# Representation of an Abstract Perceptual Decision in Macaque Superior Colliculus

# Gregory D. Horwitz, Aaron P. Batista, and William T. Newsome

Howard Hughes Medical Institute and Department of Neurobiology, Stanford University School of Medicine, Stanford, California 94305

Submitted 8 September 2003; accepted in final form 25 December 2003

Horwitz, Gregory D., Aaron P. Batista, and William T. Newsome. Representation of an abstract perceptual decision in macaque superior colliculus. J Neurophysiol 91: 2281-2296, 2004. First published January 7, 2004; 10.1152/jn.00872.2003. We recorded from neurons in the intermediate and deep layers of the superior colliculus (SC) while monkeys performed a novel direction discrimination task. In contrast to the task we used previously, the new version required the monkey to dissociate perceptual judgments from preparation to execute specific operant saccades. The monkey discriminated between 2 opposed directions of motion in a random-dot motion stimulus and was required to maintain the decision in memory throughout a delay period before the target of the required operant saccade was revealed. We hypothesized that perceptual decisions made in this paradigm would be represented in an "abstract" or "categorical" form within the brain, probably in the frontal cortex, and that decision-related neural activity would be eliminated from spatially organized preoculomotor structures such as the SC. To our surprise, however, a small population of neurons in the intermediate and deep layers of the SC fired in a choice-specific manner early in the trial well before the monkey could plan the operant saccade. Furthermore, the representation of the decision during the delay period appeared to be spatial: the active region in the SC map corresponded to the region of space toward which the perceptually discriminated stimulus motion flowed. Electrical microstimulation experiments suggested that these decisionrelated SC signals were not merely related to covert saccade planning. We conclude that monkeys may employ, in part, a spatially referenced mnemonic strategy for representing perceptual decisions, even when an abstract, categorical representation might appear more likely a priori.

## INTRODUCTION

Decision making is a key cognitive process that links perception to action. Four decades of research into the neural foundations of basic sensory and motor processes has provided neurophysiologists with an opportunity to investigate empirically the decision mechanisms that link perception and action (Glimcher 2001; Gold and Shadlen 2001; Romo and Salinas 2001; Schall 2001). Neurophysiological investigations of decision mechanisms are typically conducted in behaving animals that are trained to perform sensory-discrimination tasks. Because the animal necessarily reports the outcome of the decision process with an operant movement of some sort, the neural signals underlying decision making can be closely related to, and may in some cases be identical with, "premotor" signals associated with preparation for a specific movement. In our own previous work, for example, rhesus monkeys were trained to discriminate between opposite directions of motion

Using quantitative analytic techniques based on this intuition, we demonstrated the existence of decision-related signals in the superior colliculus (SC), a midbrain structure closely associated with the generation of eye and head movements. Using the same direction-discrimination task and similar analytic techniques, Shadlen and colleagues documented nearly identical decision-related signals in area LIP (lateral intraparietal) of the parietal lobe and area 46 of the frontal lobe, both of which have been implicated in oculomotor behavior (Kim and Shadlen 1998; Roitman and Shadlen 1998; Shadlen and Newsome 2001). A central question that has arisen from these studies is whether these neural structures contribute generally to the process of making perceptual decisions, or whether these structures are involved only because the logic of our "standard" direction-discrimination task enforces a stereotyped link between the direction of stimulus motion and the vector of the operant saccade.

To address this question, we trained 2 rhesus monkeys on a new version of our direction discrimination task—the "loose stimulus–response association task" (LSRA)—in which we broke the stereotyped link between motion direction and the vector of the operant saccade. The animal performed the same perceptual discrimination as in our previous studies and continued to report its perceptual decision with a saccadic eye movement. In the LSRA task, however,  $\leq 28$  different operant saccades could be used to report a particular direction of motion, and the metrics of the required saccade were not specified until the end of the trial, well after the perceptual

in a stochastic random-dot display and report the perceived direction with a saccadic eye movement to one of 2 targets aligned with the axis of motion (Horwitz and Newsome 1999, 2001a,b). Because the saccade targets were visible throughout the trial and bore a stereotyped geometric relationship to the axis of stimulus motion, the accumulation of sensory information favoring one or the other direction of motion could activate, nearly simultaneously, the premotor circuits that prepare one or the other eye movement (Gold and Shadlen 2000, 2003). Thus an abstract representation of the decision is not necessary to account for the behavior. To distinguish decision-related neural signals from purely motor signals in such paradigms, we and others have relied heavily on the intuition that decisionrelated activity should reflect to some extent the strength of the sensory evidence that led to the decision (i.e., the "certainty" of the decision), whereas strictly motor signals should be related more directly to the metrics of the operant movement irrespective of the evidence that led to the movement.

Address for reprint requests and other correspondence: W. T. Newsome, Department of Neurobiology, Fairchild Building, Room D209, Stanford University School of Medicine, Stanford, CA 94305 (E-mail: bill@monkeybiz.stanford.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.

decision had been made. In this task we fully expected decision-related signals to be absent from the SC, probably being displaced to frontal lobe circuitry associated with "rule-based" performance (Miller et al. 2003). To our surprise, however, we found that a small population of neurons in the intermediate and deep layers of the SC exhibited modulations of activity that predicted the decision the monkey would make at the end of the trial. In each experiment this "predictive activity" was relegated to a specific region of the SC map of space—the spatial location toward which stimulus motion flowed. These observations, and the results of microstimulation experiments that we also report, suggest that the monkey remembers the direction of stimulus motion during each trial, at least in part, by referring the motion to the spatial framework of the oculomotor system.

A specific spatial framework that could be used for this purpose is an "object-centered" coordinate system in which a virtual object is considered to form a spatial link between the 2 saccade targets. The monkey would remember the direction of stimulus motion by referring it to the right or left side (for example) of the virtual object. Such object-centered representations have been demonstrated in the supplementary eye field (SEF) of the frontal cortex (Olson and Gettner 1999, 1995; Olson and Tremblay 2000; Tremblay et al. 2002), and our data are consistent with, but do not prove, the notion that objectcentered coordinates are present in some SC neurons as well. To further investigate this possibility, we recorded from SEF neurons while one monkey performed the LSRA task. Confirming prior observations of Olsen and colleagues, SEF neurons exhibited predictive activity in a manner consistent with the existence of object-centered responses. Interestingly, predictive activity arose sooner in SEF, on average, than in the SC, suggesting that SEF may occupy an intermediate position in the processing chain that links sensory activity in occipital visual areas to the decision-related signals we have observed in the SC.

## METHODS

## Surgical procedures

Two rhesus monkeys (*Macaca mulatta*) served as subjects in these experiments. Before data collection, each monkey underwent a pair of surgical procedures that were performed under aseptic conditions and general anesthesia. In an initial surgery, the animal was implanted with a head-restraint device and scleral search coil (Judge et al. 1980). After several months of behavioral training, a second surgery was performed to implant a stainless steel recording cylinder that permitted microelectrode access to the SC. One monkey was implanted with an additional recording cylinder over the SEFs. All experimental procedures conformed to the standards established by the National Institutes of Health and were approved in advance by the Institutional Animal Care and Use Committee of Stanford University.

## Visual stimuli

Monkeys discriminated the direction of coherent motion in a stochastic random-dot stimulus that has been used extensively in this laboratory (Britten et al. 1992, 1996; Salzman et al. 1992). Stimulus movies lasted for 1 s and were presented on a CRT monitor with a 60-Hz refresh rate. Each frame of each stimulus movie consisted of a field of white dots presented within a 7° diameter aperture. Motion was created by replotting, after a delay of 50 ms, a randomly selected 51.2% of the dots (called "signal" dots) at a displacement of  $0.15^\circ$  in a single direction. Signal dots therefore moved coherently in a specified direction at a speed of 3°/s. The remaining dots (the "noise" dots) were replotted in random locations and thus served to mask the motion direction of the signal dots. Monkeys were trained to discriminate between 2 opposed directions of coherent motion and to report their judgments by making a saccade to one of 2 visual targets (see *Behavioral paradigm 1: loose stimulus–response association task* below). The density of dots in the visual stimuli was 15 dots/ (deg<sup>2</sup> · s<sup>-1</sup>), but the apparent density of dots in the stimulus was much higher because of persistence in the visual system.

#### Electrophysiology

Single-unit recordings were accomplished by means of tungsten microelectrodes (Microprobe or Frederick Haer) and standard electrophysiological amplifiers (Bak Electronics). Action potentials were detected using a time–amplitude window discriminator (Bak Electronics), time stamped (1-ms resolution), and stored on a magnetic disk for subsequent analysis.

Electrical microstimulation was accomplished using a pulse generator (Bak Electronics or Master 8 CP) in series with an optical stimulus isolation device (Frederick Haer). Stimulation pulses were biphasic, cathodal pulse leading. Each pulse was 0.2 ms in duration with a 0.1-ms pause between pulses. Stimulation currents ranged from 10 to 50  $\mu$ A, frequencies ranged from 200 to 500 Hz, and train durations ranged from 70 to 400 ms. A saccade was classified as stimulation evoked if it was initiated within 50 ms of the onset of the pulse train.

# Behavioral paradigm 1: loose stimulus-response association task

Figure 1 illustrates the LSRA task. This task is very similar to the direction-discrimination task used previously by Newsome and colleagues (Horwitz and Newsome 1999; Shadlen et al. 1996) with one critical difference: in the LSRA task the monkey does not know the metrics of, and therefore cannot plan, the operant saccade until late in the trial. Decision-related activity during the visual stimulus and delay periods is therefore unlikely to result from planning of a specific operant saccade.

