

Target Selection for Saccadic Eye Movements: Prelude Activity in the Superior Colliculus During a Direction-Discrimination Task

GREGORY D. HORWITZ AND WILLIAM T. NEWSOME

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

Received 14 July 2000; accepted in final form 13 April 2001

Horwitz, Gregory D. and William T. Newsome. Target selection for saccadic eye movements: prelude activity in the superior colliculus during a direction-discrimination task. *J Neurophysiol* 86: 2543–2558, 2001. We investigated the role of the superior colliculus (SC) in saccade target selection while macaque monkeys performed a direction-discrimination task. The monkeys selected one of two possible saccade targets based on the direction of motion in a stochastic random-dot display; the difficulty of the task was varied by adjusting the strength of the motion signal in the display. One of the two saccade targets was positioned within the movement field of the SC neuron under study while the other target was positioned well outside the movement field. Approximately 30% of the neurons in the intermediate and deep layers of the SC discharged target-specific preludes of activity that “predicted” target choices well before execution of the saccadic eye movement. Across the population of neurons, the strength of the motion signal in the display influenced the intensity of this “predictive” prelude activity: SC activity signaled the impending saccade more reliably when the motion signal was strong than when it was weak. The dependence of neural activity on motion strength could not be explained by small variations in the metrics of the saccadic eye movements. Predictive activity was particularly strong in a subpopulation of neurons with directional visual responses that we have described previously. For a subset of SC neurons, therefore, prelude activity reflects the difficulty of the direction discrimination in addition to the target of the impending saccade. These results are consistent with the notion that a restricted network of SC neurons plays a role in the process of saccade target selection.

INTRODUCTION

In the preceding paper, we showed that a subpopulation of prelude neurons in the superior colliculus (SC) exhibited direction-selective visual responses within large receptive fields that included the center of gaze. Directional visual responses were strong in three monkeys that had been extensively trained to associate specific directions of motion with saccades of specific vectors. For each neuron, the preferred direction of the visual responses pointed toward the spatial location of the cell’s movement field. As described in the preceding paper, these results suggest that two populations of prelude neurons are present in the SC of our extensively trained monkeys: one that plays a role in selecting saccade targets on the basis of visual motion cues and another that is more concerned with the specification of saccade metrics. The latter population may be

strongly involved in programming saccades to salient visual stimuli within the movement field, consistent with the traditional model of collicular function (for a review, see Wurtz and Albano 1980).

To gain further insight into the mechanisms of saccade target selection, we now compare responses recorded from these two neuronal populations while monkeys performed the direction-discrimination task. The monkeys were trained to discriminate between two opposite directions of coherent motion in a stochastic random-dot display. Following a brief delay period, the monkey revealed its decision by making a saccadic eye movement to one of two possible saccade targets, one of which was placed inside the neuron’s movement field while the other was placed well outside the movement field. Thus the monkeys selected one of two possible saccade targets based on the direction of motion in the visual stimulus.

Central to our strategy is the use of near-threshold motion signals that compel the monkey to accumulate motion information gradually (Britten et al. 1992; Roitman and Shadlen 1998). On some trials, the motion stimulus was strong and the monkey was quickly certain of the correct answer; on other trials, the motion stimulus was weak and the monkey had to guess the correct answer. For both strong and weak motion signals, however, the saccades to the same pair of targets served to report the monkey’s decisions.

We recorded exclusively from neurons whose activity early in the trial correlated with the monkeys’ target choices (Horwitz and Newsome 1999). This presaccadic prelude activity differed qualitatively between the two groups of cells. Neurons that exhibited direction-selective visual responses during a passive-fixation task also reflected the effect of the cue stimulus during the discrimination task: the intensity of prelude activity correlated both with the saccade target selected by the monkey and with the strength of the motion signal that guided decision making and target selection. In contrast, neurons that lacked directional responses during the fixation task were relatively insensitive to the cue stimulus during the direction-discrimination task. The results are consistent with the suggestion advanced in the preceding paper that one population of SC prelude neurons appears well-suited for participating in the process of target selection while the other population appears more tightly linked to programming the saccade to the selected

Address for reprint requests: W. T. Newsome, Dept. of Neurobiology, Stanford University School of Medicine, Fairchild Building, Rm. D209, Stanford, CA 94305.

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.

target. Our results are qualitatively similar to those obtained in frontal and parietal cortical neurons using the same experimental paradigm (Kim and Shadlen 1998; Shadlen and Newsome 1996, 2001).

METHODS

Three monkeys (*Macaca mulatta*) served as subjects in the experiments reported here and in the companion paper (Horwitz and Newsome 2001). Methods employed in the two studies overlapped to a large extent. Here we emphasize methods particular to the current study.

Behavioral paradigms and stimuli

The behavioral paradigm employed in this study was a two-alternative, forced-choice direction-discrimination task used previously by Shadlen and Newsome (1996). Monkeys were trained to determine which of two opposed directions of motion was dominant in a stochastic visual stimulus. On each trial, the monkey expressed the direction judgment by making a saccade to a visual target lying in the perceived direction of motion. Visual stimuli and saccade targets were generated on a special-purpose graphics board (Number Nine Computer or Cambridge Research Systems) in an IBM-compatible personal computer and presented on a CRT monitor.

The visual stimulus was a random-dot motion display, which has been used extensively in this laboratory (Britten et al. 1992, 1993; Salzman et al. 1992; Shadlen and Newsome 1996). The random-dot patterns appeared within a circular aperture that subtended 7° of visual angle. Stimulus movies were generated by updating dot positions in successive 60-Hz frames. A variable proportion of dots was replotted at a displacement of 0.15° with respect to their original positions after a delay of 50 ms. Thus these dots appeared to move at a speed of $3^\circ/\text{s}$ in a common direction. These dots, which we will refer to as "signal" dots, served as the basis for the direction discrimination. The remaining dots in the display, called "noise" dots, were replotted in random locations and thus appeared to move in random directions with random speeds. The density of dots in the visual stimuli was 15 dots/($\text{deg}^2 \cdot \text{s}$), but the apparent density of dots in the stimulus was much higher because of persistence in the visual system.

The proportion of signal dots in the display will be referred to as the coherence of the motion stimulus. High-coherence stimuli contain a large proportion of signal dots, and their direction of motion is easily discriminated. Low-coherence stimuli contain only a few signal dots, making the discrimination more difficult. Note that the completely ambiguous stimulus containing no signal dots (the 0% coherence condition) lies on the stimulus continuum. Although direction discrimination for this stimulus is impossible by definition, we routinely included it in our stimulus set, and rewarded choices randomly on these trials. Other stimulus coherences routinely included were 3.2, 6.4, 12.8, 25.6, and 51.2%. Stimulus coherences and directions were pseudorandomly varied from trial to trial by the method of constant stimuli.

Figure 1 illustrates the geometry of the visual display and the timing of events in each trial. Each trial began when the monkey fixated a small point of light subtending 0.2° of visual angle at a distance of 57 cm. Three hundred milliseconds after visual fixation was achieved, two target disks (each subtending 0.7° of visual angle) appeared, flanking the fixation point and collinear with it. Five hundred to 900 ms after the targets appeared, a 2-s-long motion stimulus movie was presented, usually at the center of gaze. After an enforced delay period of randomized length (1–1.5 s), the fixation point disappeared, cueing the monkey to make a saccade to one of the two visual targets. A saccade to the target in the direction of stimulus motion counted as a correct response and was reinforced with a liquid reward. For each cell isolated, the geometry of the display was adjusted for




FIG. 1. Two-alternative, forced-choice direction-discrimination task. The geometry of the display (A) and timing of events (B) are shown. Three hundred milliseconds after the monkey fixated a fixation point, 2 saccade targets were illuminated. Five hundred to 900 ms later, a stochastic motion stimulus was shown at the center of gaze for 2 s and was followed by a delay period lasting from 1 to 1.5 s. After the delay period, the fixation point was extinguished, whereupon the monkey had 500 ms to shift gaze to the target in the direction of stimulus motion. For each cell studied, 1 of the saccade targets (T1) was presented inside the movement field and the other (T2) was presented outside.

each experiment so that one of the targets, hereafter referred to as "T1," lay inside the movement field and the other, "T2," lay outside.

During fixation, the monkey's eye position was required to be within a $3 \times 3^\circ$ electronically defined window surrounding the fixation point. If the monkey broke fixation while the fixation point was lit, the trial was aborted. Trials were also aborted if the monkey failed to make a saccade within 500 ms of fixation point offset or if the saccade failed to land within an analogous electronic window surrounding the target. Target windows were square in shape and varied in size depending on the eccentricity of the target. Saccades landing in this window tended to be quite accurate. Seventy-two percent of saccades landed within 2° of the nominal target location; 96% landed within 5° .

Each monkey also performed a simple delayed-saccade task. In this condition, a single eccentric target appeared 300 ms after the monkey acquired the fixation point. After a delay period lasting 2,500–4,000 ms, the fixation light disappeared, and the monkey made a saccade to the target within 500 ms to acquire a reward. The only exception to this is the cell shown in Fig. 17 for which the delay period was randomized between 1,200 and 1,700 ms.

Data analysis

Data were analyzed with custom software written in Matlab (The MathWorks). Unless specified otherwise, neural responses are shown for correctly answered trials only. One general exception to this rule is that both rewarded and unrewarded trials are shown for 0% coherence because in this condition correctness is arbitrary.

After each experimental session, the percentage of correct choices was plotted as a function of the log coherence and fitted by Quick (cumulative Weibull) functions. The function was of the form

$$\% \text{correct} = 100 \cdot (1 - 0.5e^{-(\text{coherence}/\text{threshold})^{\text{slope}}})$$

Parameters were estimated via the method of maximum likelihood assuming binomially distributed errors. The *threshold* parameter corresponds to the coherence at which the monkey makes 82% correct choices.

Predictive activity

A central goal of this study was to measure the time course of target-specific prelude activity in the SC during performance of our direction-discrimination task. This activity allows an ideal observer (or an experimenter) to “predict” the monkey’s target choices well in advance of saccade execution. We compute the target specificity of neural responses, or “predictive activity,” using a technique based in signal-detection theory (Green and Swets 1966). This procedure has been used in previous studies of saccade target selection (Shadlen and Newsome 1996; Thompson et al. 1997).

