieved by a common mechanism that involves adaptive interaction between cortex and striatum. Stable patterns of activity are established in cortex, arising from direct sensory input, or from a combination of sensory and motor efferent inputs and internal state (i.e., the previous stable pattern of activity in cortex). These patterns of activity in cortex influence the oculomotor saccade circuit via adaptive corticostriatal projections. A reward-based learning mechanism modifies corticostriatal synapses to link these patterns of cortical activity to the correct saccade responses during learning. Learning proceeds on a trial-and-error basis, where initial "guesses" are provided by low levels of random activity that are injected into the network to break the "symmetry" between choices. Two promising aspects of this modeling approach are that, indeed, the same principles used in the CVD model successfully apply to the rather different task of sequence reproduction, and that in both cases, not only external behavior, but also single unit activity matches well with experimental data. While we focus on the "oculomotor loop," this model could be extended to address context-dependent behavior in the other parallel circuits linking basal ganglia and cortex (Alexander et al., 1986).

It is clear that in the behaving primate, there are brain mechanisms involved in solving the CVD and sequence learning tasks that are not captured in these models. However, data from striatal lesion (Reading et al., 1991) and dopamine depletion (Robbins et al., 1990) studies demonstrate that the intact frontostriatal system is necessary for normal learning of this kind.

Dopamine Function

In our models, we suggest two roles for the diffuse release of dopamine in the striatum, consistent with the view that neuromodulation in basal ganglia operates on multiple time scales (Graybiel, 1990). First, DA modulates cortical inputs to the striatum, ensuring that these inputs remain in a range to which striatum is sensitive. This agrees with the view that DA increases the signal-to-noise ratio in a population of cells, as suggested by Servan-Schreiber, Printz, and Cohen, (1990). Insufficient DA augments the excitability of striatum (Calabresi et al., 1993), leaving it potentially insensitive to the relative differences between strong cortical inputs. This insensitivity will impair learning and recall of learned associations, as seen in the striatal DA depletion studies of Robbins et al. (1990). In constructing this regulatory circuit, we suggest, as do Carlsson and Carlsson (1990), that the striatal input to the SNc dopamine-producing cells undergoes a sign change via influence from SNr. The functional result is a negative feedback control on corticostriatal activity, ensuring that striatum is not saturated by cortical inputs. In addition, by its down-regulation of striatal activity, this circuit will attenuate thalamocortical

activity, thus bringing cortex itself back into a nominal range of activation.

In addition, the phasic release of DA following a liquid reward during task learning (Ljungberg et al., 1991; Schultz et al., 1993) participates in an adaptive mechanism for strengthening corticostriatal synapses following reward [see Eq. (11)]. These synaptic changes construct causal associations between cortical patterns of activity, representing behavioral contexts, and the correct saccades for those contexts. Our models demonstrate that, as implemented here, the regulatory and adaptive functions of DA can coexist and cooperate.

Ljungberg, Apicella, and Schultz, (1992) found that during learning of simple tasks, nigrostriatal dopamine-producing cells are activated by the reward, and cues associated with the reward, and that these responses diminish after the behavior is overlearned. During performance of a more complex delayed alternation task, however, the same investigators (Ljungberg et al., 1991) found that responses in DA cells to trigger and reward were qualitatively preserved in conditioning, postconditioning, and overtraining. These selective responses to salient stimuli only during learning of a simple task suggest that DA neurons participate in behavioral adaptation, and are no longer required once a behavior becomes a habit. This kind of conditional response associated with important or reinforced stimuli and behavior is typical of neurons found in the primate ventral striatum (Rolls & Williams, 1987). Cells in this region project to the SNc, and can thus influence the modulation of behavior by the release of nigral dopamine in the striatum (Smith & Bolam, 1990), as in our Eqs. (10) and (11).

While we only model "forgetting" as a negative value for the reinforcement signal, we note that depletion of dopamine from nucleus accumbens septi impairs extinction of associations in CVD tasks (Robbins et al., 1990), suggesting its role in the use of error information. In tests of temporal order and conditional associative learning, PD patients were impaired only when learning by trial-and-error was required (Vriezen & Moscovitch, 1990), but not when corrections were provided on error trials, indicating that the reinforcement pathway is intact, while the error processing pathway is not. We can observe this effect by setting our penalty signal to zero, after which both CVD and sequence learning are significantly impaired.