Each trial began with onset of a fixation point (Fig. 1A, cross; Fig. 1B, top trace). Three hundred milliseconds after the monkey achieved fixation, a stochastic motion stimulus appeared for 1 s at the center of gaze flowing in one of 2 opposed directions. The visual stimulus was followed by a delay period that was a constant 1-s duration in some experiments and was randomized from 1 to 2 s in others. At the conclusion of the delay period, a pair of saccade targets appeared at one of several possible locations. The 2 targets were always separated by 5° of visual angle and were always positioned along an axis parallel to the direction of coherent motion in the visual stimulus. A second delay period followed the appearance of the saccade targets, the duration of which was 1 s in some experiments and was randomized between 1 and 1.5 s in others. Disappearance of the fixation point cued the monkey to make a saccade to the target that lay in the direction of motion relative to the location of the other target. In Fig. 1A, for example, coherent motion was leftward, and the correct target would be the leftmost member of the pair regardless of its position on the retina. Thus leftward motion could be reported with a saccade to the left or to the right, and with an upward or a downward component.

The monkey received a liquid reward for initiating a saccade within 500 ms of fixation point offset that terminated in a  $5 \times 5^{\circ}$  electronic window surrounding the correct target. Saccades that terminated in a  $5 \times 5^{\circ}$  window surrounding the distractor target were counted as incorrect choices. Fixation breaks, failures to initiate the saccade within 500 ms of fixation point offset, or saccades to locations outside either target window were aborted immediately without reward.

Target locations were defined on a grid as shown in Fig. 1C.



FIG. 1. Methods for the loose stimulus–response association (LSRA) task. Example displays (*A*) and event timings (*B*) are shown. Once the monkey achieved visual fixation, the motion stimulus was displayed at the center of gaze for 1 s. During a subsequent delay period of 1- to 2-s duration, the monkey was required to remember the direction of motion, but did not yet know the metrics of the saccade that would be required to report it. A pair of saccade targets then appeared in one of several locations. After an additional 1-1.5 s, the fixation light was extinguished, and the monkey was rewarded for making a saccade to the target that, relative to the other, lay in the direction of motion. *C*: display geometry for the LSRA task. Motion flowed toward and away from the neuron's movement field (dotted line) on randomly interleaved trials. Perceptual judgments of motion direction were reported by saccades to an array of potential target positions (disks).

Spacing between target locations was fixed at  $5^{\circ}$  along the axis of motion and varied, between experimental sessions, from 4 to  $8^{\circ}$  in the perpendicular dimension. A target pair was defined as 2 target locations that were adjacent along the dimension parallel to the motion axis. Each target was thus a member of 2 target pairs, one in which it lay to one side of its mate and one in which it lay to the other side. In the context of the task this means that a saccade to a given target location could report motion in either direction, depending on where the other target appeared.

To prevent inaccurate saccades from being misclassified as incorrect choices we constrained target positions in 2 ways. First, we never used target locations  $>22^\circ$  eccentric and rarely  $>15^\circ$  eccentric. Second, we never used target pairs for which the 2 operant saccades were very similar in direction to avoid choice biases resulting from the path of the saccade to one target from passing directly over the other target. In practice, this exclusion meant that no target pairs were used whose midpoints were within 45° of the motion axis, accounting for the gap between the 2 clusters of targets in Fig. 1*C*. Saccades to targets satisfying both of these constraints were extremely accurate; only 5% of the saccades that fell within one of the 2 target windows landed closer to the midpoint of the target pair than either individual target. The number of allowable target locations varied from 12 to 40 across experimental sessions.

For each neuron, we adjusted the display geometry so that coherent motion on each trial flowed toward or away from the movement field (dashed curve). Thus in the example in Fig. 1, the monkey discriminated rightward from leftward motion (arrows) because the movement field lay directly to the right of the fixation point. As indicated above, each member of a target pair could represent motion flowing toward or away from the cell's movement field. We will refer to the target that corresponded to motion *toward* the movement field as "T1" and the other target as "T2" regardless of their positions on the computer monitor.

In the design described thus far, a small but systematic association remains between the direction of stimulus motion and the metrics of the operant saccade. Each target location in Fig. 1C, for example, is the rightmost member of one pair and the leftmost member of another pair—with the exception of targets at the right and left edges of the array. A saccade to one of these edge targets always indicates that motion flowed in the direction of the target. Across an entire experiment, therefore, the collection of saccades that the monkey makes to report rightward motion would have a small rightward component on average. Conversely, saccades that indicate leftward motion would have a net leftward component. Thus delay period activity in the SC might occur for a trivial reason: if the monkey "knows" that the required saccade is likely to have a component in the direction of motion, it may prepare to make a saccade in that direction on viewing the motion stimulus. To eliminate this cue to the metrics of the upcoming operant saccade, we presented "edge" targets only when they were distractors. Thus the distribution of correct saccade vectors was identical for both directions of motion, and the monkey had no basis for anticipating the metrics of the correct operant saccade.

It is important to note that on most trials *neither* saccade target fell within the neuron's movement field (e.g., only 2 of 20 target locations in Fig. 1*C*). This procedure appears odd at first glance because SC neurons are commonly thought to respond only for saccades to targets within the movement field. We chose this procedure so that the results of the experiments could be compared directly to those of our previous studies in which coherent stimulus motion always flowed toward or away from the movement field (Horwitz and Newsome 1999, 2001a,b). Thus the geometrical relationship of the stimulus aperture, coherent motion flow, and movement field location were identical in all studies.

# Training history

In designing the LSRA task, we sought to create a behavioral context in which the monkeys made directional judgments without translating those judgments immediately into preparation for a particular saccade. Ideally, the delay period before onset of the saccade targets should provide a "window" for observing the representation of the decision in a relatively abstract form, divorced from a specific operant response. For this reason, naïve monkeys were selected for training on this task; neither animal had been employed previously in studies that used the standard version of our direction-discrimination task. In the first stage of LSRA training, however, both monkeys were briefly trained to perform the standard task along a single axis of motion (i.e., to make saccades in the direction of visual motion). As soon as each animal performed this discrimination reliably ( $\sim$ 70% correct; 1 mo for monkey A, 8 sessions for monkey T), the 2 targets were displaced away from the axis of stimulus motion during an extended period of "generalization" training. For monkey A, training continued for a single new target pair until performance exceeded approximately 80% correct choices, whereupon a new target pair was selected. Monkey T was trained on sets of 4 to 6 target pairs in short blocks until performance was reliable for all positions. Only after reliable performance was attained at a variety of target positions were multiple target pairs interleaved within a single block of trials. The number of target pairs increased progressively and the intertarget distances decreased over the course of training for both monkeys. For monkey A, the axis of stimulus motion was rotated relatively early during generalization training, while new target pairs were being introduced. For monkey T, the axis of motion was changed only during the last month of training after the monkey had fully general2284

ized to all target pairs for one axis of motion. In all, training took 14 mo for monkey A and 7 mo for monkey T.

### Behavioral paradigm 2: memory-guided saccade task

We generally screened cells on the basis of their discharge during a memory-guided saccade task. In this task, a single eccentric saccade target was displayed for 200 ms while the monkey fixated a central point. After a random-length delay period of 1–1.5 s, the fixation point was extinguished and the monkey was required to initiate a saccade to the remembered location of the target within 500 ms of fixation point offset. The monkey received a liquid reward for saccades that terminated in an electronic window, usually  $10 \times 10^{\circ}$ , surrounding the target location. If the monkey failed to initiate the saccade within the stipulated time interval, or if the saccade terminated outside of the window, the trial was aborted. At least 2 target locations, one inside and one outside of the movement field, were pseudorandomly interleaved from trial to trial.

# Data analysis

During single-unit recording, monkey T responded correctly to 88% of 7,072 LSRA trials, and monkey A responded correctly to 88% of 6,427 LSRA trials. Performance was consistent across sessions for both animals (the SD of percent correct was 6% for monkey A and 7% for monkey T). Unless specified otherwise in the text, we analyzed data from correctly performed trials only.

#### Predictive activity

We classified individual neurons as "choice-predictive" if the firing rate during the stimulus presentation and initial delay period was significantly higher preceding reports of one direction of motion than the other (Mann–Whitney U tests:  $P \leq 0.01$ ). The evolution of predictive activity across a population of neurons was assessed using a technique based in signal detection theory (Green and Swets 1966). For each choice-predictive neuron, we calculated responses preceding T1 and T2 choices in 100-ms bins. We normalized these responses by dividing each by the maximal response across all time bins and then pooled the normalized responses across neurons. For each bin, we calculated a receiver operating characteristic (ROC) curve from the distributions of normalized responses. The area beneath this curve was defined as the predictive activity value for that time bin. Related methods have been used to measure predictive activity in previous studies (Kim and Shadlen 1998; Krauzlis and Dill 2002; Shadlen and Newsome 2001; Thompson et al. 1996).

## Cell screening

While advancing our electrode in search of a neuron, we typically had the monkey perform the memory-guided saccade task, or occasionally, the LSRA task. Each time we isolated a neuron from the intermediate or deep layers of the SC, we qualitatively assessed responses in both tasks. If the cell exhibited either maintained delayperiod activity in the remembered saccade task or choice-predictive activity in the LSRA task, we recorded data for subsequent off-line analysis. We selected cells on the basis of the remembered saccade task for 2 reasons. First, at the outset of data collection, we did not know (or even expect) that we would find any choice-predictive neurons in the SC. It did not seem reasonable to select cells on the basis of a response we did not think existed. Second, we wanted to study roughly the same population of SC neurons that we had analyzed in our previous studies (Horwitz and Newsome 1999, 2001a,b). Our selection criterion in the previous studies correlated well with delay-period activity in the remembered saccade task. Saccade targets in the current study typically fell outside the SC movement field (Fig. 1C), so many cells that we encountered responded rarely or not at all during LSRA task performance. We quickly abandoned cells that failed to generate any responses during the LSRA task. Experiments were aborted if either electrical isolation or behavioral control was lost.

