

Memory Activity of LIP Neurons for Sequential Eye Movements Simulated With Neural Networks

JING XING AND RICHARD A. ANDERSEN

Division of Biology, California Institute of Technology, Pasadena, California 91125

Received 30 July 1999; accepted in final form 5 April 2000

Xing, Jing and Richard A. Andersen. Memory activity of LIP neurons for sequential eye movements simulated with neural networks. *J Neurophysiol* 84: 651–665, 2000. Many neurons in macaque lateral intraparietal cortex (LIP) maintain elevated activity induced by visual or auditory targets during tasks in which monkeys are required to withhold one or more planned eye movements. We studied the mechanisms for such memory activity with neural network modeling. Recurrent connections among simulated LIP neurons were used to model memory responses of LIP neurons. The connection weights were computed using an optimization procedure to produce desired outputs in memory-saccade tasks. One constraint for the training process is the “single-purpose” rule, which mimics the fact that once LIP neurons hold the memory activity of a saccade, they are insensitive to further stimuli until the motor action is completed. After training, excitatory connections were developed between units with similar preferred saccade directions, while inhibitory connections were formed between units with dissimilar directions. This “push-pull” mechanism enables the network to encode the next intended eye movement and is essential for programming sequential saccades. In simulating double saccades, the push-pull connections locked the on-going activity in the network for the first saccade until the saccade was made, then a new population of units became active to prepare for the second saccade. The simulated LIP neurons exhibited sensory responses and memory activities similar to those recorded in LIP neurons. We propose that push-pull recurrent connections might be the basic structure mediating the memory activity of area LIP in planning sequential eye movements.

INTRODUCTION

The lateral intraparietal cortex (LIP) is involved in programming saccadic eye movements (Andersen and Gnadt 1989; Lynch et al. 1977). Many LIP neurons exhibit sustained responses to remembered visual or auditory targets (Mazzoni et al. 1996a). During delayed-saccade tasks in which the monkey withheld a saccade to a remembered target for a short period of time, the response of LIP neurons triggered by the target was sustained until the saccade was initiated (Andersen et al. 1990a,b; Gnadt and Andersen 1988). Moreover, neurons could maintain the memory for the saccade even if the monkey was presented with new stimuli during the withholding period. Negatively correlated memory responses have also been observed in LIP, and such responses occurred when the remembered saccade was opposite the neuron's preferred saccadic direction (Barash et al. 1991a,b). Memory activity was further

characterized with the delayed-double-saccade experiments (Mazzoni et al. 1996b) in which the monkey was trained to memorize two consecutively flashed targets and to plan two saccades to the targets in the order that the targets were presented. During the delayed period, many LIP neurons whose preferred directions were in the direction of the first saccade fired continuously until the execution of the saccade. These neurons thus held the correct memory for the first saccade regardless of the flash of the second target. Neurons coding for the second saccade started to fire only after the first saccade was executed. The results indicated that memory activities for the majority of LIP neurons encode the next planned saccade. On the other hand, a small percentage of LIP neurons encode the memory of target locations instead. The sustained responses in all kinds of delayed-saccade tasks have a common feature: neurons begin to encode a new saccadic movement only after the current motor plan is disengaged. We call this the “single-purpose” feature.

Short-term memory activity has been observed in a number of cortical areas (Funahashi et al. 1989; Gnadt and Andersen 1988; Goldman-Rakic 1995; Kalaska and Grammond 1995; Quintana and Fuster 1992). Several computational studies have proposed that recurrent connections might be the mechanism for this activity (Cowan 1972; Dehaene and Changeux 1989; Fuster 1995; Zipser 1991). The purpose of this report was to study the mechanisms of saccadic-related memory activities in area LIP. Especially, we were interested in how the single-purpose feature was related to programming delayed double saccades. Based on experimental tasks, we used recurrent neural networks to simulate the memory features of LIP neurons. We first studied the mechanisms of memory saccades and then examined an extended model for planning double-saccades. Preliminary results of this report have been presented in abstract form (Xing et al. 1995).

METHODS

The model is a three-layered neural network, with a similar structure to that of Zipser and Andersen model (1988). The diagram of this model is shown in Fig. 1. The model was not designed to resemble the complex anatomy of area LIP. It is the typical classic neural network that can be trained to carry out the required sensorimotor transformations. The input layer, like area LIP, has access to visual and auditory target locations as well as eye position in the orbit. The output layer

Address for reprint requests: R. A. Andersen, Div. of Biology, 216-76, California Institute of Technology, Pasadena, CA 91125 (E-mail: andersen@vis.caltech.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked “advertisement” in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

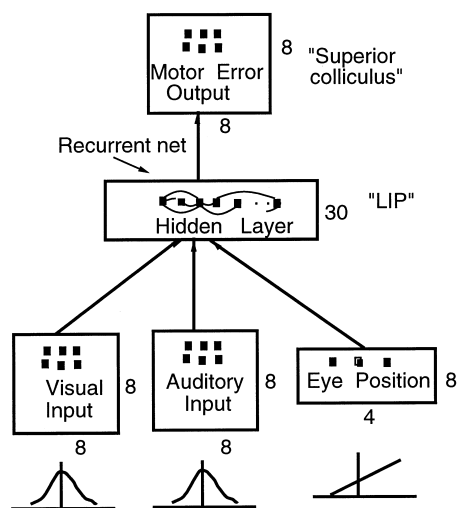


FIG. 1. The diagram of the recurrent network model for memory saccades.

is a topographic map of eye motor errors. The middle layer, or the hidden layer, is a recurrent network with every unit receiving activities from all other hidden units. Every unit in the input layer is connected to each of the hidden units, which are in turn connected to all the output units. The weights of connections vary between -1 and $+1$. They are initially set to small random values between -0.1 and 0.1 . The weights are adjusted to encode the motor errors of visual or auditory targets at the output map.

The input layer consists of a visual map in retinal coordinates, an auditory map in head-centered coordinates, and eye-position units. The visual map uses 8×8 units to model a -40° to 40° retinal space. Each of the units has a Gaussian receptive field (RF) with a $1/e$ width of 15° . The centers of the RFs were equally spaced over the 8×8 grid with 10° spacing. These units encode target locations with their activation values between 0 and 1. The auditory input is modeled using an auditory map of an 8×8 array of units, similar to the visual one. The only difference between the two input maps is that the auditory units encode target locations in head-centered coordinates and the visual units encode target locations in eye-centered coordinates. Eye positions is coded by four sets of eight units representing horizontal and vertical eye coordinates with positive and negative slopes. The activation of the units, with various intercepts and slopes, is thus an increasing function of eye positions.

The middle layer, also called the hidden layer, typically has 30 units in the simulations presented in this report. Each hidden unit receives inputs from all three input channels. In addition, each hidden unit receives recurrent projections from all other hidden units. The activation of a hidden unit is calculated by first summing all inputs and then calculating the output as a sigmoidal function of the total input. At a given simulated time step, the activation of a hidden unit can be expressed as the following: output activation = $1/[1 + \exp(-\text{net})]$ where net = sum of weighted inputs + bias.

The inputs here include the activities of the visual, auditory, and eye-position units at the current time step and the activities of other hidden units at the previous time step. The sigmoid function is chosen as the activation function because it resembles the operation performed by actual neurons that sum inputs, have a threshold, and saturate at high levels of activity. In the middle region of its dynamic range, the sigmoid approximates a linear function.

The output layer is an eye-centered map encoding eye motor errors (ME) of saccades. An 8×8 array of output units is used to represent MEs topographically. Each of the units covers a 10° space of MEs with a Gaussian $1/e$ width of 15° . The activation of the output units, like the hidden units, is a sigmoidal function of the sum of the weighted inputs from the hidden units. We use E to represent the initial eye position, V for the locations of visual targets in retinal

coordinates, and A for the locations of auditory targets in head-centered coordinates. For simple saccades, $ME = V$ for visual targets and $ME = A - E$ for auditory targets. For double saccades, we use E0 to represent the initial eye position, V1 for the location of the first visual target in retinal coordinates, and A1 for the location of the first auditory target in head-centered coordinates. E1 represents the eye position after the first saccade. V2 and A2 indicate the second visual and auditory targets, respectively. The desired ME output for the first saccade is $ME = V1$ for visual targets or $ME = A1 - E0$ for auditory targets. The ME for the second saccade is $ME = A2 - E1$ or $ME = V2 + E0 - E1$.

Training process

We use an algorithm "backpropagation-through-time" to train the network. This algorithm gradually optimizes connection weights to produce the desired output in a recurrent neural network (Munro et al. 1994; Werbos 1990; Williams and Zipser 1995). We use this algorithm simply to train the network to perform the required sensorimotor transformations with no intention to claim that the algorithm is similar to the learning mechanisms in the brain.

The backpropagation algorithm uses supervised learning. It first computes an error signal, which is the difference of the desired output (the teacher signal) and the actual output. This error signal is then used to update connection weights. The amount of weight change depends on the error signal, the activities of the two connected units, and an arbitrary learning rate. In our implementation of the algorithm, the desired activity A_{exp} for each output unit k is determined by the expected ME of a saccadic target. The actual output A_o of an output unit is computed for a given target location, eye position and the initial weights. The error signal δ_k for an output unit k is

$$\delta_k = A_{exp} - A_o$$

A connection weight W_{ho} from a hidden unit to an output unit is updated according to

$$\Delta W_{ho} = n * A_h * A_o * (1 - A_o) * \delta_k$$

where A_h is the activity of the hidden unit. The learning rate n in our simulations is 0.05.

A connection weight W_{ih} from an input unit to a hidden unit is updated according to

$$\Delta W_{ih} = n * A_i * A_h * (1 - A_h) * \sum_k (\delta_k * W_{ho_k})$$

where A_i is the activity of the input unit and δ_k is the error signal of an output unit k .

A recurrent connection weight W_{hh} from a hidden unit i to another hidden unit j is updated according to

$$\Delta W_{hh} = n * A_{h_i}(t-1) * A_{h_j}(t) * (1 - A_{h_j}(t)) * \sum_k (\delta_k * W_{ho_{jk}})$$

where $A_{h_i}(t-1)$ is the activity of the hidden unit i at the previous time step and $A_{h_j}(t)$ is the activity of the hidden unit j at the present time step.

