

# Reward Action in the Initiation of Smooth Pursuit Eye Movements

Mati Joshua<sup>2,3</sup> and Stephen G. Lisberger<sup>1,2,3</sup>

<sup>1</sup>Howard Hughes Medical Institute, <sup>2</sup>Department of Physiology, and <sup>3</sup>W.M. Keck Foundation Center for Integrative Neuroscience, University of California, San Francisco, San Francisco, California 94143

Reward has a powerful influence on motor behavior. To probe how and where reward systems alter motor behavior, we studied smooth pursuit eye movements in monkeys trained to associate the color of a visual cue with the size of the reward to be issued at the end of the target motion. When the tracking task presented two different colored targets that moved orthogonally, monkeys biased the initiation of pursuit toward the direction of motion of the target that led to larger reward. The bias was larger than expected given the modest effects of reward size on tracking of single targets. Experiments with three different reward sizes suggested that the bias afforded a given target depends mainly on the size of the larger reward. To analyze the effect of reward on directional learning in pursuit, monkeys tracked a single moving target that changed direction 250 ms after the onset of motion. Expectation of a larger reward led to a larger learned eye movement during the acquisition of the learned response and during subsequent probes of what had been learned, implying that reward influenced the expression rather than the acquisition of learning. The specific effects of reward size on learning and two-target stimuli imply that the site of reward modulation is at a level where multiple target motions compete for control of eye movement, downstream from sensory processing and learning and upstream from final motor processing.

## Introduction

Reward size and expectation are strong modulators of behavior. Reward and reward expectation are strongly represented in the basal ganglia (Schultz, 1998; Joshua et al., 2008), and effects of reward appear on the discharge of neurons in many parts of various motor systems (Platt and Glimcher, 1999; Arkadir et al., 2004; Ding and Hikosaka, 2006). However, we do not know how reward impacts the planning and execution of a motor behavior, at either the conceptual level in terms of the components of a motor behavior or the practical level of neurons and circuits. Some headway has been made for saccadic eye movements (Lawwereyns et al., 2002; Takikawa et al., 2002), but the current understanding is at best incomplete.

Smooth pursuit eye movements offer a unique opportunity to determine how reward controls motor planning and execution. Much is already known about the different components of pursuit and their interaction. Signals for pursuit flow from a sensory representation in extrastriate middle temporal area MT (Newsome et al., 1985; Groh et al., 1997) to motor processing in the brainstem and cerebellum (Fig. 1). The pursuit system is subject

to reliable forms of learning (Medina et al., 2005) that occur downstream from sensory processing (Carey et al., 2005). Still farther downstream (Kahlon and Lisberger, 1999; Recanzone and Wurtz, 1999), vector averaging is used to combine multiple visual motion signals and create a single command for smooth eye movement (Lisberger and Ferrera, 1997). Target choice alters the weight afforded a given direction of target motion in the vector averaging process (Gardner and Lisberger, 2001). Our goal was to localize the effects of reward on pursuit eye movements within the context of the conceptual structure given in Figure 1, as a step toward understanding the effects of reward at a neural level.

The skeleton of the neural circuit for pursuit and the signals carried at different levels of the circuit are known (Tian and Lynch, 1997). The basal ganglia, as the most likely source of signals related to reward expectation, interact with the pursuit circuit through the smooth eye movement region of the frontal eye field (FEF<sub>SEM</sub>). The FEF<sub>SEM</sub> plays important roles in pursuit learning (Li and Lisberger, 2011) and direction-specific modulation of the strength of visual-motor transmission for pursuit (Tanaka and Lisberger, 2001, 2002). Thus, it will be important to understand the impact of reward on these components of pursuit to reach the ultimate goal of understanding how and where reward alters the firing of specific neurons.

In the present paper, we show that reward accesses the pursuit system at the level where vector averaging occurs. Reward acts by altering the relative strength afforded different target motions in relation to how strongly they will be rewarded. Reward affects learned eye movements rather than the learning process itself. By revealing how reward fits into the conceptual organization of the essential circuit, we provide information that is critical for understanding the interaction of the reward systems and a well-understood sensory–motor circuit.

Received Sept. 13, 2011; revised Dec. 23, 2011; accepted Jan. 4, 2012.

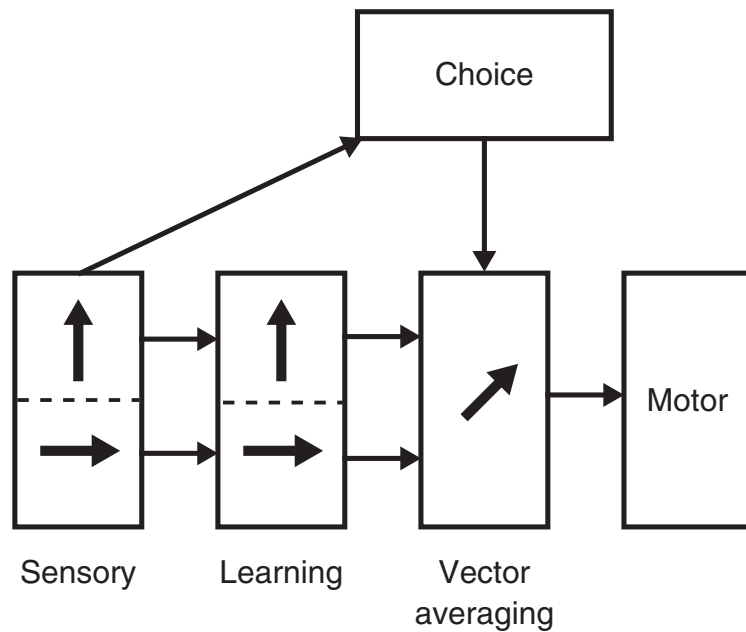
Author contributions: M.J. and S.G.L. designed research; M.J. performed research; M.J. analyzed data; M.J. and S.G.L. wrote the paper.

This research was supported by the Howard Hughes Medical Institute, National Institutes of Health Grant EY03878, and the Human Frontiers Science Program. We are grateful to Allison Doupe for helpful comments on a previous version of this manuscript. We thank K. MacLeod, E. Montgomery, S. Tokiyama, S. Ruffner, D. Kleinhesselink, D. Wolfgang-Kimball, D. Floyd, and K. McGary for technical assistance.

Correspondence should be addressed to Mati Joshua, Department of Physiology, University of California, San Francisco, Box 0444, 513 Parnassus Avenue, Room HSE-808, San Francisco, CA 94143-0444. E-mail: mati@phy.ucsf.edu.

DOI:10.1523/JNEUROSCI.4676-11.2012

Copyright © 2012 the authors 0270-6474/12/322856-12\$15.00/0



**Figure 1.** Flow of information from sensory processing to final motor processing for smooth pursuit eye movements. The rectangular blocks represent different components of pursuit behavior, and the arrows between the boxes indicate the conceptual flow of information based on previous reports. The two arrows inside the boxes for sensory processing and learning indicate the existence of two parallel representations of two targets moving in different directions. The single arrow inside the box for vector averaging indicates the combination of multiple target motions to produce a single pursuit eye movement.

## Materials and Methods

Three male rhesus monkeys (*Macaca mulatta*) served as subjects. To instrument them for experiments, each monkey was anesthetized with isoflurane, and a search coil was implanted on one eye (Ramachandran and Lisberger, 2005) so that eye position could be measured using the magnetic search coil technique. Custom-cut orthopedic stainless steel strips were attached to the monkeys' skull with 6-mm-long screws. The strips served as the foundation for dental acrylic to secure a receptacle that was used to fix the head to the primate chair. Appropriate analgesic and antibiotic treatments were administered postoperatively. After recovery from surgery, monkeys were trained to sit in a primate chair with the head restrained and to fixate and track spots of light that moved across a monitor placed in front of them. All procedures involving the monkeys had been approved in advance by the Institutional Animal Care and Use Committee at University of California, San Francisco and followed the National Institutes of Health *Guide for the Care and Use of Laboratory Animals*.

**Visual stimuli and experimental design.** Visual stimuli were displayed on a Barco monitor at a distance of 30 cm from the monkeys' eye. Targets appeared as bright  $0.5^\circ$  circles on a dark background. All experiments were performed in a dimly lit room. Sequences of target motion were controlled by a computer that also performed all the real-time operations. Signals proportional to horizontal and vertical eye position were passed through an analog circuit to create signals proportional to horizontal and vertical eye velocity. The circuit differentiated frequency content from 0 to 25 Hz and filtered higher frequencies with a roll-off of 20 dB/decade. The analog signals proportional to eye position and velocity were digitized at 1000 samples/s on each channel and stored for analysis.

Pursuit stimuli were presented in trials. At the start of a trial, a stationary white target appeared in the middle of the screen, and monkeys were required to fixate within a  $2^\circ \times 2^\circ$  window. After the monkey attained steady fixation on the white target, one or two additional targets appeared  $4\text{--}5^\circ$  eccentric along the horizontal and/or vertical meridian. The color of the eccentric targets cued the size of the reward that would be received at the end of the trial. After a variable delay of 800–1200 ms, in which the monkeys were required to continue and fixate on the central target, the fixation target disappeared and the eccentric targets started to

move toward the middle of the screen at  $30^\circ/\text{s}$ . In some of the recordings sessions of monkey S, we reduced the time before the onset of target motion to 500–700 ms to minimize anticipatory movements. Changing the fixation time did not change any of our findings.

In single-target trials, the target moved at constant velocity for 600–750 ms, and monkeys were required to keep their gaze within a  $2\text{--}4^\circ$  square window centered on the moving target. At the end of the trial, the target stopped and the monkey received a water reward. In monkey P, we used a green target for the large reward (0.2–0.4 ml) and a yellow target for small reward (0.05–0.1 ml). In monkeys I and S, we reversed the association between color and reward size.