Generation of Diversity

Associating input patterns with output patterns is facilitated if the input patterns do not have significant overlap, and can thus be easily distinguished. IT and PFC in our models ensure that visual cues and sequence execution states are represented by patterns of activity that are sufficiently nonoverlapping to be associated uniquely with the correct saccade. By using mixed excitatory and

inhibitory connections from V4 feature cells to IT we get IT cells that have preferences for different features and their conjunctions and disjunctions, and the general property that the degree of similarity or dissimilarity in cues is reflected in the corresponding IT activity. In PFC, the coded "features" are related to position, temporal order, and previous saccade activity. Again by using combined inhibitory and excitatory input from visual, saccade-related, and self-input sources, PFC cells detect these individual "features," as well as their conjunctions and disjunctions. By allowing the output of the network to influence its state, we produce a sequence of connected states that can be used to store and reproduce behavioral sequences. It is likely that intrinsic cortical plasticity plays a role in these primate tasks, but by allowing the required information "features" relevant to the sequencing task to combine in our PFC via *fixed* connections, the resulting cellular activity is surprisingly similar to that in the primate PFC, and provides the required representational diversity for learning the task.

Related Models

Many recurrent models (e.g., Jordan, 1990, Elman, 1990) use network output as part of the next input to generate sequences. These models learn by computing an error measure from the network output and the desired output and adjusting connections strengths to minimize this error. Our approach differs in that it is based on a global reinforcement/penalty signal. This approach has been applied to the control of dynamic systems (see Barto, 1990), and to learning sensorimotor and coordinate transformations (Fagg & Arbib, 1992; Mazzoni, Anderson, & Jordan, 1991). The Fagg-Arbib model learns stimulus-response mappings using reinforcement to form feature detectors that activate voting units that become associated, again by reinforcement, with the correct response. In our models, the PFC and IT cells are hardwired feature detectors whose votes are determined by reinforcement of corticostriatal synapses.

Dehaene, Changeux, and Nadal, (1987) developed a model of sequence learning and reproduction in which prerepresented sequence detectors are selected by training to accommodate the sensory sequence input using a local Hebbian rule. Our PFC also prerepresents sequences via the existence of excitatory and inhibitory connection from sensory and self inputs, so that patterns of activity PFC encode sequential states, and subsequently become associated with production of the correct saccades. Wang and Arbib (1990) note that the selection model of Dehaene et al. relies on prerepresentations that would have to be immense to learn arbitrary nontrivial sequences. They address this problem in a model of sequence recognition and reproduction based on short term memory [see our Eq. (13)], establishing sequence sensitivity by learning rather than by selecting existing prerepresentations.

Dehaene and Changeux have also modeled prefrontal cortical function in solving delayed-response tasks (1989) and the Wisconsin Card Sorting Test (1991). In these models, reinforcement is used to stabilize correct performance and de-stabilize incorrect performance, and to modify "rule coding" activity as a form of search for correct state. An important aspect of these tasks is the need to rapidly respond to a change in reinforcement schedule. This constraint is addressed by coding "rules of behavior" as patterns of activity, capable of rapid change, rather than as slowly adapting synaptic weights. The trade-off for this rapid rule changing capability is a limitation on the number of available rules that are built into the system as prerepresentations.

An Alternate View and Final Prediction

We have taken a specific, refutable position on a mechanism for sensorimotor associations that is implemented by synaptic changes between cortical inputs to striatal cells that participate in motor loops. An alternative mechanism could involve synaptic changes between corticocortical sensory inputs from sensory to motor-related cortices, possibly facilitated by striatal activity. These two mechanisms may be resolved by examination of unit activity in the caudate during the learning of a complex stimulus-response habit. If the caudate only participates transiently, by facilitating corticocortical learning, then a transitory change will be observed in caudate cells during learning. If corticostriatal synapses are permanently changed in SR formation, then a permanent change will be seen in the caudate cells, as predicted by our model.