In a recurrent network, the output of the network accounts for both the current inputs and the activities at earlier times. We run the network in 13 discrete time steps for each training cycle. To compare with experimental recordings, one time step can be viewed as a duration of 100 ms. The time lag of the recurrent connection is one time step. The input of a visual or auditory target location lasts for one time step while an eye-position signal sustains until a saccade is made. The teacher signal, which is the expected ME in the output layer, appears several steps after the onset of a target simulation and lasts for one time step. This signal mimics the command to make a saccade. The weights of the feedforward connections and recurrent connections

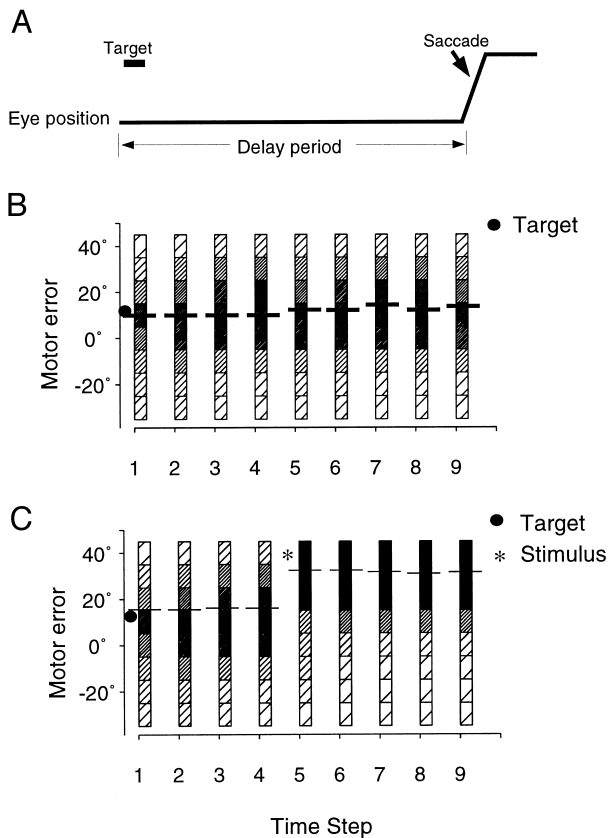


FIG. 2. The training pattern and the performance of the simple memory-saccade model. *A*: the typical training pattern. Eye movement is indicated with the lines and the short bar indicates the timing of the target. *B*: the 1-dimensional (1-D) ME output (*y* axis) through time (*x* axis). Each square indicates 1 output unit, with the gray level of the square representing the responsiveness. The short bars indicate the averaged response center of the 1-D output. The black dot indicates the time of the presentation (*x* axis) and the expected ME of the target. *C*: the 1-D output when a stimulus is presented during the delay period. The output is shifted by the stimulus.

are updated at the time of this saccade command. Note that we did not simulate the shut-off of the neuronal activity after a saccade is made (i.e., the postsaccadic suppression). Therefore the recurrent activity in the network may sustain indefinitely unless it is turned off by other mechanisms, as detailed later in the extended double-saccade model. Since different training patterns are employed for models of single- and double-memory saccades, details about the training patterns will be described in each section as needed.

RESULTS

Model of memory saccades

MODEL TRAINED WITH SINGLE MEMORY SACCADES. We first trained the model to perform single memory saccades. Twenty-five target locations across the input space and 25 eye positions were chosen as training samples. For each training cycle, a visual or an auditory target at a chosen location was presented at the first time step. A saccadic target was simulated as a dot stimulus with the amplitude of 1. The saccade was made randomly between the fifth to ninth time steps. The model was trained to encode the ME of a saccade at the time step when the saccade was made. The paradigm is illustrated in Fig. 2A. After approximately 3,000 training cycles, the network learned to produce and memorize saccadic MEs correctly to any input

pairs of target location and eye position. The performance of the trained network was evaluated by comparing the expected ME for a given target location and eye position with the produced ME at the output layer. We tested 100 random input pairs of eye position and target location for the trained network. The standard deviation of the actual ME outputs from the expected MEs was 2.62° . Figure 2B shows one example of the model output. For simplification, only eight units (which include the one with the maximum response) along one dimension (1-D) of the two-dimensional (2-D) output map are shown. The vertical axis is the 1-D ME and the horizontal axis indicates time steps. The gray level of squares is proportional to the responses of the output units. The horizontal bar indicates the gravity center of the responses. “T” indicates the expected ME of the target. Figure 2B shows that the model produces the correct output and the activity sustains throughout the delay period.

One important feature of LIP neurons is that the memory activity sustains even when new stimuli are presented during the memory period. A stimulus that appears at a different location from the target during the delay period is called an irrelevant stimulus. The memory activity of LIP neurons is resistant to irrelevant stimuli (Mazzoni et al. 1996b). However, the preceding trained network failed to produce this feature. When a new stimulus was presented during the delay period, the output pattern of the network shifted away from the expected ME, as shown in Fig. 2C. The final motor command for the saccade was thus incorrect. Correspondingly, the memory activity of hidden units was disturbed with the presentation of the irrelevant stimulus. Therefore although this network can perform simple memory saccades, it is insufficient to model the memory properties of LIP neurons.

The network also failed to produce the inhibitory activity observed in many LIP neurons. By examining the weights of the recurrent connections, we found that connections between units with similar preferred saccade directions (PD) became stronger with the progress of training. Strong excitatory connections mostly occurred between units with similar PD at the end of the training. The responses to targets were sustained through the circulation of the activity using these connections. On the other hand, inhibitory connections were rarely observed in the network.

MEMORY-SACCADE MODEL TRAINED WITH THE SINGLE-PURPOSE FEATURE. *Training and network performance.* We retrained the same model in Fig. 1 by applying the single-purpose feature to the training procedure as a constraint. Figure 3A shows a typical training pattern. The target was flashed for one time step as before. In addition to the target, an irrelevant stimulus was presented at a random location during the delay period. The irrelevant stimulus was a dot stimulus lasting for one time step. The network was required to yield the correct ME of the saccade to the target. Thus the activation of the irrelevant stimulus was to be ignored. At the time of the saccade command, the difference between the expected ME and the actual output was computed for each output unit; the weights of connections were adjusted accordingly.

In the beginning of training, the output of the network was shifted by the presentation of the irrelevant stimulus. Gradually, the effect of the irrelevant stimulus became less. Eventually, after 4,000–5,000 training cycles, the network learned to

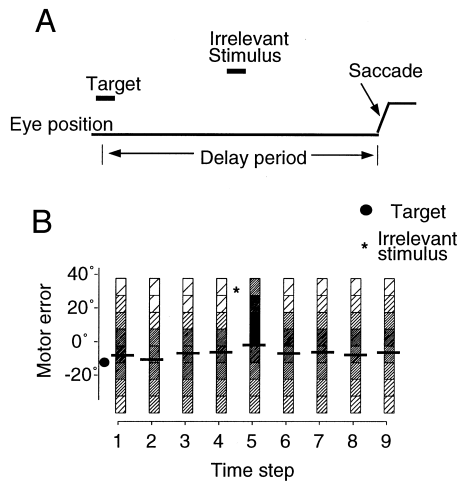


FIG. 3. The training pattern and the performance of the model with the memory-saccade feature. *A*: the training pattern. The timing of the target and the irrelevant stimulus are indicated with short bars and the eye movements are indicated with lines. *B*: the 1-D output of the model after training, illustrated in the same way as in Fig. 2. *, the presentation of the irrelevant stimulus. Compared to Fig. 2C, the stimulus did not shift the output center.

hold the correct ME memory of the first target at the end of the delay period, irrespective of the presentation of the irrelevant stimuli at any location and any time during the delay period. Figure 3B shows an example of such a response. The standard deviation of the actual output from the expected output is 2.91° .

Through training the hidden units acquired localized RFs for both visual and auditory inputs. The visual and auditory responses to targets were modulated by eye position. The map of this modulation over different eye positions is called a “gain field” (Zipser and Andersen 1988). Most hidden units were also tuned to saccadic movement directions. These properties are similar to those observed in LIP neurons and to those obtained from a similar model without recurrent connections (Xing et al. 1994). In this report, we are more interested in the sustained response patterns of the hidden units.

Figure 4 shows two typical response patterns of a hidden unit. The *left panel* indicates the RF of the unit as well as the locations of the targets and irrelevant stimuli. The *right panel* shows the responses through 13 time steps. The timing of inputs and the expected saccades is indicated at the top of the figure. In Fig. 4A, the target falls onto the unit’s RF, and the saccade is in the unit’s preferred direction. The unit responds to the target and activity sustains throughout the delay period. Notice that the brief presentation of the irrelevant stimulus during the delay period does not affect the memory activity. In Fig. 4B, the target is opposite to the preferred saccadic direction, while the irrelevant stimulus falls in the center of the RF. The unit does not respond to the target. It has a brief response to the irrelevant stimulus, but the response is immediately suppressed. Some units do not respond to the irrelevant stimuli at all. This kind of activity is often observed in double-saccade experiments in which neurons have little or no response to the brief flash of the second target during the delay period.

Through the use of various test patterns, we find that there are different types of hidden units. A small portion of the hidden units only have sensory responses but no sustained activity during the delay period—the units respond to a target,

and the responses die away soon after the target disappears. Detailed examination of these units reveals that the weights of inward recurrent connections to them are very weak. These units are merely the result of the random process of training. The majority of the hidden units exhibit different types of response patterns, depending on the tasks. A unit may have sensory responses and memory activity to a saccadic target presented in its RF as shown in Fig. 4A. Alternatively, if the stimulus in the RF is an irrelevant stimulus, the unit may only show a weak, brief responses or no response at all (Fig. 4B). More importantly, the irrelevant stimulus does not shift the firing activity away from the response evoked by the first stimulus. When the target is in a unit’s preferred direction but does not fall in the center of the RF, the unit has a weak response to the flash of the target but its elevated activity is sustained during the delay period. This memory activity is due to the excitatory inputs from other units with similar PDs (as will be explained in the next section). These response patterns are exactly what were found in LIP neurons in memory-saccade experiments (Mazzoni et al. 1996b).

Structure of the recurrent network. To understand the underlying mechanisms of saccadic memory activity, we examined the connectivity developed in the recurrent network. Figure 5A shows the weights of recurrent connections between the hidden units. The weights are plotted against the difference of preferred directions of the connected units with each dot for one connection. Compared to the recurrent connectivity in the model trained without the single-purpose constraint, strong inhibitory connections were developed between the units with dissimilar PDs in addition to the excitatory connections between the units with similar PDs. The distribution of all connection weights is relatively continuous, varying between -1 and $+1$. Figure 5B summarizes the data in Fig. 5A. The units with similar PDs have the strongest excitatory connections, and the excitatory connections become weaker as the PD difference increases. With the PDs further apart, the connections between the units become inhibitory. The strongest inhibitory connections occur to units with opposite PDs.