## RESULTS

## Predictive activity

We recorded from 36 neurons from monkey A and 41 neurons from monkey T during performance of the LSRA task, selected according to the criteria described in METHODS (*Cell screening*). Firing rates were analyzed quantitatively off-line. Of the total sample, 17 neurons from monkey A and 18 neurons from monkey T were significantly more active preceding T1 choices (defined in METHODS) than T2 choices during the stimulus-presentation period and the initial delay period (Mann–Whitney U test:  $P \leq 0.01$ ). Recall that the saccade targets had not yet appeared during these epochs, and that most targets would ultimately fall outside the movement field. Thus the difference in firing rate preceding T1 and T2 saccades appears to reflect the monkey's perceptual decision per se, not preparation to report the decision with a particular saccade.

Two cells that were significantly more active preceding T2 choices than T1 choices were excluded from analysis. One cell whose response was clearly time-locked to small-amplitude saccades confined to the fixation window was likewise excluded.

The proportion of choice-predictive neurons in our data set (35/77) should not be taken as an estimate of the proportion of choice-predictive neurons in the SC as a whole. We deliberately restricted our recordings to a select population of SC neurons (see METHODS), and choice-predictive neurons are certainly less frequent in the SC as a whole than in our biased sample. Although we cannot provide a precise estimate, we believe that choice-predictive neurons in this task constitute roughly 10% of the neurons we isolated in the intermediate and deep layers of the SC.

Figure 2 illustrates the responses of a representative choicepredictive neuron on correctly answered trials. On trials in which the monkey chose T1, the firing rate rose steadily during the stimulus presentation period and the subsequent delay period. In contrast, the firing rate remained low throughout the trials on which the monkey chose T2. Thus the discharge of this neuron during the stimulus-presentation period and the delay period predicted the monkey's choice at the end of the trial despite complete uncertainty regarding the metrics of the correct operant saccade, which was randomized across 12 locations in both the contralateral and ipsilateral hemifields in this experiment (see following text).

Surprisingly, the neuron's discharge was remarkably consistent across trials after onset of the saccade targets despite the substantial variation in target location. On trials in which the monkey chose T1, the firing rate increased transiently after onset of the saccade targets, then decreased gradually until the time of fixation point offset (the "go" signal). In contrast, the firing rate remained uniformly low during the same period for trials ending in T2 choices. Recall that the distribution of correct target locations was identical for trials in which motion flowed toward or away from the movement field. If the monkey had made no errors, each trial in Fig. 2A would have an exact counterpart in Fig. 2B

# DECISION-RELATED ACTIVITY IN SUPERIOR COLLICULUS



FIG. 2. Discharge of a single predictive superior colliculus (SC) neuron in the LSRA task. All trials were aligned to stimulus onset (*left*) and target onset (*right*). Vertical lines correspond to stimulus onset, stimulus offset, target onset, and fixation point offset (from *left to right*). Break between *left* and *right panels* reflects random epoch durations. This cell discharged more vigorously preceding T1 choices (*A*) than T2 choices (*B*) throughout nearly the entire trial, although the required saccade could not be planned until the target presentation.

with respect to required metrics of the operant saccade (91% of the trials depicted in Fig. 2 were in fact members of such a matched pair). Despite the metrical similarity of T1- and T2-directed saccades, the firing rate of this neuron differed substantially on T1 and T2 choice trials, even well after the saccade targets appeared. (Responses at the time of the saccade will be considered in detail below.)

Figure 3 shows the responses of the same neuron, segregated by perceptual choice and by target location, during 3 epochs of the trial (correct choices only). *Left* and *right panels* illustrate responses preceding T1 and T2 choices, respectively. The *top panels* depict the responses from 50 to 125 ms after the visual stimulus was extinguished. The firing rate of the cell was higher preceding T1 choices than T2 choices during this epoch (as noted in Fig. 2), and this relationship, of course, was independent of the (future) location of the targets.

The *middle panels* show the responses from 50 to 125 ms after the saccade targets appeared. By this point in the trial, firing rates preceding T1 choices have increased dramatically, whereas firing rates preceding T2 choices have remained relatively low. The firing rate of the neuron is still largely unaffected by the target location, suggesting that the firing rate is independent of any spatially localized visual or premovement response to the saccade targets.

The *bottom panels* illustrate the responses during a 75-ms epoch starting 400 ms before saccade initiation. Three aspects of the responses during this epoch are noteworthy. First, the response is considerably lower at this time than in the previous epoch, as indicated in Fig. 2. This is true preceding saccades to either T1 or T2 but is clearest preceding saccades to T1. Second, the response is only weakly related to the metrics of the impending saccade. Third, the firing rate is still significantly higher preceding T1 choices than T2 choices. Thus the signal carried by the cell at this relatively late point in the trial continues to reflect the perceptual decision made by the monkey and is inappropriate for specifying the metrics of the operant saccade.

# Predictive activity: population analysis

To examine decision-related signals quantitatively, we normalized the responses for each choice-predictive neuron, pooled the data across the population, and performed the signal-detection analysis described in METHODS. We analyzed data from correct and error trials separately. Figure 4, A and B depicts the results of this analysis separately for the 2 monkeys.



FIG. 3. Response as a function of saccade endpoint and epoch. Firing rate preceding T1 choices (*left column*) and T2 choices (*right column*) is shown as a function of saccade endpoint. Firing rate was calculated during 3 epochs: a "stimulus OFF" period (50 to 125 ms after stimulus offset; *top*), a "targets ON" period (50 to 125 ms after targets onset; *middle*), and a "presaccade" period (400 to 325 ms preceding the saccade; *bottom*).

The development of predictive activity within the population of SC neurons followed closely the firing rate dynamics shown for the single example cell in Fig. 2. Predictive activity increased during the stimulus-presentation period and the delay period, peaked around the time of target presentation, and declined subsequently. The *rightmost panels*, aligned on the saccades, show that predictive activity was maintained through the time of saccade initiation. Comparison of *A* and *B* in Fig. 4 reveals that the time course of predictive activity was quite similar for the 2 monkeys, although the magnitude differed, reaching higher levels on average in monkey A.

SC neurons accurately predicted erroneous as well as correct choices (*gray traces*, Fig. 4). This observation has 2 important implications. First, the responses of choice-predictive neurons were not strictly a function of the visual stimulus; responses accurately predicted the monkeys' choices irrespective of the direction of motion in the visual stimulus. Second, although



monkey T (B). Trials were aligned to stimulus onset (*left*), targets onset (*middle*), and saccade initiation (*right*). Solid lines indicate significant predictive activity (randomization test:  $P \le 0.05$ ) and dashed lines indicate nonsignificant predictive activity (P > 0.05). On both correct and error trials, the firing rate of choice-predictive SC neurons reflects the monkey's choice before and after the presentation of the targets.

predictive activity on error trials lagged that on correct trials, the genesis of erroneous choices occurred relatively early in the trial. Predictive activity on error trials reached statistical significance during the delay period in both monkeys, several hundred milliseconds before onset of the saccade targets. Note that the visual motion stimulus used in these experiments was 51.2% coherent, well above the psychophysical threshold, and the aperture was centered on the fovea, an ideal location for performing the discrimination. It thus seems unlikely that errors resulted from faulty perceptual processes. Rather, attentional and memory-related processes during the delay period seem a more likely source of erroneous performance.

Almost all (35 of 37) choice-predictive neurons preferred T1 choices to T2 choices. This demonstrates a close relationship between the choice preference of individual neurons and their movement field locations: choice-predictive SC neurons discharged more vigorously during trials in which the monkey reported visual motion flowing toward the movement field. We shall return to this point in the DISCUSSION.

# Saccade-related discharge

The single-neuron data in Figs. 2 and 3, and the population data in Fig. 4, suggest that choice-predictive neurons in the SC carry signals related to the monkey's perceptual decision up to and including the time of the saccade, irrespective of target location. The question naturally arises whether these neurons *ever* exhibit the spatially localized movements fields that are typical of SC neurons.

Figure 5 illustrates the responses of the example cell of Fig. 2 during a perisaccadic epoch, defined as 50 ms preceding saccade initiation until 25 ms afterward. Responses are segregated by perceptual decision and target location as in Fig. 3. To determine whether these 2 factors significantly influenced perisaccadic firing rate, we performed a randomization test analogous to an ANOVA (Edgington 1995). Unlike a conventional ANOVA, this randomization test does not assume normally distributed firing rates. The data were permuted 10,000 times for each analysis.

Visual inspection of Fig. 5 shows that perisaccadic activity was higher for saccades with upward and leftward components than for saccades with downward and rightward components. This tuning was roughly consistent with the location of the neuron in the collicular map: the mean direction of saccades evoked by electrical stimulation at the recording site was 174°  $(6^{\circ}$  clockwise from straight left) and the mean amplitude was 7.4°. A randomization ANOVA confirmed that the effect of target location was statistically significant (P < 0.01). As is evident in Fig. 5, however, perisaccadic activity of this neuron was also modulated by the monkey's perceptual decision (P <0.01). Thus at the time of the saccade this cell carried signals related to the metrics of the operant saccade, as would be expected from a typical SC neuron, and to the perceptual decision as well. Comparison of Fig. 5, A and B indicates that perisaccadic activity, like predictive activity earlier in the trial, was higher when motion flowed toward the movement field.