Figure 2 illustrates the calculation of predictive activity. Trials are sorted into those that resulted in T1 choices and those that resulted in T2 choices; spike trains are aligned either to the onset of the motion stimulus or to the initiation of the saccade (Fig. 2A). We divide time into nonoverlapping 100-ms bins and count the number of spikes occurring in each bin. Pooling across trials yields two distributions of spike counts (preceding T1 and T2 choices) at each time point (Fig. 2B). From each pair of spike count distributions, we calculate a receiver operating characteristic (ROC curve, Fig. 2C) (Britten et al. 1992; Green and Swets 1966). Each point in the ROC curve is the proportion of T1 spike counts exceeding an arbitrary criterion value as a function of the proportion of T2 spike counts exceeding the same criterion. Entire curves are obtained by sweeping the criterion value through the range of the data. The integrated area beneath each ROC curve is the predictive activity at each time point. Predictive activity is then plotted for

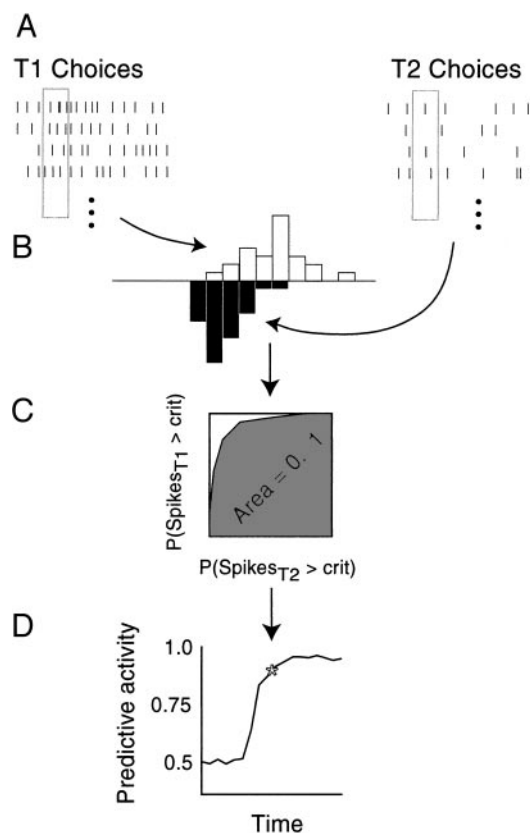


FIG. 2. Predictive activity calculation. Rasters preceding T1 and T2 choices were compiled and spikes were counted in 100-ms bins (A). For a given time bin, distributions of spike counts preceding T1 and T2 saccades were compared (B). A receiver operating characteristic (ROC) curve was calculated from this pair of distributions (C, see METHODS for details). The area beneath this curve is defined as the predictive activity for that time bin (D). This value can be interpreted as the probability with which an ideal observer can predict the monkey’s choice based on the spike counts (Green and Swets 1966).

sequential time bins to visualize its development during single trials (Fig. 2D).

This metric of predictive activity has a number of attractive properties. First, it can be interpreted as the probability with which an ideal observer can correctly predict the monkey’s choice based on spike counts occurring at different times in the trial (Green and Swets 1966). This interpretation is intuitive because during the data collection, we rated neurons qualitatively on the basis of how well we could predict the monkey’s choices from the discharge heard on the audio monitor. Because the predictive activity metric is a probability, it is bounded between 0 and 1. The midpoint of this range, 0.5, is the value expected if no relationship exists between the firing rate of the neuron and the monkey’s choice. Second, this metric does not require assumptions about the parametric form of the distributions of spike counts. Third, it measures only target-specific (or movement-specific) changes in firing rate; nonspecific increases in firing rate related to arousal or anticipation of the “go” signal would not contribute to predictive activity.

For some analyses in this paper, we computed predictive activity across populations of neurons that were recorded sequentially. To accomplish this analysis, we normalized spike counts for each cell as follows. First, as described in the preceding text, we divided time into 100-ms bins and counted the number of spikes occurring in each bin. For each neuron, these counts were averaged across trials of a common stimulus coherence to compute the mean response as a function of coherence and time in the trial. The maximum of these average responses, across all time bins and coherence levels, was used to divide the individual spike counts. We define the result of this operation as the normalized spike count. Normalized spike counts were then pooled across neurons and subjected to the predictive activity calculation described in the preceding text.

We assessed the statistical significance of individual predictive activity values via a permutation test (Britten et al. 1996). We calculated an ROC curve from the distributions of spike counts preceding T1 and T2 saccades, integrated its area, and recorded the result. We then randomly reassigned spike counts to two groups and recalculated the ROC area. This procedure was repeated 2,000 times to generate a reasonable estimate of the ROC area distribution under the null hypothesis. If <100 ($0.05 \times 2,000$) ROC areas equaled or exceeded the one calculated from the actual (nonpermuted) data, we rejected the null hypothesis at the 0.05 level.

Saccade parameter model

The firing rate of some collicular neurons covaries with saccade end-point, velocity, and latency (Dorris and Munoz 1998; Rohrer et al. 1987; Sparks et al. 1976; Wurtz and Goldberg 1972). We used linear regression analysis to determine the degree to which the observed differences in firing rate across trials was attributable to subtle differences in these saccade parameters. For each recorded cell, we fit a model relating the firing rate during one of several temporal epochs to the end-point, velocity, and latency of a subsequent T1-directed saccade. The model assumed the relationship was of the form

$$\text{response} = \beta_0 + \beta_1 x + \beta_2 y + \beta_3 x^2 + \beta_4 y^2 + \beta_5 xy + \beta_6 \text{vel} + \beta_7 x \text{vel} + \beta_8 y \text{vel} + \beta_9 \text{lat}$$

where x and y are saccade end-point coordinates (adjusted for differences in initial position), vel is the saccade peak velocity, and lat is saccadic latency. Coefficients were estimated by the method of least-squares. In individual experiments, saccades to single targets were generally very stereotyped with little variation in precise parameters. Thus our use of a linear model is justified; while the relationship between these saccade parameters and the

firing rate of SC neurons is not linear in general, it is nearly so over narrow ranges (Dorris et al. 1997; Sparks and Mays 1980).

Our main goal in this analysis was to determine whether stimulus coherence accounts for variance in firing rate after eye movement parameters are taken into account. To accomplish this, we fit the neural data to a second model that included the log of the stimulus coherence as an additional factor. Naturally, inclusion of an extra parameter improved the model fit. The significance of this improvement was assessed by partial F tests on the mean squared residual errors in the two models (Draper and Smith 1998). Significance in this test indicates that the coherence of the motion stimulus affects the firing rate of the cell even after saccade parameters have been taken into account.

Our model of saccade metrics includes more terms than necessary. Indeed, in none of the fitted models were all nine coefficients deemed significant at the 0.05 level. This is a conservative procedure, however, because our interest is in the *additional* effect of stimulus coherence. The more terms we include in the model, the less residual variance is available for coherence to explain.

Raster “streak” index

The temporal structure of prelude spike trains varies across trials of identical stimuli. Here, we describe a metric (“streak index”) that quantifies this variability with respect to a Poisson process, a common benchmark of random firing.

Figure 3 schematizes the steps involved in the streak index calculation. We divide the stimulus presentation into eighty 25-ms bins and count the number of spikes occurring in each bin in each trial. Then we calculate the median number of spikes occurring in each time bin (across trials) thereby obtaining 80 spike-count medians. Each spike count in each trial is then compared with the median spike count for its time bin, and individual spike counts are converted into ones or zeros depending on whether they exceed or

fall below the median, respectively.¹ We then concatenate the ones and zeros into a single string and count the number of runs of ones or zeros (e.g., the string “0100011” contains 4 runs). The expected number of runs under the hypothesis of randomly mixed ones and zeros can be shown to be

$$\mu = 1 + \left(\frac{2nm}{n+m} \right)$$

where n and m are the number of ones and zeros in the sequence. The standard deviation, under the same hypothesis, can be shown to be (Zar 1984)

$$\sigma = \sqrt{\frac{(2nm) \cdot (2nm - n - m)}{(n+m)^2 \cdot (n+m-1)}}$$

The streak index is the difference between the observed and the expected number of runs divided by the standard deviation. Notice that this unit-less quantity is independent of the absolute firing rate of the cell.

A neuron that fires spikes according to a Poisson process with a constant underlying rate will have a streak index near 0. In fact, it can be shown that for this hypothetical neuron, the streak index (asymptotically) has a normal distribution with mean 0 and variance 1. This is the large-sample version of the well-established runs test (Zar 1984).

A neuron that fires spikes according to an inhomogeneous Poisson process (a Poisson process with a time-varying rate) will have a streak index drawn from this normal distribution as well, provided the changes in firing rate are identical on every trial. For instance, a Poisson neuron whose firing rate ramps up over the course of each trial, identically on each trial, falls in this category. On the other hand, a neuron that fires spikes according to an inhomogeneous Poisson process whose underlying firing rate differs across trials will not, in general, have a streak index drawn from this distribution.

A positive streak index indicates that a neuron, having fired above the median in one time bin, has a smaller than 50% chance of firing above the median in the next time bin. A neuron with an extremely long refractory period, for instance, might have a positive streak index. Conversely, a large negative index is a signature of a neuron whose firing rate tends to be above or below the median for many bins in a row before switching state. The one free parameter in this procedure, the 25-ms sampling period, was chosen to be long with respect to the refractory period.

Two potential limitations of the streak index warrant explicit mention. First, the 25-ms sampling period causes firing rate fluctuations over 20 Hz (the Nyquist limit) to be aliased to lower frequencies and could thus lead to artifactually low streak indices. Firing rates during very brief epochs are notoriously difficult to assess, and we can make no firm assertions about their presence or absence in our data set. For the purposes of this analysis, therefore we assume that firing rate transitions occur predominantly below 20 Hz and note that our conclusions are conditional on this assumption. Second, the timing of firing rate transitions across trials affects the streak index: randomly timed, rare firing rate transitions will drive the streak index below zero, whereas transitions that are consistent across trials (irrespective of frequency) will not. Thus the streak index can be influenced both by the frequency of firing rate transitions as well as by the timing of




FIG. 3. Streak index calculation. See METHODS for details.

¹ If an individual spike count equaled the median exactly, we randomly converted it to a one with probability p and a zero with probability $(1 - p)$. The value of p was calculated as

$$p = \frac{1}{2} \cdot \frac{[\# \text{ equal} - (\# \text{ greater} - \# \text{ less})]}{\# \text{ equal}}$$

which ensures that roughly equal numbers of ones and zeros are expected in each time bin.

these transitions across trials and should not be interpreted as reflecting of either one alone.

