

Vector Averaging Occurs Downstream from Learning in Smooth Pursuit Eye Movements of Monkeys

Maninder Kahlon and Stephen G. Lisberger

Howard Hughes Medical Institute, Department of Physiology, Neuroscience Graduate Program, and W. M. Keck Foundation Center for Integrative Neuroscience, University of California, San Francisco, California 94143

How are sensory-motor transformations organized in a cortical motor system? In general, sensory information is transformed through a variety of signal processing operations in the context of distinct coordinate frameworks. We studied the interaction of two distinct operations in pursuit eye movements, learning and vector-averaging, to gain insight into their underlying coordinate frameworks and their sequence in sensory-motor processing. Learning was induced in the initiation of pursuit eye movements by targets that moved initially at one speed for 100 msec and then increased or decreased to a sustained final speed. Vector averaging was studied by comparing the initial eye acceleration evoked by the simultaneous motion of two targets with that evoked by each target singly. Learning caused specific effects on the direction of the vector-averaged responses to two-target stimuli that included one target moving in the direction used to induce learning. Learned increases or de-

creases in eye acceleration caused the direction of the responses to two-targets to rotate toward or away from the learning direction. Learning also caused nonspecific changes in the responses to two-target stimuli. After any learning protocol, two-target responses usually became smaller, and their directions rotated away from the axis of the target motion used for learning. Quantitative analysis showed that the specific effects of learning were predicted most closely by a model in which vector averaging occurs downstream from the site(s) of learning. We suggest that the pursuit system creates parallel commands for potential movements to each of the targets in two-target stimuli, and that learning occurs in the coordinates of the potential movements.

Key words: oculomotor system; sensory-motor transformation; coordinate system; population code; visual-motor processing; learning

Conversion of sensory inputs into motor outputs involves a complex series of neural transformations. These transformations can be described in terms of the coordinate system of the representation at each stage of processing. For visual-motor processing, sensory signals encoded in retinal coordinates are converted into intermediate sensorimotor coordinates (Andersen et al., 1993) and then into motor coordinates that specify muscle contractions. Superimposed on these coordinate transformations are a variety of signal processing operations that compute the metrics, kinematics, and dynamics of the movement. What are the coordinate frameworks for sensory-motor transformations, and how are they related to the signal processing operations that create them?

Pursuit eye movements provide an ideal system to study the coordinate transformations underlying signal processing operations of the brain. First, something is known about the coordinate systems at the different sites in the neural pursuit system. The specifics of the representation of sensory information in retinal coordinates (Lisberger and Movshon, 1999) and the encoding of commands for eye movement in motor coordinates (Skavenski and Robinson, 1973; Shidara et al., 1993; Krauzlis and Lisberger, 1994; Van der Steen et al., 1994) have been described quantita-

tively. A number of recent behavioral observations have suggested that much of pursuit processing is done in an intermediate, world-centered coordinate framework (Grasse and Lisberger, 1992; Kahlon and Lisberger, 1996; Kiorpes et al., 1996). Second, several studies have revealed diverse signal processing operations that are part of the generation of pursuit. These include predictive pursuit (Kowler, 1990), on-line gain control (Goldreich et al., 1992; Schwartz and Lisberger, 1994), target selection (Ferrera and Lisberger, 1995), learning (Optican et al., 1985; Kahlon and Lisberger, 1996; Ogawa and Fujita, 1997), and most recently vector averaging (Lisberger and Ferrera, 1997).

We have now used behavioral approaches to determine the coordinate systems and the relative placements of two of these operations: learning and vector averaging. Learning occurs in the initial pursuit response during repeated presentation of targets that move at one speed for 100 msec and then change to a higher or lower speed (Optican et al., 1985; van Donkelaar et al., 1994; Kahlon and Lisberger, 1996; Ogawa and Fujita, 1997). Previous behavioral analysis revealed that learning in pursuit is expressed in coordinates related to eye or target motion in the world, rather than to image motion on the retina (Kahlon and Lisberger, 1996). Vector-averaged pursuit responses occur when a monkey is presented simultaneously with two potential targets (Lisberger and Ferrera, 1997). Although available evidence does not favor any hypothesis for the coordinate system of vector averaging, we and others have favored tacitly the hypothesis that it occurs in retinal coordinates as an operation on image motion signals (Groh et al., 1997; Lisberger and Ferrera, 1997; Recanzone et al., 1997).

We evaluated the relative placement of vector averaging and learning in the neural circuitry of pursuit by comparing vector-

Received Dec. 31, 1998; revised Aug. 5, 1999; accepted Aug. 5, 1999.

This work was supported by National Institutes of Health Grant P01-NS34835 (S.G.L.) and by stipend support from the Boyer Fund and the University of California Regents (M.K.). S.G.L. is an investigator of the Howard Hughes Medical Institute. We thank Stefanie Tokiyama for excellent technical assistance, Mark Kvale for helpful discussions, and Allan Basbaum, Michael Stryker, Ken Miller, and Fred Miles for thoughtful comments on this manuscript.

Correspondence should be addressed to Stephen G. Lisberger, Department of Physiology, Box 0444, 513 Parnassus Avenue, Room S-762, University of California, San Francisco, CA 94143-0444.