In two-target trials, the monkeys initiated a smooth eye movement and then typically made a saccade to one of the targets. To allow the monkey to choose a pursuit target freely, we suspended fixation requirements from the time of disappearance of the fixation target to the end of the first saccade. We detected the first saccade by sensing when eye velocity exceeded and then crossed back downward through  $80^\circ/\text{s}$ , at which time we extinguished the target that was farther from the endpoint of the saccade. The monkey received a reward after successful tracking of the remaining target. The size of the reward was

linked to the color of the target chosen for tracking and followed the same rules used to determine the reward size during single-trial experiments. In rare trials when the monkey did not execute a saccade, the reward size was determined according to the color of the target that was closer to eye position at the end of the target movement, but only if the eye was within  $4^\circ$  of at least one of the targets at the end of the trial.

In some experiments, we delivered only single-target trials and studied the effect of reward size on the response to target motion. In other experiments, we delivered both single- and two-target trials to study the effect of reward size on the weighting of the target motions. The directions of the single-target motions were selected to provide appropriate controls for assessing the responses in two-target trials. Different pairs of two-target motion directions were interleaved and compared with single-target trials collected in the same daily experiment. To control for the influence of motivational changes during a recording day, blocks of single-target trials could precede or follow blocks with two trials. The order of the blocks did not change any of our conclusions.

**Directional learning paradigm.** Learning experiments followed the directional learning paradigm established by Medina et al. (2005). Learning trials began with the same sequence of target presentations as the single-target trials. After the colored target had moved at  $30^\circ/\text{s}$  for 250 ms, we changed target direction by adding an orthogonal component of motion at  $30^\circ/\text{s}$ . We considered the change in target direction as an “instruction” that teaches the pursuit system to emit eye motion orthogonal to the direction of target motion (Medina et al., 2005).

Each learning block of 300 trials comprised 45% learning trials with target motion that started in a given direction and changed direction 250 ms later, 10% probe trials with target motion that started in the same direction as the learning trials and did not change direction, and 45% control trials that started with motion in the opposite direction from the learning and probe trials and also did not change direction. To equalize the overall reward that could be attained in each block of trials, the control trials in a learning block provided the opposite reward size from the learning trials. Blocks with large rewards for learning trials provided small rewards for the control trials and vice versa. The different trials were interleaved in random order within blocks in which the learning direction, reward size, and time of the instructive change in target direction were fixed. To control for the influence of motivational changes

during a recording day, blocks of learning with small and large reward for learning trials were alternated and the reward size of the first block of learning alternated across days.

Probe trials always were interleaved with learning trials and led to rewards consistent with the color of the target in the probe trial. In some blocks, probe trials cued and delivered the opposite reward from the learning trials. In other blocks, probe trials cued and delivered rewards of an intermediate size that was the same for blocks that provided large or small rewards for learning trials. Because monkeys did not always complete the full 300 trials of the learning block (containing 135 learning trials), we analyzed learning blocks that contained at least 75 learning trials. Each learning block was preceded by a baseline block of at least 50 probe trials in the same direction used for the start of the subsequent learning trials, again without changing the direction of the target.

**Data analysis.** We used eye velocity thresholds to detect saccades automatically and then verified the automatic decisions by visual inspection of the traces from each trial. In the one-target and two-target experiments, we excluded from analysis any trials that had saccades within an analysis window from 0 to 200 ms after the onset of target motion. We averaged the eye velocity across responses to identical target motions and reward contingencies, to obtain the mean eye velocity vector as a function of time. Statistics were performed by testing for significance across sessions. We set the significance level at 0.05. We obtained similar results using both parametric and nonparametric statistics and with tests using more complicated designs with the session number as one factor (Friedman's test or two-way ANOVA).

In learning trials, we were interested in the time course of the trial-by-trial progress of learning and hence we did not simply discard trials with saccades. Instead, we treated the eye velocity between the start and end of each saccade as missing data. To confirm that including the 2.5% of trials with saccades in the analysis window did not affect the results, we verified that the conclusions did not change if we discarded all trials with saccades. To quantify the learned response, we took the average of the eye velocity in the learned direction in the interval from 200 to 300 ms after the onset of target motion, which corresponds to the interval from 50 ms before to 50 ms after the time of the instructive change in target motion used to induce learning. Comparison between different conditions (e.g., probe vs learning) was performed by nonparametric statistical tests (paired Wilcoxon's and Mann–Whitney tests) on the average learned responses.

We estimated the latency of average smooth eye velocity responses as the time when eye velocity became significantly different from zero and remained different until the end of the analysis window. This approach avoids early detections of response attributable to noise and exploits the fact that the eye velocity response increases with time and hence tends to remain significant after the genuine response onset.

## Results

Our study used tasks designed to probe three different facets of pursuit behavior to test how reward interacts with sensory–motor transformations for smooth pursuit eye movements. First, we evaluate the relative weighting of different visual stimuli for the initiation of pursuit using the simultaneous motion of two targets: the initial 100 ms of pursuit is biased toward the direction of motion of the target that yields a larger reward. Second, we show that larger rewards lead to faster initial eye accelerations even for the motion of single targets but that the effects are too small to explain the bias toward the larger reward in the two-target experiment. Third, we evaluate the effect of reward on motor learning in pursuit: reward size modulates mainly the expression of learning.

### Reward size biases responses to two-target stimuli

When two targets start to move simultaneously in orthogonal directions under conditions in which the targets are identical and lead to equal rewards, the speeds of the horizontal and vertical eye components are approximately half of those evoked by each tar-

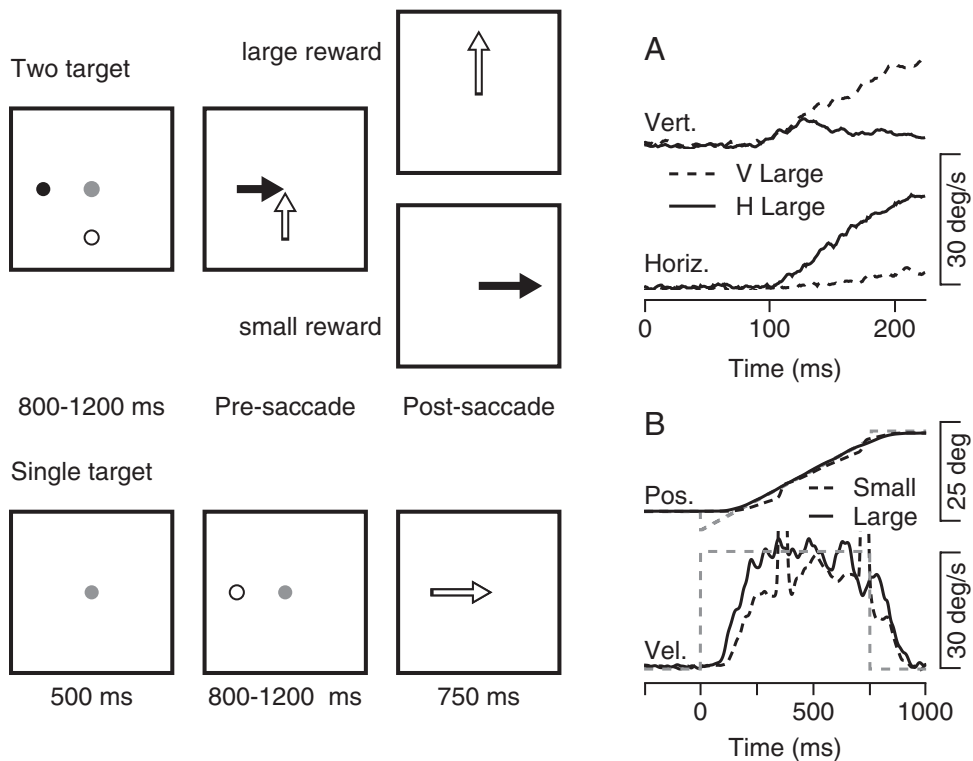
get moving singly. The direction of the net eye movement is approximately halfway between the directions of the responses to the two target motions singly. Quantitative analysis has shown that the pursuit direction and speed of two equally salient targets can be accounted for in terms of vector averaging of the orthogonal eye movements evoked by each target singly (Lisberger and Ferrera, 1997). We now show that previous knowledge of the reward size associated with each target can bias the initiation of pursuit strongly toward one of the two targets.

In the experiment illustrated in the top schematic of Figure 2, in which the colors of the two targets are represented by open and filled spots, the monkey knew in advance that tracking the filled spot would lead to a small reward and tracking the open spot would lead to a large reward. In 99% of the two-target trials, the monkey's first saccade took the eye to the target that would lead to the larger reward. We analyzed the eye velocity that preceded the first saccade, in an interval of presaccadic pursuit that is influenced only very weakly by target choice for the subsequent saccade when the targets lead to equal rewards (Gardner and Lisberger, 2001).