Acknowledgments

This work was supported by a grant from the Human Frontier Research Organization and by Grant NOOO14-92-J-4026 from ONR to M.A. Arbib, and by a Chateaubriand Fellowship, an NIH Multidisciplinary Training Grant in Gerontology, and a grant from the Ministère de l'Éducation Nationale, D.R.E.D, to P. F. Dominey. The authors thank the anonymous reviewers for substantial and useful comments on a previous version of the manuscript.

Reprint requests should be sent to Dr. Peter Dominey, INSERM Unité 94, 16 av. du Doyen Lepine, 69500 Bron, France.

Notes

1. Abbreviations used: CD, caudate nucleus; CVD, conditional visual discrimination; DA, dopamine; FEF, frontal eye field; FON, fovea on—active during visual fixation; IT, inferior temporal cortex; LTD, long-term depression; LTP, long-term potentiation; MD, mediodorsal thalamus; NMDA, *N*-methyl-*D*-aspartate; NSL, neural simulation language; PFC, prefrontal cortex; PP, posterior parietal cortex; RLU, right, left, up—target sequence; SC, superior colliculus; SG, brainstem saccade generator; SNc, substantia nigra pars compacta; SNr, substantia nigra pars reticulata; SR,

stimulus response; VT, visual tonic cells—recorded in PFC; WTA, winner-take-all.

2. Dynamic spatial remapping of target representations for double-step saccades developed in Dominey and Arbib (1992) will not be addressed in the present article.
3. Compare the symmetry breaking in the Fagg and Arbib (1992) model.
4. If the animal makes a wrong move at any point in the sequence, then the trial is terminated. Thus we can consider that when he makes a correct move and the trial is not terminated, this will be a form of "internal reward," which will, in the model, lead to reinforcement of the correct move, even though the actual juice reward has not yet been received.

REFERENCES

- Alexander, G. E., DeLong, M.R., & Strick, P.L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience*, 9, 357-381.
- Arbib, M. A. (1969). Memory limitations of stimulus-response models. *Psychological Review*, 76(5), 507-510.
- Barone, P., & Joseph, J.-P. (1989a). Prefrontal cortex and spatial sequencing in macaque monkey. *Experimental Brain Research*, 78, 447-464.
- Barone, P., & Joseph, J.-P. (1989b). Role of dorsolateral prefrontal cortex in organizing visually guided behavior. *Brain Behavior Evolution*, 33, 132-135.
- Barto A.G. (1990). Connectionist learning for control: An overview. In W. T. Miller, R. S. Sutton, & P. J. Werbos (Eds.), *Neural networks for control* (pp. 1-58). Cambridge, MA: MIT Press.
- Boussaoud, D., & Joseph, J. P. (1985). Role of the cat substantia nigra pars reticulata in eye and head movements. II. Effects of local pharmacological injections. *Experimental Brain Research*, 57, 297-304.
- Bruce C. J., & Goldberg, M. E. (1984). Physiology of the frontal eye fields. *Trends in Neuroscience*, 7:436-441
- Calabresi, P., Maj, R., Pisani, A., Mercuri, N. B., & Bernardi, G. (1992). Long-term synaptic depression in the striatum: Physiological and pharmacological characterization. *Journal of Neuroscience*, 12(11), 4224-4233.
- Calabresi, P., Mercuri, N. B., Sancesario, G., & Bernardi, G. (1993). Electrophysiology of dopamine-denervated striatal neurons: Implications for Parkinson's disease. *Brain*, 116, 433-452.
- Calabresi, P., Pisani, A., Mercuri, N. B., & Bernardi, G. (1992). Long-term potentiation in the striatum is unmasked by removing the voltage-dependent magnesium block of NMDA receptor channels. *European Journal of Neuroscience*, 4, 929-935.
- Carlsson, M., & Carlsson, A. (1990). Interactions between glutamatergic and monoaminergic systems within the basal ganglia—implications for schizophrenia and Parkinson's disease. *Trends in Neuroscience*, 13(7), 272-276.
- Chevalier, G., Vacher, S., Deniau, J. M., & Desban, M. (1985). Disinhibition as a basic process in the expression of striatal functions. I. The striato-nigral influence on the tecto-spinal/tecto-diencephalic neurons. *Brain Research*, 334, 215-226.
- Dehaene, S., & Changeux, J.-P. (1991). The Wisconsin Card Sorting Test: Theoretical analysis and modeling in a neuronal network. *Cerebral Cortex*, 1, 62-79.
- Dehaene, S., Changeux, J.-P., & Nadal, J. P. (1987). Neural networks that learn temporal sequences by selection. *Proceedings of the National Academy of Sciences U.S.A.*, 84, 2727-2731.