We therefore propose a recurrent model for memory activ-

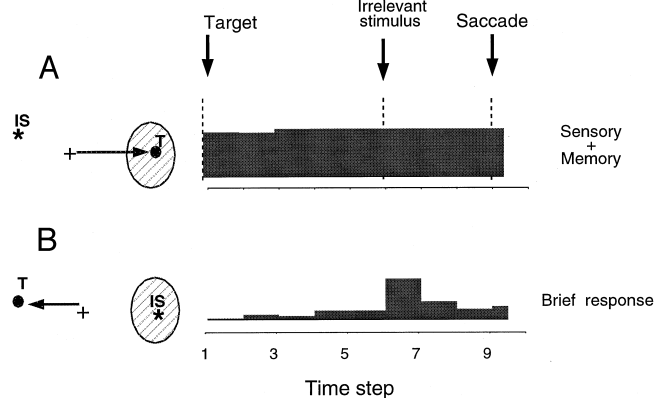


FIG. 4. Typical response patterns of a hidden unit in the memory-saccade model. *Left*: spatial arrangement of the tasks. The receptive field (RF) of the unit is indicated with the dashed area. The star symbol indicates the irrelevant stimulus (IS), and the target (T) is represented with the black dot. The arrow line shows the saccade. *Right*: each graph shows the response of the hidden unit through time. The height of the bars corresponds to the responsiveness. *A*: the sensory and memory activity to a target in the unit’s RF. *B*: the brief response to an irrelevant stimulus in the RF.

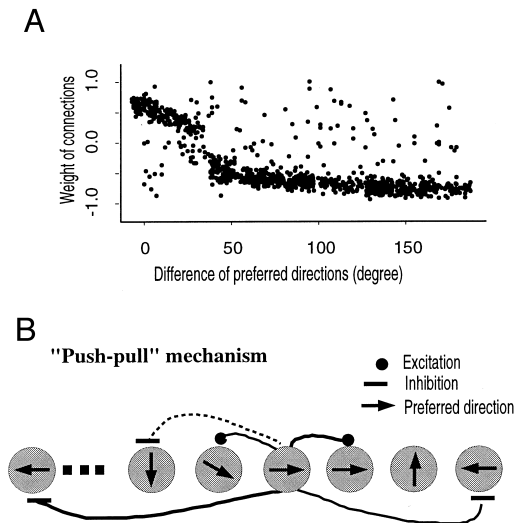


FIG. 5. The weights of recurrent connections. *A*: the weight of recurrent connections (y axis) are plotted against the difference of the preferred directions of the 2 connected hidden units. Each dot is for 1 connection. *B*: the diagram of the push-pull mechanism. Shaded circles represent hidden units and lines represent recurrent connections.

ity: lateral excitation pulls responses together from units with similar PDs to maintain the activity over a period of time, while lateral inhibition pushes away any response in units with dissimilar PDs so that their responses do not disturb the ongoing memory activity. Such a push-pull structure could be the basic architecture for the single-purpose feature of memory activity in area LIP. Recurrent excitation may invoke a set of neurons with similar PDs to maintain the memory activity. Once this neuron population is engaged, those neurons with dissimilar PDs are suppressed due to the inhibition. Thus when a new stimulus is presented at a different location, the neurons tuned to that direction are inhibited. Even if some of these neurons may respond weakly, as shown in Fig. 4*B*, the activity is immediately suppressed by the existing cooperative activity of the first population. Therefore the push-pull structure can lock the ongoing activity to prevent it from being disturbed. Only after the remembered saccade is made and the cooperative activity is turned off, can the network perform a new task.

Notice that the inhibition is a training result of ignoring irrelevant stimuli, i.e., a result of the single-purpose feature. The weights of feedforward and recurrent connections were adjusted such that the hidden unit activity evoked by the one-time-step presentation of the dot stimulus was not strong enough to override the inhibition. Since the ability of the network to resist irrelevant stimuli depends on the training stimuli used, a strong sustained irrelevant stimulus or simultaneously presented multiple stimuli could override the recurrent activity of the network trained here. Similarly, a strong sustained inhibitory input to the hidden layer could override the recurrent activity maintained by excitatory recurrent connections. This allows the network to be reset quickly.

Model of double saccades

LIP neurons participate in planning sequential eye movements. This has been typically studied with double-saccade experiments. In this section, we first summarize the neurophysiological data and then extend the memory-saccade model to make a sequence of two saccades.

PHYSIOLOGICAL RESULTS TO BE MODELED. The delayed double-saccade tasks by Mazzoni et al. (1996b) were designed to test whether LIP neurons encoded sensory locations or motor plans of saccades in sequential eye movements. The monkeys were required to memorize two targets briefly flashed in succession during a delay period and to make a sequence of two saccades to the two targets after the fixation light went off. The memory activities during the delay period (before the 1st saccade) and during the intersaccadic interval (after the 1st saccade and before the 2nd saccade) were examined. Extracellular recordings showed that during the delay period, many neurons whose movement fields were in the direction of the first saccade fired continuously until the first saccade was made, whereas neurons coding for the direction of the second saccade started to fire only after the first saccade was performed. Figure 6 shows the responses of a typical LIP neuron in different double-saccade tasks. The *left panel* shows the two saccades made toward the two remembered targets. The dashed curve indicates the neuron's RF. This neuron preferred saccades in the down-left direction. The saccadic targets are indicated with black dots and labeled as T1 and T2. Responses of the neuron are shown in the *right panel*. The delay period is labeled as M1. The horizontal and vertical eye positions are plotted under the responses. The first deflection in these eye traces corresponds to the first saccade and the second deflection corresponds to the second saccade. In Fig. 6*A*, both targets fall in the RF, and only the first saccade is in the neuron's PD. The neuron fires during the delay period. The sustained activity goes off after the first saccade is made. In Fig. 6*B*, the first target is outside the RF and the second target falls in the RF. The second saccade is in the neuron's PD. The neuron has a brief response following the flash of the second target, and this activity does not sustain during the delay period. After the first saccade is completed, the neuron begins to fire and the activity sustains until the monkey makes the second saccade. Thus the activity is related to the second saccade. In Fig. 6*C*, no targets fall in RF, but the second saccade is in the neuron's PD. The neuron has no response to the flash of either target. However, it fires during the intersaccadic interval and thus codes for the second saccade. Therefore this neuron encodes a preferred impending movement regardless of target locations. As shown in Fig. 6*C*, the activity does not even depend on sensory stimulations. Seventy-seven percent of LIP neurons recorded encode the impending saccade. It is concluded that the memory activity of the majority of LIP neurons encodes the next planned saccade. On the other hand, 16% of neurons encode target locations instead. These neurons begin to fire after the flash of the second target, which falls in their RFs, and the activity lasts through the delay period and the intersaccadic interval. These neurons may participate in programming subsequent saccades because information about the second target needs to be held until the first saccade is performed. The remaining neurons, approximately 7%, were difficult to classify into one or the other of the two categories.

From a large amount of experimental data, we generalized three basic features about LIP neurons in double-saccade tasks. *Feature 1*: Single purpose—The sustained activity for the first saccade is only minimally transiently affected by the brief presentation of the second target. *Feature 2*: Postsaccadic suppression—The sustained activity is sharply turned off after the saccade is performed. This turning off is also seen in simple

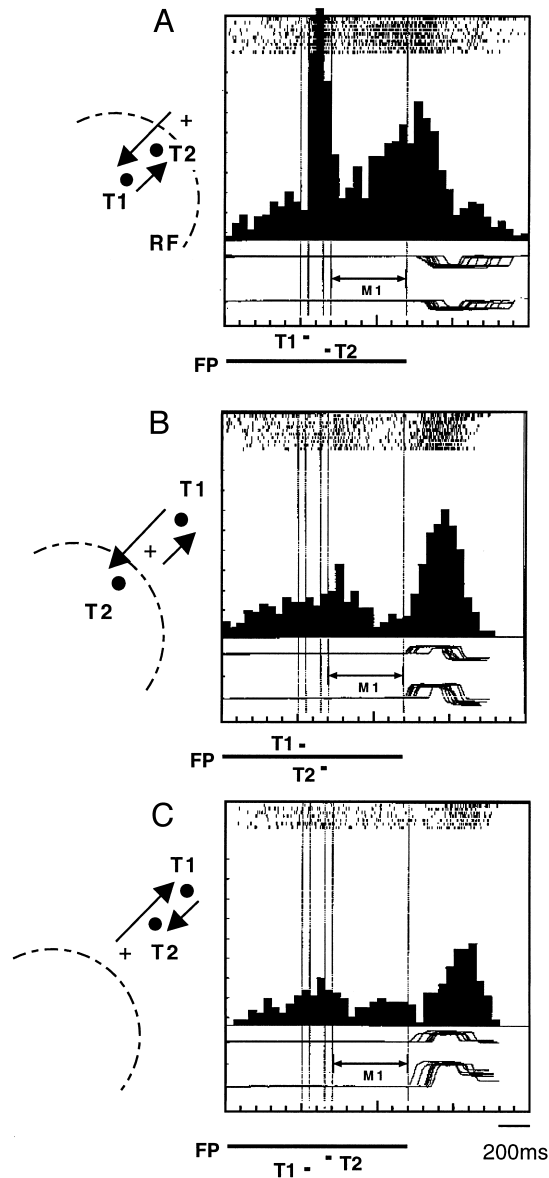


FIG. 6. Activity of a lateral intraparietal cortex (LIP) neuron coding for motor intention. The responses of a typical cell encoding impending ME in 3 double-saccade tasks. Each panel has a plot that includes, from top to bottom, the spike rasters for each trial, the time histogram of the firing rate, and the horizontal and vertical eye positions ($30^\circ/\text{division}$) (abscissa: $100 \text{ ms}/\text{division}$). The vertical dotted lines and the thick horizontal lines below each panel show the onset and offset of the visual stimuli. The deflections in the eye traces correspond to the first and the second saccades in sequence. The diagrams to the left of each panel show the spatial arrangement of the 1st and 2nd target (T1 and T2, respectively), the 1st and 2nd saccades (arrows), and the neuron's RF. This figure is modified from Mazzoni et al. (1996b).

memory saccades. *Feature 3: Memory buffer*—A separate population of neurons hold information about the second target. This population should project to those LIP neurons which, in turn, project to other motor/premotor areas.

MODEL. Based on the preceding experimental observations, we extended the memory-saccade model to simulate double-saccade tasks. Figure 7 is a diagram of the extended model. Besides the recurrent network in the original memory-saccade model, called recurrent net I (RN-I) here, the extended model has an additional recurrent net in the hidden layer (RN-II). This

population of units also receive visual, auditory and eye-position inputs. Its output projects to the primary hidden net (RN-I). Every RN-II unit projects to all RN-I units. Like RN-I units, RN-II units are fully interconnected. The postsaccadic suppression is also built into the model. It artificially resets the activity of RN-I units to the initial state after the first saccade is made. The push-pull structure of the RN-I network is capable of carrying out *feature 1*, the single-purpose feature; postsaccadic suppression serves *feature 2*, i.e., turning off the memory activity in RN-I after a saccade is made; and the RN-II network serves the memory buffer for *feature 3*. This model is expected to produce the following response patterns: RN-I units encode the first saccade, and the activity is sustained while a brief presentation of the second target does not affect the on-going activity in RN-I due to the push-pull mechanism; information about the second target and the initial eye position is maintained in RN-II; the postsaccadic suppression turns off RN-I activity after the first saccade is made; and after the first saccade is performed, RN-I combines the new eye-position information with the input from RN-II and produces a new ME for the second saccade.