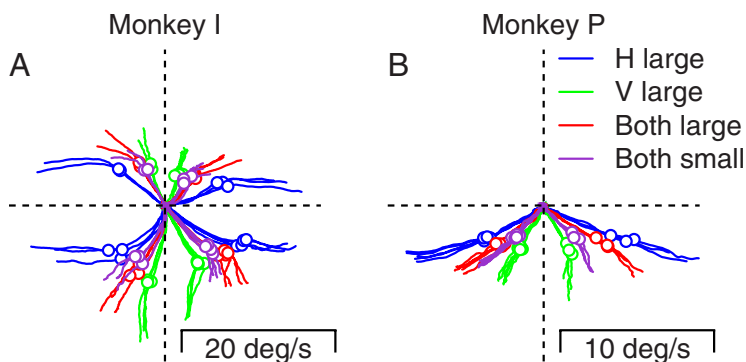
At least 100 ms before the first saccade, monkeys initiated a smooth eye movement biased toward the direction of motion of the target that would yield a large reward, henceforth called the "highly rewarded direction." Figure 2A, for example, uses records from two trials to illustrate that each eye velocity component was larger when it coincided with the highly rewarded direction. The horizontal eye velocity was large when the horizontal target led to a large reward (continuous trace) and quite small when the vertical target led to a large reward (dashed trace). The vertical eye velocity showed the opposite relationship. At least for the horizontal eye velocity, the traces for the different highly rewarded directions diverged starting from the onset of pursuit. For vertical eye velocity, the divergence occurred ~20–30 ms after the onset of pursuit.

The effect of future reward on the bias toward one target or the other was present across experimental days and monkeys. In Figure 3, each trace plots the averages of eye velocity for multiple repetitions of each experiment in two monkeys. Here, and in some of our other graphs, we have used two-dimensional plots that show vertical eye velocity as a function of horizontal eye velocity for the first 200 ms after the onset of target movement. Time is not represented explicitly in Figure 3. Instead, time begins with eye velocity at (0,0), defined by the intersection of the two dashed axes. As time proceeds toward 200 ms after the onset of target motion, eye velocity becomes non-zero and moves outward along each trace in the graph. We have chosen to use the two-dimensional representation of eye velocity because it enables a clear comparison of the directions of eye movement in different reward conditions. To help the reader appreciate the passage of time, we also have placed open circular symbols on the graph to indicate 150 ms after the onset of target motion.

We obtained the same general picture for all six pairs of orthogonal target directions that we were able to analyze in two monkeys, shown by data drawn in six quadrants in Figure 3. Thus, the same pair of target directions was used for all traces within an individual quadrant, and the traces differed only in the monkey's expectations of reward size for each target. When the reward assigned to the two moving targets was unequal (blue and green traces), the trajectory of eye velocity was biased toward the highly rewarded target almost from the initiation of pursuit. By the end of the first 100 ms of pursuit, the bias toward the highly rewarded target grew quite large. Eye velocity was more strongly horizontal when the horizontal target motion was highly re-



**Figure 2.** Example data showing the effect of reward size on the initiation of pursuit. The sequence of snapshots illustrates the structure of the behavioral task with two targets (top) or single targets (bottom). We show targets as filled and open spots for clarity of presentation, although they actually were green and yellow in our experiments. **A**, Examples of the horizontal (Horiz.) and vertical (Vert.) eye velocity in two different trials. The continuous and dashed traces show responses when the highly rewarded target moved horizontally or vertically. **B**, The black solid and dashed traces show examples of the horizontal eye position (Pos.) and velocity (Vel.) in trials that led to large or small rewards. The gray dashed lines are the target position and velocity.



**Figure 3.** The effect of reward size on the initiation of pursuit for two-target stimuli across multiple experiments. Each trace plots average vertical (V) eye velocity versus average horizontal (H) eye velocity for each millisecond in the first 200 ms after the onset of target motion. Different colors show data for different reward-size paradigms. Traces in different quadrants show responses for different pairs of orthogonal directions of target motions. The multiple traces of the same color show results from identical experiments performed on different days. **A**, **B**, the two graphs show data for two monkeys. Open circles on the traces show the eye velocities 150 ms after the onset of target motion. In monkey P we did not study the pursuit evoked by combinations of horizontal and upward target motion because pursuit was poor for the upward motion of single targets.

warded (blue traces) and more strongly vertical when the vertical target motion was highly rewarded (green traces). When the reward assigned to the two moving targets was equal (red and purple traces), the trajectory of eye velocity was intermediate and depended only slightly on whether the rewards were both large or both small.

For comparison with the effects of reward on the responses to two-target stimuli, we also assessed the effect of reward size on tracking of a single target. In single-target trials, we provided a colored cue about the future reward size and then delivered the

advertised reward after successful completion of the trial (Fig. 2, bottom schematic). Even in interleaved trials in which the color of the target (and the reward size) varied randomly from trial to trial, the monkey's smooth eye velocity was modestly larger in the trial with the large reward and his pursuit was punctuated by fewer corrective saccades. For the two examples illustrated in Figure 2B, the eye velocity generated in anticipation of a larger reward (continuous trace) had a more rapid eye acceleration, more accurate steady-state eye velocity, and no corrective saccades, all in contrast to the features of the eye velocity generated in anticipation of a smaller reward (dashed trace). To understand the mechanisms of reward action on pursuit, we will ask whether the eye movements evoked by pairs of unequally rewarded targets can be predicted by the linear combination of the

responses to the two targets presented singly with the same reward expectations used in the two-target trials.

**Relative sites of reward modulation and vector averaging**

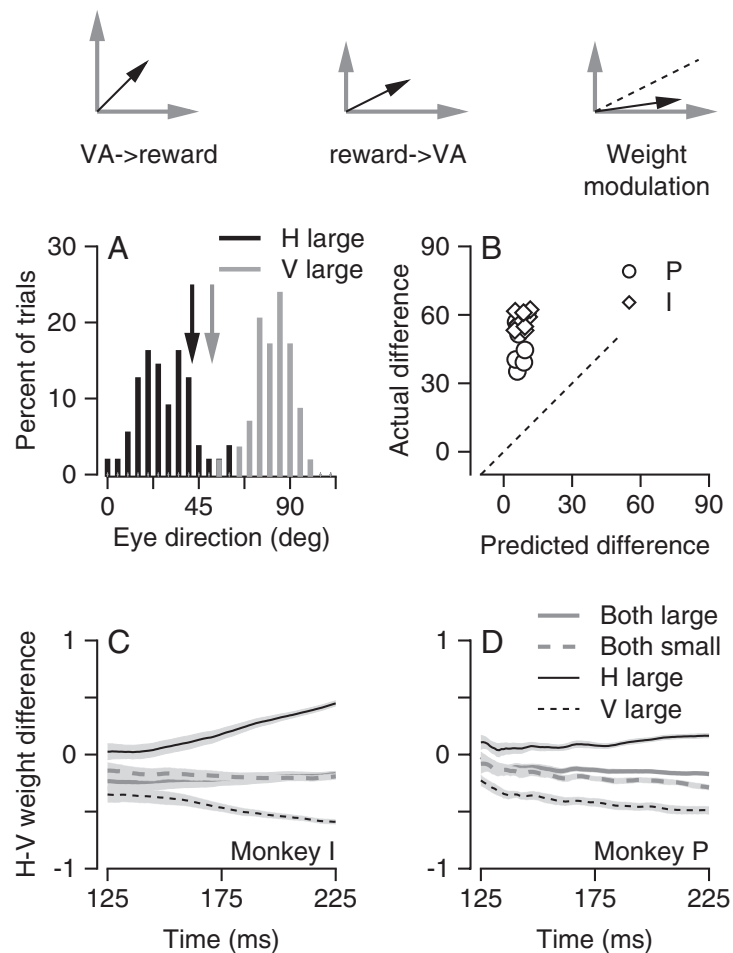
The schematic diagrams at the top of Figure 4 illustrate three conceptual possibilities for the interaction of reward size and the initiation of pursuit. Each diagram predicts the eye movements evoked by two-target stimuli in relation to the eye movements evoked by each of the target motions singly. Because we delivered

both two-target and single-target motions in daily experiments, we were able to test these three models with our data.

The left schematic (VA → reward) illustrates the possibility that reward modulation occurs downstream from the site of vector averaging in pursuit circuitry. Then, reward would act after the direction of the smooth eye movement had been determined. The direction of the eye velocity would be determined by the vector average of the responses to the two targets singly, and the reward would act only on the magnitude of eye velocity. The data in Figure 3 contradict the hypothesis that reward modulation is downstream from vector averaging, because they show an effect of reward size on the direction of eye velocity for two-target stimuli.

The middle schematic at the top of Figure 4 (reward → VA) illustrates the possibility that modulation by reward size of the responses to each of the single targets occurs upstream from the site of vector averaging. In this case, the direction of the eye movement evoked in two-target trials should be the vector average of the responses to each target presented singly with the appropriate reward sizes: we would say that reward had affected the strength of the pursuit evoked by one target direction but not that it truly had biased target choice. The rightmost schematic (Weight modulation) illustrates the results predicted if reward size can modulate the weighting of the two targets and lead to presaccadic pursuit that truly is biased toward the direction of motion of one of the two targets, even after taking into account the effects of reward on the responses to each target singly. A competitive circuit with reciprocal inhibition and excitation would be one mechanism for using reward to bias the weighting of two differently rewarded targets (Ferrera, 2000). However, our focus here is on the conceptual organizations indicated by the three models at the top of Figure 4 rather than on the possible mechanisms for implementing each model.