RESULTS

Behavioral performance

Figure 4A shows psychometric data averaged across experimental sessions for each of the three monkeys. Several aspects of these data indicate that the monkeys' behavior was well controlled during data collection. First, the percentage of correct choices increases monotonically with stimulus coherence for each monkey, forming a sigmoidal relationship. Second, at the highest stimulus coherence employed, 51.2% coherently moving dots, all of the monkeys performed the task nearly perfectly. Finally, while monkeys differed in psychometric threshold, as evidenced by the position of the curves on the coherence axis, they had similar sensitivities, as indicated by the similarity in slope. Lack of motivation is generally reflected in reduction of percent correct responses at high coherences and in psychometric functions of shallow slope.

Figure 4B shows a histogram of psychophysical thresholds, on an experiment-by-experiment basis, for the three monkeys combined. An ANOVA on the log-transformed thresholds confirmed that threshold differed significantly across the three monkeys ($P < 0.0001$). This was primarily due to *monkey T*, whose geometric mean threshold (19.25% coherence) was considerably higher than either *monkey E* (10.98) or *monkey D* (11.29).

Cell selection

We isolated 704 SC neurons from the superior colliculi of three monkeys (*monkey E*: 222, *monkey T*: 237, *monkey D*:

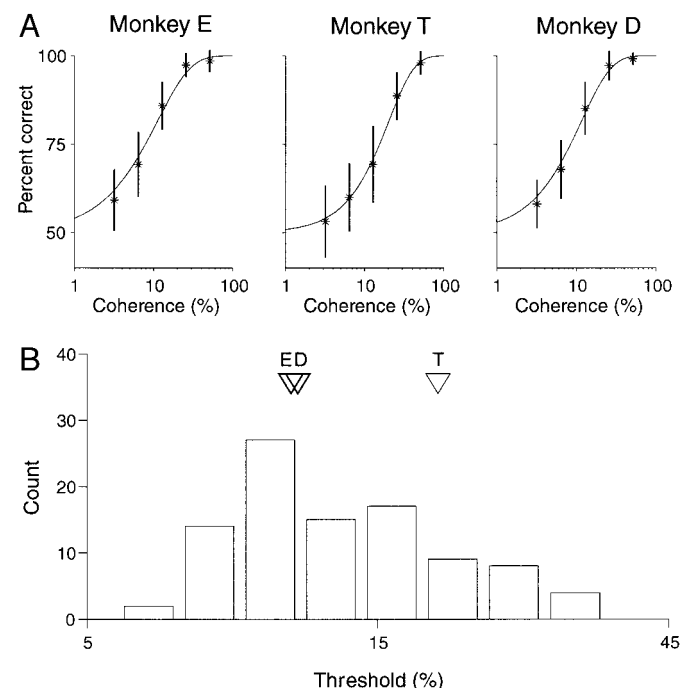


FIG. 4. Behavioral performance summary. *A*: percent correct choices across experimental sessions for each of the 3 monkeys. *, means; vertical lines extend ± 1 SD. The curves are maximum likelihood fits of cumulative Weibull functions. *B*: histogram of the psychophysical threshold values estimated from each monkey's psychophysical performance during each recording session. ∇ , geometric mean thresholds for monkeys *E*, *D*, and *T*.

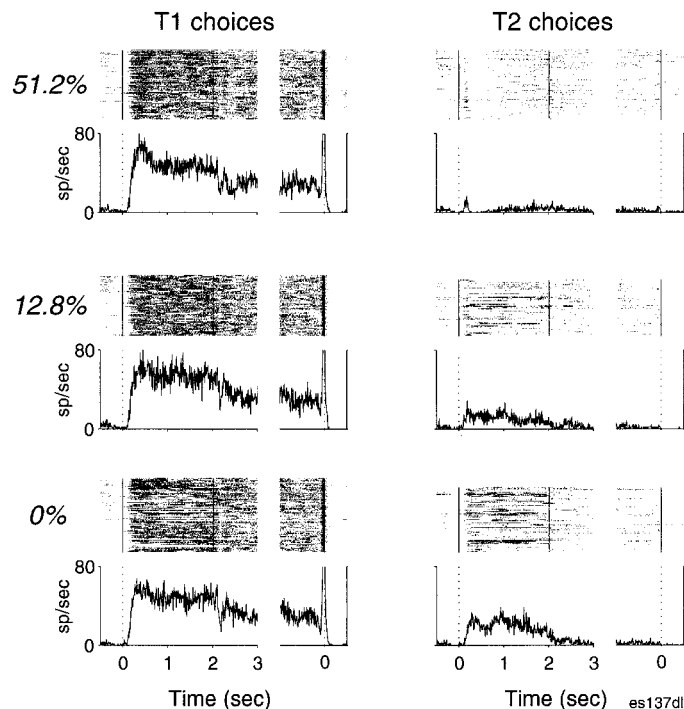


FIG. 5. Responses of a single SC neuron during performance of the direction-discrimination task. Trials are aligned to the onset of the motion stimulus and to saccade initiation. Time is expressed relative to the corresponding alignment event. Stimulus onsets, stimulus offsets, and saccade initiations are marked by vertical bars in the rasters. *Left* and *right*: responses preceding T1 and T2 choices are displayed, respectively. Data from 3 coherence levels are shown: 51.2% (top), 12.8% (middle), and 0% (bottom).

245) while they performed the direction-discrimination task described in the preceding text. We qualitatively assessed each cell's firing rate preceding saccades to T1 and T2. If prelude firing (after onset of the random dots but well before saccade execution) seemed to covary with ("predict") the target choice, we studied the neuron quantitatively. Complete direction-discrimination data sets (≥ 5 coherence levels and ≥ 30 trials per coherence level) were acquired for 96 cells that had significantly higher firing rates preceding saccades to T1 than saccades to T2 (*monkey E*: 33, *monkey T*: 29, *monkey D*: 34).² This ratio, 96/704, should not be taken as an estimate of the proportion of choice predictive neurons we encountered in the SC because many neurons were either lost part way through the experimental session or were studied with different experimental protocols. We estimate the actual proportion to be closer to 1/3.

Target-specific preludes during direction discrimination

Figure 5 shows the responses of a single SC neuron recorded from *monkey E* during performance of the direction-discrimi-

² For each cell, we counted spikes during the presentation of the motion stimulus and the first second of the delay period. The counts were normalized within each stimulus coherence and compiled into two distributions according to the monkey's choice in the discrimination task (T1 or T2). For each cell, the difference between these two distributions was evaluated by a Mann-Whitney U test with a criterion level of $P < 0.01$. By this criterion, 103 of 127 cells exhibited predictive activity. Seven of these cells were significantly more active preceding saccades to T2 than to T1 and were therefore eliminated from further analysis (for a further description, see *Cells with reversed preludes*). Thus our final database consisted of 96 choice-predicting SC neurons.

nation task. Data from three coherence levels are shown: a supra-threshold coherence (51.2%), a near-threshold coherence (12.8%), and a subthreshold coherence (0%). Within each panel, all trials have been aligned to the onset of the visual motion stimulus (*left*) and to saccade initiation (*right*). The prestimulus firing rate of this cell was a modest 1.5 spikes/s. Approximately 120 ms after the onset of the visual stimulus, the firing rate increased dramatically on trials that ended in a T1 choice (*left*). On trials ending in T2 choices (*right*), the firing rate either did not change or increased only modestly over the baseline rate. The high-frequency motor burst immediately preceding T1-directed saccades was typical of most neurons we studied.

The coherence of the motion stimulus influenced the responses of this cell, particularly during the stimulus presentation interval preceding T2 choices. On trials in which a high coherence stimulus led to a T2 choice, the cell discharged only weakly (*top right*). In contrast, the cell fired moderately when the coherence of the stimulus was low (*bottom right*). The firing rate was extremely variable under some conditions, resulting in “streaky” rasters with periods of high-frequency discharge intermixed with periods of low-frequency discharge (e.g., *bottom right*). We will consider this variability in detail in *Firing rate dynamics*.

Figure 6A shows the average firing rate of this cell as a set of superimposed peristimulus time histograms (PSTHs). For visual clarity, only the three coherence levels shown in Fig. 5 are illustrated, although six were used in the experiment. Solid curves show the activity measured on T1 choice trials, and dashed curves show the activity on T2 choice trials. This cell is “choice-predictive” because its firing rate early in the trial reveals the target that the monkey will choose at the end of the trial. The histograms confirm the impression from Fig. 5 that the firing rate of this cell varied with coherence for T2, but not T1, choices.

The target-specific prelude activity shown in Fig. 6A permits an experimenter to predict which target the monkey will select and, by extension, the outcome of the monkey’s perceptual decision process. As described in METHODS, we used techniques derived from signal detection theory to compute a “predictive activity” metric that reflects how well an ideal observer could predict the monkey’s decision based on the differential activity of a single SC neuron on T1 and T2 choice trials. Figure 6B illustrates predictive activity for the cell in Fig. 6A. Prior to appearance of the motion stimulus, predictive activity is close to a value of 0.5 for all coherences, indicating no systematic relationship between firing rate and the monkey’s eventual choice (i.e., random performance for the ideal observer). Within a few hundred milliseconds of stimulus onset, however, predictive activity rises sharply, remaining roughly constant during the delay period and returning to chance levels only after the saccadic eye movement that ends the trial. Throughout most of the trial, therefore, the firing rate of this neuron predicts the choice that the monkey will ultimately express.