Copyright © 1999 Society for Neuroscience 0270-6474/99/199039-15\$05.00/0

averaged pursuit responses before and after learning. The learning-related changes in the direction and magnitude of vector-averaged responses were predicted best by the hypothesis that learning occurs upstream from vector averaging. Our data imply that vector averaging occurs quite late in pursuit processing and operates in a directional coordinate framework. The inputs to vector averaging appear to represent commands for two potential movements, already transformed by learning.

MATERIALS AND METHODS

Behavioral experiments were conducted on four rhesus monkeys. Surgical and behavioral methods have been described previously (Lisberger and Westbrook, 1985; Kahlon and Lisberger, 1996). Briefly, monkeys were first trained to attend to spots of light in a bar press task for liquid reinforcements. Using isoflurane anesthesia and aseptic conditions, head holders were implanted on the skull of each monkey. At the same time, a coil of wire was implanted in one eye to measure voltages proportional to eye position with the scleral search coil technique (Judge et al., 1980). After postsurgical recovery, monkeys were trained to track the slow movement of small spots of light. The animals used in the experiments presented in this paper were overtrained on such tracking tasks. In previous experiments, two had tracked targets in pursuit learning paradigms (Kahlon and Lisberger, 1996), and two had tracked two-target stimuli used to reveal vector averaging (Lisberger and Ferrera, 1997). Each daily session consisted of one learning experiment and lasted ~2 hr.

Visual stimuli. Visual targets were generated by a digital signal processing board on a Pentium computer and displayed on a 12-inch diagonal oscilloscope (1304A, P-4 phosphor; Hewlett-Packard, Palo Alto, CA). The screen was 40 cm from the monkey and provided a $32 \times 26^\circ$ visual display. The system provided a spatial resolution of $65,536 \times 65,536$ pixels and a temporal resolution of 4 msec. Pursuit targets were 0.4° squares, had a luminance of 3.5 cd/m^2 , and were presented on a uniform gray background. All experiments were conducted in a moderately lit room.

Experimental paradigm. Animals tracked targets in a series of trials. Each trial had the basic structure illustrated in Figure 1*A*. A fixation target appeared in the center of the monitor for a random interval of 500–900 msec. When the fixation target was extinguished, one or two pursuit targets appeared at 3° left, right, up, or down relative to the fixation target and began moving toward the center of the monitor. The monkey was given a grace period of 350 msec to let his eye catch up with the target, after which he had to keep eye position within a $\pm 4^\circ$ square window centered on the target. If the monkey completed the trial successfully, he received fluid reinforcement.

Single-target trials provided either single or double steps of target speed. Single steps (Fig. 1*A*) were used in “control trials” that delivered target motion at $20^\circ/\text{sec}$. Double steps of target speed (Fig. 1*B*) were used as “learning trials.” In learning trials, the target moved at $20^\circ/\text{sec}$ for 100 msec before undergoing a step change to another speed. For the target shown in Figure 1*B*, target speed stepped down to $5^\circ/\text{sec}$ to cause learned decreases in the eye acceleration at the initiation of pursuit. In trials designed to increase eye acceleration (results not shown), the target speed was $20^\circ/\text{sec}$ for 100 msec and then stepped up to $40^\circ/\text{sec}$. Learning trials providing leftward or rightward target motion were always interleaved with control trials in the opposite direction. Thus, there were four possible combinations of learning: increases and decreases in eye acceleration for leftward or rightward target motion. As shorthand, we will refer to these experiments as left-increase, left-decrease, right-increase, or right-decrease experiments.

In two-target trials, two identical targets appeared simultaneously at two different locations that were 3° eccentric: left, right, up, or down relative to the fixation light. For 148 msec, both targets moved at $20^\circ/\text{sec}$ toward the position of fixation. With equal probability, one of the targets then disappeared, and the other became the tracking target for the monkey. Therefore, each two-target combination was repeated in two separate trials, with two distinct final pursuit targets but with identical initial motion of two targets. In the schematic diagram of Figure 3*A*, the two targets started at 3° left and 3° down relative to the fixation light and moved to the right and up. The upward target disappeared after 148 msec, and the rightward target continued to move at $20^\circ/\text{sec}$ for at least another 600 msec. The five other combinations of two-target motions provided three other pairs that interacted orthogonal directions of target motion (up and left, left and down, and down and right) and two pairs