To discriminate between (1) reward modulation precedes vector averaging and (2) reward modulates the weight of vector averaging, we asked whether the direction of pursuit evoked by two-target motions coincided with the direction predicted by the vector average of the responses in trials that had presented the two target motions singly. Consider, for example, an experiment in which a two-target stimulus included upward and rightward target motion and the cues indicated that upward was the highly rewarded direction. The experiment also would have included single-target trials with upward target motions that cued large rewards or rightward target motions that cued small rewards. We used the responses to these two “appropriately rewarded” single-target trials to form the predictions based on equally weighted vector averaging. If vector averaging of the two appropriately rewarded single-trial responses predicted the actual direction for two-target responses, then we would conclude that



**Figure 4.** Evaluation of three models for the relative locations of vector averaging (VA) and reward modulation. The schematic diagrams at the top of the figure show the three models. Gray arrows show the responses to each target motion presented singly, modulated by reward size in the second and third models. Black arrows show the eye velocity predicted by each model. The dashed line shows the prediction of equally weighted vector averaging, which differs from the predicted eye velocity only in the third model. **A**, The distribution of eye direction 200 ms after target onset for 1 experimental day. Black and gray histograms indicate the results when the highly rewarded target moved horizontally (H; at 0°) or vertically (V; at 90°). The vertical arrows indicate the predictions for vector averaging of the appropriately rewarded targets presented singly. **B**, Comparison of the average direction of the initiation of pursuit in the data with the predictions for vector averaging of the appropriately rewarded targets presented singly. Each symbol shows data from a different experiment ( $n = 8$  and 11 for monkeys P and I). Diamonds and circles represent data for monkeys I and P. The dashed line is the unity line. **C, D**, The difference between the weighting of the horizontal and vertical targets for two-target stimuli plotted as a function of time from the onset of target motion. Different line styles indicate different reward-size paradigms.

reward modulation is upstream to vector averaging. If, instead, the actual direction of pursuit was biased toward the highly rewarded target motion by comparison with the predictions of vector averaging, then we would conclude that reward size was affecting the relative weighting of the two targets above and beyond its effect on the responses to the two targets singly.

The critical test of the location of reward modulation relative to vector averaging appears in Figure 4A. Here, the differently shaded distributions illustrate the pursuit directions from a single experiment for many interleaved two-target trials when rightward (black) or upward (gray) target motion comprised the highly rewarded direction. The actual distributions of pursuit direction were separated by much more than the predictions for vector averaging of the pursuit in appropriately rewarded, single-target trials, shown by the black and gray arrows. To summarize the results across experiments, we computed the mean directions for each pair of highly rewarded directions in each experiment

and represented the effect of reward as the difference between the two means. In Figure 4B, the points for all experiments lie well above the unity line, indicating that the actual effect of reward was considerably larger than that predicted by the vector average of the responses to the appropriately rewarded targets presented singly. We conclude that the locus of reward modulation is not upstream from vector averaging.

Our analysis so far indicates that neither the VA  $\rightarrow$  reward hypothesis nor the reward  $\rightarrow$  VA hypothesis is supported by our data, implying that reward size truly biased the responses to two-target motions. Therefore, we next used the approach of Ferrera and Lisberger (1997) and Garbutt and Lisberger (2006) to quantify the weights used to combine the motion signals from two targets under different reward conditions:

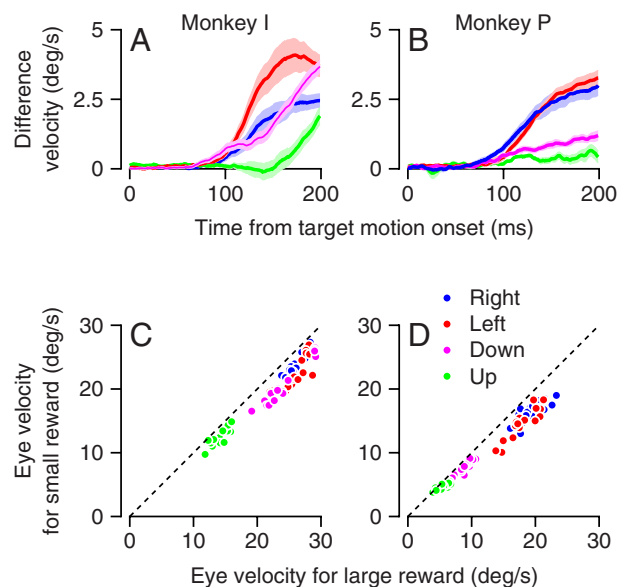
$$\bar{\dot{E}}(t) = w_H(t)\dot{E}_H(t) + w_V(t)\dot{E}_V(t).$$

We calculated the weights [ $w_H(t)$  and  $w_V(t)$ ] assigned to each target so that the net smooth eye velocity vector ( $\dot{E}$ ) is accounted for in terms of the averages of eye velocity evoked by the appropriately rewarded targets presented singly ( $\dot{E}_H$  and  $\dot{E}_V$ ). The use of the appropriately rewarded responses on the right side of the equation above assigns the full effect of reward to the values of the weights. The analysis considers each time separately and reports the weighting of each target motion as a function of time and reward size. When both coefficients are equal to 0.5 in the equation, the behavior is predicted by vector averaging of the single target responses. Biases in eye movement toward one of the targets would be reflected in larger weights for that direction of motion. To assess the bias toward a target, we calculated the difference between the horizontal and vertical weights. Thus, positive or negative weight differences indicate biases toward the horizontal or vertical target motion.

For both monkeys, reward biased the weighting of the two targets from the earliest times in the response when eye velocity was large enough to allow us to perform the weight analysis (Fig. 4C,D). When the horizontal or vertical target motion was more highly rewarded (black continuous and dashed lines), the weight difference was positive or negative and the difference between conditions was statistically different from zero starting at the earliest time in the analysis ( $p < 0.01$ , paired Wilcoxon's test). The weight difference was increasingly positive or negative as time passed up to 225 ms after the onset of target motion. In contrast, when both directions of target motion received equal large or small rewards (gray continuous and dashed lines), the weight difference did not change as a function of time and differed from zero by small amounts only because of small idiosyncratic biases in the individual monkeys (Kahlon and Lisberger, 1999). We obtained almost the same values of weights if we predicted the responses to the two-target motions using the responses either when both targets singly were highly rewarded or when neither target singly was highly rewarded. Thus, our conclusion, that reward biases the weight assigned to different targets toward the highly rewarded target, does not depend on the exact details of the model used to test the hypotheses.

### Reward size enhances smooth pursuit of a single target

In Figure 4A, the predicted effects of reward based on vector averaging of the responses to the appropriately rewarded single targets, signified by the small horizontal shift of the two arrows, were relatively small because of the modest effects of reward size on the responses to single-target motions. For each direction of single target motion, we computed the time course of the effect of



**Figure 5.** Quantitative analysis of the effect of reward size on the initiation of pursuit for single targets in two monkeys. **A, B**, Time course of the difference in the eye velocity during the initiation of pursuit between trials that led to large versus small rewards. Different colors indicate different directions of target motion, solid curves are the averages, and the lighter shading shows the SEM across experiments. Positive values indicate that eye velocity was larger in trials that led to large rewards. **C, D**, Scatter plots comparing the average velocity 200 ms after the onset of target motion for trials that led to small versus large rewards. Each point shows data for a single experimental day. Different colors correspond to the colors used to indicate different directions of target motion in **A** and **B**. The oblique dashed line indicates equal eye velocity for large versus small rewards; points that lie below the line indicate that eye velocity was smaller in trials that delivered smaller rewards. The number of sessions was 15 for monkey P in each direction and 16 and 17 sessions for the vertical and horizontal directions for monkey I.

reward size by subtracting the initial eye velocity at each time during the initiation of pursuit for small-reward targets from that for large-reward targets. The difference was always zero or positive (Fig. 5A, B) and became non-zero within 100 ms of the onset of target motion for seven of eight traces from the two monkeys.

We quantified the effect of anticipated reward size by measuring eye velocity 200 ms after the onset of target motion for each direction of target motion and each daily experiment and plotting the eye velocity evoked in anticipation of a small reward as a function of that evoked in anticipation of a larger reward (Fig. 5C,D). Essentially all the points lie below the unity line, indicating larger eye velocities during pursuit initiation for large rewards. The data were consistent across different sessions (for each monkey,  $p < 0.001$ , paired Wilcoxon's test) and across initial pursuit eye velocities, which varied widely because of the individual directional anisotropies in the two monkeys. The effect of reward size on pursuit of single targets amounted to  $< 16\%$  of the total response (estimated by regression through the origin; slopes of 0.87 and 0.84 in monkey I and monkey P). The small effects of reward size on the responses to single targets explains why the responses to two-target stimulus are accounted for best if reward controls the weighting of different targets as part of vector averaging.

### Dependence of target bias on the size of the larger reward

To understand better how reward size modulates the bias toward the highly rewarded target in a two-target trial, we performed two-target experiments with three reward sizes of 0.1, 0.2, and 0.4 ml. In each experiment, we interleaved all three combinations of

reward size [small (S), medium (M), and large (L)]: 0.4 versus 0.1 (L vs S), 0.4 versus 0.2 (L vs M), and 0.2 versus 0.1 (M vs S).

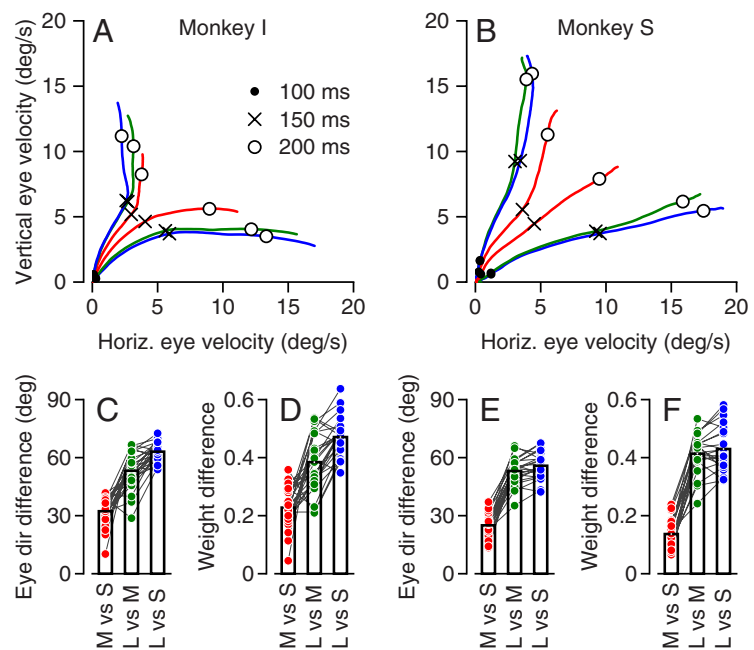
Figure 6 shows that modulation of the weights of vector averaging depends mainly on the size of the larger reward, with a possible small influence of the difference in reward size. In both monkeys (Fig. 6*A,B*), the eye velocity traces followed separate trajectories starting from the first reliable estimation after the onset of pursuit, depending on whether the highly rewarded direction was vertical or horizontal. The large versus small reward comparison (blue traces) yielded the largest difference and the large versus medium reward comparison (green traces) yielded a slightly smaller difference, although the size of the larger reward was the same. The medium versus small reward comparison (red traces) yielded the smallest difference in eye velocity directions, although the reward size ratio was the same as for the large versus medium comparison.