The RN-II network acts as a memory buffer for the second target. In a delayed double-saccade task, different populations of neurons must be involved to hold the information about each target and thus a memory buffer is necessary. This memory buffer could correspond to the 16% of LIP neurons coding for target locations (Mazzoni et al. 1996), or it may come from some brain areas outside area LIP, such as area 7a or the frontal lobe. We do not specify which of the two possibilities correspond to the RN-II network since there is currently not sufficient experimental evidence to make this determination. The RN-II network loads the target that is retained in a memory buffer, i.e., the input lines of the RN-II network are open only after the onset of the second target.

An important control structure of the model is the postsaccadic suppression, which turns off the activity of RN-I units after the saccade is made. Such a turning-off action is necessary for neurons to encode a new saccade. During single-memory-saccade and double-saccade tasks, sharp turning-off of LIP neuronal activity is often observed right before or after the saccade. One possible source of such suppression is the efferency copy of the eye movement command. However, in a change-plan experiment (Bracewell et al. 1996), where the

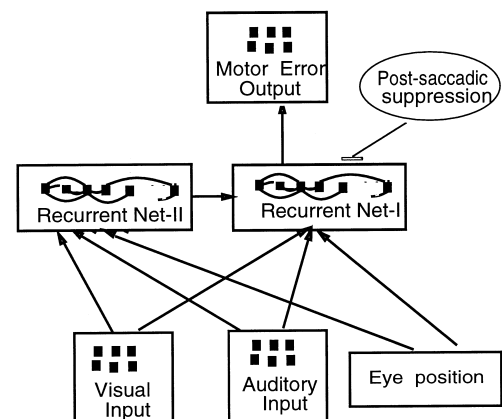


FIG. 7. The diagram of the extended model for double saccades. The recurrent net-II (RN-II) and the postsaccadic suppression are added to the simple memory-saccade model shown in Fig. 1.

monkey was required to prepare a saccade to a new target during the fixation period, the memory activity of LIP neurons for the previous planned saccade was turned off sharply even though no eye movement was made. Thus eye movement information could not be the only source for the suppression. A high-level signal that changes the memorized saccadic plan may terminate the activity of the neurons. The suppression thus could be due to strong inhibitory inputs from other high-order cortical areas, such as the frontal eye field (FEF). Many neurons in the FEF exhibit postsaccadic activities (Bruce and Goldberg 1985; Goldberg and Bruce 1990). Given that the FEF has feedback connections to area LIP, it is possible that those FEF neurons send a damping signal to LIP to provide the postsaccadic suppression. The generation of such inhibitory inputs is beyond the scope of the model. We mimicked this postsaccadic suppression by simply resetting the network artificially.

TRAINING PROCEDURE. In each training cycle, two targets (T1 and T2), either visual or auditory, are randomly selected for position and modality and presented to the network for a duration of one time step. With E0 representing the initial eye position, V1 for the location of the first visual target in retinal coordinates, and A1 for the location of the auditory target in head-centered coordinates, the desired ME output at the end of the delay period is $ME = V1$ for visual targets or $ME = A1 - E0$ for auditory targets. After the first saccade is made, the eye is moved to the new position E1. The desired output at the time of the second saccade is $ME = V2 + E0 - E1$ or $ME = A2 - E1$. Here V2 and A2 indicate the second visual and auditory targets.

Figure 8A illustrates the training protocol. Target T1 is flashed at the beginning of a training cycle and T2 is flashed at a randomly selected later time step. The two saccade commands, labeled as S1 and S2, are made for the two targets, separated by two time steps in the intersaccadic interval. The length of this interval is arbitrary, simply mimicking the approximately 100- to 200-ms time lag between two consecutive saccades in double-saccade tasks. The initial eye-position signal E0 lasts until the time of S1; the new eye-position signal E1 starts after S1. The RN-I network is open all the time except for the reset at the time of S1. The RN-II network begins to open only after the onset of target T2. The error signals for learning are computed at the time of S1 and S2. The connection weights are adjusted accordingly.

Figure 8B shows an example of the model performance at the output layer after training. After the training is completed, the connections are fixed. Like in the training period, the network is run for 13 time steps for a given combination of the two target locations and the initial eye position. The ME outputs of the network are plotted along the vertical axis. The two black dots indicate the expected MEs of the two saccades. The output of the model encodes the first saccadic ME before the first saccade is made, and then encodes the second saccade. Thus the model produced two saccade commands in sequence. The push-pull structure assures that the model carries out multiple saccade plans sequentially.

RESPONSES OF THE HIDDEN UNITS. After training, most hidden units in RN-I and RN-II developed localized RFs for both visual and auditory inputs. When eye position was centered in the orbit, the visual and auditory RFs of a given unit were

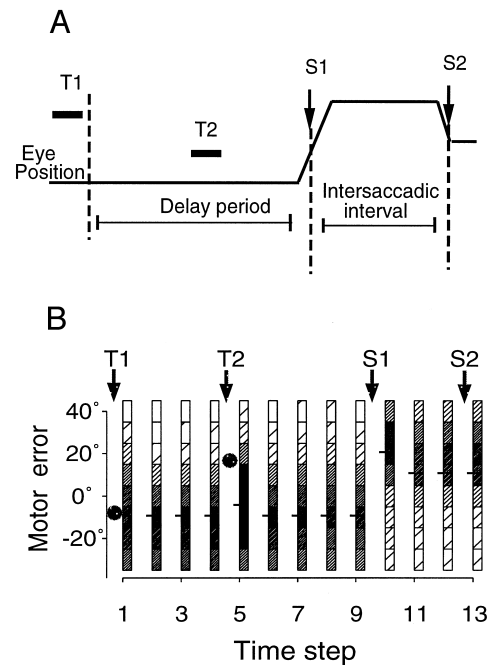


FIG. 8. Training pattern and the performance of the double-saccade model. A: the training pattern includes 2 targets, T1 and T2, and 2 saccades, S1 and S2. The time lag between S1 and S2 is the intersaccadic interval. B: the 1-D motor error output of the model illustrated in the same way as in Fig. 2. The timing of T1 and T2 and S1 and S2 are indicated on the *top*.

usually aligned. The RFs were very large; some of them even occupied up to half of the input space. The responses of the hidden units to visual or auditory targets were gain modulated by initial eye position. Later we will discuss how these gain fields are essential for coordinate transformations.

As in the single-memory-saccade model, most hidden units in the double-saccade model exhibit sustained memory activity to visual and auditory targets. Here we show the typical response patterns of the hidden units to visual targets to make direct comparisons with the experimental data in Fig. 6. Figure 9 illustrates five response patterns of two typical hidden units. Figure 9, A–C, shows the responses of a unit in the RN-I network; Fig. 9, D and E, shows the responses of a unit in the RN-II network. The *left panel* shows the spatial arrangements of the saccades. The RFs of the units are outlined with the dashed areas. The initial eye positions are indicated with + symbols. The two targets are labeled as T1 and T2, and the two saccades are labeled as S1 and S2. The responses of the units are shown on the *right panel* with the height of the vertical bars proportional to the responsiveness. The targets are flashed sequentially on the first and fourth steps and the model produces the first saccade (S1) at the 10th time step and the second saccade (S2) at the 13th time step, as indicated (*top right*).

The double-saccade arrangements in Fig. 9, A–C, are the same as those in Fig. 6, A–C. In Fig. 9A, S1 is in the unit's preferred direction. The unit responds to the target and the activity sustains until the postsaccadic suppression turns it off at the time of S1. In Fig. 9B, only T2 falls in the RF and S2 is in the unit's preferred direction. The unit has a brief response when T2 is presented, and this activity is suppressed by other hidden units that encode S1 during the delay period. This unit begins to fire after S1. In Fig. 9C, no targets fall in the RF, but S2 is in the unit's preferred direction. The unit still fires during

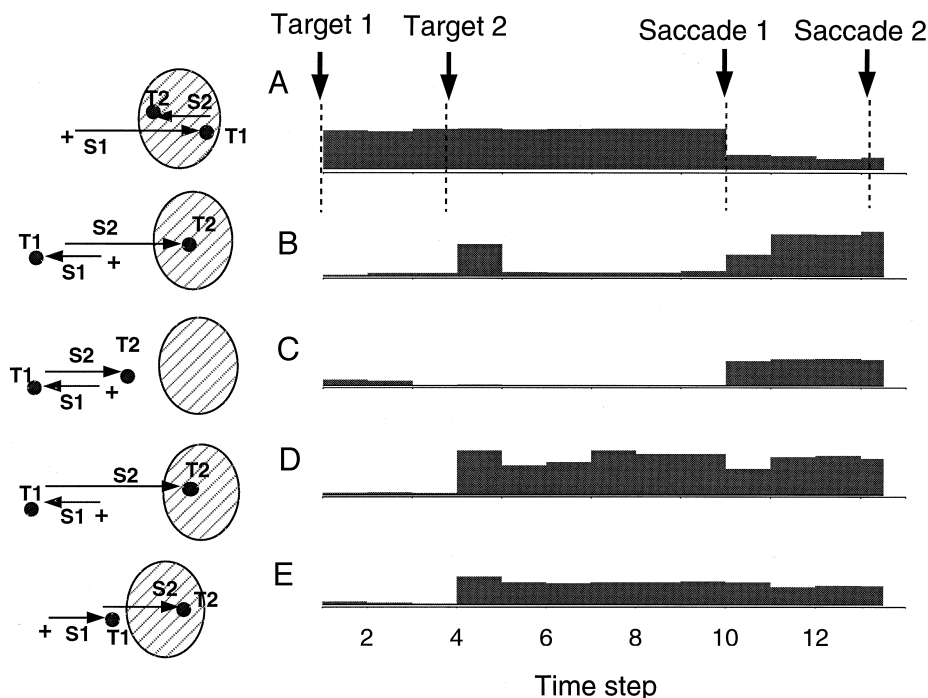


FIG. 9. Typical responses of 2 hidden units in the double-saccade model. The timing of targets and saccades are shown on the top of the figure. *Left:* the spatial arrangements, with the dashed area for the RFs, T1 and T2 for the 2 targets, and S1, S2 for the 2 saccades. *Right:* the response patterns. A–C: responses of a RN-I unit in 3 double-saccade tasks. D and E: the responses of a RN-II unit in 2 tasks.

the intersaccadic interval, coding for S2. Like the neuron shown in Fig. 6, the sustained responses of this model unit code the upcoming saccade. The result of Fig. 9C is intriguing in that a unit can be activated without a target in its RF.

Notice that the neuronal responses shown in Fig. 6 exhibited complex dynamic patterns. For example, the activity in Fig. 6A had a dip between the offset of the second target and the onset of the saccade. This dip might correspond either to the second target or to the saccade onset. The activities in Fig. 6, B and C, also had similar dips. The model responses in Fig. 9, A–C, did not capture these dynamic response patterns. The model units updated their activities at a time step of 100 ms while neurons updated their activities at an order of 1 ms. To capture those neuronal dynamics requires a network with realistic model neurons and stochastic processing.