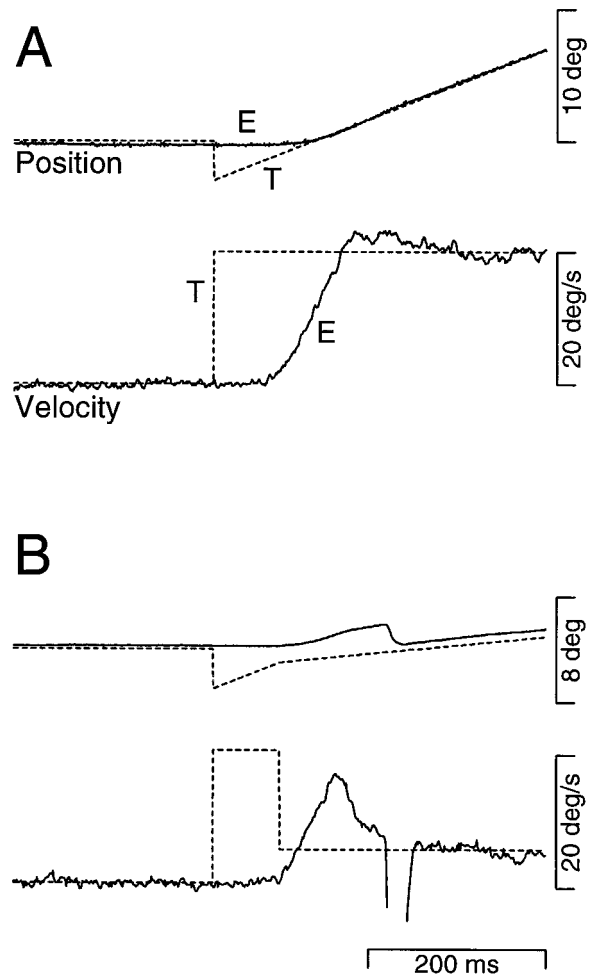


Figure 1. Basic structure of trials used to evoke pursuit. *A*, Single-step trial. The pursuit target moved at $20^\circ/\text{sec}$ for the entire duration of the trial. *B*, Double-step learning trial that caused decreases in eye acceleration. The pursuit target began to move at $20^\circ/\text{sec}$. After 100 msec, velocity was stepped down to $5^\circ/\text{sec}$. In *A* and *B*, the top traces show superimposed eye (*E*) and target (*T*) position, and the bottom traces show eye and target velocity, respectively. Dashed traces show target position and velocity. Solid traces show eye position and velocity. Data are shown starting 300 msec before the pursuit target began to move.

that interacted opposite directions of target motion (left and right and up and down).

Each experiment consisted of 20 blocks of prelearning tests, followed by 500 blocks of learning and control trials and finally another 20 blocks of postlearning tests. In three monkeys we performed two experiments in two separate configurations. One configuration was designed to determine the effect of pursuit learning on vector averaging for two-target stimuli. Prelearning and postlearning tests provided target motion in single-target and two-target trials. Each block of trials provided target motion in 12 two-target trials (6 combinations of two initial targets \times 2 final tracking targets) and 4 single-target trials in the four orthogonal directions. Animals repeated this experiment on 16 d: 4 d for each combination of leftward versus rightward learning trials and learned increases versus decreases in eye acceleration. A second configuration was designed to evaluate the generalization of pursuit learning across directions of single-target motion. Prelearning and postlearning tests provided target motion only in single-target trials, in which each trial provided motion in one of 12 directions sampled at 30° intervals. Pursuit targets were presented in step-ramp motion, configured so that the ramp took the target from 3° eccentric back through the position of fixation. Animals repeated this experiment on 12 d: 3 d for each combination of rightward versus leftward learning trials and learned increases versus decreases in eye acceleration. In one animal we combined the two

experimental configurations. Prelearning and postlearning tests provided 12 combinations of two-target motion and 12 directions of single-target motion. Monkey N completed this experiment in 16 d.

Data acquisition and analysis. Experiments were run and data acquired with a 90 MHz Pentium-based computer. This computer communicated over the local area network with a UNIX workstation that provided a user interface for determining experimental parameters. Voltages proportional to eye position, obtained from the magnetic search coil electronics, were differentiated by an analog circuit (bandpass DC to 25 Hz; -20 dB/decade) to generate eye velocity signals. Voltages proportional to horizontal and vertical eye position and velocity were sampled at 1000 Hz/channel and saved on disk along with codes representing the commands sent to the display oscilloscope. The codes were used to reconstruct horizontal and vertical target position and velocity for data analysis.

For data analysis, horizontal and vertical eye position and velocity were displayed on the video monitor and marked using software that ran on a UNIX workstation. The first and last 20 learning trials were analyzed first. Trials that lacked saccades in the first 200 msec after the onset of target motion were aligned on the onset of target motion and averaged. The averages of eye velocity in the first and last 20 pursuit learning trials were then superimposed and compared to select the analysis-interval that displayed the greatest effects of learning. These intervals were 128–176 msec (monkey I) or 138–186 msec (monkeys K, E, and N) after the onset of target motion. All the intervals corresponded approximately to the second 48 msec of pursuit eye movements, which has been shown previously to express the greatest changes in eye acceleration after a sequence of learning trials (Kahlon and Lisberger, 1996).

Once the analysis interval had been chosen, data in prelearning and postlearning tests were viewed individually for each trial. Horizontal and vertical eye velocity traces were marked and discarded if saccades occurred in or before the analysis intervals. In all but one case, this allowed us to retain at least 85% of two-target trials and 95% of single target trials but still to analyze only presaccadic smooth eye velocity. The one exception was monkey N, whose upward pursuit included many early saccades so that 50% of the responses to upward moving single targets had to be discarded. Prelearning and postlearning test trials were then analyzed separately by dividing the trials into groups that presented the same target or targets, aligning on the onset of target motion, and computing averages of horizontal and vertical eye velocity. For two-target trials, this grouping included the two types of trials that started with the same pair of target motions but ended with either of the two stimuli as the tracking target. It was legitimate to group the two trials for each pair of two-target stimuli, because we quantified responses only during intervals that were driven by the motion of the two targets, before the disappearance of one of the targets. We then used the averaged traces to measure average horizontal and vertical eye acceleration as the change in eye velocity across the 48 msec period of analysis, divided by 0.048. Most of the data displayed in the paper show the mean of these measurements across multiple repeats of the same experiment (usually four). We show SDs in Figure 4 and data taken from multiple daily experiments for one monkey in Figure 5 to provide estimates of trial-to-trial and day-to-day variability in our measurements.