To quantify the effect of the relative reward, we calculated the difference between the eye directions (Fig. 6*C,E*) and the weights assigned to the two targets (Fig. 6*D,F*) when the highly rewarded direction was horizontal versus vertical. We used a time point 200 ms after target onset. Repeated ANOVA revealed a significant difference among the different combinations of reward sizes ( $p < 0.0001$ , in Fig. 6*C–F*). Tukey's *post hoc* comparison revealed all comparisons to be significant for monkey I ( $p < 0.001$  for directions and  $p < 0.01$  for weights). In monkey S, the effects on both direction and weights were significantly larger for L versus S and L versus M than for M versus S ( $p < 0.001$ ) but were statistically identical for L versus S and L versus M ( $p > 0.05$ ).

Saccade behavior also depended on the size of the larger reward. The monkey almost always chose the target with the largest reward without regard for the size of the other reward (for monkey I, 99.4 and 99.5% in the L vs M and L vs S trials; for monkey S, 97.3 and 98.7% in the L vs M and L vs S trials). In trials that did not include the largest reward (M vs S), monkeys I and S chose the highly rewarded target in only 90 and 84.7% of the trials.

### Timing of pursuit planning for target choice

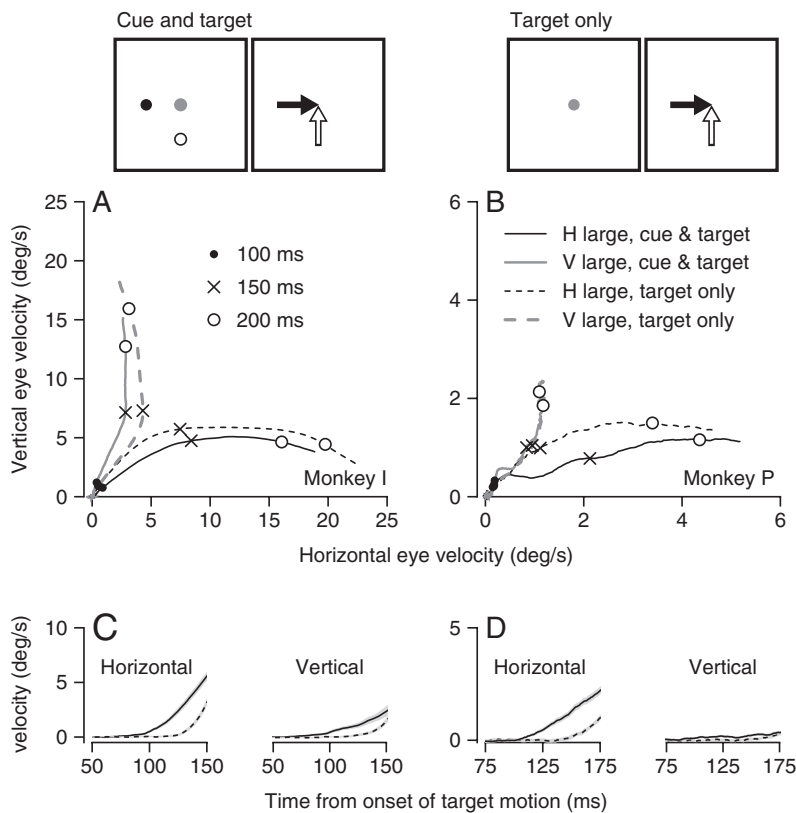
To examine whether the monkey is planning the future pursuit movement and to test the timing of the effects of reward on pursuit, we modified the two-target task so that it provided information about which target would be rewarded only at the onset of target motion (Ferrera, 2000; Case and Ferrera, 2007). In interleaved trials, we presented the two conditions illustrated in the schematic diagrams at the top of Figure 7. In our standard "cue-and-target" condition, the cues appeared well before the onset of target motion so that the monkey knew about the relationship between target motion direction and reward size in advance of the pursuit segments of the trial. In a "target-only" condition, the targets did not appear until they started to move, so that the monkey did not receive an advance cue about the highly rewarded direction.



**Figure 6.** Graded effects of different reward sizes. *A, B*, Two-dimensional plots for the two monkeys of average vertical eye velocity versus horizontal (Horiz.) eye velocity for each millisecond in the first 225 ms after the onset of target motion with different reward levels. Traces that travel closer to the horizontal and vertical axis show data for trials with horizontal or vertical highly rewarded directions. For simplicity of presentation, all directions were transformed to plot in the first quadrant. Symbols on the traces indicate different times after the onset of target motion. Blue, green, and red traces represent experiments with reward sizes of large versus small (L vs S), large versus medium (L vs M), and medium versus small (M vs S). *C, E*, Eye direction 200 ms after the onset of target motion when vertical was the highly rewarded direction minus that when horizontal was the highly rewarded direction. *D, F*, The weight afforded the highly rewarded direction 200 ms after the onset of target motion minus that afforded the other direction. In *C–F*, each symbol plots the average from one experimental session. Colors have the same significance as in *A* and *B*. Lines connect the three symbols for conditions that were interleaved in the same session. Bars represent the average within each group.

Availability of the cue before the onset of target motion led to an earlier bias of eye velocity toward the highly rewarded direction. As before, the traces followed separate trajectories starting from the onset of pursuit (monkey I) or very soon after the onset of pursuit (monkey P) when the highly rewarded direction was cued 800–1200 ms before the onset of target motion (Fig. 7*A,B*, solid traces). When the cue was withheld until the onset of target motion, the effect of reward was delayed. The eye velocity for the horizontal versus vertical highly rewarded directions did not separate until almost 150 ms after the onset of target motion (Fig. 7*A,B*, dashed traces).

To assess the time course of the development of the bias toward the highly rewarded direction of target motion, we plotted the difference between the mean horizontal and vertical eye velocities when the large reward was in the horizontal versus the vertical direction, as a function of time (Fig. 7*C,D*). The difference traces show that the bias toward the highly rewarded direction appeared earlier in the cue and target condition than in the target-only condition. In the cue-and-target condition (continuous curves), the effect of reward was statistically significant in the horizontal eye velocity 72 and 106 ms after the onset of target motion for monkeys I and P ( $p < 0.05$ , paired Wilcoxon's test). For the target-only condition (dashed curves), the bias did not reach statistical significance in the horizontal eye velocity traces until 123 and 150 ms after the onset of target motion for the two monkeys. Reward size had little or no effect on the vertical pursuit in monkey P in Figure 7. However, the control of horizontal eye velocity by reward was clear in the monkey's



**Figure 7.** Latency of modulation of pursuit initiation by reward size. The schematic diagrams at the top of the figure illustrate the structure of the two-target trials. Cue and target, The two tracking targets appeared during fixation, as described in Results, and provided cues for future reward. Target only, The tracking targets appeared and started to move when the fixation spot was extinguished, providing cues for future reward at the same time as motion signals for pursuit initiation. **A, B**, Two-dimensional plots of average vertical (V) eye velocity versus horizontal (H) eye velocity for each millisecond in the first 225 ms after the onset of target motion in different task configurations. Solid and dashed traces show results for the cue-and-target and target-only conditions. Black and gray traces show data when horizontal or vertical was the highly rewarded direction. For simplicity of presentation, all directions were transformed to plot in the first quadrant. Symbols on the traces indicate different times after the onset of target motion. **C, D**, For each component of eye velocity, the curves show the difference in each millisecond between the average component of eye velocity when it was in the highly rewarded versus less-rewarded direction. Left and right columns of the figure show data for two monkeys. Monkey P had lower eye speeds partly because we used lower target speeds of 20%/s, instead of 30%/s and partly because it had generally lower eye velocities in the first 100 ms of pursuit than did monkey I.

data, and the effect of the time of the cue was qualitatively the same in the two monkeys.

### Effect of reward size on directional learning in pursuit

In pursuit eye movements, learning is caused by sensory instructive signals that indicate errors in tracking. Thus, pursuit provides an example of error-correcting learning rather than a form of reinforcement learning. Still, the size of the reward may have an effect on either the expression of the learned response or the speed and strength of the acquisition of learning. It follows that analysis of the relationship between reward size and learning should reveal how the two processes interact in the conceptual organization of pursuit. Previous experiments have shown that pursuit learning of a change in target speed occurs upstream from vector averaging (Kahlon and Lisberger, 1999). Therefore, our finding here that reward modulation occurs at the site of vector averaging suggests that that learning would be upstream from reward modulation, as well.