Of the 17 choice-predictive neurons in monkey A, 14 exhibited significant spatial tuning at the time of the saccade and 13 were modulated by the perceptual decision during the same interval (randomization ANOVA, P < 0.01). Of the 18 choice predictive neurons in monkey T, 15 were spatially tuned and 3



FIG. 5. Perisaccadic discharge as a function of saccade endpoint and motion direction. Spikes were counted from 50 ms before saccade initiation until 25 ms afterward. Correct reports of up-left motion (T1 choices) appear in A, and correct reports of down-right motion (T2 choices) appear in B. Perisaccadic discharge is related to both saccade endpoint and motion direction.

were modulated by the decision. Thus at the time of the saccade the large majority of choice-predictive neurons indeed exhibited conventional spatially tuned movement fields, and roughly half of these neurons carried signals related to the perceptual decision as well. Neurons without significant spatial tuning may have had small movement fields that did not overlap with any saccade target in the array.

## Saccade metrics

The effect of the perceptual decision on perisaccadic activity might be artifactual if the saccades to specific target locations varied systematically with the direction of stimulus motion that instructed the saccade. We therefore asked whether saccade endpoints varied with motion direction. Only experiments in which  $\geq 3$  saccades for each direction of motion were made to  $\geq 3$  targets were considered in this analysis. For each target, we calculated the signed distance between the mean endpoints of T1 and T2 saccades, projected onto the axis of motion (positive values indicate displacements in the direction of motion, and negative values indicate displacements in the opposite direction). Under the null hypothesis, the signed distance should be distributed symmetrically around zero. One-sample *t*-tests (P < 0.05) showed that this was not the case for 16/35 experiments from monkey A and 13/33 experiments from monkey T. The mean endpoint displacement across targets and experiments was  $0.15^{\circ}$  for monkey A and  $0.11^{\circ}$  for monkey T in the direction of motion. Saccade endpoints were thus weakly biased in the direction of motion. We obtained qualitatively similar results irrespective of whether we subtracted initial eye position.

This small bias in saccade metrics clearly cannot account for the effect of choice on perisaccadic activity in Fig. 5 ( $\sim$ 20 spikes/s for several target locations), but it could in principle account for some of the small effects we observed. The randomization ANOVA described in the preceding section was blind to this dependency because firing rates were regressed on *target location*, not on precise saccade metrics. We therefore reanalyzed the effect of choice on perisaccadic activity using a second randomization test, which is similar to the ANOVA, with the important distinction that it takes precise saccade metrics into account.

For each cell, we fit 2 separate regressions, one for T1 choices and one for T2 choices, modeling perisaccadic firing rate as a third-order polynomial of saccade endpoint coordinates. The fitted equation was

$$z = b_0 + b_1 x + b_2 y + b_3 x^2 + b_4 y^2 + b_5 x y + b_6 x^3 + b_7 y^3 + b_8 x y^2 + b_8 y x^2$$

where x and y are saccade endpoint coordinates and z is the perisaccadic firing rate. The higher-order terms in this regression model allow for curvature of the fitted movement field. Having obtained movement field fits for the 2 decision states considered independently, we then randomized T1 and T2 choices within each target location and refit the models 10,000 times. The sum of squared residuals (SSR) provides a measure of the quality of the fits. If perisaccadic firing rates are determined by single underlying movement field, SSRs should be similar irrespective of whether the data were randomized across choice. Instead, we found that SSRs from the unrandomized data were significantly smaller (P < 0.05) than SSRs from the randomized data for 14 of 16 cells. Thus we could reject the hypothesis of a single underlying movement field for all but 2 cells (one from monkey A and one from monkey T). We thus conclude that, although saccades were biased weakly in the direction of motion, these biases did not account for the relationship between motion direction and perisaccadic discharge that we observed in many cells.

Finally, we considered the possibility that small saccades made within the fixation window during the stimulus presentation and delay periods were biased by the motion direction. Such a bias could produce artifactual predictive activity. A circular analog of the 2-sample unpaired *t*-test, the Mardia– Watson–Wheeler test (Batschelet 1981), revealed significant differences in fixational saccades (P < 0.05) in only 2/36 experiments in monkey A and 3/41 experiments in monkey T. The directions of fixational saccades thus do not appear to be strongly affected by the motion direction.

# Electrical stimulation

We considered the possibility that predictive activity in the SC may be related to saccade planning even though it cannot,

by virtue of the task design, be related to the planning of the operant saccade. For example, the monkeys may have planned saccades during the delay period that were congruent with the direction of motion during the stimulus period. Such saccade plans might even have served as a mnemonic device for remembering the direction of stimulus motion. This explanation accounts for the general pattern of selectivity we observed (T1 choice preference) because T1 choices were consistently associated with motion *toward* the movement field.

We used an electrical-stimulation technique to determine whether the monkeys formed saccade plans before target presentation. To a first approximation, saccades elicited by SC stimulation depend only on the position of the stimulating electrode within the collicular map (Robinson 1972). Several groups have shown, however, that the exact metrics of an electrically elicited saccade can be modified when the monkey simultaneously prepares a saccade to a different region of space (Gold and Shadlen 2001; Kustov and Robinson 1996; Sparks and Mays 1983). For example, electrical stimulation of an SC site that normally generates leftward saccades will produce a saccade with an upward component (in addition to the standard leftward component) if the monkey is preparing an upward saccade at the time of electrical stimulation, and a downward component if the monkey is preparing a downward saccade. If our monkeys consistently plan to make saccades in the direction of stimulus motion in the LSRA task, we thus expect systematic deviations in stimulation-evoked saccade vectors to reflect this fact. On the other hand, if the signals we have reported are not related to saccade planning, we expect stimulation-evoked saccades to be minimally affected by the direction of stimulus motion.

Figure 6 illustrates the sequence of events during these experiments. To ensure that the monkey would continue to perform the perceptual discrimination task, only one-ninth of the trials were interrupted with electrical stimulation. On these trials, the fixation point disappeared at the moment that the targets would normally have appeared, and 80 ms later a train of electrical stimulating pulses into the SC evoked a saccade. In half of the stimulation trials visual motion flowed in one direction; in the other half motion flowed in the opposite direction. The axis of motion was selected to be roughly orthogonal to the direction of the stimulation-evoked saccades. Rewards were given 200 ms after the onset of electrical stimulation on all stimulation trials and for correct decisions, as usual, on nonstimulated trials.

Figure 7 illustrates the results of a single microstimulation experiment. Electrical stimulation at this site evoked saccades with a mean amplitude of  $6^{\circ}$ , leftward and slightly downward. In this experiment, the monkey discriminated between upward and downward directions of motion, and saccade target pairs were distributed accordingly. Target-directed saccades on non-stimulated trials (*pale trajectories* in *A*) were quite accurate, as evidenced by the consistent trajectories to the 8 individual target locations. The electrically evoked saccades were also consistent in direction and amplitude (*saturated traces*), whether the monkey saw upward or downward motion during the visual stimulus period (*red* and *green traces*, respectively).

Figure 7B shows the endpoint of each stimulation-evoked saccade, corrected for small differences in initial fixation position (all saccade endpoints henceforth are likewise corrected). The mean endpoint of saccades evoked after an up-



FIG. 6. A: methods for LSRA electrical stimulation experiment. Most trials did not involve electrical stimulation but rather were conventional LSRA trials illustrated in Fig. 1. However, on one-ninth of the trials, selected randomly, the display was extinguished when the targets would have been illuminated and, 80 ms later, a train of electrical stimulation pulses was delivered to the SC, evoking a saccade. *B*: event timings in nonstimulated (solid lines) and stimulated trials (dashed lines).

ward motion stimulus (red asterisk) was slightly above the mean saccade endpoint after downward motion (green asterisk). The distance between these 2 mean endpoints, or "centroids," was  $0.47^{\circ}$ , or 8% of the average saccade amplitude. The distance is not significantly greater than expected by chance (randomization test, exchanging saccade endpoints randomly between groups: P > 0.2). However, the *direction* of the shift is correctly predicted by the saccade-planning hypothesis: when motion flowed upward, the stimulation-evoked saccade had a larger upward component than it did when motion flowed downward.

We performed this analysis for 14 stimulation experiments from monkey A and 23 from monkey T in which we obtained at least 10 stimulation-evoked saccades for each direction of visual motion. For each experiment, we conducted a randomization test to determine whether the separation between saccade endpoint centroids was statistically significant. Only 4 experiments yielded significant effects (P < 0.05). This number of significant effects is, itself, not significantly greater than the false-positive rate expected by chance. Thus any effect of predictive activity on saccade-generating circuitry in the SC was too weak to be measured within individual experiments.