- Deniau, J. M., & Chevalier, G. (1985). Disinhibition as a basic process in the expression of striatal functions. II. The striato-nigral influence on thalamocortical cells of the ventromedial thalamic nucleus. *Brain Research*, 334, 227-233.
- Desimone, R. (1991). Face-selective cells in the temporal cortex of monkeys. *Journal of Cognitive Neuroscience*, 3(1), 1-8.
- Desimone, R., Albright, T. D., Gross, C. G., & Bruce, C. (1984). Stimulus-selective properties of inferior temporal neurons in the macaque. *Journal of Neuroscience*, 4(8), 2051-2062.
- Didday, R. L. (1976). A model of visuomotor mechanisms in the frog optic tectum. *Math Bioscience*, 30, 169-180.
- Dominey, P. F., & Arbib, M. A. (1992). A cortico-subcortical model for generation of spatially accurate sequential saccades. *Cerebral Cortex*, 2, 153-175.
- Dominey, P. F., Joseph, J.-P., & Arbib, M. A. (1992). A model of striatal function in habit formation. *Society of Neuroscience Abstracts*, 357, 17.
- Dominey, P. F., Joseph, J.-P., & Arbib, M. A. (1993). Caudate activity in the execution of oculomotor habits. *Society of Neuroscience Abstracts*, 411, 1.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14, 179-211.
- Fagg, A. H., & Arbib, M. A. (1992). A model of primate visual-motor conditional learning. *Journal of Adaptive Behavior*, 1(1), 3-37.
- Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology*, 61, 331-349.
- Fuster, J. M. (1990). Inferotemporal units in selective visual attention and short-term memory. *Journal of Neurophysiology*, 64(2), 681-697.
- Garcia-Munoz, M., Young, S. J., & Groves, P. M. (1991). Terminal excitability of the corticostriatal pathway. I. Regulation by dopamine receptor stimulation. *Brain Research*, 551(1-2), 195-206.
- Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In Vernon B. Mountcastle (Ed.), *Handbook of Physiology*, 5, 373-417.
- Graybiel, A. (1990). Neurotransmitters and neuromodulators in the basal ganglia. *Trends in Neuroscience*, 13(7), 244-254.
- Hikosaka, O., Sakamoto, M., & Usui, S. (1989a). Functional properties of monkey caudate neurons. I. Activities related to saccadic eye movements. *Journal of Neurophysiology*, 61, 780-798.
- Hikosaka, O., Sakamoto, M., & Usui, S. (1989b). Functional properties of monkey caudate neurons. II. Visual and auditory responses. *Journal of Neurophysiology*, 61, 799-813.
- Hikosaka, O., Sakamoto, M., & Usui, S. (1989c). Functional properties of monkey caudate neurons. III. Activities related to expectation of target and reward. *Journal of Neurophysiology*, 61, 814-832.
- Hikosaka, O., & Wurtz, R. (1983a). Visual and oculomotor functions of monkey substantia nigra pars reticulata. I. Relation of visual and auditory responses to saccades. *Journal of Neurophysiology*, 49, 1230-1253.
- Hikosaka, O., & Wurtz, R. (1983b). Visual and oculomotor functions of monkey substantia nigra pars reticulata. II. Visual responses related to fixation of gaze. *Journal of Neurophysiology*, 49, 1254-1267.
- Hikosaka, O., & Wurtz, R. (1983c). Visual and oculomotor functions of monkey substantia nigra pars reticulata. III. Memory-contingent visual and saccade responses. *Journal of Neurophysiology*, 49, 1268-1284.
- Hikosaka, O., & Wurtz, R. (1983d). Visual and oculomotor functions of monkey substantia nigra pars reticulata. IV. Relation of substantia nigra to superior colliculus. *Journal of Neurophysiology*, 49, 1285-1301.