RESULTS

Hypotheses for the effects of learning on responses to two-target stimuli

Figure 2 illustrates two predictions for the effect of learning on the responses to two-target stimuli, depending on whether the neural sites of learning are upstream or downstream from those for vector averaging in the flow of signals that guide pursuit eye movements. Each prediction is for experiments that cause a learned decrease in rightward eye acceleration, under the assumption that learning does not cause changes in the direction of responses to the upward motion of single targets.

(1) If vector averaging occurs upstream from learning (Fig. 2A), then learning may cause changes in the amplitude but not in the direction of the responses to two-target stimuli. In Figure 2A, *leftmost diagram*, the vectors simulate the responses to upward or rightward motion of single targets. The result of vector averaging before learning (Fig. 2A, *middle diagram*, *Average*) is a response

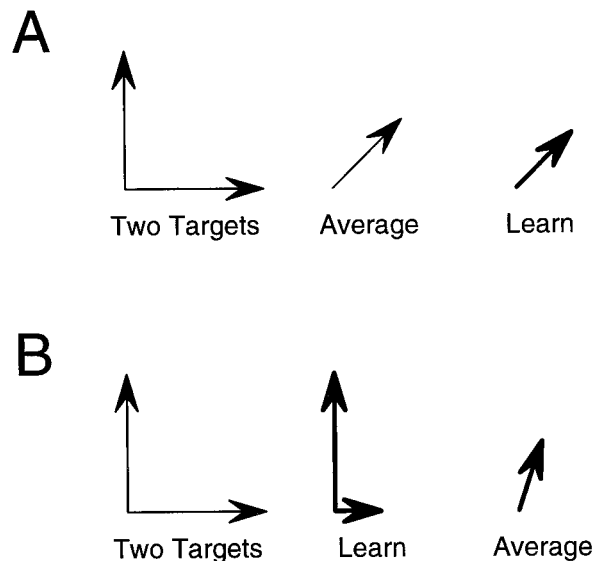


Figure 2. Two hypotheses for the relative placement of vector averaging and learning. *A*, Learning is downstream of vector averaging. *B*, Learning is upstream from vector averaging. Each arrow scheme uses polar notation to summarize the steps in converting the visual inputs from a two-target stimulus consisting of rightward and upward target motions into a command for a single pursuit eye acceleration. The *arrows* labeled *Two Targets* at the *left* are the same in *A* and *B* and show equal amplitude signals related to rightward and upward target motion. *Bold arrows* indicate steps that are after learning-induced decreases in eye acceleration for rightward target motion. Schemes that contain *one arrow* are after vector averaging; schemes that contain *two arrows* are before vector averaging.

that is oblique upward and rightward with an amplitude equal to half the amplitude of the sum of the rightward and upward vectors. Because subsequent learning operates on the averaged response, it causes only small changes in the amplitude of the response to two-target stimuli without changes in the direction of the response (Fig. 2A, *rightmost diagram*, *Learn*).

(2) If vector averaging occurs downstream from learning (Fig. 2B), then learning should cause a change in the direction of the responses to two-target stimuli. The leftmost, prelearning diagram of Figure 2B is the same as that in Figure 2A and simulates responses to rightward and upward single-target motions. If learning causes a decrease in the size of the response to a rightward target with no change in the size or direction of the response to an upward moving target (Fig. 2B, *middle diagram*, *Learn*), then subsequent vector averaging predicts a change in the direction of smooth eye movement. The simulated postlearning response (Fig. 2B, *rightmost diagram*, *Average*) rotates away from rightward (i.e., toward upward) relative to the response predicted for the same two-target stimulus if averaging occurs before learning (Fig. 2A, *rightmost diagram*).

We will discriminate between these two alternatives by showing that learning causes consistent and specific changes in the direction of the responses to two-target stimuli.

Two additional hypotheses seemed plausible but are not illustrated in Figure 2. In one, the mechanism of learning is a change in the weights used for vector averaging. This hypothesis takes the same general computational form but has different predictions from the two outlined above. It will be tested and rejected in the section of the paper that considers quantitatively the order of learning and vector averaging. In the other, which we call the




Figure 3. Vector-averaged responses to two-target stimuli. *A*, Schematic description of presentation of two-targets. The *dashed vectors* labeled *Target* and *Distractor* indicate rightward and upward moving targets that start 3° left and down, respectively, and move toward the fixation point, shown by \times . The *solid vector* labeled *Eye* shows an initial eye movement response that is a vector average of the two target motions. *B*, Average horizontal (*H*) and vertical (*V*) eye velocities for the stimuli shown in *A*. *Solid traces* describe the response to the two-target stimulus. *Short dashed traces* show the eye velocity response to a single rightward moving target. *Long dashed traces* show the eye velocity response to a single upward moving target. The *vertical arrows* indicate the time that was 70 msec after the distractor disappeared and coincides with the end of the interval in which pursuit is influenced by the 148 msec of two-target motion. The traces in *B* start 100 msec before the onset of target motion.