We then performed a population analysis to gain additional statistical power. In this analysis we asked whether the direction of the stimulation effect agreed with the hypothesized effect of predictive activity on saccade preparation (as in Fig. 7), irrespective of whether the effect was significant in individual experiments. For each experiment, we arbitrarily de-



FIG. 7. Data from a single LSRA stimulation experiment. A: saccade trajectories during interleaved voluntary target-directed saccades (*pale traces*) and saccades evoked by electrical stimulation (*saturated traces*). Targetdirected saccades reported either upward (red) or downward (green) motion. Stimulation-evoked saccades were elicited after upward (red) or downward (green) stimulus presentations. Square represents the region of space depicted in *B*. *B*: saccade endpoints (dots) and endpoint centroids (asterisks) of stimulation-evoked saccades. Distance between the centroids,  $0.47^\circ$ , is not statistically significant but the shift is in the predicted direction.

fined one of the 2 motion directions as the "reference direction," the other being the "nonreference direction." We then calculated a vector from the centroid of saccade endpoints obtained from nonreference trials to that obtained from the reference trials. We will refer to the direction of this vector as the "direction of the stimulation effect." In the example of Fig. 7, upward was the reference direction, and the direction of the stimulation effect (the vector from the green centroid to the red centroid in Fig. 7B) was very close to the reference direction. Figure 8 plots the angular difference between the reference direction and the direction of the stimulation effect for each experiment for monkey A (8A) and for monkey T (8B). Under the null hypothesis that no relationship exists between the 2 quantities, these differences should be uniformly distributed about the circle. Visual inspection shows, however, that they tend to cluster near 0°, as we would expect if saccades were deviated in the direction of visual motion. Treating each data point as a vector of unit length, we calculated an average vector for each monkey. The length of each average vector exceeded the upper 95% confidence bound (dashed circle) computed from a Rayleigh test (Batschelet 1981), verifying that the data points are not distributed uniformly. The population effect illustrated in Fig. 8 is consistent with the saccade-planning hypothesis, but the individual effects were very small and rarely significant, raising doubt about their role in a hypothesized saccade-planning process. We therefore compared these data to similar microstimulation effects obtained when the monkey performed a delayed-saccade task—a task that explicitly encourages saccade planning. In the delayed-saccade task, a single saccade target appeared 400 ms after the monkey achieved fixation and remained visible throughout a 2- to 3-s delay period. The saccade target appeared at one of 2 locations arranged symmetrically about the fixation point. On 8 out of every 9 trials, the monkey obtained a reward by making a saccade to the target within 500 ms of fixation point offset. On the remaining one-ninth of the trials, the targets were extinguished along with the fixation point, and after an 80-ms pause, electrical stimulation was applied to elicit a saccadic eye movement.

Importantly, the axis defined by a line linking the 2 saccade targets was identical to the axis of motion in a matched LSRA experiment, and the electrical stimulation parameters were identical in the matched experiments. This allowed us to compare, at individual SC sites, the amplitude of the saccade deviation in LSRA task relative to that in the delayed-saccade task. We ran delayed-saccade experiments matched to 16 of the LSRA experiments (5 from monkey A and 11 from monkey T).

To compare stimulation effect magnitudes, we sorted saccade endpoints by stimulus condition (motion direction in the LSRA task and target position in the delayed-saccade task), and calculated the centroid of each group of endpoints. Figure



FIG. 8. Angular differences between the reference direction of motion and the direction of the stimulation effect. *A* and *B*: data from monkeys A and T, respectively. Differences tend to cluster near  $0^{\circ}$  (black radius), indicating that saccade endpoints tend to be shifted, relative to each other, in the remembered direction of motion. Treating each data point as a vector of unit length permits the calculation of vector averages (arrows). Amplitude of each vector exceeds the upper 95% critical value (dashed circle) from a Rayleigh test of circular uniformity.

#### 2290

9A is a scatterplot of intercentroid distances (in normalized units; see APPENDIX). Every point lies above the identity line, indicating that stimulation effects were consistently larger in the delayed-saccade task than in the LSRA task. We thus conclude that choice-predictive activity in the SC during the LSRA task does not constitute a saccade "plan," at least not in the standard sense that applies in the delayed-saccade task. We will return to this issue in the DISCUSSION.

# *Relationship between stimulation effect magnitude and task performance*

Since the stimulation effects during LSRA task performance do not appear to reflect saccade plans, we must ask whether the effects are related to task performance at all. It is possible, for example, that the effects result solely from the presence of the visual stimulus at the point of fixation, perhaps related to very slow pursuit or optokinetic nystagmus responses elicited by the foveal motion stimulus. Figure 9*B* shows that the stimulation effects are in fact related to task performance. Here we plot the monkey's performance in each experiment (in percentage of



FIG. 9. A: stimulation effect sizes in matched LSRA and delayed-saccade blocks. Stimulation effects were quantified as the distance between saccade endpoint centroids in normalized distance units,  $d_{norm}$  (see APPENDIX). Effects were consistently larger in delayed-saccade blocks than in matched LSRA blocks. *B*: percentage of correct choices on nonstimulated LSRA trials as a function of intercentroid distance on stimulated trials. Correlation coefficient between these 2 variables, 0.49, is highly significant (P < 0.005), indicating that stimulation effects tend to be larger in sessions when the monkey is performing the task well.

correct choices) as a function of the size of the microstimulation effect. Across the database of microstimulation experiments, the monkey's performance varied from 71 to 100% correct, with an average of 89% correct. (These percentages indicate performance on nonstimulated trials only.) The size of the stimulation effect was positively correlated with task performance. The correlation is statistically significant when the data are pooled across the 2 monkeys (r = 0.49, P < 0.005), suggesting that the observed deviations in the metrics of stimulation-evoked saccades, although very small, are linked to LSRA task performance. When considered separately, the correlation was considerably more robust in monkey A (r = 0.68, P < 0.005) than in monkey T (r = 0.42, P < 0.05), consistent with the fact that choice-predictive activity, although present in both animals, was stronger in monkey A than in monkey T.

## Stimulus-response compatibility affects saccade latency

The microstimulation experiments described above showed the monkeys' perceptual decisions influenced the precise metrics of saccades evoked with electrical stimulation. We now show that decisions were likewise manifest in the latencies of the operant saccades made to the targets on interleaved (nonstimulated) trials. Figure 10A illustrates the basic finding presented in the preceding sections: stimulus motion in a particular direction excited a population of SC neurons whose movement fields (dashed circles) lay in the direction of stimulus motion (arrows). For a given direction of motion, however, the correct operant saccade may be directionally compatible or incompatible depending on the location of the pair of saccade targets on any given trial. In Fig. 10A, for example, motion down and to the right, which elicited predictive activity in a population of neurons with movement fields near the dashed black circle, could be followed by an operant saccade in the same general direction (target locations C and D) or in the opposite direction (target locations A and B).

Saccade latencies were shorter when stimulus motion and the operant saccade were directionally compatible. Figure 10B shows the relationship between saccade latency and target position in the experimental session whose spatial layout is illustrated in Fig. 10A. Target position is defined relative to the fixation point along the axis of stimulus motion. Because of the symmetry of our display, we arbitrarily designated one of the 2 motion directions as the "reference" direction (45° clockwise from rightward in the experiment of Fig. 10A, black arrow). By convention, positive target positions lie in the reference direction, and negative target positions lie in the opposite direction. Black points in Fig. 10B represent saccade latencies on reference trials and gray points represent saccade latencies on nonreference direction trials. Black and gray lines are leastsquares fits to the data points. The slope of the black line is negative, indicating that latencies were shorter when the direction of the operant saccade was consistent with the direction of motion flow (downward and rightward). In contrast, the slope of the gray line is positive, indicating that latencies were shorter for saccades upward and leftward when motion flowed in that direction.

To assess this effect across a population of behavioral experiments, we performed a multiple regression analysis of each data set, modeling the latency of correct saccades as a weighted sum of target position, the direction of stimulus motion, and the



FIG. 10. A: geometry of a single LSRA experiment. Motion in the stimulus aperture (central circle) flowed down and right (black arrow) or up and left (gray arrow), exciting SC neurons with movement fields in the direction of motion (dashed circles). Positions of correct targets are labeled A, B, C, and D. Pairs of targets with the same identifying letter were considered equivalent for this analysis because they were equidistant from the fixation point along the motion axis. *B*: scatterplot of saccade latency and target position. Saccades reported down-right motion (black points) or up-left motion (gray points). Each point has been shifted horizontally by a small random amount for visual clarity. Black and gray lines are regression fits (to unshifted points). Saccades in the direction of motion have shorter latencies than saccades in the opposite direction.

interaction between these 2 variables. The interaction term is of particular interest because it captures differences in the slope of the relationship between target position and saccade latency that depend on the motion direction. The data in Fig. 10*B* clearly illustrate such a slope difference, and the interaction term for these data was indeed significant (P < 0.05), indicating that the slope difference is not attributable to random variability.

Figure 11 shows a scatterplot of the 2 regression slopes (e.g., the black and gray lines in Fig. 10*B*) obtained in each of the 86 experiments we performed. Data from monkey A appear in Fig.11*A* and data from monkey T appear in Fig. 11*B*. Experiments with significant interactions, which constitute a majority, are indicated with filled symbols. The relationship between slopes was generally consistent across sessions: motion flowing in the reference direction yielded negative slopes, whereas motion in the opposite direction yielded positive slopes. In other words, for either direction of motion, saccade latencies decreased as the target was presented with increasing eccenters.

tricity in the direction of motion. On average, a displacement of the target by  $1^{\circ}$  in the direction of motion decreased saccade latency by an average of 1.2 ms for monkey A and by 1.0 ms for monkey T. Thus on correct trials saccades in the direction of motion have shorter latencies than saccades in the opposite direction.

These results are equally consistent with 2 interpretations: short latency might be a property of saccades in the direction of stimulus motion or in the direction of the monkey's judgment. These 2 possibilities cannot be distinguished from the previous analysis of correct trials because the monkey's judgment is, by definition, in the direction of stimulus motion on these trials. We thus analyzed error trials to determine whether latency decreased in the direction of stimulus motion of in the direction of the judgment. We included in this analysis only the



FIG. 11. Scatterplot of regression slopes (saccade latency vs. target position) calculated from reference direction and nonreference direction trials. Filled symbols indicate significant interaction terms (target position  $\times$  motion direction). Open symbols indicate nonsignificant interactions. Identity line is shown for reference (diagonal line). In almost all experiments, the slope of the relationship between saccade latency and target position is more negative when motion flowed in the reference direction than when it flowed in the opposite direction. Data from monkey A appear in *A*, and data from monkey T appear in *B*.