We subjected monkeys to tracking conditions that cause a learned change in the direction of pursuit eye movements (Medina et al., 2005; Yang and Lisberger, 2010) under different reward conditions. The schematic at the top of Figure 8 illustrates

the configuration of the “learning” and “probe” trials used for these experiments. While the monkey fixated a stationary target, a second target appeared several degrees eccentric and indicated, by its color, whether successful completion of the trial would result in a large or small reward. After 800–1200 ms, the fixation point disappeared and the colored spot moved at constant speed for 250 ms in a fixed direction that we will call the “probe direction.” In learning trials, we provided an instructive change in target direction by having the target acquire an orthogonal component of motion in what we will call the “learning direction”; net target motion was in an oblique direction for the remainder of the trial. In probe trials, the target simply continued to move in the probe direction throughout the duration of the trial.

We characterize a learning experiment with the colored image in Figure 8A, in which color indicates eye velocity in the learning direction, each horizontal line of pixels indicates the time course of eye velocity in the learning direction for one trial, and learning progresses trial-by-trial from the bottom to the top of the image. In the first several learning trials, eye velocity was initiated in the direction of the probe target motion but there was no eye movement in the learning direction, as shown by the entirely blue color of the bottom lines of Figure 8A. As the learning trials progressed, monkeys first acquired a small eye velocity in the learning direction that appeared before the time when the target started to move in the learning direction (250 ms). Thus, all except the bottom two or three horizontal lines in Figure

8A are light blue starting ~150 ms after the onset of target motion in the probe direction. As learning progressed further, the size of eye velocity in the learning direction increased until it was close to 5%/s by the time of the instructive change in target direction, 250 ms after the onset of target motion in the probe direction. Figure 8A shows only the eye velocity between the onset of target motion in the probe direction and 50 ms after the instructive change in target direction. Any change in the eye movement in this interval must reflect the learned component of the movement, because it is too early to be driven by the instructive target motion in the learning direction. At later times, learning would be obscured by the direct eye velocity response to the instructive change in target motion. We quantified the size of the learned eye movement on the basis of the eye velocity in the interval from 200 to 300 ms after the onset of target motion.

The size of the learned eye velocity reached a larger value when the learning trials led to a larger reward. Figure 8B shows that the eye velocity in the learning direction, averaged over the 50th to the last learning trial, was consistently larger for learning blocks that used the large reward (blue curves) versus trials in learning blocks that used the small reward (red curves). The difference appeared over the entire duration of the eye velocity response in

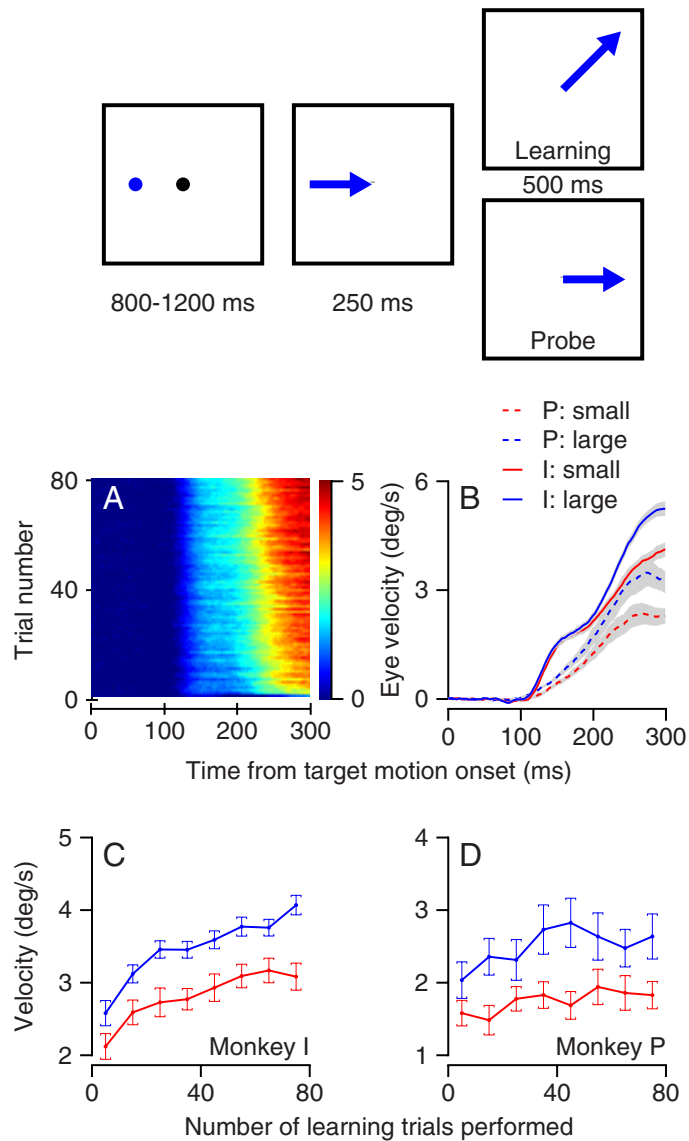


the learning direction. Changing the order of the blocks with the large and small rewards in learning trials confirmed that this was an effect of reward size and not fatigue.

The effect of reward on the size of the learned eye movement was present across the entire learning block (Fig. 8*C,D*). We binned the learning trials into groups of 10 for each experiment and took the average of the mean eye velocity in the interval from 200 to 300 ms after the target movement onset across experiments. The shapes of the resulting learning curves did not depend strongly on the size of the reward, but the amplitude of the learned response was larger in the high-reward blocks for almost every bin in both monkeys.

The effect of reward size on learning could be attributable to either (1) faster or bigger acquisition of learning under conditions in which the monkey receives large versus small reward or (2) modulation of the expression of learning without any effect on the learning process itself. We test these possibilities by measuring the learned eye movement in probe trials. Suppose, for example, that the learned eye movement is the same size during a probe trial with a target that cues a large reward as during a learning trial with a target that cues a small reward. Then, we would conclude that reward size modulates the time course or amplitude of the learning process but not the expression of a previously learned response. Conversely, suppose that the learned eye movement is larger during a probe trial with a target that cues a large reward than during a learning trial with a target that cues a small reward. Then we can conclude that reward modulates the expression of learning, at least to some degree. If reward modulates the expression of learning, then the locus of learning may be upstream from the site of reward modulation. The latter possibility is similar to findings that learning can occur without any explicit reward (Blodgett, 1929; Tolman, 1948).

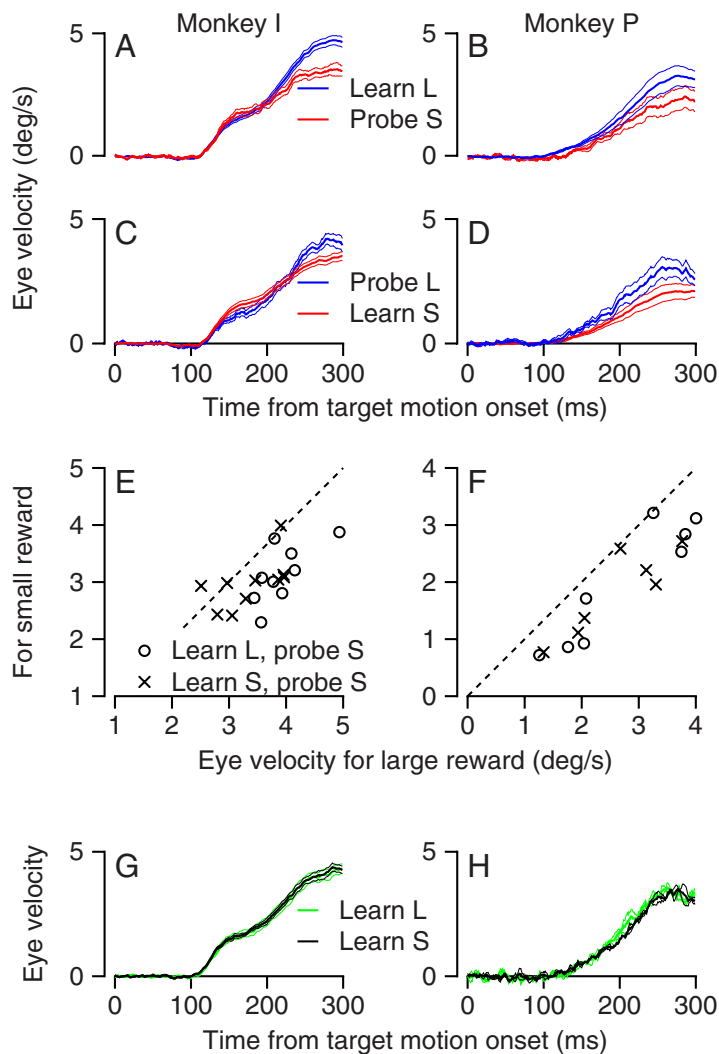
To discriminate between effects of reward size on acquisition versus expression of learning, we probed the learned response with targets that cued large or small rewards during and after learning induced with targets that cued the opposite, small or large rewards. As before, probes comprised target motions that started in the same direction as learning target motions but did not contain an instructive change in the direction of target motion. Any expression of the learned eye movement in probe trials must be attributed to previous learning. Note that the use of probes with the opposite reward size from the learning trials is an important element of the experimental design and ensures that the probe target color never was associated with the learning direction and hence did not contribute to learning.



**Figure 8.** Effect of reward size on the expression of directional learning in pursuit. The schematic at the top of the figure shows the structure of the learning paradigm, as described in Results. **A**, The progression of the learned eye velocity averaged across all learning experiments of monkey I. Colors represent eye velocity in the learning direction, and each horizontal line of the image shows eye velocity as a function of time for a single trial. The trials in a learning block progress from the bottom to the top of the image. **B**, The time course of the average eye velocity in the learning direction for the 50th trial to the end of learning block. Blue and red traces show results when the learning target motion cued and delivered large versus small rewards. Solid and dashed traces show data from monkeys I and P. The ribbons along the traces show the SEM across experiments. **C, D**, Learning curves showing the acquisition of learned eye velocity in bins of 10 learning trials. Blue and red show experiments in which the learning target received a large versus a small reward. Symbols and error bars indicate the mean and SE across experiments for the eye velocity averaged across the interval from 200 to 300 ms after the onset of target motion.