- Hikosaka, O., & Wurtz, R. H. (1985). Modification of saccadic eye movements by GABA-related substances. II. Effects of muscimol in monkey substantia nigra pars reticulata. *Journal of Neurophysiology*, 53, 292-308.
- Ilinsky, I. A., Jouandret, M. L., & Goldman-Rakic, P. S. (1985). Organization of the nigrothalamocortical system in the rhesus monkey. *Journal of Comparative Neurology*, 236, 315-330.
- Jordan, M. I. (1990). Learning to articulate: Sequential networks and distal constraints. In M. Jeannerod (Ed.), *Attention and Performance XIII*. Hillsdale, N.Y.: Lawrence Erlbaum.
- Lee, T. D., Magill, R. A., & Weeks, D. J. (1985). Influence of practice schedule on testing schema theory predictions in adults. *Journal of Motor Behavior*, 17(3), 283-299.
- Ljungberg, T., Apicella, P., & Schultz, W. (1991). Responses of monkey midbrain dopamine neurons during delayed alternation performance. *Brain Research*, 567, 337-341.
- Ljungberg, T., Apicella, P., & Schultz, W. (1992). Responses of monkey dopamine neurons during learning of behavioral reactions. *Journal of Neurophysiology*, 67(1), 145-163.
- Mazzoni, P., Andersen, R. A., & Jordan, M. I. (1991). A more biologically plausible learning rule for neural networks. *Proceedings of the National Academy of Sciences, U.S.A.*, 88, 4433-4437.
- Mercuri, M., Bernardi, G., Calabresi, P., Cotugno, A., Levi, G., & Stanzione, P. (1985). Dopamine decreases cell excitability in rat striatal neurons by pre- and postsynaptic mechanisms. *Brain Research*, 385, 110-121.
- Mishkin, M. (1990). Cerebral memory circuits. In *Proceedings of Yakult International Symposium: Perception, cognition and brain*. Yakult Honsha Co., Ltd.
- Mishkin, M., Malamut, B., & Bachevalier, J. (1984). Memories and habits: Two neural systems. In G. Lynch, J. L. McGaugh, & N. M. Weinberger (Eds.), *Neurobiology of learning and memory* (pp. 65-77). New York: Guilford Press.
- Mitz, A. R., Godshalk, M., & Wise, S. P. (1991). Learning-dependent neuronal activity in the premotor cortex. Activity during the acquisition of conditional motor associations. *Journal of Neuroscience*, 6, 1855-1872.
- Managhan, D. T., & Cotman, C. W. (1985). Distribution of N-methyl-D-aspartate-sensitive L-[3]glutamate-binding sites in rat brain. *Journal of Neuroscience*, 5(11), 2909-2919.
- Niki, H. (1974). Prefrontal unit activity during delayed alternation in the monkey. II. Relation to absolute versus relative direction of response. *Brain Research*, 68, 197-204.
- Nisenbaum, E. S., & Berger, T. W. (1992). Functionally distinct subpopulations of striatal neurons are differentially regulated by GABAergic and Dopaminergic inputs-I. An in vivo analysis. *Neuroscience*, 48(3), 561-578.
- Reading, P. J., Dunnet, S. B., & Robbins, T. W. (1991). Dissociable roles of the ventral, medial and lateral striatum on the acquisition and performance of a complex visual stimulus-response habit. *Behavioral Brain Research*, 45, 147-161.
- Robbins, T. W., Giardini, V., Jones, G. H., Reading, P. J., & Sahakian, B. J. (1990). Effects of dopamine depletion from the caudate-putamen and nucleus accumbens septi on the acquisition and performance of a conditional discrimination task. *Behavioral Brain Research*, 38, 243-261.
- Rolls, E. T., Thorpe, S. J., & Maddison, S. P. (1983). Responses of striatal neurons in the behaving monkey. I. Head of the caudate nucleus. *Behavioral Brain Research*, 7, 179-210.
- Rolls, E. T., & Williams, G. V. (1987). Sensory and movement-