Figure 9, D and E, shows the responses of a typical unit in the RN-II network. The unit begins to respond after the onset of T2 in its RF, and the activity is sustained. The locations of target T2 in Fig. 9, D and E, are the same, but the initial eye positions in the two graphs are different. The unit responds to T2 in both cases. However, the responses are strongly modulated by the eye position. The responsiveness in Fig. 9E is weaker than that in Fig. 9D as the eye position moves in the opposite direction to the unit's RF from Fig. 9, D to E. The information about the eye position is thus combined with the target's retinal location through this modulation. Therefore the information about head-centered representation is implicitly carried in the activity of a set of RN-II units.

COORDINATE TRANSFORMATIONS. One traditional question about double-saccade tasks is how the motor vector for the second saccade is computed, given that eye position at the time of the second saccade is different from the time when the visual target was flashed. How are the spatial transformations required for double-saccades carried out? To answer this question, we examined how eye-position information is utilized by the hidden units in the model.

We first examined the hidden units in the RN-II network. The RF of a unit was first measured at the central eye position. Next, for 8×8 eye positions, the responses to a target presented in the RF were measured. Results showed that the responses of most hidden units were modulated by eye position. The 2-D plot of responses against different eye positions is called gain field (GF) as reported by Andersen et al. (1985). The GFs of RN-II units monotonically increase in particular directions. Figure 10, A and B, shows the RF and the GF of a typical unit. In Fig. 10A, the gray levels of the small squares correspond to the responses of the unit to the target presented at different locations of the input map while the eye position is pointed at the central fixation. In Fig. 10B, the sizes of the squares indicate the responses to a target presented in the center of the RF for different eye positions. Notice that the unit's RF and its GF are in the same direction. This is typical for the majority of the RN-II units. Figure 10C shows the direction differences of the GF and the RF for every RN-II unit. Most units have an aligned RF-GF structure. The two units whose RF-GF direction differences are close to 180° have weak responses. Therefore they have little contribution to the network. A group of units with the aligned RF-GF structure is well suited for the transformation from eye- to head-centered coordinates since this transformation requires addition of eye position and retinal position. Previously we have demonstrated that a population of units with aligned RF-GF contains an implicit representation of target locations in head-centered coordinates (Xing et al. 1994; Zipser and Andersen 1988).

Next we examined RFs and GFs of the units in the RN-I network. After training, most RN-I units developed RFs for visual and auditory inputs and monotonic GFs for eye position. Figure 11, A and B, shows the RF and GF of a typical RN-I unit. Unlike the one in Fig. 10, this RF and GF of the unit are in opposite directions. Figure 11C shows the RF-GF direction differences for all the RN-I units. The result indicates that most

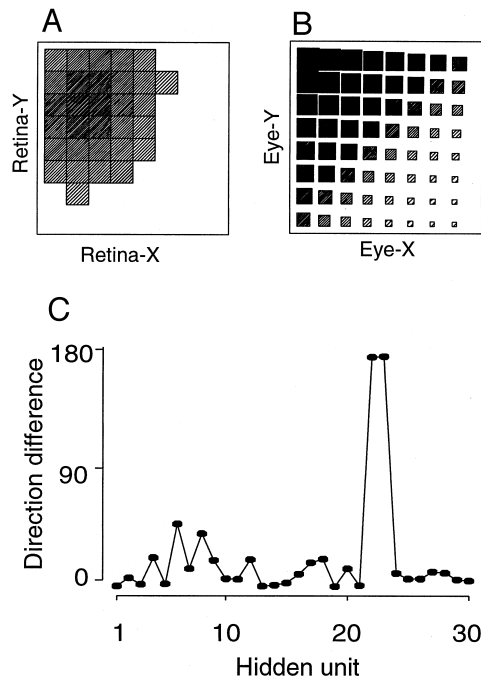


FIG. 10. RF and GF of a RN-II unit. *A*: the visual RF of a typical hidden unit measured at the central eye position. The gray level of the dashed squares is proportional to the evoked response in the hidden unit. *B*: the spatial gain field (GF) of the unit. The GFs are the responsiveness of the unit to a target within its RF plotted against an 8×8 grid of eye positions spaced by 10° . The gray level and the size of small squares in the graphs correspond to the activation of the unit. *C*: the RF-GF direction differences for every hidden unit in the RN-II network. The direction of a RF was calculated as the vector direction from the center of the input map to the center of the RF. The difference between tuning direction of a GF and the RF direction was computed for every hidden unit and shown in the figure. Hidden units are listed along the horizontal axis; the vertical axis indicates the corresponding absolute value of direction difference. Most hidden units have direction differences close to 0° , i.e., the aligned RF-GF structure.

RN-I units have an opposite RF-GF structure. The opposite RF-GF structure is well suited to carry out the transformation from head-centered coordinates to eye-centered coordinates. This transformation requires subtraction of eye position from a head-centered target location. The opposite signs for changes of eye position and head-centered location, due to the opposite RF-GF structure, meet the requirement of the subtractive operation (Xing et al. 1994). This operation is required for computing the ME of the second saccade from a distributed head-centered representation in the RN-II network.

We further found that the visual and auditory RFs (VRFs and ARFs) of RN-I units differed in two aspects: 1) although the VRF and ARF of a given unit usually aligned roughly, the VRF was smaller than the ARF. Most ARFs were planar and spread toward the edges of the auditory input map. 2) The GFs for ARFs were much stronger (with steeper slopes) than those for the VRFs. Further examining the connection weights we found that the weights to the RN-I units from the visual inputs were on the average stronger than the weights from the auditory inputs and the RN-II inputs. With the sigmoidal integration between the signals of eye position and target location, the stronger connection weights to visual inputs resulted in a weak effect of eye position on the visual responses. Due to this weak gain modulation, no coordinate transformation occurred to visual inputs of single visual saccades or the first visual sac-

cade in a double-saccade task. This resulted in one of the model functions: $ME = V1$.

The results of Figs. 10 and 11 show that the majority of RN-I units have an opposite RF-GF structure and the majority of RN-II units have an aligned RF-GF structure. This segregation of RF-GF types is associated with the output function of the double-saccade model. The output layer in the present model is a single map of eye MEs. Hence the only task for the hidden layer is to compute MEs. There might be other types of coordinate transformations occurring in area LIP as well. The double-saccade model here may only reflect a part of the more complicated LIP functional structures for different sensorimotor integrations. When we modified this model by having multiple output maps in different coordinates, for example, a ME map and a head-centered spatial map, the distribution of RF-GF types in the hidden layer changed; the units in both RN-I and RN-II networks exhibited aligned, opposite, and intermediate RF-GF structures (Andersen et al. 1997).

The results in the preceding text outline the coordinate frames used by the hidden units to encode saccadic targets in double-saccade tasks. We further looked into the coordinates of RN-I and RN-II network. For RN-II units, the tuning curves to target retinal locations were plotted for different initial eye positions. A diagram showing how the tuning curves are computed is shown in Fig. 12A. The — and - - - represent retinotopic frames at the two eye positions E and E'; the trajectories indicate the saccades to eight target locations. The retinal target locations in the two eye-position frames are identical. If a unit encodes saccadic targets in retinal coordinates, the retinotopic tuning curves for different eye positions should align with each other. In contrast, if the unit encodes targets in motor coordinates, the tuning curves should shift with eye position. Figure 12B shows the retinotopic tuning curves for a typical RN-II

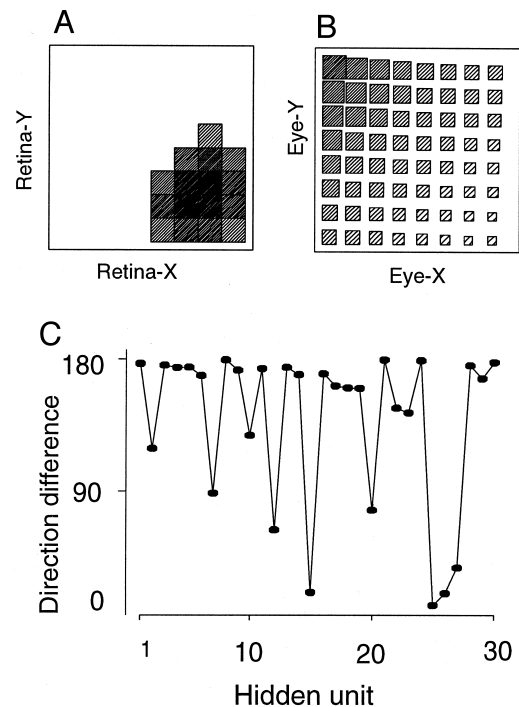


FIG. 11. RF and GF of a RN-I unit. The illustrations are the same as in Fig. 10. *A*: the RF of a typical RN-I unit. *B*: the GF. *C*: the direction differences. Notice that most RN-I units have GF and RF in the opposite direction.

unit. In Fig. 12B, the vertical axis represents the responses and the horizontal axis indicates retinal locations. The — and - - - are for the two eye positions. Although the responsiveness for a given retinal location is modulated by eye position, the two tuning curves align well with each other. Therefore RN-II units encode inputs of saccadic targets in retinal coordinates. These units may correspond to the small portion of LIP neurons that encode the memory of target locations (Mazzoni et al. 1996b).

Next we investigated the coordinates of the RN-I network. Figure 13A shows the diagram for computing tuning curves of RN-I units. The x and y axes are in head-centered coordinates. The target positions of the first saccade are indicated by \odot , i.e., the initial eye positions of the second saccade. The trajectories represent the second saccades. In one test, the second saccades are made outward to eight targets, as indicated with —. In the other test, the second saccades are made inward from a different set of initial eye positions to the same sets of targets as in the first test. Thus the target locations are the same for the two tests although the directions of the second saccades toward a given target are different in the two tests. The tuning curves of RN-I units are plotted for each test. Figure 13B shows the tuning curves for a typical RN-I unit. The responsiveness of the unit is plotted against the retinal location of the second saccadic targets; the — and - - - are for the two tests. The results show that the two tuning curves do not align in retinotopic coordinates. In Fig. 13C, the same set of data in Fig. 13B is plotted against eye MEs. The two ME tuning curves align well with each other. Therefore RN-I units encode saccadic targets in motor coordinates. These units may correspond to the ma-

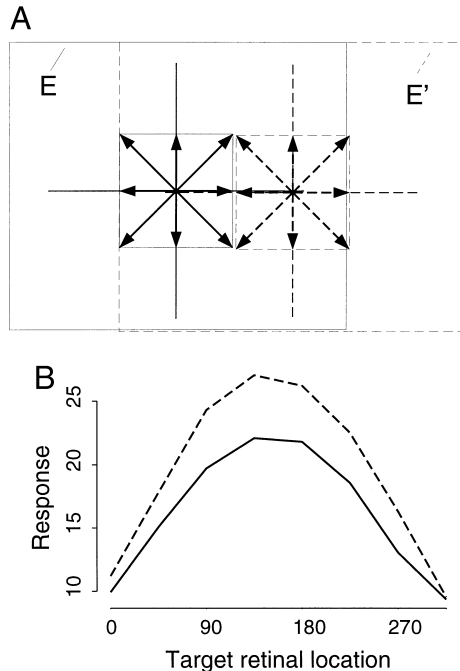


FIG. 12. Tuning curves of a typical RN-II unit. A: the method for computing retinal location tuning curves for different eye positions. — and - - -, retinotopic frames at the 2 eye positions, E and E'. For each eye position, the arrays indicate the saccades to 8 retinal locations. The retinal locations are identical for the 2 eye positions. B: the tuning curves of a typical RN-II unit. The responsiveness of the unit (the vertical axis) is plotted against target retinal locations (the horizontal axis). — and - - -, the two initial eye positions, respectively. The 2 tuning curves align well with each other.