Predictive activity during the visual stimulus period varied strongly with motion coherence, as we would expect given the differential effects of motion coherence on neural activity on T1 and T2 choice trials (Fig. 6A). For high-coherence stimuli, predictive activity develops rapidly, reaching levels near unity, the theoretical maximum. Predictive activity during the visual stimulus interval is weaker when the coherence is lower. The

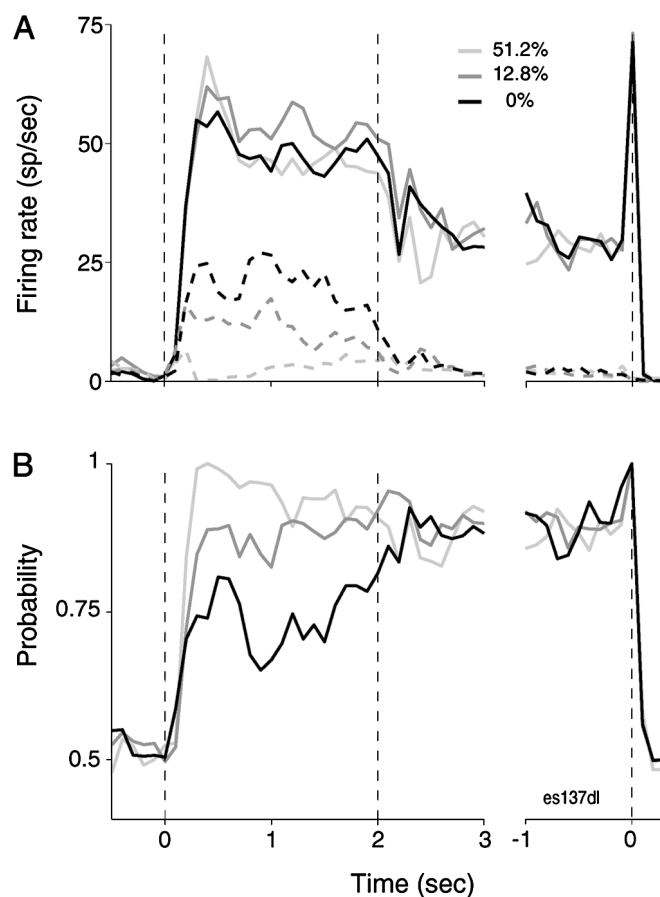


FIG. 6. Average firing rate (A) and predictive activity (B) computed from the rasters in Fig. 5. Trials have been aligned both on the stimulus presentation (*left*) and saccade initiation (*right*). Vertical dashed lines indicate the times of stimulus onset, stimulus offset, and saccade initiation. Gray-level corresponds to stimulus coherence (black: 0%, dark gray: 12.8%, light gray: 51.2%). Solid and dashed curves in A illustrate responses preceding correct T1 and T2 choices, respectively.

activity of this neuron thus tends to predict the monkey’s choice more quickly and accurately when the stimulus motion is strong than when it is weak.

The predictive responses shown in Figs. 5 and 6 were among the strongest we recorded. Figure 7 shows spike rasters from a cell that exhibited weak prelude activity. Despite its relatively subdued discharge, this neuron, like the one in Figs. 5 and 6, fired more spikes preceding T1 choices than T2 choices, thus qualifying as a choice-predictive neuron. Superimposed PSTHs for this cell appear in Fig. 8A. The average firing rates are somewhat noisy due to the low overall level of responsiveness and the modest number of trials collected per condition.

Motion coherence influenced the activity of this neuron but differently than for the cell illustrated in Figs. 5 and 6. Activity varied inversely with coherence for *both* T1 and T2 choice trials; only T2 choice trials exhibited this effect in Figs. 5 and 6. For the neuron in Figs. 7 and 8, the influence of coherence can be described as a modulation of overall response gain: activity increases as stimulus coherence decreases, irrespective of the monkey’s choice. Despite its clear effect on firing rate, motion coherence exerted no influence at all on predictive activity (Fig. 8B). This outcome is expected from the data in Fig. 8A: predictive activity is a *differential* measure of re-

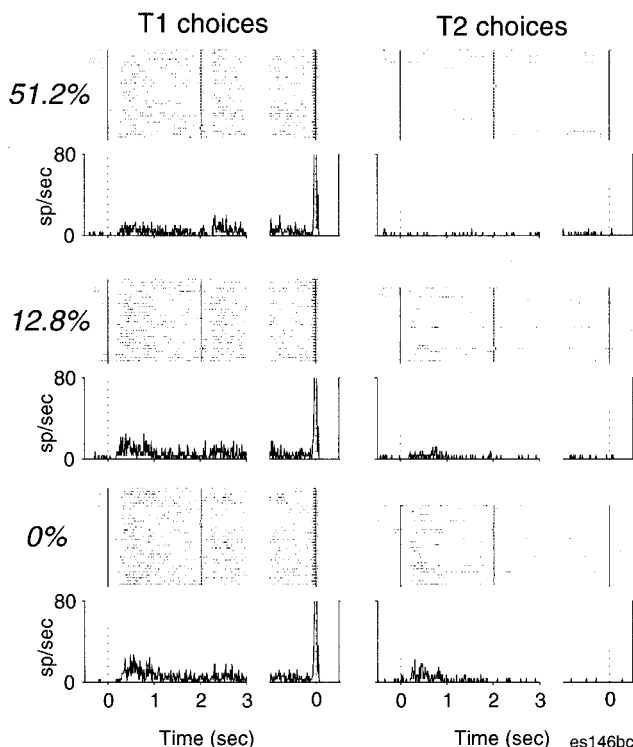


FIG. 7. Responses of a single SC neuron during performance of the direction-discrimination task. Conventions are as in Fig. 5.

sponses on T1 and T2 choice trials, but coherence affected T1 and T2 choice trials similarly. If the SC was populated only with neurons like this one, predicting the monkey's perceptual decision and eventual saccadic eye movement would be equally easy (or difficult) irrespective of the motion coherence.

In the preceding paper, we showed that choice-predicting prelude neurons in our monkeys can be divided into two subpopulations based on the presence or absence of direction-selective visual inputs revealed during a passive-fixation task. Direction-selective cells responded significantly more to visual motion flowing toward, than away from, their movement fields (permutation tests: $P < 0.05$) (see Fig. 13 of Horwitz and Newsome 2001). The cell shown in Figs. 5 and 6, for example, is a member of the direction-selective population, whereas the cell shown in Figs. 7 and 8 is a non-direction-selective cell (data not shown).

To visualize the evolution of predictive activity across each population, we pooled and analyzed the normalized spike counts (see METHODS) from 44 direction-selective cells and from 52 non-direction-selective cells. Six predictive activity curves, one for each motion coherence level tested, appear in both panels of Fig. 9. Recall that each predictive activity curve is derived from firing rates preceding both T1 and T2 choices. Across the population of direction-selective cells (Fig. 9A), predictive activity follows a pattern similar to the data shown for the single neuron in Fig. 6B. Predictive activity evolves rapidly and attains higher levels for high-coherence than for low-coherence motion stimuli. In contrast, predictive activity across the population of non-direction-selective cells (Fig. 9B) is more similar to the single neuron example shown in Fig. 8B. Non-direction-selective neurons exhibit a modest level of predictive activity that is not strongly influenced by the coherence of the motion stimulus.

We compared predictive activity in the two neuronal populations via permutation tests (randomly shuffling cells 5,000 times between groups to assess statistical significance). Specifically, we asked whether predictive activity in the two populations differed in latency, time course, and magnitude. For these analyses, we combined responses across coherence levels and calculated a single predictive activity curve for each cell population (see METHODS). We then measured the latency (time point at which predictive activity first exceeded the baseline level by 3 SD), time course (time to reach half of the maximum predictive activity during the stimulus presentation), and magnitude (average level from stimulus onset through the 1st second of the delay period) for each population. Predictive activity was calculated in 100-ms-wide bins for all analyses except for the latency analysis in which we used 10-ms-wide bins to improve temporal resolution. Predictive activity across the direction-selective population developed at a shorter latency ($P < 0.025$), evolved with a more rapid time course ($P < 0.01$), and was higher on average ($P < 0.01$) than predictive activity across the non-direction-selective population.

Finally, we assessed the statistical significance of the effect of coherence on predictive activity. We calculated the correlation between stimulus coherence (transformed to ranks ranging from 1 to 6) and predictive activity (over the stimulus presentation) for each cell. Mean correlation coefficients for direction-selective cells ($r = 0.58$) and non-direction-selective cells ($r = 0.18$) were significantly greater than zero (t -tests: $P < 0.0001$ and $P < 0.05$, respectively). The difference be-

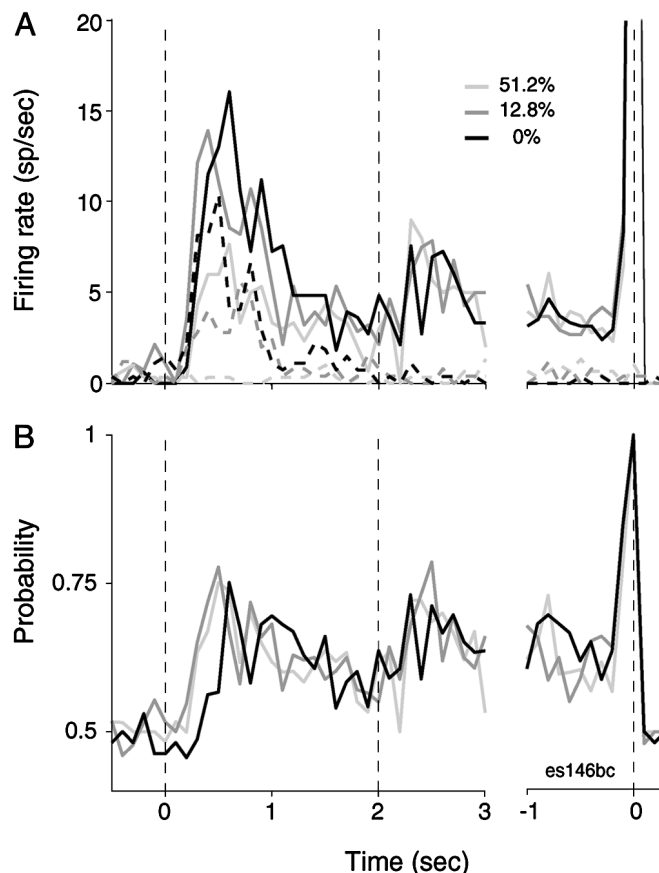


FIG. 8. Average responses and predictive activity computed from the rasters in Fig. 7. Conventions are as in Fig. 6.

tween these groups was also significant (t -test: $P < 0.0001$). Thus while predictive activity in both populations varies with stimulus coherence, this effect is greater in the direction-selective population.

The relationship between coherence and predictive activity was qualitatively similar across the three monkeys in our study, but the relationship between coherence and prelude firing rates varied across monkeys. The three panels in Fig. 10 depict data for each monkey, averaged across all neurons recorded from that monkey. Responses are aligned both to stimulus onset (*left*) and to saccade initiation (*right*). Solid and dashed lines illustrate firing rates preceding T1 and T2 choices, respectively; gray-level indicates stimulus coherence. We combined data from direction- and non-direction-selective cells in this analysis because both populations exhibited similar relationships between stimulus coherence and prelude firing rate.