- related neuronal activity in different regions of the primate striatum. In J. S. Schneider & T. I. Lidsky (Eds.), *Basal ganglia and behavior: Sensory aspects of motor functioning* (pp. 37-60). Toronto: Hans Huber.
- Schultz, W. (1989). Neurophysiology of basal ganglia. In D. B. Caine (Ed.), *Handbook of experimental pharmacology* (Vol. 88, pp. 1-45). Berlin: Springer Verlag.
- Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *Journal of Neuroscience*, 13(3), 900-913.
- Scudder, C. A. (1988). A new local feedback model of the saccadic burst generator. *Journal of Neurophysiology*, 59, 1455-1475.
- Segraves, M., & Goldberg, M. E. (1987). Functional properties of corticotectal neurons in the monkey's frontal eye field. *Journal of Neurophysiology*, 58, 1387-1419.
- Selemon, L. D., & Goldman-Rakic, P. S. (1985). Longitudinal topography and interdigitation of corticostriatal projections in the rhesus monkey. *Journal of Neuroscience*, 5, 776-794.
- Servan-Schreiber, D., Printz, H., & Cohen, J. D. (1990). A network model of catecholamine effects: Gain, signal-to-noise ratio, and behavior. *Science*, 249, 892-895.
- Smith, D. A., & Bolam, J. P. (1990). The neural network of the basal ganglia as revealed by the study of synaptic connections of identified neurons. *Trends in Neuroscience*, 13(7), 259-265.
- Sparks, D. L. (1986). Translation of sensory signals into commands for control of saccadic eye movements: Role of primate superior colliculus. *Physiological Review*, 66, 118-171.
- Stanton, G. B., Goldberg, M. E., & Bruce, C. J. (1988a). Frontal eye field efferents in the Macaque monkey. I. Subcortical pathways and topography of striatal and thalamic terminal fields. *Journal of Comparative Neurology*, 271, 473-492.
- Stanton, G. B., Goldberg, M. E., & Bruce, C. J. (1988b). Frontal eye field efferents in the Macaque monkey. II. Topography of terminal fields in midbrain and pons. *Journal of Comparative Neurology*, 271, 493-506.
- Suzuki, H., & Azuma, M. (1977). Prefrontal neural activity during gazing at a light spot in the monkey. *Brain Research*, 126, 497-508.
- Tanaka, K., Saito, H.-A., Fukada, Y., & Mori, M. (1991). Coding visual images of objects in the inferotemporal cortex of the Macaque Monkey. *Journal of Neurophysiology*, 66(1), 170-189.
- Vriezen, E. R., & Moscovitch, M. (1990). Memory for temporal order and conditional-associative learning in patients with Parkinson's Disease. *Neuropsychologia*, 28(12), 1283-1293.
- Walsh, J.P., & Dunia, R. (1993). Synaptic activation of N-methyl-D-aspartate receptors induces short-term potentiation at excitatory synapses in the striatum of the rat. *Neuroscience*, 57(2), 241-248.
- Wang, D., & Arbib, M. A. (1990). Complex temporal sequence learning based on short-term memory. *Proceedings of the IEEE*, 78(9), 1536-1543.
- Weihmuller, F. B., Ulas, J., Nguyen, L., Cotman, C. W., & Marshall, J. F. (1992). Elevated NMDA receptors in Parkinsonian striatum. *Neuroreport*, 3, 977-980.
- Weitzenfeld, A. (1991). NSL Neural Simulation Language Version 2.1. Center for Neural Engineering, University of Southern California Technical Report 91-5.
- Wilson, C. J., & Groves, P. M. (1980). Fine structure and synaptic connections of the common spiny neuron of the rat neostriatum. *Journal of Comparative Neurology*, 194, 599-615.
- Wise, R. A., & Rompre, P. P. (1989). Brain dopamine and reward. *Annual Review of Psychology*, 40, 191-225.
- Yeterian, E. H., & Pandya, D. (1991). Prefrontostriatal connections in relation to cortical architectonic organization in rhesus monkeys. *Journal of Comparative Neurology*, 312, 43-67.