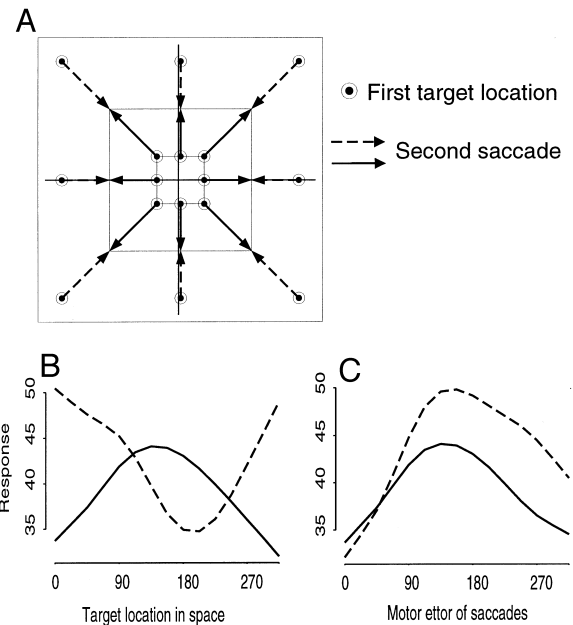


FIG. 13. Tuning curves of a typical RN-I unit. A: the diagram for computing tuning curves in double-saccade tests. \odot , retinal locations of the 1st target, i.e., the initial eye positions of the 2nd saccade. The arrays represent the 2nd saccades. —, saccades to 8 target locations in 1 test; - - -, the saccades that are in the opposite direction but end at the same locations as in the 1st test. B: the tuning curves for a typical RN-I unit. The responsiveness of the unit is plotted against the retinal location of the 2nd saccadic targets; — and - - -, for the 2 tests, respectively. The 2 tuning curves do not align on the retinotopic space. C: the same set of data as in Fig. 13B is plotted against saccadic motor errors (the horizontal axis). The 2 motor error tuning curves align well with each other. Thus the unit encodes saccades in motor coordinates.

majority of LIP neurons that encode MEs of saccades (Mazzoni et al. 1996b).

Theoretically, eye-position modulation with the aligned RF-GF structure yields a head-centered representation of target locations in the distributed activity of RN-II units. This head-centered representation is fed into the RN-I network after the first saccade. With the opposite RF-GF structure in the RN-I network, the new eye-position information is subtracted from the head-centered representation to yield a ME of the second saccade. Hence RN-I units encode saccadic targets in motor coordinates. Therefore the model carries out coordinate transformations required for double saccades in the following steps: the RN-I network transforms target locations into the representation of MEs through the opposite RF-GF structure (for auditory targets), and the RN-II network provide the RN-I network with a head-centered representation of the second target through the aligned RF-GF structure. Thus the RN-I network encodes motor plans of saccades, while the RN-II network represents head-centered information implicitly through distributed coding of RN-II units.

DISCUSSION

The mechanisms for programming memory saccades and sequential saccades remain unclear to neurophysiologists. A number of computational models of saccade generation have been proposed. Dominey and Arbib (1992) proposed a cortical-subcortical model of the control of saccadic eye movement and suggested that the parietal cortex may dynamically remap the

target locations in saccade ME maps to program double saccades. The network model developed by Droulez and Berthoz (1991) showed that target position could be memorized in a sensory map and updated with eye-movement signals. Krommenhoek et al. (1993) trained a neural network to compute MEs using information about eye position. These computational approaches yield valuable insights into memory saccades. On the other hand, the frameworks in these models did not correspond well to known neurophysiological data. Given that LIP neurons can withhold their saccade-related activity and participate in programming double saccades, the network model in this report studied the memory activity in area LIP for saccadic eye movements. With the implementation of the single-purpose rule in the training process, the network developed lateral excitation-inhibition (the push-pull structure) that was essential to memory and sequential saccades. The simulated neurons in our model exhibited properties similar to those recorded in area LIP. After training to make double saccades, the model carried out the coordinate transformations required to program double saccades by the means of gain modulations. In our model, one group of neurons maintain the sensory memory of saccadic targets, while the other group of neurons encode the motor plan of an impending saccade. Thus coding the motor commands of double saccades is achieved by different neuronal populations rather than by dynamically remapping the same neuronal population.

One prediction of the model is that neurons corresponding to the memory buffer RN-II respond to the second target but not the first one. Mazzoni et al. (1996b) found that 16% of LIP neurons encoded the location of the second target in a memory-saccade task. These cells were referred as the “sensory memory” cells. It would be interesting to test experimentally whether these cells encode only the second target or any visual stimuli within their RFs. Furthermore the model predicts that the responses of RN-II neurons are gain modulated by the initial eye position. Due to the push-pull structure, the activity of RN-II neurons is not affected by the new eye position. This remains to be tested experimentally.

Push-pull structure

Examining data from various kinds of delayed-saccade experiments, we found a common feature in the response patterns of LIP neurons—once a neuron is engaged in a saccade command, its activity is maintained irrespective of further stimuli; the neuron starts to respond to another stimulus only after the saccade being encoded is completed or the intention of the saccade is dismissed by some high level command. We called this feature single-purpose. This feature is essential for the behaviors of a motor system. The eyes can never make saccades simultaneously to two different spots.

The single-purpose feature is used as a constraint for the training process of our networks. This constraint results in the push-pull structure, i.e., excitatory connections between units with similar preferred saccadic directions and inhibitory connections between units with dissimilar preferred directions. Such an excitation-inhibition structure is the neuronal basis for the single-purpose feature. In the extended model of the double-saccade system, the push-pull structure allows the model to program two saccadic commands sequentially, rather than mixing the two commands into one. Therefore the excitation-

inhibition connections ensure that LIP encodes the next planned saccade.

Ideally, a push-pull structure suppresses any irrelevant stimuli that differ from the target location. However, through simulation we found that responses to stimuli close to the target were often sustained rather than suppressed. The minimal distance for an irrelevant stimulus to be suppressed varied from unit to unit but was roughly in an order of about 10° . Within this distance, the output memory activity represented a ME that was a weighted average between the irrelevant stimulus and the target. The function of the connection weights in Fig. 5A reflected this inaccuracy: Excitatory connections could occur to PDs that are $10\text{--}20^\circ$ apart. Several reasons contribute to this inaccuracy: the tuning of the hidden units is broad; the limited number of the hidden units prevents precise excitatory connections; and the training samples of the stimulus locations are often more than 10° apart. We expect that using a large set of hidden units and finer spaced training stimuli, or an attention mechanism, would improve the accuracy of the push-pull structure.

A number of neural network studies have used a push-pull structure as a memory-storage mechanism (Grossberg and Levine 1975; Seung 1996; Zhang 1995). Typically, adjacent units in these networks excite each other while distant units inhibit each other. Such an arrangement could prevent recurrent activities from spreading to the whole network. Thus a push-pull mechanism also enforces the stability of a recurrent network. Salinas and Abbott (1996) recently proposed another functional role of the push-pull structure in the parietal cortex. They found that neurons in a recurrently connected network with push-pull connections could perform a product operation on additive synaptic inputs. The resulting multiplicative gain modulation is important for coordinate transformations in the parietal cortex. In our model, the push-pull structure emerged as the result of the single-purpose feature. Moreover, the excitation and inhibition were organized according to the preferred directions of the units, rather than the geometric positions. Goldman-Rakic (1995) observed a similar lateral inhibition structure in the opponent memory field of neurons in the frontal cortex. Schlag et al. (1998) found that, in the FEF, cells that encoded similar eye movements mutually excited each other while silencing those that would produce conflicting eye movements. Since the single-purpose feature might be common for cortical areas involved in motor planning, it is likely that the push-pull structure is a principle applicable to these cortical areas.

An analogy to the single-purpose feature is the winner-take-all mechanism. The latter has been widely applied to the models of visual search processes (Braddick 1997; Ferrera and Lisberger 1995; Lee et al. 1999). In a visual search task, a target is searched among a number of distractors. A winner-take-all mechanism allows the neurons representing the target and the distractors to compete against one another. Attention serves to bias the outcome of this competition toward the direction of the selected target. As a result, the neuronal response to the target remains and the response to the distractors is suppressed. Salzman and Newsome (1994) also proposed that a winner-take-all mechanism existed in the motion cortex (area MT and MST). When more than one motion cue was presented, monkeys chose the direction encoded by the largest signal in the representation of motion direction. Braddick

(1997) suggested that local motion detectors use winner-take-all interactions in global motion analysis.

The single-purpose feature and the winner-take-all mechanism are similar in that both generate only one single output representation. The latter evokes neuronal competition based on the context of stimuli and enhances the response to the target stimulus through attention. Such a mechanism is not suitable for area LIP because LIP neurons are generally insensitive to stimulus context and thus do not support a competition process. The target to be represented in LIP is chosen by motor intention and is not the result of an attention-biased competition. The winner-take-all mechanism handles spatial conflicts in visual selection. The single-purpose feature assures no conflicts in a temporal sequence of motor plans. Neurophysiological data support our assumption that a single-purpose feature exists in area LIP. It would be interesting to test this assumption further by recording the responses of LIP neurons to a target and many distractors presented simultaneously.

Coordinate transformations

A traditional question about planning double saccades is how the motor command of the second saccade is computed. Two hypotheses have been proposed. One hypothesis is head-centered coding (Robinson 1975; Sparks and Mays 1983): the absolute target location in head-centered coordinates is computed and stored, and then the new eye position after the first saccade is subtracted. With this hypothesis, one would expect to find neurons that encode visual targets explicitly in head-centered coordinates. However, physiological studies have largely failed to find such neurons. Most LIP neurons have retinal RFs with their responses modulated by eye position. The other hypothesis is retinotopic coding, also called vector subtraction (Bruce and Goldberg 1985; Scudder 1988): the retinal location of the target is stored and then the change of eye position is subtracted. This hypothesis requires neurons that explicitly encode the change of eye position.