For T2 choices, average firing rate varied inversely with coherence for all three monkeys. For T1 choices, however, the influence of coherence on firing rate differed among the animals. For *monkey E* (*top*), and to a lesser extent *monkey D* (*middle*), average firing rate varied inversely with coherence for T1 choices as well as for T2 choices, as in the single neuron data of Figs. 7 and 8. Data from these two monkeys suggest a strong inverse effect of coherence on response gain. For *monkey T* (*bottom*), however, average firing rates on T1 trials were positively correlated with coherence (in the 1st second of the stimulus presentation), following a pattern similar to that of the single neuron of Figs. 5 and 6. These differences were not

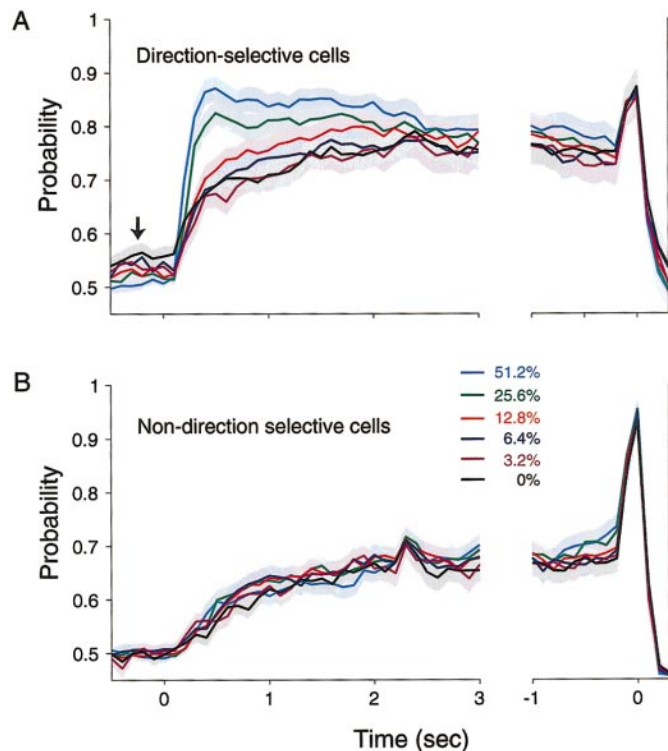


FIG. 9. Predictive activity for the population of direction-selective cells (*A*) and non-direction-selective cells (*B*) for each stimulus coherence. All trials were aligned on the stimulus onset (*left*) and saccade initiation (*right*). Standard errors were estimated by bootstrap (2,000 resamples per point) and are indicated *above* and *below* the predictive activity values by extent of the light-colored swaths. Direction-selective, but not non-direction-selective, cells predict choices weakly before the onset of the stimulus (arrow).

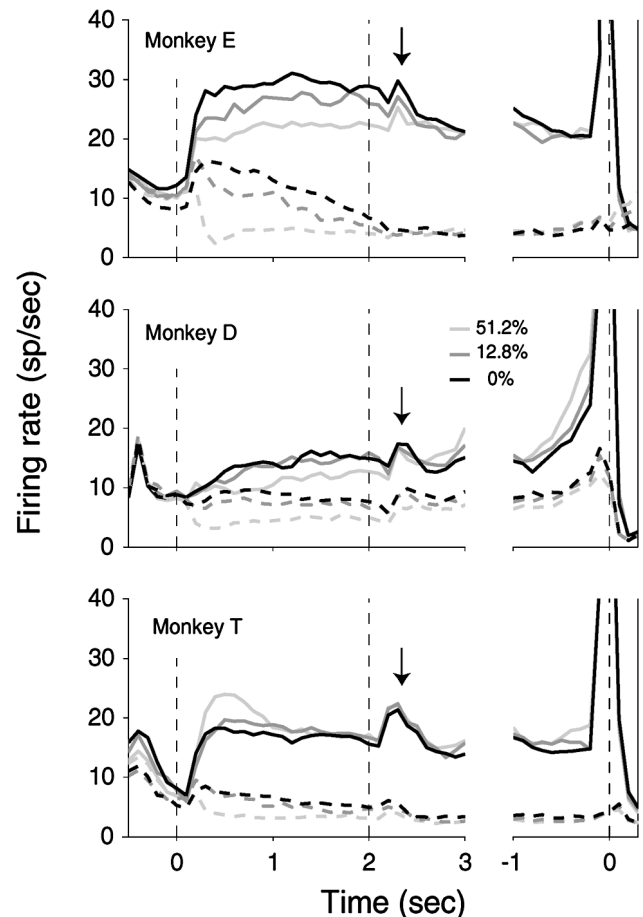


FIG. 10. Population average responses preceding T1 and T2 choices. Responses were averaged across trials and then across neurons. Conventions are as in Fig. 6*A*. Arrows indicate the transient increase in activity associated with stimulus offset.

easily accounted for by difference in proportion of direction- and non-direction-selective cells.

In all three monkeys, some neurons exhibited a transient increase in firing rate ~ 200 ms after the disappearance of the motion stimulus (arrows, Fig. 10). We considered the possibility that this response might be related to small saccades confined to the fixation window. The frequency of fixational saccades, however, did not increase near the time of the transient discharge. Saccade frequency *decreased* briefly ~ 200 ms after the stimulus presentation in some experiments; but this seems unlikely to account for the transient. In several experiments, we found robust transient discharges without concomitant changes in saccade frequency of either polarity.

Choice bias

Surprisingly, predictive activity in direction-selective cells is significantly greater than 0.5 even *before* the onset of the motion stimulus on 0% coherence trials (Fig. 9*A*, arrow; permutation tests: $P < 0.01$ for all time bins). In other words, the neuronal discharge at the very beginning of the trial is slightly higher on trials that end in T1 choices than on trials that end in T2 choices. This was not true in the non-direction-selective population ($P > 0.1$ for all time bins).

The existence of this early predictive activity does not mean




FIG. 11. Peristimulus time histograms for a single neuron showing bias signals. Time 0 is stimulus onset. Responses are shown preceding T1 choices (thick lines) and T2 choices (thin lines) at both 51.2% coherence (A) and 0% coherence (B). At 0% coherence, target choice varies with the prestimulus neuronal response.

that the monkey is precognizant of the direction of impending motion stimuli, which are randomly chosen from trial to trial. Rather we propose that the monkey enters some trials with a bias to one of the two targets, and that this choice bias is manifest in the discharge of collicular neurons. Naturally, this bias is *not* predictive of the actual direction of motion in the upcoming trial and is overridden so long as the motion signal is sufficiently strong. For high coherences, therefore, prestimulus activity does not correlate with the direction judgment expressed at the end of the trial. On trials in which the motion signal is weak, however, the monkey may ultimately choose the target to which it was originally predisposed. In this case, prestimulus activity becomes correlated with the monkey's choice by virtue of the fact that the motion stimulus provides no additional information to override the initial bias.

For *monkey T*, these bias-related signals were strong enough to be detectable within single experiments. Figure 11 shows the firing rate of a single SC neuron in this animal during the 1-s interval spanning onset of the motion stimulus. A shows responses during high-coherence (51.2%) trials, and B shows responses during low-coherence (0%) trials. The thick and thin curves in each panel were calculated from trials ending in T1 and T2 choices, respectively. In the 51.2% coherence condition, the two curves are largely overlapping until ~150 ms after the onset of the visual stimulus, indicating that the average firing rate during the prestimulus period is not related to the target choice. In the 0% coherence condition, on the other

hand, the firing rate preceding T1 choices exceeds the firing rate preceding T2 choices over this entire interval.

Figure 12 shows the magnitude of bias-related activity for each animal. Bias-related activity (quantified by our "predictive activity" metric—see METHODS) was calculated from normalized responses over the 500 ms preceding presentation of a 0% coherence stimulus. Values significantly >0.5 are indicated (*; permutation tests: $P < 0.05$). Significant bias-related signals occurred in *monkeys E* and *T*, but not in *monkey D*. For both *monkeys E* and *T*, direction-selective neurons carried significantly greater bias-related signals than non-direction-selective neurons (permutation tests: $P < 0.05$). Interestingly, *monkey T*, the animal with the strongest bias signals, was also the animal with the highest psychophysical thresholds, as would be expected if internal variables exert greater influence on this animal's choices.

Influence of variation in saccade parameters

In our task, each monkey must make accurate, target-directed saccades to obtain rewards irrespective of stimulus coherence. We thus expect that motor signals preceding saccades to a particular target should be substantially independent of stimulus coherence. Recall that our monkeys make saccades only after a delay period of 1–1.5 s and should therefore be relatively immune to the well known effect of task difficulty on response latency as measured in reaction time tasks. Nevertheless, even saccades made to single targets in our task vary slightly from trial to trial in latency, end point, and velocity. We therefore considered the possibility that the effects of motion coherence on prelude activity might be a neural correlate of subtle variations in motor output. This possibility is credible because saccade latency, end point, and velocity have been shown to covary with the discharge of SC neurons (Dorris and Munoz 1998; Rohrer et al. 1987; Sparks et al. 1976; Wurtz and Goldberg 1972). We approached this problem by imple-




FIG. 12. Predictive activity during the 500 ms preceding stimulus presentation. Each monkey was considered separately. *, predictive activity values that were significantly >0.5 (by permutation test, $P < 0.05$). Error bars indicate SEs estimated by bootstrap (2,000 resamples).

menting linear regression models (see METHODS) and testing hypotheses with partial F tests. This technique determines how much of the variance in the neuronal response can be accounted for by saccade parameters and then asks how much additional variance can be accounted for by stimulus coherence.

Regressions were performed for each cell individually using trials that ended in correct T1 choices only (i.e., rewarded saccades into the movement field). We calculated firing rates in five different epochs: the first 1 s of the stimulus presentation, the second 1 s of the stimulus presentation, the first 500 ms of the delay period, the last 500 ms of the delay period, and a perisaccadic interval defined as 50 ms before saccade initiation until 25 ms after it. For each of these epochs, we regressed firing rate onto the measured saccade parameters. The proportion of regressions that attained statistical significance ($P < 0.05$) is illustrated by the dark bars in Fig. 13. Few cells (6/96) yielded significant regressions for firing rates calculated during the first 1 s of the stimulus presentation, indicating little if any relationship between firing rate and saccade parameters early in the trial. For spikes occurring during the early delay interval, however, a greater proportion of regressions achieved significance (11/96). The greatest proportion of significant regressions occurred for spikes counted during the perisaccadic interval (30/96). Thus the relationship between neuronal discharge and saccade parameters becomes more pronounced as the time of the saccade approaches.