27 experiments (15 from monkey A and 12 from monkey T) in which  $\geq 10$  errors were made in each direction.

This analysis revealed that saccade latencies were related to the direction judgments, not to the direction of motion in the stimulus. Thus for error trials saccade latency *increased* with displacement of the target in the direction of motion (1.1 ms/°). This effect was significant for monkey T (Wilcoxon test on regression slopes: P < 0.01) and nearly so for monkey A (P <0.1). This effect is a further indication that the monkey's perceptual judgment influenced the state of the oculomotor system.

### Predictive signals in SEF

Where do the predictive signals that we observed in the SC originate? The supplementary eye field (SEF) is a good candidate for providing such signals, given that Olson and colleagues have documented predictive activity in the SEF in a task that is similar to our LSRA task (Olson and Gettner 1999, 1995; Olson and Tremblay 2000; Tremblay et al. 2002). To explore this idea, we recorded from SEF neurons in monkey T while he performed the LSRA task.

The SEF was identified using conventional criteria: the medial surface of the frontal lobe from which saccadic eye movements could be elicited with electrical stimulation currents  $\leq 60 \ \mu$ A (Mitz and Godschalk 1989). Neurons were isolated for study while the monkey performed either the LSRA task or the memory-guided saccade task. For each neuron isolated, the preferred saccade direction was estimated from response during a block of delayed-saccade trials or from the vector of saccades elicited with microstimulation. Each neuron was then tested for responsiveness in the LSRA task, first using an axis of dot motion aligned with the cell's estimated preferred direction, and second, with dot motion along an orthogonal axis.

Figure 12 illustrates predictive activity in a population of 25 choice-predictive SEF neurons (thick curve). We included only correct trials in this analysis because the animal generally performed the task very well, generating too few errors for analysis. For comparison, Fig. 12 also shows predictive activity from the population of 18 SC neurons from the same animal (thin curve). Early in the trial, SEF neurons were more predictive than SC neurons. The difference in predictive activity was statistically significant from 300 ms after stimulus onset until 300 ms into the delay period (randomization tests: P < 0.01; shaded area, Fig. 12). The population of SC neurons was never significantly more predictive than the population of SEF neurons. Nine of 23 SEF neurons tested also exhibited statistically significant delay activity in the memory-guided saccade task (Mann–Whitney U test, P < 0.05). In every case, there was reasonable agreement between preferred directions in the delayed-saccade task and in the LSRA task.

These observations are consistent with the ideas that: 1) processes occurring in the frontal lobe might drive predictive activity in the SC, and 2) the predictive activity that we observed in the SC may be very similar to that reported in SEF by Olson and colleagues during studies of object-centered coordinates. We will return to these issues in the DISCUSSION.



FIG. 12. Predictive activity from 25 supplementary eye field (SEF) neurons (heavy line) and 18 SC neurons (thin line) recorded from monkey T. Data from correct trials only are shown. Difference between the 2 predictive activity curves becomes statistically significant (randomization test: P < 0.01) 300 ms into the stimulus presentation and stays significant until 300 ms into the delay period (gray patch). Evolution of predictive activity is qualitatively similar between the SEF and SC, but reaches higher levels more quickly in the SEF.

### DISCUSSION

Our primary goal in this study was to characterize the responses of decision-related SC neurons in a direction-discrimination task in which the perceptual decision was dissociated from planning of the operant response. In a previous study, we identified decision-related SC responses in the standard version of our direction-discrimination task in which the vector of the operant saccade bore a stereotyped spatial relationship to the motion signal in the random-dot stimulus (Horwitz and Newsome 1999, 2001a,b). In that paradigm the monkey could prepare the operant saccade nearly simultaneously with the perceptual accumulation of motion signals originating in the visual cortex. Thus in the standard task it is reasonable to expect decision-related signals to appear in oculomotor structures, and this is indeed the case (Gold and Shadlen 2000; Horwitz and Newsome 2001b; Kim and Shadlen 1998; Roitman and Shadlen 2002; Shadlen and Newsome 1996).

In contrast to the standard task, the LSRA task used in the current study did not permit the monkey to know the metrics of the operant saccade during formation of the perceptual decision. We therefore expected that decision-related activity would be eliminated from the SC (and probably from other oculomotor structures as well) during performance of this task. To our surprise, however, a population of intermediate and deep layer SC neurons, as well as a population of SEF neurons, discharged in a choice-specific manner during the LSRA task. This decision-related discharge, or "predictive activity," began up to several seconds before the monkey could plan the operant saccade, allowing an experimenter to predict the monkey's perceptual report irrespective of the direction and amplitude of the operant saccade. Thus the activity of a subpopulation of SC and SEF neurons during this task appears to be related to the

perceptual judgment per se, and not simply to the operant saccade.

Predictive activity rose steadily after the motion stimulus was extinguished, peaking around the time that the targets were illuminated. This observation is inconsistent with the idea that the SC buildup activity reflects the integration of motion information (Gold and Shadlen 2001), but rather points to more complex dynamics (e.g., Wang 2001). Remarkably, predictive activity persisted after the presentation of the 2 saccade targets and even extended through the time of the saccade itself for some neurons (Fig. 5). Thus even after the monkey was able to plan and execute the correct saccade, the discharge of some SC neurons was not strictly related to saccade metrics, but rather, was strikingly modulated by behavioral context (i.e., the perceptual decision).

A possible (uninteresting) explanation for our data is that choice-predictive activity reflects covert saccade planning. In this scenario, the monkey covertly prepares a saccade in the direction of stimulus motion, accounting for our observation of choice-predictive activity in an oculomotor structure. The saccade plan would then be changed to the correct operant saccade on appearance of the pair of targets. Three lines of evidence militate against this interpretation. First, the LSRA task was specifically designed to discourage the monkey from planning saccades in advance of the target presentation; at the time in the trial that predictive activity evolved, the monkey could not predict the metrics of the required saccade because the distribution of correct target locations was identical for both directions of stimulus motion (see METHODS). Second, the proportion of neurons that predicted choices in the LSRA task (about 10%) was considerably smaller than the proportion of neurons exhibiting predictive activity in our standard direction-discrimination task or exhibiting delay-period activity in a memoryguided saccade task (both about 30%). If a saccade plan can be thought of as a well-defined, unitary entity, we would expect similar numbers of neurons to contribute to a plan in all 3 tasks. Finally, stimulation-evoked saccades were influenced only weakly by the direction of stimulus motion in the LSRA task (Figs. 7, 8, 9A). In contrast, saccade plans formed during a delayed-saccade task exerted striking effects on the direction of stimulation-evoked saccades. If the monkey actively planned saccades in the LSRA task, we would expect stimulationevoked saccade vectors to be affected similarly. Together, these observations build a strong case that decision-related activity in the LSRA task does not result from covert saccade planning.

## Spatial representation of the perceptual decision

A priori, perceptual decisions made during the LSRA task might be represented in a completely "abstract" fashion, that is, unrelated to a particular motor system or to particular regions of visual space. For example, the monkey's decisions might be represented in categorical "rightward" or "leftward" neurons in a nontopographic region of the frontal cortex. That decisionrelated signals are present in a spatially organized structure like the SC raises the alternative possibility that decisions in the LSRA task are in fact coded with respect to space. This notion is supported by 3 different observations presented in this paper: electrophysiological, microstimulation, and behavioral. First, decision-related signals recorded electrophysiologically were systematically related to the spatial map in the SC: most choice-predictive SC neurons discharged most vigorously when the monkey reported motion flowing *toward* their movement fields. Thus stimulus motion and the accompanying decision selectively excited a topographic region of the SC lying in the direction of motion (i.e., stimulus motion "pointed" toward the neuron's spatial movement field). Second, small deviations in electrically evoked saccades were systematically related to the direction of stimulus motion. On trials in which rightward or leftward motion was presented, for example, the endpoints of evoked saccades were displaced slightly to the right and left, respectively, showing that decision-related signals influenced the oculomotor system in a spatially specific manner. This relationship was particularly strong during sessions of good performance (Fig. 9B). Finally, the behavioral effects of stimulus-response compatibility (Figs. 10 and 11) support the idea that perceptual decisions are coded spatially during the LSRA task. If the neural representation of decisions was completely abstract, we would not expect to find latency differences when the direction of the operant saccade is compatible or incompatible with the direction of perceived motion. The direction-specific bias on oculomotor behavior again suggests that the representation of the decision is at least in part spatially coded.

One can imagine several possible mechanisms for storing the perceptual decision in a spatially mapped structure. Most simply, perhaps, the monkey could remember (or attend to) the location in space toward which motion flowed as a proxy for (or in addition to) remembering the motion direction abstractly. Such a "spatial mnemonic" could span the delay period and provide a basis for selecting the appropriate saccade after the target pair appears. Monkeys have a well-developed ability to remember the spatial location of an object even when the object is not visible. This type of memory, termed "spatial working memory," is characterized by the active maintenance of behaviorally relevant spatial information for durations of several seconds (Assad and Maunsell 1995; Chafee and Goldman-Rakic 1998, 2000; Funahashi et al. 1989; Funahashi and Takeda 2002; Logie 1995; Pesaran et al. 2002; Umeno and Goldberg 2001). Although the LSRA task does not explicitly require the storage of spatial locations, the monkeys may have exploited this memory system to encode the motion direction.