“motor hypothesis,” learning occurs after vector averaging but also after the pursuit signals have been divided into separate commands for the horizontal and vertical extraocular muscles. *A priori* the motor hypothesis seemed unlikely to be true, because it was incompatible with some of our behavioral findings (Kahlon and Lisberger, 1996). It also predicts agreement we did not observe between the effects of learning on the responses to single-target and two-target stimuli in the experiments of the present paper (see discussion of Figs. 11, 12 below).

Vector averaging for two-target stimuli

When two moving, potential targets were presented simultaneously to an animal trained to pursue single targets, the initial smooth eye velocity was between the responses that would have been evoked by either target separately (Lisberger and Ferrera, 1997). In the example shown schematically in Figure 3*A*, two targets moved up and right for 148 msec before the upward moving target disappeared and the rightward target became the tracking target. The data traces in Figure 3*B* show the time courses of the average horizontal and vertical eye velocities evoked by these stimuli for a typical experiment. For single targets, the responses to both the upward (*long dashed traces*) and the rightward (*short dashed traces*) target motion consisted of large smooth eye velocities in the direction of motion and small changes in eye velocity in the orthogonal direction. The response to simultaneous motion of two targets, rightward and upward, was intermediate to the response to either the rightward- or the upward moving single target (Fig. 3*B*, *solid traces*). In the first 148 msec after the onset of pursuit (Fig. 3*B*, interval ending at the *arrow*), the responses to two targets included both horizontal and vertical components that were smaller than those evoked by the motion of either target alone but much larger than those evoked when a single target moved in the orthogonal direction. Thereafter, vertical eye velocity began to decrease toward zero, and horizontal eye velocity increased, because the upward moving target had disappeared, and only the rightward moving target

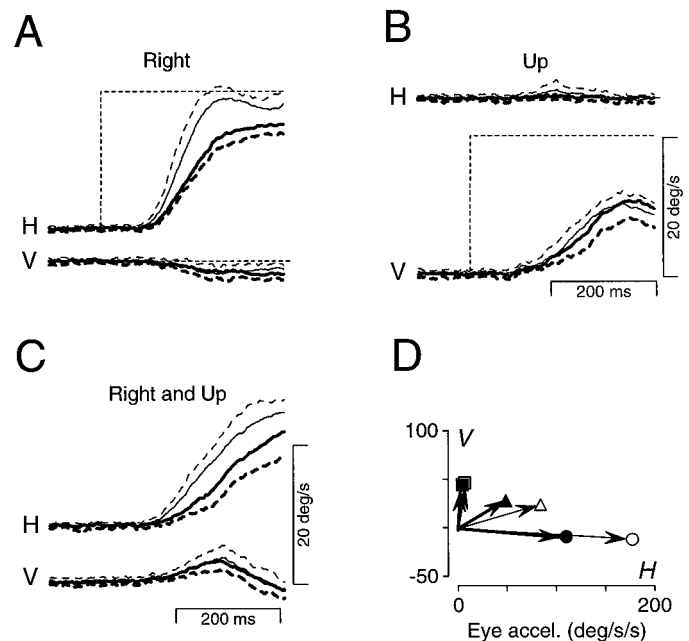


Figure 4. Example showing the effect of learning on responses to single-target and two-target stimuli in one experiment. The two targets moved rightward and upward, and learning caused a decrease in rightward eye acceleration. *A*, Horizontal (*H*) and vertical (*V*) eye velocity responses to rightward moving targets. *B*, Horizontal and vertical eye velocity responses to upward moving targets. *C*, Horizontal and vertical eye velocity responses to right and up two-target stimuli. In *A* and *B*, the *dashed traces* show steps of target velocity. In *A–C*, the *fine* and *bold solid traces* show prelearning and postlearning averages of eye velocity. The *dashed traces* that parallel the eye velocity traces show SDs: for clearer viewing, the SDs are shown above the average trace for prelearning data (*fine traces*) and below the average trace for postlearning data (*bold traces*). *D*, Polar plot showing quantification of average eye acceleration in the interval from 138 to 186 msec after the onset of target motion. *Open symbols* show eye acceleration measured from prelearning data; *filled symbols* show eye acceleration measured from postlearning data. *Circles, squares, and triangles* show the end points of vectors summarizing responses to single rightward moving targets, single upward moving targets, and right and up two-target stimuli, respectively. *Fine arrows* give the magnitude and direction of prelearning eye acceleration. *Bold arrows* give the magnitude and direction of postlearning eye acceleration.

remained visible. If the traces had been extended longer, then horizontal eye velocity would have attained final target velocity.

Lisberger and Ferrera (1997) showed previously that the intermediate eye velocities in the first 100 msec of the responses to two-target stimuli could be described as a weighted vector average of the responses to single-target motion in each of the two directions. Our data on the responses to two-target stimuli were entirely consistent with theirs. For each animal, the two-target stimuli could be used to determine weights associated with each direction of target motion that best described the pursuit response to interactions of that direction of target motion with all others. These weights were different for target motions in different directions but were consistent over repeats of the two-target paradigm in separate experimental sessions.