Our results contradict the alternative that the larger learned response is solely attributable to modulation of the acquisition of learning. In the experiment summarized by Figure 9, *A* and *B*, probe trials with a target that cued and led to a small reward (red curves) caused somewhat smaller eye velocities than interleaved learning trials with targets that cued and led to larger rewards (blue curves). In Figure 9, *C* and *D*, probe trials with a large reward (blue curves) caused somewhat larger eye velocities than did interleaved learning trials with smaller rewards.

We quantified the effect of reward size on the expression of learning by averaging the eye velocity in the learning direction in the 200–300 ms interval after the onset of target motion for learning and probe trials and plotting the eye velocity evoked for trials



**Figure 9.** Effect of reward size on the learned eye velocity in probe trials. *A–D*, Time course of the average eye velocity in the learning direction for interleaved learning and probe trials. Blue and red traces indicate responses to targets that cued and led to large or small rewards. Ribbons along the traces indicate SEM. In *A* and *B*, the learning trials provided large rewards and the probe trials provided small rewards. In *C* and *D*, the learning trials provided small rewards and the probe trials provided large rewards. *E*, *F*, Quantitative analysis of the effect of reward on the expression of the learning. Each symbol shows the average eye velocity from 200 to 300 ms after target movement during learning and probe trials in one learning block. Data are plotted so that the eye velocity for the targets with small (S) versus large (L) rewards always is on the y-axis versus x-axis. *G*, *H*, Time course of the average eye velocity in the learning direction for probe trials with intermediate reward size. Green and black traces indicate responses to probe trials that delivered medium-sized rewards in blocks in which learning trials delivered large or small rewards.

that led to a small reward versus that for trials that led to a large reward (Fig. 9*E,F*). Whether we induced learning with a large reward and probed with a small reward (open circles) or vice versa (× symbols), almost all the points lay below the unity line. Therefore, both of our monkeys emitted larger eye velocities for targets that cued larger rewards without regard for whether the target motion comprised a probe trial or a learning trial. We conclude that at least part of the effect of reward on the learned eye velocity in the learning trials (Fig. 8*B*) is attributable to modulation of the expression of the learned eye movement. Given that reward appears to act at the site of vector averaging, our results on learning are compatible with the previous conclusion (Kahlon and Lisberger, 1999) that learning occurs upstream from vector averaging.

In an independent test of whether reward affected the acquisition or expression of learning, we ran a separate set of learning experiments with probe trials that were delivered after the learning block, always used a cue with a different color (red), and led to

a medium size reward. In different experiments, the reward for the probe trials always was the same, and learning trials led to large or small rewards. If reward modulates the acquisition of learning, then we would expect to find responses in probe trials that depended on whether learning was induced with a small or large reward, although the probe trials always delivered the same, medium-sized reward. We did not find any significant differences in the average eye velocity in probe trials that delivered medium reward after learning with large or small rewards (Fig. 9*G,H*,  $p > 0.89$ , Mann–Whitney test, 300 ms after target movement onset). We conclude that reward modulates primarily the expression of learning and that it does so in both learning and probe trials. Still, additional experiments will be needed to fully exclude the possibility that reward size also modulates the size or speed of the acquisition of the learned response.

## Discussion

The motivation provided by reward has a powerful effect on human and animal behavior. The past 15 years have seen considerable interest in the mechanisms of action of reward, with the primary insight that the basal ganglia contain a flexible representation of the actual and expected reward. However, identification of a “code” for reward leaves open the question of how and where that code alters behavior. Our premise in studying the effect of reward size on smooth pursuit eye movements was that we could capitalize on the extensive knowledge of the neural basis for pursuit and its relation to different components of pursuit behavior. Thus, our experiments were designed to determine which components of the pursuit behavior are affected by reward.

We have taken advantage of two attributes of the pursuit system. First, we have restricted our attention to the earliest parts of the response before there has been time for visual feedback. For the initiation of pursuit, this has meant the first 100 ms of the response, up to 200 ms after the onset of target motion. For the learned responses, it has meant the component of eye velocity up to 50 ms after the instructive change in target direction. As a result, we can interpret our data in the context of the flow of neural signals from sensory to motor parts of the circuit in that brief interval. Second, we have asked how pursuit is affected by reward under different conditions that exercise different features of the pursuit response: in the presence of a single target or two targets, and in a learning task. Because we already know much about the flow of neural signals under each of these conditions, our results allow us to localize reward modulation within the functional signal flow of the pursuit circuit.

### The location of reward modulation

Our data indicate that reward modulation occurs downstream from sensory processing, downstream from the site of pursuit learning, and upstream from final motor processing. If we assume a single site of reward modulation, then our experiments with two-target trials, single target trials, and learning all are consistent with reward modulation at the site at which vector averaging occurs and multiple targets compete for control of pursuit eye movement.

Our finding of a bias in the weighting of the highly rewarded target in pursuit initiation for two-target stimuli argues against the hypothesis that the site of reward modulation is in the sensory representation of the highly rewarded target. Additional evidence against a sensory locus comes from the effect of reward on a learned change in target direction. Because the learned eye movement is orthogonal to the sensory stimulus, it probably is not driven directly by the sensory stimulus: reward must have an effect outside of sensory processing. Thus, it seems unlikely that reward modulation occurs in the middle temporal visual area, a key sensory structure in the processing of visual motion information for pursuit (Newsome and Paré, 1988; Ferrera and Lisberger, 1997; Groh et al., 1997). In agreement with our conclusion that reward modulation occurs downstream from sensory processing, previous papers have shown that the first 100 ms of the responses of MT neurons are spared modulation by extra-retinal signals such as attention (Lee and Maunsell, 2010). Furthermore, Ferrera and Lisberger (1997) showed that the response biases in MT and the medial superior temporal area MST are far too small to account for the high degree of selectivity shown by pursuit behavior when monkeys use a color cue to choose between two targets moving in different directions.

Our data on the effects of reward size on the mean weighting of two targets might be explained by a competitive neural circuit based on reciprocal inhibition and excitation to implement weighted vector averaging (Ferrera, 2000). In principle, a competitive network could amplify the small effects of reward size on the sensory signals for two targets, leading to the large and nonlinear biases seen in our data, even if the locus of reward action was upstream from vector averaging. However, three other pieces of evidence contradict predictions of the competitive network. (1) Targets of different speeds are equally weighted when they are used in a two-target stimulus (Lisberger and Ferrera, 1997); a competitive network might be expected to weight faster targets more strongly. (2) Target luminance has a larger effect on the responses to single targets than we found here for reward but biases the responses to two-target stimuli more weakly than does reward size when the targets appear at different spatial locations as in our design here (Niu and Lisberger, 2011); a competitive network based on imbalances in the sensory inputs might be expected to weight brighter targets much more strongly. (3) Because of the trial-by-trial variation in the strength of sensory signals, a competitive network should cause a large increase in the variance for two-target stimuli, relative to the variance predicted by simple vector averaging of the responses to the two targets singly. Niu and Lisberger (2011) as well as an analysis of the variation in our data fail to support this prediction. Thus, the weight of the evidence appears to contradict the possibility that reward acts entirely upstream from a competitive mechanism for combining multiple motion signals.

The effects of reward on the pursuit evoked by motion of two targets argue against a purely motor locus for reward modulation. If the site of reward modulation were in the motor system, downstream from the site where two target motions compete to control pursuit, then reward size should alter the magnitude of eye velocity but would not be able to bias the direction of eye

velocity in two-target trials toward the direction of the highly rewarded target. If we assume that there is a single site of reward modulation, then that site cannot be in the final motor processing. If, however, we allow multiple sites of reward modulation, then a site in the motor system would be consistent with the modest effects of reward size on pursuit of a single target and on the expression of pursuit learning.

### Loci in which two moving targets compete for the initiation of pursuit

There are at least two loci within the pursuit circuit in which two targets compete for control of the initiation of pursuit. At one locus, hypothesized to be in the primary visual cortex or extrastriate area MT, stimuli compete if they occupy overlapping regions of the visual field, and the relative luminance of the two moving stimuli is important (Niu and Lisberger, 2011). If one stimulus is eight times brighter than the other, then the initiation of pursuit is dominated by the brighter stimulus, although the dimmer stimulus is a very effective pursuit target if presented singly. The second locus for target competition seems to be downstream from sensory processing, dominates the pursuit response when the targets moved across separate parts of the visual field, and does not support strong competition between two targets based on their relative luminance (Niu and Lisberger, 2011).

The precedence of a more highly rewarded target in our present data resembles, as a behavioral phenomenon, the dominance of the brighter of two targets. However, the competition on the basis of luminance appears to occur in early visual processing, and we have argued above that the locus of the reward modulation is downstream from area MT. In addition, in the present experiments, we used pairs of targets that were separated by at least 5° in the visual field so that the size of receptive fields in MT would not allow direct interaction of the two target motions. Therefore, we think that the mechanism of reward action is different from the mechanism for modulation by relative luminance.