In a somewhat more sophisticated model, the perceptual decision could be encoded as a gradient of enhanced activity across a spatially organized structure such as the SC. A rightward decision, for example, would generate maximally or minimally enhanced delay period activity for SC locations corresponding to large rightward or leftward eye movements, respectively, with activity varying monotonically across intermediate locations. This gradient of activity, once established in a subpopulation of SC neurons (e.g., the predictive neurons identified in the present study), could provide the memory trace necessary for spanning the delay period. After the targets appeared, it would be a computationally simple matter to select the target that lies closest to the maximally active end of the gradient. A model of this nature has, in fact, been proposed to explain the perceptual deficits associated with hemi-neglect in some human patients (Pouget and Sejnowski 1997).

Ironically, then, the available evidence suggests that perceptual decisions in the LSRA task may be represented, at least in part, relative to spatial maps within the oculomotor system, even though the task design eliminates any correspondence between the direction of stimulus motion and the metrics of the operant saccade. Gold and Shadlen (2003) recently designed a third version of the direction-discrimination task in which the relationship between stimulus motion and the operant saccade is even more abstract than that in the LSRA task. Motion with a rightward component required an operant saccade to a red target whereas motion with a leftward component required a saccade to a green target. From trial to trial, the red and green target locations were varied randomly within an array of possible target locations without any consistent spatial relationship with each other. Using microstimulation techniques nearly identical to those we used in Figs. 6-9, Gold and Shadlen observed no effect of motion direction on the metrics of electrically evoked saccades and concluded that perceptual decisions in the red/green version of the task are represented in a genuinely abstract, nonspatial manner within the nervous system. It may well be the case that the additional level of spatial abstraction provided by the red/green task (abolishing the spatial relationship between the 2 saccade targets) gives rise to a completely abstract decision representation in the brain; we certainly agree that Gold and Shadlen's microstimuation results support this interpretation, although we are hesitant to draw firm conclusions on this issue until the other 2 lines of evidence presented in this paper are evaluated in the context of the red/green task. Specifically, 1) neural activity should be recorded in a spatially organized structure [such as the SC or frontal eye field (FEF)] during performance of the task to determine whether decision-related signals might be present in a subpopulation of neurons as in the present study, and 2) the behavioral data recorded in the red/green task should be analyzed to determine whether stimulus-compatibility effects on saccade latency are present in the data.

## **Object-centered** coordinates?

The LSRA task is very similar to tasks used by Olson and colleagues to investigate object-centered coordinate frames in the brain (Olson and Gettner 1999, 1995; Olson and Tremblay 2000; Tremblay et al. 2002). In these tasks a monkey is cued to make a saccade to one side of an object (e.g., the left end of a horizontal bar) irrespective of the location of the bar in oculocentric coordinates. Remarkably, Olson and colleagues found that a subpopulation of SEF neurons responds for all saccades that conform to the instruction, regardless of the actual oculocentric vector of the saccade. By analogy, target pairs in our task can be considered the ends of a virtual object, and our motion stimulus instructs a saccade to one end of the virtual object. Thus neurons, like the one illustrated in Figs. 2 and 3, can be thought of as signaling the upcoming saccade in object-centered coordinates. The neuron fires vigorously preceding saccades to the upper-left end of the virtual object irrespective of its position in the visual field (Fig. 3, left column). Conversely, it responds only weakly when the monkey makes identical saccades to the lower-right end of the virtual object (Fig. 3, right panels). From this perspective, the decision-related neurons we have studied in the SC and SEF appear virtually identical to many of the SEF neurons reported by Olson and colleagues. Indeed, Olson and colleagues have shown directly that target pairs like those used in the present study can elicit object-centered responses from many SEF neurons (Olson and Tremblay 2000).

Should the SC and SEF neurons we have studied be regarded as coding saccades in object-centered coordinates? This is certainly a plausible hypothesis that is consistent with our current data. It is important to note, however, that an objectcentered representation is but one of several spatial strategies that could enable the monkey to solve the LSRA task, and we are unable to distinguish among these strategies based on current data. We proposed 2 equally plausible "spatial mnemonic" strategies, for example, in an earlier section of this DISCUSSION. Although the object-centered account remains a viable explanation of our data, further studies appear necessary to test this model incisively.

## High-level signals in the SC

The most surprising result of this study is the presence in the SC of signals related to the outcome of a perceptual decision in the absence of a statistical association between motion direction and the vector of the operant saccade. What does this finding imply about the function of the SC? One possibility is that these are "stray" signals that "leak" into the SC from cortical structures that play a primary role in the formation of perceptual decisions (perhaps SEF, LIP, FEF, and area 46). In this conception the SC signals described in this paper have no functional significance in terms of behaviorally significant output from the SC. We should therefore continue to think of the SC according to the standard model in which the function of the intermediate and deep layers is simply to prepare and execute eye movements to the next target location.

Alternatively, the SC may play a heretofore underappreciated role in higher-level processing leading to the selection of targets for saccades. In this conception different cell types in the SC, like those in cortical areas, occupy different levels in the hierarchichal processes that result in target selection, coordinate system transformation, saccade preparation, and saccade execution. Cells that occupy an equivalent hierarchical level in different brain structures would interact cooperatively to implement the computations appropriate to each level. Thus we should not expect a hierarchy of brain areas to correspond to a hierarchical sequence of computations. Rather, specific cell types in a network of brain areas, including midbrain structures such as the SC, would interact cooperatively to process signals at each computational level. (A hybrid model between these extremes is certainly plausible: a hierarchical sequence of computations could be performed cooperatively by a network of brain areas amid diversity in the *relative* contributions of individual areas to different computations.)

The "network" model is consistent with a growing body of evidence acquired while monkeys perform a standard series of oculomotor tasks including short-latency saccades to visible targets, delayed saccades to visible targets, saccades to the remembered location of briefly flashed targets, and "doublejump" saccade tasks. A variety of neural signals is revealed during such tasks, including visual responses, memory-related activity, motor "burst" responses, and responses reflecting coordinate computations after the first saccade in a double-jump pair (Colby et al. 1995; Goldberg and Segraves 1990; Mays and Sparks 1980). Interestingly, all of these signals can be observed in a variety of structures related to oculomotor behavior, including structures in the cerebral cortex, thalamus, midbrain, and basal ganglia. In addition, a series of studies carried out by Wurtz, Pare, and Sommer indicates that the output from many individual areas incorporates a similarly diverse range of neural signals (Ferraina et al. 2002; Pare and Wurtz 1997, 2001; Sommer and Wurtz 1998, 2001, 2002; Wurtz et al. 2001). Such data reinforce the notion that multiple processing levels are incorporated into anatomically diverse neural structures, and suggest that "lower" structures such as the SC may participate in a broader range of cognitively relevant computations than has been recognized previously. Critical tests of this hypothesis must ultimately include experimental approaches that can establish whether the SC actually plays a causal role in the genesis of higher-level cognitive functions.

# APPENDIX: NORMALIZATION PROCEDURE FOR ELECTRICALLY EVOKED SACCADE VECTORS

The measurement of the distance between the mean position (centroids) of 2 sets of saccade endpoint distributions is highly dependent on endpoint variability. Here we describe a normalized distance metric that compensates for differences in saccade endpoint variability and thus allows for stimulation effects to be compared across experiments. Consider 2 sets of random draws from a single distribution of endpoints. If the endpoint distribution is spread over many degrees in both dimensions, the centroids will tend to be far apart (even though both sets of endpoints came from the same distribution). Conversely, centroids will lie close together if the endpoint distribution is restricted, say, to a few minutes of arc. The expected distance between centroids depends on the amount of data collected, but for a finite number of data points, the expected distance between the centroids increases with the variance of the endpoints. The variance of saccade endpoints depends on saccade amplitude, the location of a stimulation site in the SC, and the precise parameters of a stimulation train (Becker 1989; Stanford et al. 1996).

We assume that, for a given stimulus condition, saccade endpoints are independent samples from a single bivariate distribution. We also assume that different directions of stimulus motion may shift this distribution spatially, but do not change its shape. Our goal is to measure the distance between mean endpoints in units that are related to the width of these distributions.