The simulation results of this report suggest a third possibility: instead of computing explicit head-centered target locations or the change of eye position, LIP neurons utilize eye position with the use of GFs to carry out coordinate transformation through the distributed activity of many neurons. In the double-saccade model, information about the second target location is combined with the current eye-position signal through aligned RF-GF gain modulation to form a distributed head-centered representation. After the first saccade, the new eye position comes in and is combined with the head-centered representation through the opposite RF-GF structure so that the ME of the second saccade is computed. This model does not require individual neurons to encode target locations in explicit head-centered coordinates. The presence of GFs could account for the computation of double saccades. Moreover, the experimental results by Li et al. (1995) suggested that a distributed head-centered representation of targets might be maintained in LIP for programming sequences of eye movements. Using reversible lesions of LIP, Li et al. found that the monkeys depended on the new eye position more than the retinal vectors to make the second saccade. Thus this model fits current data well.

Theoretically, coordinate transformations suggested by the first two hypotheses above can be carried out by shift circuits.

Quaia et al. (1998) proposed a shift circuit to simulate RF remapping in LIP, in which the FEF neurons shifted the RFs of the LIP neurons. However, the large RFs and the distributed coding feature of parietal neurons make it difficult for a precise shift circuit to work. The modeling results in this report show that the gain modulation is essential to carry out the coordinate transformations in area LIP. Using this strategy, neurons may remain in retinotopic coordinates for visual stimuli. With eye-position modulation, the distributed activity of these neurons can represent the stimuli in other coordinates. Varying RF-GF structures carries out different kinds of transformations. Hence the gain modulation along with distributed coding is an efficient way to achieve sensorimotor transformations without using complex shift circuits. Other theoretic studies also revealed the importance of GF properties in coordinate transformations. Goodman and Andersen (1990) analytically demonstrated that an aligned GF and RF relationship was required for transformations from oculocentric to craniocentric coordinates. A similar mechanism of eye-position modulation in the saccadic system was studied by Krommenhoek et al. (1993, 1996). They developed a neural network in which retinal signals and an efference copy of eye position could be remapped to a ME map in two steps: distributed coding of head-centered target position at one level and of ME in eye-centered coordinates at another stage.

RF remapping versus ME coding

Experimental data demonstrate that the memory activity of LIP neurons encodes saccadic eye movements (Snyder et al. 1997). Furthermore it has been shown that LIP neurons encode motor intention, irrespective of the actual execution of the planned movements (Bracewell et al. 1996; Snyder et al. 1997). The simulated LIP neurons in our models indeed encode the impending saccade. On the other hand, Duhamel et al. (1992) proposed that LIP neurons encoded sensory stimuli instead of saccades. In their experiment, as illustrated in Fig. 14A, the monkey was required to make a saccade to a remembered target, and this saccade would bring a stimulus onto the RF of the LIP neuron being recorded. It was found that the neuron responded to the stimulus outside its classic RF when an impending saccade brought the stimuli into the RF. Some neurons became active before the stimulus was brought into the neurons' RFs by the saccades. Duhamel et al. thus concluded that the RF of the neuron transiently shifted with the eyes to the retinal location at which the stimulus could excite the neuron. This hypothesis is diagrammatically illustrated in Fig. 14B. During fixation, the representation of the visual scene was stable (*left*). Immediately before or during the saccade the cortical representation shifted into the destination of the intended saccade. The neuron thus began to respond to the stimulus at a new retinal location (*middle*). After the eye movement, the cortical representation shifted back to match the visual inputs so that the neuron continued to respond to the stimulus (*right*).

Quaia et al. (1998) proposed a model to explain the observed shifts of RFs. In their model, a group FEF neurons carry the signal about impending saccades; LIP phasic-tonic neurons have stable local RFs and LIP phasic cells have shifting RFs. If a FEF neuron and a LIP phasic-tonic neuron are active at the same time, a LIP phasic neuron, whose RF is equal to the

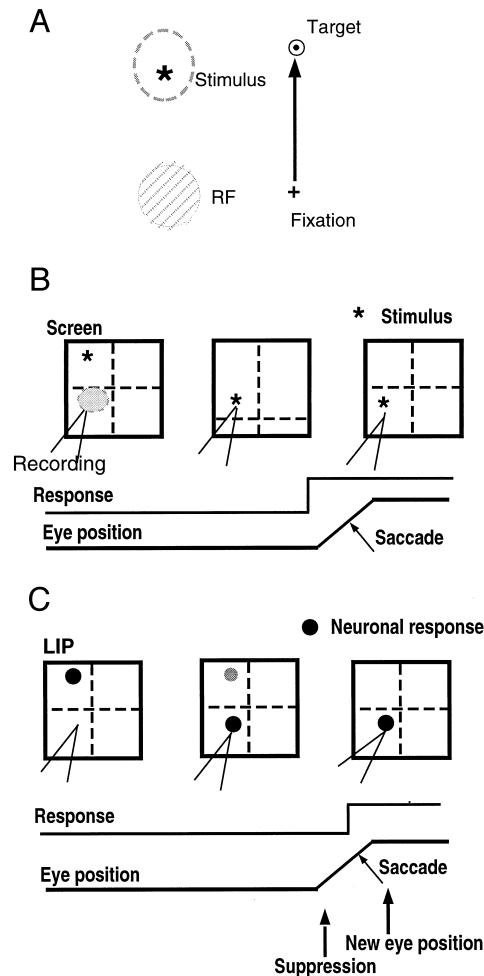


FIG. 14. Illustrations of predictive responses before saccades. *A*: the experimental paradigm used by Duhamel et al. (1992) to test the predictable activity in LIP. The monkey was asked to make a saccade to the target after the fixation light went off. A stimulus outside the RF of the recorded neuron was presented during the fixation, and the saccade would bring the stimulus onto the neuron's RF. The neuron predictably responded to the stimulus before the saccade. *B*: the diagram of transient RF shift accounting for predictable responses. The solid-lined box corresponds the cortical representation of the space, the broken lines indicate the retinal coordinates. The RF of the neuron being-recorded is illustrated with the dashed area. The black dot indicates the response to the stimulus. *Left*: the neuron has no response because the stimulus is outside the RF. *Middle*: the neurons begin to respond before the saccade because the RF shifts to overlap with the stimulus. *Right*: the situation after the saccade. The RF moves back and the stimulus is inside the RF due to the eye movements. Thus the neuron continues responding to the stimulus. The neuron's response and the eye movement through time are illustrated on the lower part of the graph. *C*: the diagram of neurons coding for the intended saccades. The black dot represents the response of the neurons. *Left*: other neurons fire for the intended saccade, the neuron being-recorded has no response. *Middle*: the command for the saccade is issued so that the on-going activity is suppressed by postsaccadic suppression. *Right*: the 1st saccade is completed and the neuron continues to fire for the next intended saccade, although this saccade may not actually be executed. The dashed rectangle indicates that the information about the new eye position can reach LIP before, during, or after the saccade is made. In both models, the neuron can predictably fire before the saccade is made. The 1st one requires remapping RFs; the 2nd one requires predictably updating eye positions.

difference of the RF of the LIP phasic-tonic neuron and the motor field of the FEF neuron, is activated. All pairs of LIP phasic-tonic neurons and FEF neurons, whose RF/motor field difference is equal, must be connected to the same LIP phasic

neuron. Therefore the RFs of the LIP phasic neurons are shifted with impending saccades. Such a model requires a specific connectivity: precise pairings between LIP and FEF neurons. It also requires specific computations at the dendritic level, i.e., a multiplication between cells in a pair and a logic OR computation between different pairs to the same LIP phasic cell. Both the connectivity and the computations are biologically difficult to implement. Moreover, although the model explained RF remapping, it did not account for the coordinate transformations in sequential saccades.

It is interesting to see how our model responds to the paradigm of Fig. 14A. In Fig. 9C, the hidden unit responded to the second target even though the target never appeared in the unit's RF. This response appeared as if the RF of the unit shifted to capture the second target, while in fact there was no RF shift and the response was merely encoding the impending preferred saccade. Figure 14C illustrates the model results in the same experiment. After the target onset, some neurons fire to the first intended saccade, and the neuron being recorded has no response (*left*) since its RF is not in the preferred direction. Next, after the command to make the saccade is issued, the on-going activity is suppressed by postsaccadic suppression. The network computes the ME of an intended saccade to the stimulus based on the inputs of the new eye position and the information about the stimulus location. As a result, the neuron under recording becomes active since its PD is in the direction of the next intended saccade (*middle*). Finally, the first saccade is completed, and the neuron continues to fire for the next intended saccade, although this saccade may not actually be executed. Therefore the hidden units in the model can encode stimuli outside the unit's RF using dynamically updated information about eye position. In this model, the cortical representation does not shift toward the stimulus and then shift back. Instead, the activity of one group of cells goes up while the others come down for a new saccade plan. Thus different populations of neurons are engaged and disengaged rather than individual neurons shifting their retinal RFs back and forth.

Duhamel et al. (1992) reported that 44% of LIP neurons became active before the saccade brought the stimuli into the neurons' RFs. Our model can account for these predictive responses. Before the first saccade is made, the RN-I network may already begin to compute the ME of the second saccade using the information about the new eye position. Therefore the units coding for the second saccade could become active before the first saccade. Thus the observed predictive remapping could be the result of the sequential activation of different populations of LIP neurons rather than jumping RFs. There is experimental evidence that signals for new eye positions appear in some LIP neurons before the beginning of a saccade (C. Li, B. Breznen, and R. A. Andersen, unpublished data). In addition, psychophysical studies by Dassonville et al. (1995) and Schlag and Schlag-Rey (1995) indicated that spatial localization during saccades was largely based on updating of the internal representation of eye position.

Modeling multiple sequential saccades

The present model only simulated single and double saccades. How would the model handle more than two saccadic targets in a sequence? Our double-saccade model can be viewed as a schematic version of a model of multiple-sequen-

tial-saccades. In this report, we focused our model on how the coordinate transformations of two sequential saccades were carried out. We could extend the model to handle multiple saccades in the following two ways: the model has more memory buffers each holding the memory of every additional saccade and the input layer of the model could correspond either to sensory inputs or to inputs from a memory system. The first possibility is perhaps too rigid and the architecture is difficult for brain to implement. In the second possibility, the target locations and the orders of the presentation are held in the memory system while RN-I and RN-II networks carry out the coordinate transformations for the impending saccade. Behavioral and physiological evidence supports this possibility. Training monkeys to perform more than two sequential saccades is difficult. Barone and Joseph (1989) were able to train monkeys to make sequential saccades to three fixed target locations. However, they only observed prefrontal neurons that responded to the first target or the second target, but no neurons responded to the third target. The result suggested that the memory for more than two sequential targets was not directly handled by the parietal or the prefrontal cortex.