Effects of learning on responses to single-target and two-target stimuli

Figure 4 contains averages of eye velocity that illustrate the effect of learning-induced decreases in rightward eye acceleration on responses to a two-target stimulus consisting of rightward and upward target motion. Before learning (*fine traces*), a single

rightward moving target (Fig. 4A) evoked smooth eye movements that began ~100 msec after the onset of target motion and that consisted of a brisk increase in rightward eye velocity toward the final target velocity of 20°/sec and a slight downward eye acceleration. After the monkey had completed 500 rightward learning trials and 500 leftward control trials, there was a large decrease in the initial rightward eye acceleration evoked by rightward target motion and no change in the small downward response to the same target motion (Fig. 4A, *bold traces*). In contrast to their large effects on responses to rightward single-target stimuli, learning trials along the horizontal axis caused little or no change in the response to single upward moving targets (Fig. 4B). In this monkey, the initiation of upward pursuit had a low gain and failed to reach target velocity in the part of the response that is illustrated, but the low-gain upward pursuit was the same before (Fig. 4B, *fine traces*) and after (Fig. 4B, *bold traces*) learning.

Learned decreases in eye acceleration for rightward single target motion caused a change in the direction of the initial pursuit response to two-target stimuli that combined a rightward moving target with an upward moving target. Figure 4C shows this finding by plotting average eye velocity as a function of time. Comparison of the responses to two-target stimuli before learning (*fine traces*) and after learning (*bold traces*) reveals a reduction in the horizontal component of eye velocity with no change in the vertical component. We quantified the effect of learning on the responses to two-target stimuli by measuring horizontal and vertical eye acceleration in 48 msec intervals, starting either 128 msec (monkey N) or 138 msec (monkeys E, I, and K) after the onset of target motion. Figure 4D illustrates a polar plot that graphs these measurements for the prelearning data (*open symbols*) and postlearning data (*filled symbols*) from the single experiment documented in Figure 4A–C. For rightward target motion (*circles*), learning caused a decrease in the magnitude of eye acceleration without a change in direction. For upward target motion (*squares*), learning caused only a very small change in magnitude and direction. The prelearning response to the simultaneous presentation of the rightward- and upward moving target (*open triangle*) was intermediate to the responses to the rightward moving target and the upward moving target presented singly. The postlearning response to the two-target stimulus (*filled triangle*) included a smaller horizontal component relative to the prelearning response. Thus, learned decreases in rightward eye acceleration caused the vector for the response to the rightward and upward two-target stimulus (*bold arrow, filled triangle*) to rotate away from the rightward direction and decrease slightly in magnitude. In subsequent figures, the responses to single- and two-target stimuli will be represented as vectors that describe the direction and magnitude of eye acceleration, as in Figure 4D.

Although the data shown in Figure 4D were measured from the averaged eye velocity traces shown in Figures 4A–C, we use this example to describe the trial-by-trial variance of eye acceleration in single- and two-target trials. For the averages in Figure 4A, the SDs are demonstrated by the *dashed traces* that follow above or below the averages, which are shown as *solid traces*. In a separate trial-by-trial analysis, the SDs of prelearning eye accelerations to single upward or rightward moving targets ranged from 14.59 to 30.25°/sec², and those of postlearning eye accelerations ranged from 10.26 to 21.15°/sec². SDs of horizontal and vertical eye acceleration in right-up two-target trials were 26.67 and 22.83°/sec² in prelearning trials, and 31.12 and 25.84°/sec² in postlearning trials, respectively. Tests of statistical significance on the responses to all stimuli shown in Figure 4D revealed that learning

induced significant changes ($p < 0.05$) only in the horizontal eye acceleration components of the responses to rightward single-target stimuli and right-up two-target stimuli (unpaired t tests, $p = 0.0001$ for both). Unfortunately, it would not have been meaningful to perform similar statistical tests on most experiments. We will show in the following sections that learning caused both specific and nonspecific effects on the responses to two-target stimuli in almost all experiments. The nonspecific effects of learning could have created statistical significance in the specific effects, even when they were actually not significant. In the experiment illustrated in Figure 4, the nonspecific effects were not present, and statistical evaluation was feasible. We will resort to other controls to support our contention about the consistency and veracity of the specific effects of learning on responses to two-target stimuli.

Specific effects of learning on responses to two-target stimuli

Learning had both specific and nonspecific effects on the initiation of pursuit evoked by two-target stimuli. Specific effects of learning depended systematically on the learning paradigm, whereas nonspecific effects of learning were the same for a given two-target stimulus, regardless of the learning paradigm. We begin by describing the specific effects for a monkey that showed a relatively small nonspecific effect. Consider first Figure 5A, which shows the results of “left-increase” experiments, in which the responses to two-target stimuli were measured before and after learning-induced increases in leftward eye acceleration. Results are shown for the four two-target pairs that consisted of one horizontal and one vertical target motion, and the data are plotted in the quadrant that would reflect the vector average of each pair of target motions. Thus, responses to two-target stimuli consisting of leftward and upward target motion are plotted in the left, top quadrant of the graph. The vectors show the mean eye acceleration during the initiation of pursuit to two-target stimuli before and after learning, and the points show the same data from four individual experiments.