For each regression, we then tested whether incorporating stimulus coherence as an added predictor significantly improved the model fit. The proportion of cells for which this was the case is shown by the light gray bars in Fig. 13 alongside the results of the original regressions. For the early stimulus pre-

sensation interval (Stim1), the inclusion of coherence was significant many times more often than expected by chance (33/96). Thus the effect of coherence on the firing rate early in the trial cannot be accounted for by small parametric differences in the saccadic eye movements made on each trial. Coherence influenced the firing rate progressively less later in the trial. During the later part of the delay and at saccade initiation, stimulus coherence does not account for any detectable additional variance.

Error trials

When the monkey is performing the direction-discrimination task correctly, the firing rate of collicular prelude neurons varies in accordance with the monkey's choice. On correctly answered trials, of course, the target choice and the direction of the stimulus motion are perfectly correlated. Thus it is not possible to determine whether the neural response is more closely related to the target choice or the stimulus direction. On error trials, however, the direction of motion and the monkey's choice are opposed. A comparison of neural responses between correct and error trials revealed that target choice exerts more influence on the activity of SC cells than does motion direction.

Each panel in Fig. 14 displays average firing rates during correctly and incorrectly answered trials (solid and dashed, respectively) which ended in a T1 choice (black) or a T2 choice (gray). Because reasonable numbers of errors were made only when the stimulus coherence was relatively low, we have restricted our attention to coherences of 3.2, 6.4, and 12.8% in this analysis.

High firing rates preceded T1 choices and low firing rates preceded T2 choices, for all three monkeys, irrespective of whether the choice was correct or not. Thus the solid and dashed lines of a common gray-level tend to lie near each other throughout the trial. For solid and dashed lines of a common gray-level, stimulus motion is in opposite directions, but the saccade is made to the same target. On average, therefore SC firing rates are more closely related to the target that the monkey selects than to the direction of stimulus motion that instructed the choice. This was true for both the direction- and non-direction-selective cells.

Firing rate dynamics

Figures 9 and 10 show that firing rate and predictive activity, averaged across many trials and many neurons, evolve smoothly over time. For analyses of time course, however, averaged data can be deceptive. The gradual increase in average predictive activity could indeed reflect the fact that firing rates change smoothly over time *during* individual trials and that this pattern is consistent *across* trials. Alternatively, however, the firing rate on individual trials could change abruptly but at different times on different trials. Both scenarios could lead to the same average data. Visual inspection of individual rasters revealed that firing rates were highly variable for some SC neurons. In Fig. 5, for example, neural activity appears to jump between low and high firing rate states, particularly during low-coherence trials.

To examine this issue quantitatively, we calculated a "streak index" (see METHODS) that reflects the number of times that the firing rate of a neuron changed from a "high" level to a "low"




FIG. 13. Regression summary barchart. Bar height corresponds to the proportion of significant regressions ($P < 0.05$). Asterisks indicate bars whose height is significantly greater than expected by chance ($*P < 0.01$, $**P < 0.0001$). Spikes were counted during 5 temporal epochs spanning the trial (see text). For regressions of firing rate on saccade parameters (dark gray), significance was determined by standard F tests. The proportions of significant regressions (from left to right) are: 0.06, 0.08, 0.11, 0.17, and 0.31. Partial F tests were used to assess the additional contribution of stimulus coherence (light gray). The proportions of significant partial F tests (from left to right) are: 0.34, 0.29, 0.21, 0.10, and 0.07.




FIG. 14. Average neuronal responses preceding T1 choices (black) and T2 choices (gray) on correct (solid) and error (dashed) trials. Trials are aligned on stimulus onset (*left*) and saccade initiation (*right*). Responses are more closely related to the target choice than to the direction of motion in the stimulus.

level (or vice versa), where high and low are defined relative to the median spike count observed in each time bin. A hypothetical neuron that discharges spikes in accordance with a Poisson process has a 50–50 chance of firing more or fewer spikes than the median in any given time bin, irrespective of the spike counts in other bins. This is true irrespective of how the firing rate modulates over time provided that the firing rate modulations are the same across trials (i.e., the calculation is equally valid for an inhomogeneous Poisson process). For the hypothetical Poisson neuron, the concatenated thresholded firing rates (see METHODS) can be thought of as the outcomes of a series of independent coin flips.³ The asymptotic distribution of the streak index, under this hypothesis, is normal with a mean of 0 and a variance of 1 so that ~95% of the observations should lie in the range ± 2 .

For each neuron in our database, we calculated streak indices during the presentation of 0% coherence stimuli and 51.2%

coherence stimuli for trials ending in T1 choices. Figure 15 illustrates several example spike rasters and their associated streak indices. Negative indices indicate that periods of high- or low-frequency discharge are longer than expected under the Poisson model, whereas positive indices indicate that such periods are shorter than expected (see METHODS). The rasters in the *top panels* had streak indices that were among the most negative we calculated and, correspondingly, reflected the clearest variations in firing rate across trials. These responses appeared previously in Fig. 5. Rasters in the *bottom panels* yielded positive streak indices due to remarkably regular interspike intervals and moderate firing rate. All but one of the streak indices in Fig. 15 lie outside of the range of ± 2 , so we can reject the Poisson firing hypothesis for these rasters with 95% confidence. Indeed, 74% of the indices calculated across our entire data set lay outside the range ± 2 . The observed streakiness in neuronal discharge is thus not simply a product of Poisson randomness. Streak indices tended to be negative for both coherence levels (*t*-tests: $P < 0.0001$), indicating that the neurons in our database tend to fire above or below the median rate for several bins in a row before switching states. Note that the streak indices would tend to be positive if the only departure from the Poisson model were that imposed by a neuron's refractory period.

An intriguing possibility is that the abrupt changes in firing rate may correspond to changes in an internal decision variable. For example, the monkey may waver between the alter-

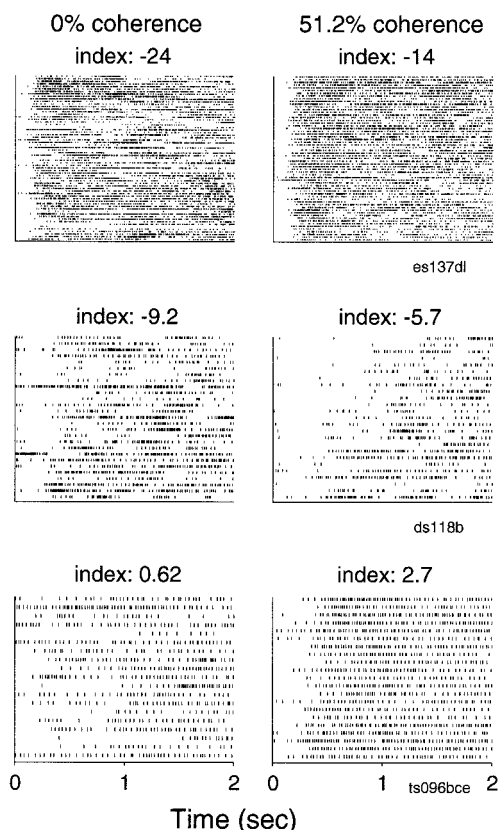


FIG. 15. Example spike rasters and corresponding streak indices. Responses preceding T1 choices at 0 and 51.2% coherence are shown for 3 neurons. Streak indices range from large negative values (indicating extremely variable firing rates) to modest positive values (indicating very regular firing rates).

³ The fact that we estimate the median firing rate from the data complicates matters slightly. As a consequence of this fact, the number of ones and zeros in a column (Fig. 3, middle panel) are not independent, which undermines the coin flip analogy. However, because many trials typically contribute to the calculation, the dependence is negligible.

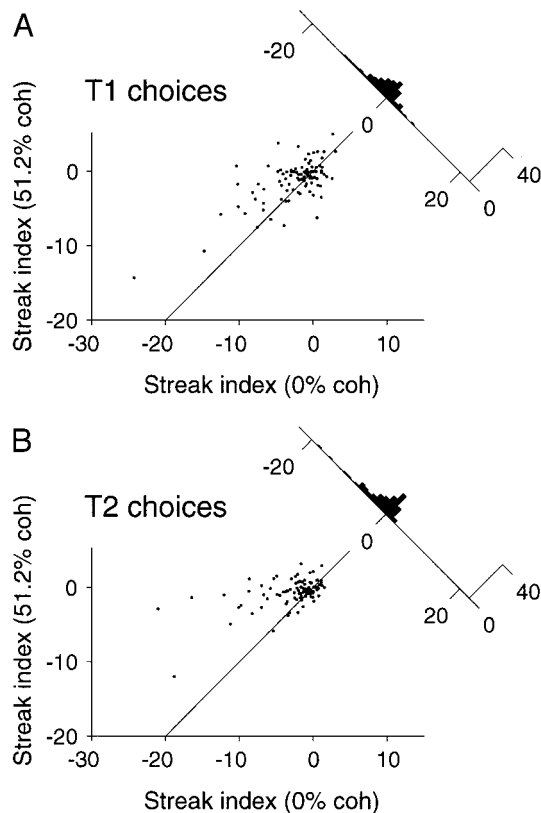


FIG. 16. Scatterplots of streak indices calculated from 51.2 and 0% coherence trials. Histograms show the distributions of streak index differences between stimulus coherence conditions. Indices calculated from correct T1 choice trials and correct T2 choice trials are presented in A and B, respectively.

native choices (T1 or T2) during individual trials, and this vacillation may correspond to abrupt changes in prelude discharge (such states of indecision are certainly familiar to experienced psychophysical observers). Consistent with the data, this hypothesis predicts that streak indices should be predominantly negative because the length of time during which the monkey “leans” toward one choice or the other is presumably long with respect to the 25-ms time window used in this analysis and because choice vacillations presumably occur at different times on different trials. Interestingly, this hypothesis also predicts that the streak index should vary with stimulus coherence. On low-coherence trials, choice vacillations are likely to be more common, resulting in firing rate changes that vary within single trials and are unsynchronized across trials. Both of these response dynamics tend to drive the streak index negative. Choice vacillations on high coherence trials are presumably uncommon, so we expect relatively consistent firing rates on these trials and streak indices that are more positive.

To test this prediction, we compared streak indices at 0 and 51.2% coherence. T1 and T2 choices were analyzed separately. The scatterplot in Fig. 16 shows that streak indices indeed tended to be lower at 0% coherence than at 51.2% coherence for both T1 and T2 choices, in agreement with the prediction. The difference in streak index between coherence levels (histograms) was statistically significant for both T1 and T2 choices (t -tests: $P < 0.0001$).