We estimate the variances and covariance of the underlying endpoint distributions by the usual formulas

$$\begin{split} \text{Var} \left( x_{j} \right) &= \frac{\sum_{i=1}^{n} \left( x_{ij} - \bar{x}_{j} \right)^{2}}{n_{j} - 1} \\ \text{Var} \left( y_{j} \right) &= \frac{\sum_{i=1}^{n} \left( y_{ij} - \bar{y}_{j} \right)^{2}}{n_{j} - 1} \\ \text{Cov} \left( x_{j}, \, y_{j} \right) &= \frac{\sum_{i=1}^{n} \left( x_{ij} - \bar{x}_{j} \right) (y_{ij} - \bar{y}_{j})}{n_{i} - 1} \end{split}$$

where x and y refer to endpoint coordinates, n is the number of saccades, i is a trial number index, and j is an index relating to the motion direction. These estimates can be pooled across stimulus conditions according to

$$\begin{aligned} &\operatorname{Var}_{\text{pooled}}\left(x\right) = \frac{n_{1} - 1}{n_{1} + n_{2} - 2} \operatorname{Var}\left(x_{1}\right) + \frac{n_{2} - 1}{n_{1} + n_{2} - 2} \operatorname{Var}\left(x_{2}\right) \\ &\operatorname{Var}_{\text{pooled}}\left(y\right) = \frac{n_{1} - 1}{n_{1} + n_{2} - 2} \operatorname{Var}\left(y_{1}\right) + \frac{n_{2} - 1}{n_{1} + n_{2} - 2} \operatorname{Var}\left(y_{2}\right) \end{aligned}$$

$$Cov_{pooled}(x, y) = \frac{n_1 - 1}{n_1 + n_2 - 2} Cov(x_1, y_1) + \frac{n_2 - 1}{n_1 + n_2 - 2} Cov(x_2, y_2)$$

This results in the covariance matrix, S

$$\begin{bmatrix} Var_{pooled}(x) & Cov_{pooled}(x, y) \\ Cov_{pooled}(x, y) & Var_{pooled}(y) \end{bmatrix}$$

In general, there exists a linear transformation of the variables (X, Y) such that the covariance matrix of the transformed variables (X', Y') is the identity matrix (Picinbono 1993). This transformation can be estimated by

$$\begin{bmatrix} \mathbf{x}' & \mathbf{y}' \end{bmatrix} = \begin{bmatrix} sqrt(\mathbf{S}) \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}$$

where *sqrt* indicates the matrix square root. Geometrically, the transformation is equivalent to rotating and scaling the coordinate axes so that each cluster of endpoints has approximately variance 1 in both dimensions and covariance 0. The distance between the centroids in this space,  $d_{norm}$ , can then be fairly compared across stimulation sites. In the one-dimensional case, this procedure yields d', the distance between the means in units of SD.

## ACKNOWLEDGMENT

We thank J. Riehl, J. Stein, and M. Warden for technical assistance during the course of the study. Drs. Gregory DeAngelis, M. James Nichols, and Eyal Seideman contributed to the design of the LSRA task. Drs. Mazyar Fallah and James Muller provided valuable comments on an earlier version of the manuscript.

Present address of G. D. Horwitz: The Salk Institute, P.O. Box 85800, San Diego, CA 92186.

### G R A N T S

G. D. Horwitz was supported by a training grant from the National Institute of Mental Health (MH-17047). A. P. Batista was supported by the Howard Hughes Medical Institute. W. T. Newsome is an investigator at the Howard Hughes Medical Institute, and received additional support from the National Eye Institute (EY-05603).

#### REFERENCES

- Assad J and Maunsell J. Neuronal correlates of inferred motion in primate posterior parietal cortex. *Nature* 373: 518–521, 1995.
- Batschelet E. Circular Statistics in Biology. New York: Academic Press, 1981.
- Britten KH, Newsome WT, Shadlen MN, Celebrini S, and Movshon JA. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13: 87–100, 1996.
- Britten KH, Shadlen MN, Newsome WT, and Movshon JA. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12: 4745–4765, 1992.
- Chafee MV and Goldman-Rakic PS. Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *J Neurophysiol* 79: 2919–2940, 1998.
- Chafee MV and Goldman-Rakic PS. Inactivation of parietal and prefrontal cortex reveals interdependence of neural activity during memory-guided saccades. J Neurophysiol 83: 1550–1566, 2000.
- Colby CL, Duhamel JR, and Goldberg ME. Oculocentric spatial representation in parietal cortex. *Cereb Cortex* 5: 470–481, 1995.
- Edgington ES. Randomization Tests. New York: Marcel Dekker, 1995.
- Ferraina S, Pare M, and Wurtz RH. Comparison of cortico-cortical and cortico-collicular signals for the generation of saccadic eye movements. *J Neurophysiol* 87: 845–858, 2002.
- Funahashi S, Bruce C, and Goldman-Rakic P. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. J Neurophysiol 61: 331–349, 1989.

Funahashi S and Takeda K. Information processes in the primate prefrontal cortex in relation to working memory processes. *Rev Neurosci* 13: 313–345, 2002.

**Glimcher PW.** Making choices: the neurophysiology of visual-saccadic decision making. *Trends Neurosci* 24: 654–659, 2001.

Gold JI and Shadlen MN. Representation of a perceptual decision in developing oculomotor commands. *Nature* 404: 390–394, 2000.

- **Gold JI and Shadlen MN.** Neural computations that underlie decisions about sensory stimuli. *Trends Cogn Sci* 5: 10–16, 2001.
- **Gold JI and Shadlen MN.** The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *J Neurosci* 23: 632–651, 2003.
- **Goldberg ME and Segraves MA.** The role of the frontal eye field and its corticotectal projection in the generation of eye movements. *Res Publ Assoc Res Nerv Ment Dis* 67: 195–209, 1990.
- Green DM and Swets JA. Signal Detection Theory and Psychophysics. New York: Wiley, 1966.
- Horwitz GD and Newsome WT. Separate signals for target selection and movement specification in the superior colliculus. *Science* 284: 1158–1161, 1999.
- Horwitz GD and Newsome WT. Target selection for saccadic eye movements: direction-selective visual responses in the superior colliculus. J Neurophysiol 86: 2527–2542, 2001a.
- Horwitz GD and Newsome WT. Target selection for saccadic eye movements: prelude activity in the superior colliculus during a direction-discrimination task. J Neurophysiol 86: 2543–2558, 2001b.
- Judge SJ, Richmond BJ, and Chu FC. Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res* 20: 535–538, 1980.
- Kim JN and Shadlen MN. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2: 176–185, 1998.
- Krauzlis R and Dill N. Neural correlates of target choice for pursuit and saccades in the primate superior colliculus. *Neuron* 35: 355–363, 2002.
- Kustov AA and Robinson DL. Shared neural control of attentional shifts and eye movements. *Nature* 384: 74–77, 1996.
- Logie RH. Visuo-Spatial Working Memory. East Sussex, UK: Lawrence Erlbaum, 1995.
- Mays LE and Sparks DL. Dissociation of visual and saccade-related responses in superior colliculus neurons. J Neurophysiol 43: 207–232, 1980.
- Miller EK, Nieder A, Freedman DJ, and Wallis JD. Neural correlates of categories and concepts. Curr Opin Neurobiol 13: 198–203, 2003.
- Mitz AR and Godschalk M. Eye-movement representation in the frontal lobe of rhesus monkeys. *Neurosci Lett* 106: 157–162, 1989.
- **Olson CR and Gettner SN.** Object-centered direction selectivity in the macaque supplementary eye field. *Science* 269: 985–988, 1995.
- Olson CR and Gettner SN. Macaque SEF neurons encode object-centered directions of eye movements regardless of the visual attributes of instructional cues. J Neurophysiol 81: 2340–2346, 1999.
- **Olson CR and Tremblay L.** Macaque supplementary eye field neurons encode object-centered locations relative to both continuous and discontinuous objects. *J Neurophysiol* 83: 2392–2411, 2000.
- Pare M and Wurtz RH. Monkey posterior parietal cortex neurons antidromically activated from superior colliculus. J Neurophysiol 78: 3493–3497, 1997.
- Pare M and Wurtz RH. Progression in neuronal processing for saccadic eye movements from parietal cortex area lip to superior colliculus. J Neurophysiol 85: 2545–2562, 2001.

- **Pesaran B, Pezaris JS, Sahani M, Mitra PP, and Andersen RA.** Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci* 5: 805–811, 2002.
- Pouget A and Sejnowski TJ. A new view of hemineglect based on the response properties of parietal neurones. *Philos Trans R Soc Lond B Biol Sci* 352: 1449–1459, 1997.
- Roitman JD and Shadlen MN. Response of neurons in area LIP during a reaction-time direction discrimination task. Soc Neurosci Abstr 24: 262, 1998.
- Roitman JD and Shadlen MN. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22: 9475–9489, 2002.
- Romo R and Salinas E. Touch and go: decision-making mechanisms in somatosensation. Annu Rev Neurosci 24: 107–137, 2001.
- Salzman CD, Murasugi CM, Britten KH, and Newsome WT. Microstimulation in visual area MT: effects on direction discrimination performance. *J Neurosci* 12: 2331–2355, 1992.
- Schall JD. Neural basis of deciding, choosing and acting. *Nat Rev Neurosci* 2: 33–42, 2001.
- Shadlen MN, Britten KH, Newsome WT, and Movshon JA. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. J Neurosci 16: 1486–1510, 1996.
- Shadlen MN and Newsome WT. Motion perception: seeing and deciding. Proc Natl Acad Sci USA 93: 628–633, 1996.
- Shadlen MN and Newsome WT. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J Neurophysiol 86: 1916– 1936, 2001.
- Sommer MA and Wurtz RH. Frontal eye field neurons orthodromically activated from the superior colliculus. J Neurophysiol 80: 3331–3335, 1998.
- Sommer MA and Wurtz RH. Frontal eye field sends delay activity related to movement, memory, and vision to the superior colliculus. J Neurophysiol 85: 1673–1685, 2001.
- Sommer MA and Wurtz RH. A pathway in primate brain for internal monitoring of movements. *Science* 296: 1480–1482, 2002.
- 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.
- Thompson KG, Hanes DP, Bichot NP, and Schall JD. Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. J Neurophysiol 76: 4040–4055, 1996.
- Tremblay L, Gettner SN, and Olson CR. Neurons with object-centered spatial selectivity in macaque SEF: do they represent locations or rules? *J Neurophysiol* 87: 333–350, 2002.
- Umeno MM and Goldberg ME. Spatial processing in the monkey frontal eye field. II. Memory responses. J Neurophysiol 86: 2344–2352, 2001.
- Wang XJ. Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci* 24: 455–463, 2001.
- Wurtz RH, Sommer MA, Pare M, and Ferraina S. Signal transformations from cerebral cortex to superior colliculus for the generation of saccades. *Vision Res* 41: 3399–3412, 2001.