Left-increase learning (Fig. 5A) caused a change in both the magnitude and direction of the responses to two-target stimuli that paired upward or downward target motion with leftward target motion. Both before learning (*fine vectors* and *open symbols*) and after learning (*bold vectors* and *filled symbols*), the responses were intermediate between the directions of the two-targets. Comparison of the responses before and after learning reveals that learning caused the responses to rotate in the direction of the learned increase in leftward eye acceleration. The change in direction reflected by the two vectors is also evident in the results from the four individual experiments (*symbols*) that were averaged to obtain the vectors. In contrast, left-increase learning had very little effect on the responses to two-target stimuli that paired rightward target motion with upward or downward target motion. There may have been a small decrease in the magnitude of the vectors after learning, but there was no change in the direction.

Right-increase learning caused complementary effects (Fig. 5B). There were changes in the direction and magnitude of the responses to two-target stimuli that paired upward or downward target motion with rightward target motion. After learning, the responses were rotated in the direction of the learned increase in rightward eye acceleration. The effects again are visible in the results of the four individual experiments (*symbols*) as well as in the averages across experiments (*vectors*). There were only small

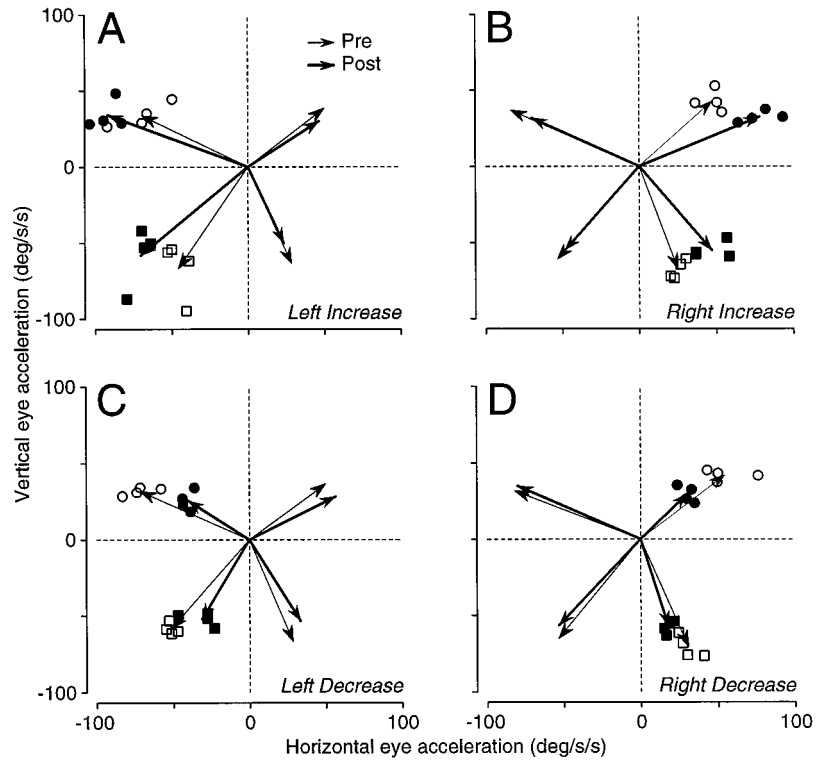


Figure 5. Specific effects of all four learning conditions on two-target responses for orthogonal target motions. Each plot summarizes responses averaged over four experiments for one learning protocol. *A*, Left-increase learning. *B*, Right-increase learning. *C*, Left-decrease learning. *D*, Right-decrease learning. Within each graph, the *arrows* in the four oblique directions show the effect of the given learning condition for two-target stimuli with vector averages in that direction. *Fine* and *bold arrows* show the vectors for pre-learning and postlearning eye acceleration. *Open symbols* show eye acceleration from the four individual experiments that make up each prelearning average. *Filled symbols* show eye acceleration from the four individual experiments that make up each postlearning average. Data are from monkey K.

changes in the magnitude and no change in the direction of the responses to two-target stimuli that paired upward or downward target motion with leftward target motion.

The results were similar but a little more complex for experiments that measured the effect of learning-induced decreases in eye acceleration. In Figure 5*C*, left-decrease learning caused changes in the direction of the responses to all pairs of two-target stimuli. When one of the two targets provided leftward target motion, learning caused the responses to be smaller and to rotate away from the left. When one of the two targets provided rightward target motion, an equivalent change occurred. Postlearning responses changed direction, in this case rotating toward the right (away from the left), with only a small change in magnitude. In Figure 5*D*, right-decrease learning caused small changes in the direction of the responses to all pairs of two-target stimuli. When one of the two targets provided rightward motion, the postlearning response rotated consistently away from the right and decreased in magnitude. When one of the two targets provided leftward motion, the changes in direction were smaller and inconsistent.

The 16 daily experiments summarized in Figure 5 show that pursuit learning had consistent effects on the directions of responses to two-target stimuli in the monkey we have chosen to illustrate our general findings. Learned increases in eye acceleration for a given direction of horizontal target motion caused responses to two targets to be rotated toward that direction. Learned decreases in eye acceleration for a given direction of horizontal target motion caused responses to two targets to be rotated away from that direction. Changes in the direction of eye acceleration were always seen when one of the targets in a two-target stimulus moved in the learning direction. Changes were sometimes seen when one of the targets moved in the control direction.