In the preceding paper, we reported that small amplitude saccades within the fixation window modulate the activity of some choice predictive SC neurons (Fig. 15 of Horwitz and

Newsome 2001). We considered the possibility that such eye movements might contribute to the streaks we observed in the spike rasters. Cross-correlation analysis revealed significant coincidence of saccade occurrences and streak terminations (transitions from high to low firing rates) in 30 of 384 rasters (Fisher exact tests: $P < 0.05$). Streak indices for these rasters, however, were unremarkable and did not differ statistically from the 354 rasters lacking significant coincidences (unpaired t -tests: $P > 0.05$). While saccades within the fixation window modulate the ongoing firing rate of some neurons, this effect cannot account for the streaks in the majority of spike rasters.

Cells with reversed preludes

Some SC neurons exhibited higher frequency preludes preceding saccades to T2, the target outside of the movement field, than to T1, the target inside the movement field. An example of a neuron with such a “reversed” prelude appears in Fig. 17. This cell, like most SC neurons, fired a brief burst of action potentials preceding saccades to T1 in accordance with its position in the collicular map. On the other hand, prelude activity preceding T2 choices actually exceeded prelude activity preceding T1 choices. Rasters and PSTHs in Fig. 17A are

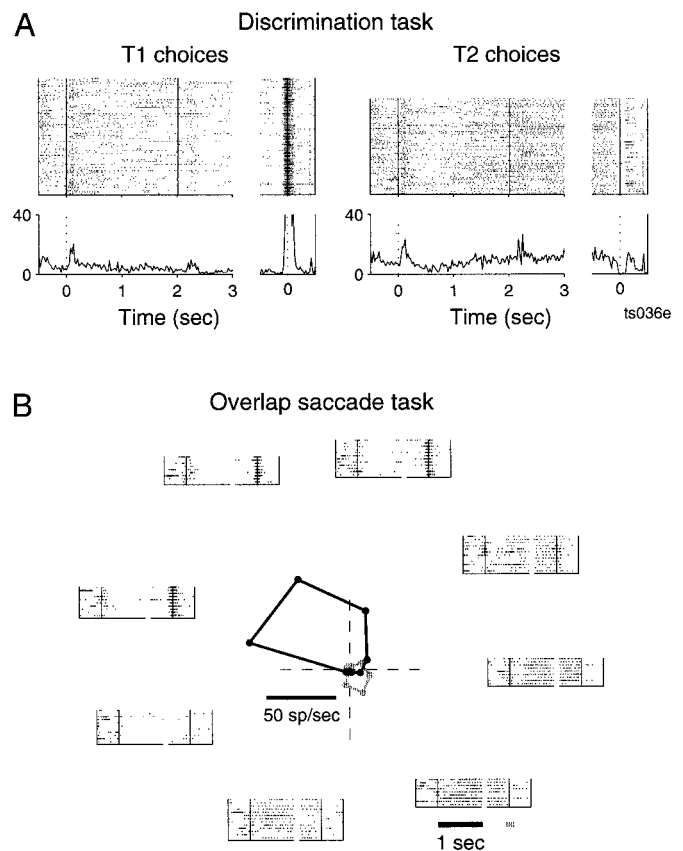


FIG. 17. Responses of a single SC neuron with reversed prelude activity. A: responses preceding T1 (left) and T2 (right) choices in the direction-discrimination task. Rasters have been aligned both to stimulus onset and saccade initiation. B: firing rate as a function of saccade direction in the delayed saccade task. Firing rates were calculated during the first second of the delay period (black) and perisaccadic period (gray); these epochs are shown as horizontal bars below the bottom-right raster. Rasters, showing the raw response for each saccade direction, are aligned both to target onset and saccade initiation.

aligned both to stimulus onset as well as to the saccade initiation to illustrate the difference between prelude and the burst discharges.

We also recorded from this cell while the monkey performed a delayed saccade task. Eight possible target locations lay on a circle around the fixation point; the rasters in Fig. 17B show neural responses obtained for each location. The cell exhibited reversed prelude activity in this task as well. We calculated the mean firing rate of this cell during two epochs: a prelude interval and a perisaccadic interval. The former was defined as the first 1 s following target presentation and the latter as 50 ms before until 25 ms after saccade initiation. The polar plots at the center of this figure show the mean firing rate during these two epochs as a function of saccade direction. The cell exhibited the highest frequency prelude when the target appeared down and to the right of the fovea. The greatest peri-saccadic discharge, on the other hand, accompanied movements to the upper left. Neurons with this type of “reversed prelude” activity have been reported in the frontal eye fields (Friedman et al. 1998), but we are unaware of any previous reports documenting their existence in the SC.

We defined a prelude index contrasting the mean firing rates preceding T1 saccades and T2 saccades: $(T1 - T2)/(T1 + T2)$. This metric assumes positive values if the prelude preceding T1 saccades exceeds the prelude preceding T2 saccades, negative values if the reverse is true, and is constrained to lie between 1 and -1 . For this analysis, we counted spikes from the onset of the motion stimulus (discrimination trials) or target onset (delayed saccade trials) until fixation point offset. Figure 18 shows a comparison of prelude indices for the 18 cells that were tested in both the delayed saccade and direction-discrimination tasks and had significant reversed prelude activity (1-tailed t -test, $P < 0.05$) in at least one. Prelude indices in the two tasks were significantly correlated, that is, neurons with strongly reversed preludes in one task tend to have strongly reversed preludes in the other task as well ($r = 0.46$, $P < 0.05$).




FIG. 18. Scatterplot of prelude indices calculated from the delayed saccade task (abscissa) and discrimination task (ordinate). Firing rates were calculated from target onset (delayed saccade task) or stimulus onset (discrimination task) until fixation point offset.

Additionally, preludes tended to be more strongly reversed in the delayed saccade task than in the discrimination task (paired t -test, $P < 0.0001$).

DISCUSSION

We investigated the activity of intermediate and deep layer SC neurons while monkeys performed a two-alternative, forced-choice direction-discrimination task. Approximately one-third of the cells that we screened exhibited target-specific preludes of activity, allowing an experimenter to “predict” target choices well in advance of the operant saccade. For some cells, predictive activity was positively correlated with motion coherence, consistent with the notion that these neurons bear a signature of the visual stimulus that instructed the perceptual decision. This finding might be explained trivially, however, if saccade metrics varied systematically with stimulus coherence. In this case, purely motor cells could easily exhibit an influence of stimulus coherence. Regression analyses demonstrated that variation in saccade metrics does not account for our results. The effect of stimulus coherence on neuronal firing rate early in the trial was statistically significant even after accounting for differences in saccade metrics, showing that the coherence effect is related to the strength of the sensory evidence that instructs the saccade.

Relationship between firing rate and stimulus coherence

The relationship between stimulus coherence and predictive activity was qualitatively similar across all three monkeys. However, substantial inter-monkey differences were observed in the relationship between coherence and firing rate (Fig. 10). In *monkey E* and to a lesser extent in *monkey D*, firing rate during the visual stimulus presentation was negatively correlated with coherence; this was true preceding a saccade to either target. In *monkey T*, this relationship was reversed for T1 choices: firing rate was positively correlated with coherence.

This difference among monkeys may reflect idiosyncrasies in task strategy. In *monkeys E* and *D*, for example, low-coherence stimuli may engage arousal mechanisms that increase activity throughout the SC. Interestingly, *monkeys E* and *D* had significantly lower psychophysical thresholds than did *monkey T*, consistent with the notion that heightened arousal may improve task performance.

We were surprised by the negative correlation between coherence and the firing rate for T1 choice trials in *monkeys E* and *D*. This pattern of responses contrasts markedly with the responses of LIP and frontal lobe neurons tested using the same behavioral paradigm (Kim and Shadlen 1998; Shadlen and Newsome 1996, 2001). The vast majority of neurons in these structures discharged vigorously when a high-coherence stimulus instructed a T1 choice and weakly when a low-coherence stimulus instructed the same choice. In addition, several previous studies of the SC have documented a positive correlation between neural firing rate and the probability that a saccade will be made into the movement field (Basso and Wurtz 1998; Dorris and Munoz 1998; Everling et al. 1998; Glimcher and Sparks 1992; Sparks 1978). This relationship is present in our data as well: prelude discharges on T1 choice trials typically exceed those on T2 choice trials. Only by restricting our attention to T1 choice trials do we find, in some monkeys, an

inverse relationship between SC prelude activity and the strength of the sensory evidence favoring that choice. Despite this inverted relationship, *predictive activity* increased with stimulus coherence. Thus activity at a single location in the colliculus may be greater at low coherences than at high coherences, but the *difference* in activity between two collicular locations representing the two competing targets may be greatest at high coherences. Thus the level of activity in SC prelude neurons may encode the readiness to make a saccade into the movement field relative to the level of activity at other points in the collicular map.

Comparison among SC, LIP, and frontal lobe discharges during task performance

To a first approximation, the responses of SC neurons we have described closely resemble the responses of neurons in the parietal (LIP) and frontal lobes (frontal eye field and area 46) studied in the same experimental paradigm (Kim and Shadlen 1998; Shadlen and Newsome 1996, 2001). In all three areas, a population of neurons fires in a choice-specific manner soon after the onset of the visual motion stimulus. Predictive activity in each area is greatest when the stimulus coherence is high and is lowest when the stimulus coherence is low. Response latency is similar across areas although methodological differences across studies preclude rigorous comparison.

Collicular neurons appear to differ from neurons in the parietal and frontal lobes in how stimulus coherence affects firing rate. Specifically, prelude discharge prior to T1-directed saccades was, on average, negatively correlated with coherence (for 2 monkeys), in contrast to the positive correlation observed in LIP and in the frontal lobe (and for the SC in 1 monkey). Each of the three relevant studies employed a small number of animals, and only one animal was common to more than one study. Inter-animal differences may thus account for some or all of the difference between studies. *Monkey E* participated in both the LIP study and the current study, and this monkey yielded opposite results in the two studies. However, the LIP recordings and the SC recordings were performed several months apart, and it remains possible that the observed response difference stems from a change in task strategy over time, not from a difference in the neural structure studied.

Neuronal correlates of choice bias

Predictive activity actually *preceded* the onset of the motion stimulus when the stimulus coherence was sufficiently low. We interpret these signals as a neural correlate of an internal bias that influences target selection when relevant sensory information is weak or absent. The magnitude of these bias signals varied significantly across animals, suggesting that it may reflect individual differences in task strategy.