In summary, the models in this report capture the important characteristics of LIP neurons and provide insights into the mechanisms of LIP in programming eye movements. By optimizing the network to implement various saccadic tasks, two important properties emerge from the model: push-pull recurrent connections and opposite/aligned GF structures. These properties are the basics for programming memory saccades and sequential saccades. The consistency of simulated results and current experimental data suggests that the models are well suited to describe the sensorimotor processing in area LIP and thus can be used as a framework to guide future experiments in understanding the neural functions of LIP.

We thank P. Mazzoni for providing experimental data, C. Li for valuable discussions, and K. Shenoy for valuable comments on the manuscript.

This work was supported by National Eye Institute Grant EY-05522. J. Xing was supported by the Del Webb Foundation fellowship.

REFERENCES

- ANDERSEN RA. Visual and visual-motor functions of the posterior parietal cortex. In: *Neurobiology of Neocortex*, edited by Rakic P and Singer W. New York: Wiley, 1988, p. 285–295.
- ANDERSEN RA, ASANUMA C, ESSICK C, AND SIEGEL RMM. Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *J Comp Neurol* 296: 65–113, 1990a.
- ANDERSEN RA, BRACEWELL RM, BARASH S, GNADT JW, AND FOGASSI L. Eye position effects on visual, memory, and saccade-related activity in areas LIP and 7a of macaque. *J Neurosci* 10: 1176–1196, 1990b.
- ANDERSEN RA, ESSICK GK, AND SIEGEL RM. The encoding of spatial location by posterior parietal neurons. *Science* 230: 456–458, 1985.
- ANDERSEN RA AND GNADT JW. Posterior parietal cortex. In: *The Neurobiology of Saccadic Eye Movements*, edited by Wurtz RH and Goldberg ME. Amsterdam: Elsevier, 1989, p. 315–335.
- ANDERSEN RA, SYNDER LH, BRADLEY DC, AND XING J. Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annu Rev Neurosci* 20: 303–330, 1997.
- BARASH S, BRACEWELL RM, FOGASSI L, GNADT JW, AND ANDERSEN RA. Saccade related activity in the lateral intraparietal area. II. Temporal properties; comparison with area 7a. *J Neurophysiol* 66: 1095–1108, 1991a.
- BARASH S, BRACEWELL RM, FOGASSI L, GNADT JW, AND ANDERSEN RA. Saccade related activity in the lateral intraparietal area. II. Spatial properties. *J Neurophysiol* 66: 1109–1124, 1991b.
- BARONE P AND JOSEPH JP. Prefrontal cortex and spatial sequencing in macaque monkey. *Exp Brain Res* 78: 447–464, 1989.
- BEISER DG AND HOUK JC. Model of cortical-basal ganglionic processing: encoding the serial order of sensory events. *J Neurophysiol* 79: 3168–3188, 1998.
- BRACEWELL RM, MAZZONI P, BARASH S, AND ANDERSEN RA. Motor intention activity in the macaque's posterior parietal area LIP. II. Changes of motor plan. *J Neurophysiol* 76: 1457–1463, 1996.
- BRADDICK O. Local and global representations of velocity: transparency, opponency, and global direction perception. *Perception* 26: 995–1010, 1997.
- BRUCE CJ AND GOLDBERG ME. Primate frontal eye fields. I. Single neurons discharging before saccades. *J Neurophysiol* 53: 603–635, 1985.
- COWAN JD. Statistical mechanics of nervous net. In: *Neural in Ravello*, edited by Caianiello R. Berlin: Springer, 1972, p. 181–188.
- DASSONVILLE P, SCHLAG J, AND SCHLAG-REY M. The use of egocentric and exocentric location cues in saccadic programming. *Vision Res* 35: 2191–2199, 1995.
- DEHAENE S AND CHANGEUX J. A simple model of prefrontal cortex function in delayed-response tasks. *J Cognit Neurosci* 1: 244–261, 1989.
- DOMINEY PF AND ARBIB MA. A cortical-subcortical model for the generation of spatially accurate sequential saccades. *Cereb Cortex* 2: 153–175, 1992.
- DROULEZ J AND BERTHOZ A. A neural model of sensoritopic maps with predictive short-term memory properties. *Proc Natl Acad Sci USA* 88: 9653–9657, 1991.
- DUHAMEL J-R, COLBY C, AND GOLDBERG ME. The updating of the representation of the visual space in parietal cortex by intended eye movement. *Science* 255: 90–92, 1992.
- FERRERA VP AND LISBERGER SG. Neuronal responses in visual areas MT and MST during smooth pursuit target selection. *J Neurophysiol* 78: 1433–1446, 1997.
- FUNAHASHI S, BRUCE CJ, AND GOLDMAN-RAKIC PS. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J Neurophysiol* 61: 331–349, 1989.
- FUSTER JM. Not the module does memory make—but the network. *Behav Brain Sci* 18: 631–633, 1995.
- GOLDBERG ME AND BRUCE CJ. Primate frontal eye fields. III. Maintenance of a spatially accurate saccade signal. *J Neurophysiol* 64: 489–508, 1990.
- GOODMAN SJ AND ANDERSEN RA. Algorithm programmed by neural network model for coordinate transformations. In: *Proceedings of the International Joint Conference on Neural Networks, San Diego*. Ann Arbor, MI: IEEE Neural Network Council, 1990, vol. 2, p. 381–386.
- GNADT JW AND ANDERSEN RA. Memory related motor planning activity in posterior parietal cortex of macaque. *Exp Brain Res* 70: 216–220, 1988.
- GOLDMAN-RAKIC PS. Cellular basis of working memory. *Neuron* 14: 477–485, 1995.
- GROSSBERG S AND LEVINE D. Some developmental and attentional biases in the contrast enhancement and short term memory of recurrent neural networks. *J Theor Biol* 53: 341–80, 1975.
- KALASKA JF AND CRAMMOND DJ. Deciding not to go—neuronal correlates of response selection in a go/nogo task in primate premotor and parietal cortex. *Cereb Cortex* 5: 410–428, 1995.
- KROMMENHOEK KP, VAN OPSTAL AJ, GIELEN CCAM, AND VAN GISBERGEN JAM. Remapping of neural activity in the motor colliculus: a neural network study. *Vision Res* 33: 1287–1298, 1993.
- KROMMENHOEK KP, VAN OPSTAL AJ AND VAN GISBERGEN JAM. An analysis of craniocentric and oculocentric coding stages in a neural-network model of the saccadic system. *J Neural Networks* 9: 1497–1511, 1996.
- LEE DK, ITTI L, KOCH C, AND BRAUN J. Attention activates winner-take-all competition among visual filters. *Nat Neurosci* 2: 375–381, 1999.
- LI CS, MAZZONI P, AND ANDERSEN RA. Reversible inactivation of area LIP disrupts saccadic eye movements. *Soc Neurosci Abstr* 21: 281.1, 1995.
- LYNCH JC, MOUNTCASTLE VB, TALBOT WH, AND YIN TC. Parietal lobe mechanisms for directed visual attention. *J Neurophysiol* 40: 362–389, 1977.
- MAZZONI P, BRACEWELL RM, BARASH S, AND ANDERSEN RA. Spatially tuned auditory responses in area LIP of macaques performing delayed memory saccades to acoustic targets. *J Neurophysiol* 75: 1233–1241, 1996a.
- MAZZONI P, BRACEWELL RM, BARASH S, AND ANDERSEN RA. Motor intention activity in the macaque's lateral intraparietal area. I. Dissociation of motor plan from sensory memory. *J Neurophysiol* 76: 1439–1456, 1996b.
- MUNRO EE, SHUPE LE, AND FETZ EE. Integration and differentiation in dynamic recurrent neural networks. *J Neural Comput* 6: 405–419, 1994.
- QUAIA C, OPTICAN LM, AND GOLDBERG ME. The maintenance of spatial accuracy by the perisaccadic remapping of visual receptive fields. *Neural Networks* 11: 1229–1240, 1998.

- QUINTANA J AND FUSTER JM. Mnemonic and predictive functions of cortical-neurons in a memory task. *Neuroreport* 3: 721–724, 1992.
- ROBINSON DA. Oculomotor control signals. In: *Basic Mechanisms of Ocular Motility and Their Clinical Implications*, edited by Lennerstrand G and Bachy-Rita P. Oxford: Pergamon, 1975, p. 337–374.
- SALINAS E AND ABBOTT LF. A model of multiplicative neural responses in parietal cortex. *Proc Natl Acad Sci USA* 93: 11956–11961, 1996.
- SALZMAN CD AND NEWSOME WT. Neural mechanisms for forming a perceptual decision. *Science* 264: 231–237, 1994.
- SCHLAG J, DASSONVILLE P, AND SCHLAG-REY M. Interaction of the two frontal eye fields before saccade onset. *J Neurophysiol* 79: 64–72, 1998.
- SCHLAG J AND SCHLAG-REY M. Illusory localization of stimuli flashed in the dark before saccades. *Vision Res* 35: 2347–2357, 1995.
- SCUDDER CA. A new local feedback model of the saccadic burst generator. *J Neurophysiol* 59: 1455–1475, 1988.
- SEUNG HS. How the brain keeps the eyes still. *Proc Natl Acad Sci USA* 93: 13339–13344, 1996.
- SNYDER LH, BATISTA AP, AND ANDERSEN RA. Coding of intention in the posterior parietal cortex. *Nature* 386: 167–170, 1997.
- SPARKS DL AND MAYS LE. Spatial localization of saccade targets. I. Compensation for stimulation-induced perturbations in eye position. *J Neurophysiol* 49: 45–63, 1983.
- WERBOS PJ. Backpropagation through time: what it does and how to do it. *Proc IEEE* 78: 1550–1560, 1990.
- WILLIAMS RJ AND ZIPSER D. Gradient-based learning algorithms for recurrent networks. In: *Back-Propagation: Theory, Architectures, and Applications*, edited by Chauvin Y and Rumelhart DE. Hillsdale, NJ: Erlbaum, 1995, p. 433–486.
- XING J, LI C-S, AND ANDERSEN RA. The temporal-spatial properties of LIP neurons for sequential eye movements simulated in a neural network model. *Soc Neurosci Abstr* 21: 120, 1995.
- XING J, STRICANNE B, AND ANDERSEN RA. A neural network model for sensorimotor transformation in macaque area LIP. *Soc Neurosci Abstr* 20: 143, 1994.
- ZHANG K. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *J Neurosci* 16: 2112–2126, 1995.
- ZIPSER D. Recurrent network model of the neural mechanism of short-term active memory. *Neural Comput* 3: 179–193, 1991.
- ZIPSER D AND ANDERSEN RA. A backpropagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331: 679–684, 1988.
- ZIPSER D, KEHOE B, LITTLEWORT G, AND FUSTER J. A spiking network model of short-term active memory. *J Neurosci* 13: 3406–3420, 1993.