We think that the modulation of pursuit by reward size is more closely related to paradigms that bias the initial pursuit response toward one target or the other by providing cues about the direction of the future target motion (Garbutt and Lisberger, 2006; Shichinohe et al., 2009). Biases with longer latencies occur when monkeys are cued about the direction of the required eye movement at the onset of target motion (Ferrera and Lisberger, 1995; Ferrera, 2000; Case and Ferrera, 2007; Mahaffy and Krauzlis, 2011a). These different forms of modulation of the initiation of pursuit all share the property that they rely on voluntary control and that they occur with pursuit targets that occupy different regions of the visual field. We think they may occur at a single neural locus, and we propose that the use of reward size provides the simplest approach to study the neural mechanism of the modulation, because it takes so little time or effort to induce a monkey to track a highly rewarded target preferentially.

### Possible anatomical locations of reward modulation

The basal ganglia are specifically linked to the presence or absence of a reward and its effect on behavior (Schultz, 1998; Lauwereyns et al., 2002; Joshua et al., 2008), and we imagine that they also play an important role in modulation of pursuit behavior by reward size. There is a strong anatomical relationship between the basal ganglia and the smooth eye movement region of the frontal eye fields, or FEF<sub>SEM</sub> (Tian and Lynch, 1997; Cui et al., 2003). Activity in the basal ganglia is related to pursuit (O'Driscoll et al., 2000; Basso et al., 2005). Therefore, the FEF<sub>SEM</sub> could be the location of reward modulation.

The main features of the FEF<sub>SEM</sub> are as follows: (1) its neurons are direction selective for pursuit eye movements (Gottlieb et al., 1994; Tanaka and Fukushima, 1998); (2) its output appears to modulate the strength, or gain, of visual–motor transmission (Tanaka and Lisberger, 2001); and (3) different neurons in the FEF<sub>SEM</sub> contribute to pursuit and pursuit learning selectively over different narrow time segments of the total pursuit response (Schoppik et al., 2008; Li and Lisberger, 2011). The role of the FEF<sub>SEM</sub> in modulating the gain of visual–motor transmission fits well with our finding that the effects of reward modulation can be understood in terms of modulation of the weight of the visual motion signals from one or multiple targets. Other studies of the FEF<sub>SEM</sub> show that its activity mostly follows the selection of a pursuit target, in agreement with the idea that modulation by reward and other voluntary mechanisms occurs in the inputs to, or directly on, the neurons of the FEF<sub>SEM</sub> (Mahaffy and Krauzlis, 2011a,b). The FEF<sub>SEM</sub> need not act alone in modulation of pursuit eye movement, and other evidence indicates that the superior colliculus also could play a role (Nummela and Krauzlis, 2011).

Available data clearly place the locus of reward modulation of pursuit eye movements and other forms of target selection at a middle level of sensory–motor processing, both after the early sensory representation of target motion and before the final motor commands are assembled in the cerebellum. The hypothesis that a representation of reward in the basal ganglia modulates pursuit initiation through the FEF<sub>SEM</sub>, with a possible role for the superior colliculus as well, is appealing in light of available data.

## References

- Arkadir D, Morris G, Vaadia E, Bergman H (2004) Independent coding of movement direction and reward prediction by single pallidal neurons. *J Neurosci* 24:10047–10056.
- Basso MA, Pokorny JJ, Liu P (2005) Activity of substantia nigra pars reticulata neurons during smooth pursuit eye movements in monkeys. *Eur J Neurosci* 22:448–464.
- Blodgett HC (1929) The effect of the introduction of reward upon the maze performance of rats. Berkeley, CA: University of California.
- Carey MR, Medina JF, Lisberger SG (2005) Instructive signals for motor learning from visual cortical area MT. *Nat Neurosci* 8:813–819.
- Case GR, Ferrera VP (2007) Coordination of smooth pursuit and saccade target selection in monkeys. *J Neurophysiol* 98:2206–2214.
- Cui DM, Yan YJ, Lynch JC (2003) Pursuit subregion of the frontal eye field projects to the caudate nucleus in monkeys. *J Neurophysiol* 89:2678–2684.
- Ding L, Hikosaka O (2006) Comparison of reward modulation in the frontal eye field and caudate of the macaque. *J Neurosci* 26:6695–6703.
- Ferrera VP (2000) Task-dependent modulation of the sensorimotor transformation for smooth pursuit eye movements. *J Neurophysiol* 84:2725–2738.
- Ferrera VP, Lisberger SG (1995) Attention and target selection for smooth pursuit eye movements. *J Neurosci* 15:7472–7484.
- Ferrera VP, Lisberger SG (1997) Neuronal responses in visual areas MT and MST during smooth pursuit target selection. *J Neurophysiol* 78:1433–1446.
- Garbutt S, Lisberger SG (2006) Directional cuing of target choice in human smooth pursuit eye movements. *J Neurosci* 26:12479–12486.
- Gardner JL, Lisberger SG (2001) Linked target selection for saccadic and smooth pursuit eye movements. *J Neurosci* 21:2075–2084.
- Gottlieb JP, MacAvoy MG, Bruce CJ (1994) Neural responses related to smooth-pursuit eye movements and their correspondence with electrically elicited smooth eye movements in the primate frontal eye field. *J Neurophysiol* 72:1634–1653.
- Groh JM, Born RT, Newsome WT (1997) How is a sensory map read out? Effects of microstimulation in visual area MT on saccades and smooth pursuit eye movements. *J Neurosci* 17:4312–4330.
- Joshua M, Adler A, Mitelman R, Vaadia E, Bergman H (2008) Midbrain dopaminergic neurons and striatal cholinergic interneurons encode the difference between reward and aversive events at different epochs of probabilistic classical conditioning trials. *J Neurosci* 28:11673–11684.
- Kahlon M, Lisberger SG (1999) Vector averaging occurs downstream from learning in smooth pursuit eye movements of monkeys. *J Neurosci* 19:9039–9053.
- Lauwereyns J, Watanabe K, Coe B, Hikosaka O (2002) A neural correlate of response bias in monkey caudate nucleus. *Nature* 418:413–417.
- Lee J, Maunsell JH (2010) Attentional modulation of MT neurons with single or multiple stimuli in their receptive fields. *J Neurosci* 30:3058–3066.
- Li JX, Lisberger SG (2011) Learned timing of motor behavior in the smooth eye movement region of the frontal eye fields. *Neuron* 69:159–169.
- Lisberger SG, Ferrera VP (1997) Vector averaging for smooth pursuit eye movements initiated by two moving targets in monkeys. *J Neurosci* 17:7490–7502.
- Mahaffy S, Krauzlis RJ (2011a) Neural activity in the frontal pursuit area does not underlie pursuit target selection. *Vision Res* 51:853–866.
- Mahaffy S, Krauzlis RJ (2011b) Inactivation and stimulation of the frontal pursuit area change pursuit metrics without affecting pursuit target selection. *J Neurophysiol* 106:347–360.
- Medina JF, Carey MR, Lisberger SG (2005) The representation of time for motor learning. *Neuron* 45:157–167.
- Newsome WT, Paré EB (1988) A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J Neurosci* 8:2201–2211.
- Newsome WT, Wurtz RH, Dürsteler MR, Mikami A (1985) Deficits in visual motion processing following ibotenic acid lesions of the middle temporal visual area of the macaque monkey. *J Neurosci* 5:825–840.
- Niu YQ, Lisberger SG (2011) Sensory versus motor loci for integration of multiple motion signals in smooth pursuit eye movements and human motion perception. *J Neurophysiol* 106:741–753.
- Nummela SU, Krauzlis RJ (2011) Superior colliculus inactivation alters the weighted integration of visual stimuli. *J Neurosci* 31:8059–8066.
- O'Driscoll GA, Wolff AL, Benkelfat C, Florencio PS, Lal S, Evans AC (2000) Functional neuroanatomy of smooth pursuit and predictive saccades. *Neuroreport* 11:1335–1340.
- Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233–238.
- Ramachandran R, Lisberger SG (2005) Normal performance and expression of learning in the vestibulo-ocular reflex (VOR) at high frequencies. *J Neurophysiol* 93:2028–2038.
- Recanzone GH, Wurtz RH (1999) Shift in smooth pursuit initiation and MT and MST neuronal activity under different stimulus conditions. *J Neurophysiol* 82:1710–1727.
- Schoppik D, Nagel KI, Lisberger SG (2008) Cortical mechanisms of smooth eye movements revealed by dynamic covariations of neural and behavioral responses. *Neuron* 58:248–260.
- Schultz W (1998) Predictive reward signal of dopamine neurons. *J Neurophysiol* 80:1–27.
- Shichinohe N, Akao T, Kurkin S, Fukushima J, Kaneko CR, Fukushima K (2009) Memory and decision making in the frontal cortex during visual motion processing for smooth pursuit eye movements. *Neuron* 62:717–732.
- Takikawa Y, Kawagoe R, Itoh H, Nakahara H, Hikosaka O (2002) Modulation of saccadic eye movements by predicted reward outcome. *Exp Brain Res* 142:284–291.
- Tanaka M, Fukushima K (1998) Neuronal responses related to smooth pursuit eye movements in the periarculate cortical area of monkeys. *J Neurophysiol* 80:28–47.
- Tanaka M, Lisberger SG (2001) Regulation of the gain of visually guided smooth-pursuit eye movements by frontal cortex. *Nature* 409:191–194.
- Tanaka M, Lisberger SG (2002) Role of arcuate frontal cortex of monkeys in smooth pursuit eye movements. I. Basic response properties to retinal image motion and position. *J Neurophysiol* 87:2684–2699.
- Tian JR, Lynch JC (1997) Subcortical input to the smooth and saccadic eye movement subregions of the frontal eye field in Cebus monkey. *J Neurosci* 17:9233–9247.
- Tolman EC (1948) Cognitive maps in rats and men. *Psychol Rev* 55:189–208.
- Yang Y, Lisberger SG (2010) Learning on multiple timescales in smooth pursuit eye movements. *J Neurophysiol* 104:2850–2862.