The role of biases in decision-making has been studied in the fields of ethology, statistics, and psychology (Davison and McCarthy 1988; Kahneman and Tversky 1984; Treisman 1987) but has come under neurophysiological investigation only recently. When confronted with multiple choice options, animals tend to select options that are associated with large rewards or a high probability of a reward (Davison and McCarthy 1988). By manipulating reward size and probability, Platt and Glimcher (1999) altered a monkey's choice bias in a

saccade target selection task and found that the discharge of neurons in area LIP reflected the induced bias. Similar factors may contribute to bias-related signals in the SC as well.

In our direction-discrimination task, the animals' choice bias was not controlled experimentally. Post hoc analyses, however, revealed that choice bias was related to recent reward history (Seidemann 1998). Following a correct T1 choice, for instance, the animal was more likely to choose T1 than T2 if the motion stimulus on the next trial provided ambiguous information about the correct alternative (i.e., a low coherence stimulus). The behavior of our animals is probably related to that of Platt and Glimcher's monkeys, even though we did not deliberately manipulate reward probability; fluctuations in the sequence of rewarded and nonrewarded trials probably affect our monkeys' estimates of reward probability on upcoming trials even though these fluctuations are actually random.

On average, bias signals were more common and substantially stronger in the population of direction-selective SC neurons than in the population of non-direction-selective neurons. This observation is consistent with the notion that the direction-selective neurons represent a higher level of processing that integrates inputs from numerous sources that influence the choice of target for the upcoming saccade.

Temporal dynamics of the response

Predictive prelude activity in both LIP and SC appears to evolve smoothly over time. This observation is consistent with the idea that the evolution of neural activity reflects the gradual accumulation of motion information toward a perceptual decision. Our measure of predictive activity combines responses across trials, however, and thus does not reveal whether the firing rate on individual trials changes smoothly over time or whether the gradual change in average predictive activity results from combining data across heterogeneous trials in which firing rates fluctuate differently over time.

If the discharge of SC neurons reflects the accumulation of sensory information, we expect firing rate modulations to be smooth within individual trials and consistent across trials of a common stimulus type (coherence and motion direction). Alternatively, however, firing rate may increment or decrement sharply at a moment when the monkey decides to make a saccade to T1 or to T2, respectively. In this case, the firing rate on individual trials could be grossly discontinuous and variable from trial to trial.

Our analysis of streak indices confirmed that, for many cells, the time course of firing rate modulation varied across trials in a manner consistent with rare, randomly timed transitions. This was particularly true in trials with ambiguous motion stimuli (0% coherence). Vacillations in target choice, if present, should be more prevalent in this condition than at higher coherences. This raises the intriguing possibility that some of the observed fluctuations in firing rate may result from circuit-level state changes related to vacillation between behavioral choices. Experimental techniques for "reading out" saccade plans at specific times in the trial could test this hypothesis directly (Gold and Shadlen 2000; Roitman and Shadlen 1998). Simultaneous recordings from multiple SC neurons would be revealing as well. If circuit-level state changes are indeed occurring in the SC, the firing rate fluctuations should be highly correlated between neighboring neurons, and anticorre-

Two types of SC prelude neuron

Target selection	Movement specification
Directional visual response	No directional visual response
Predictive quickly	Predictive slowly
Strong effects of motion strength on predictive activity	Weak effect of motion strength on predictive activity
Strong effect of choice-bias on predictive activity	Weak effect of choice-bias on predictive activity
Weak perisaccadic bursts	Strong perisaccadic bursts
Weak relationship with saccade parameters	Strong relationship with saccade parameters
Non-"classical" visual response	"Classical" visual response

FIG. 19. Response property summary for 2 types of SC prelude neuron. Column labels indicate hypothesized functional roles: 1 population has response properties appropriate for a role in saccade target selection, the other in specifying the precise parameters of a saccadic eye movement.

lated between cells recorded at disparate points in the SC corresponding to the spatial locations of the T1 and T2 saccade targets.

Concluding remarks

A population of prelude neurons in the intermediate and deep layers of the SC exhibits choice-predictive activity during performance of a direction-discrimination task. As summarized in Fig. 19, these neurons can be divided into two subpopulations that possess constellations of physiological properties suggestive of different functional roles during task performance. As described in the companion paper, we initially assigned prelude neurons into these two subpopulations based on the existence of direction-selective visual responses during a passive-fixation task or the lack thereof. This initial division, however, was also consistent with several other physiological differences. For direction-selective neurons, the choice-predictive power of the prelude activity during the discrimination task was positively correlated with the coherence of the motion stimulus that instructed the saccade. In addition, direction-selective neurons exhibited significant effects of choice bias on firing rate prior to appearance of the motion stimulus. For non-direction-selective neurons, on the other hand, prelude activity predicts the outcome of the decision in a manner that is only weakly dependent on the strength of the motion stimulus that instructs the decision. In addition, perisaccadic activity is more intense in the non-direction-selective neurons and shows a greater tendency to vary with subtle parametric differences in saccadic eye movements.

Together, these data are consistent with the interpretation advanced in the preceding paper. One population of prelude neurons appears well suited for a role in target selection, while the other is better suited for movement preparation. The direc-

tion-selective neurons are influenced in a graded manner by the motion stimulus that instructs the perceptual decision, their preferred directions point toward the spatial location of the movement field as demanded by the logic of the task, and their activity during the discrimination task strongly predicts the outcome of the decision (i.e., the upcoming saccade). When examined in the context of our discrimination task, the non-direction-selective neurons appear to be involved only in preparation for the selected saccade.

That the SC contains neurons involved in preparation for an upcoming saccade is not controversial. More contentious is the possibility that some SC neurons play an active role in the cognitively demanding process of target selection. Whether the directional neurons we have described actually exert a causal influence on target selection is not known. Microstimulation experiments may eventually be able to address this possibility, but our initial efforts in this direction have proven inconclusive. Nevertheless, this is an important direction for future research.

We thank J. Stein and A. Milloux for technical assistance during the course of this study. Dr. Jennifer Groh provided valuable advice during the first collicular recordings, and Dr. Michael Shadlen was generous in sharing advice and software for data analysis. C. Barberini, A. Batista, J. Muller, and J. Nichols provided valuable comments on an earlier version of the manuscript. W. T. Newsome is an Investigator at the Howard Hughes Medical Institute.

G. D. Horwitz was supported by a graduate research fellowship from the Office of Naval Research and by a training grant from the National Institute of Mental Health (MH-17047). Research in this laboratory is also supported by the National Eye Institute (EY-05603).

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

REFERENCES

- BASSO MA AND WURTZ RH. Modulation of neuronal activity in superior colliculus by changes in target probability. *J Neurosci* 18: 7519–7519, 1998.
- 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. *Visual 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.
- BRITTEN KH, SHADLEN MN, NEWSOME WT, AND MOVSHON JA. Responses of neurons in macaque MT to stochastic motion signals. *Vis Neurosci* 10: 1157–1169, 1993.
- DAVISON M AND MCCARTHY D. *The Matching Law*. Hillsdale, NJ: Erlbaum, 1988.
- DORRIS MC AND MUNOZ DP. Saccadic probability influences motor preparation signals and time to saccadic initiation. *J Neurosci* 18: 7015–7026, 1998.
- DORRIS MC, PARÉ M, AND MUNOZ DP. Neuronal activity in monkey superior colliculus related to the initiation of saccadic eye movements. *J Neurosci* 17: 8566–8579, 1997.
- DRAPER NR AND SMITH H. *Applied Regression Analysis*. New York: Wiley, 1998.
- EVERLING S, DORRIS MC, AND MUNOZ DP. Reflex suppression in the anti-saccade task is dependent on prestimulus neural processes. *J Neurophysiol* 80: 1584–1598, 1998.
- FRIEDMAN HR, BURMAN DD, SHI D, AND BRUCE CJ. Independence of visual and movement response fields of visuomotor cells in monkey frontal eye field. *Soc Neurosci Abstr* 24: 522, 1998.
- GLIMCHER PW AND SPARKS DL. Movement selection in advance of action in the superior colliculus. *Nature* 355: 542–545, 1992.
- GOLD JI AND SHADLEN MN. Representation of a perceptual decision in developing oculomotor commands. *Nature* 404: 390–394, 2000.
- 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, 2001.
- KAHNEMAN D AND TVERSKY A. Choices, values, and frames. *Am Psychol* 39: 341–350, 1984.
- KIM JN AND SHADLEN MN. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2: 176–185, 1998.
- PLATT ML AND GLIMCHER PW. Neural correlates of decision variables in parietal cortex. *Nature* 400: 233–238, 1999.
- ROHRER WH, WHITE JM, AND SPARKS DL. Saccade-related burst cells in the superior colliculus: relationship of activity with saccadic velocity. *Soc Neurosci Abstr* 13: 1092, 1987.
- ROITMAN JD AND SHADLEN MN. Response of neurons in area LIP during a reaction-time direction discrimination task. *Soc Neurosci Abstr* 24: 262, 1998.
- 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.
- SEIDEMANN E. *Neuronal Mechanisms Mediating Conversion of Visual Signals into Perceptual Decisions in a Direction Discrimination Task* (Ph.D. dissertation). Stanford, CA: Stanford University, 1998.
- SHADLEN MN AND NEWSOME WT. Motion perception: seeing and deciding. *Proc Nat 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.
- SPARKS DL. Functional properties of neurons in the monkey superior colliculus: coupling of neuronal activity and saccade onset. *Brain Res* 156: 1–16, 1978.
- SPARKS DL, HOLLAND R, AND GUTHRIE BL. Size and distribution of movement fields in the monkey superior colliculus. *Brain Res* 113: 21–34, 1976.
- SPARKS DL AND MAYS LE. Movement fields of saccade-related burst neurons in the monkey superior colliculus. *Brain Res* 190: 39–50, 1980.
- THOMPSON KG, BICHOT NP, AND SCHALL JD. Dissociation of visual discrimination from saccade planning in macaque frontal eye field. *J Neurophysiol* 77: 1046–1050, 1997.
- TREISMAN M. Effects of setting and adjustment of decision criteria on psychophysical performance. *Prog Math Psych* 1: 253–297, 1987.
- WURTZ RH AND ALBANO JE. Visual-motor function of the primate superior colliculus. *Annu Rev Neurosci* 3: 189–226, 1980.
- WURTZ RH AND GOLDBERG ME. Activity of superior colliculus in behaving monkeys. III. Cells discharging before eye movements. *J Neurophysiol* 35: 575–586, 1972.
- ZAR JH. *Biostatistical Analysis* (2nd ed.). Englewood Cliffs, NJ: Prentice-Hall, 1984.