We documented the specific effects of learning on the responses

to each two-target stimulus by comparing the postlearning responses after learned increases versus decreases in eye acceleration for a given learning direction. Figure 6 summarizes results for a total of 64 experiments (16 daily experiments on each of four monkeys). Consider first the four graphs in Figure 6*A₁–A₄*, which compare the effects of left-increase and left-decrease learning on the responses to two-target stimuli that paired leftward target motion with upward or downward target motion. Of the eight quadrants available for comparison (two quadrants by four monkeys), six showed a consistent effect of learning on the direction of the eye acceleration evoked by two-target stimuli. The two exceptions are the upper-left quadrants for monkeys N and I (Fig. 6*A₁, A₄*). In general, responses after left-increase learning (*solid vectors*) were rotated toward the left, whereas responses after left-decrease learning (*dashed vectors*) were rotated away from the left. Figure 6*B₁–B₄* shows that learning-induced changes in rightward eye acceleration had complementary effects in all eight of the available quadrants (two quadrants by four monkeys). Responses after right-increase learning (*solid vectors*) were always rotated to the right relative to responses after right-decrease learning (*dashed vectors*). On top of the very consistent general trend in the data in Figure 6, there is considerable variability between different monkeys. This variability represents genuine differences in the details of the responses and can be attributed to intersubject differences in (1) the baseline weighting of different directions in two-target stimuli, (2) pursuit gain for different directions of eye motion, and (3) size of the nonspecific effects of learning.

As a control to assess the specificity of the effects illustrated in Figure 6, Figure 7 illustrates mean data from two-target trials that paired target motion in the nonlearning (control) direction with upward or downward target motion. If the small but consistent effects in Figure 6 are real and specific, then Figure 7 should reveal no effect of learning on the responses to two-target stimuli

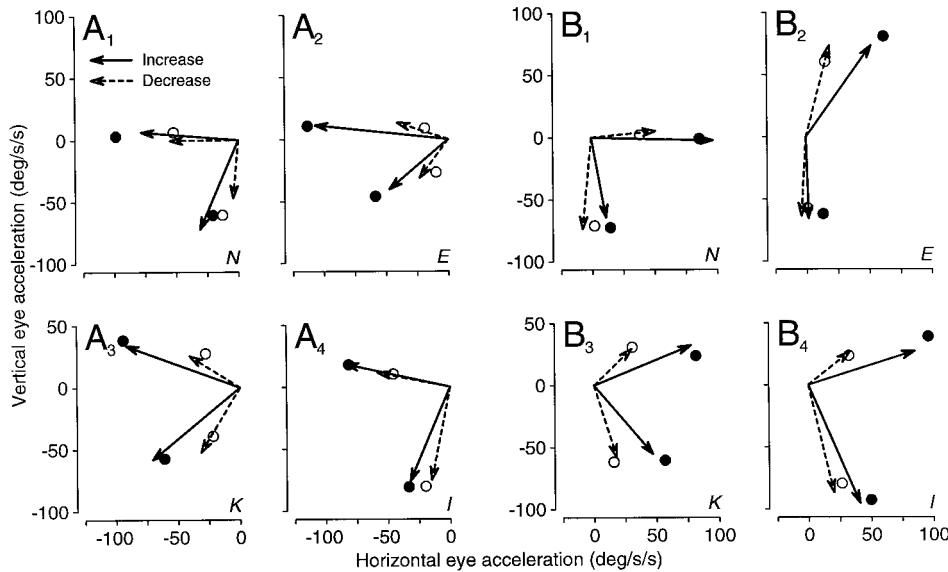


Figure 6. Summary of specific effects of learning on responses to two-target stimuli that combined motion in the learning direction with vertical target motion. In *A* and *B*, the four separate plots summarize data from four animals: monkeys N, E, K, and I. *A*₁–*A*₄, Responses to two-target stimuli that interacted leftward target motion with vertical target motions after experiments that caused learning for leftward target motion. *B*₁–*B*₄, Responses to two-target stimuli that interacted rightward target motion with vertical target motions after experiments that caused learning for rightward target motion. *Solid arrows* describe the vectors of average postlearning eye acceleration after experiments that increased eye acceleration. *Dashed arrows* plot the vectors of average postlearning eye acceleration after experiments that decreased eye acceleration. *Circles* display the predictions of model 1. *Open* and *filled circles* show predictions for experiments that decreased or increased eye acceleration, respectively.

with horizontal target motion in the control direction. In each of the 16 quadrants illustrated in Figure 7, the differences between the prelearning and postlearning vectors are small. The success of the control analysis in Figure 7 persuades us that the effects illustrated in Figures 5 and 6 are specific to the learning direction and are functionally significant.

The effect of learning on responses to two-target stimuli consisting of leftward and rightward target motion was entirely consistent with the picture described above for two-target stimuli comprising orthogonal target motions. Consider first experiments that induced learning for leftward target motion (Fig. 8, left column). After left-increase learning, two-target stimuli consisting of leftward and rightward target motion consistently caused more leftward eye acceleration (*solid vectors with filled arrowheads*) than after left-decrease learning (*dashed vectors with open arrowheads*). We obtained complementary effects for rightward learning directions (Fig. 8, right column). The responses to two-target stimuli consisting of rightward and leftward target motion

were always more rightward or less leftward after right-increase learning (*solid arrows with filled arrowheads*) than after right-decrease learning (*dashed arrows with open arrowheads*). In fact, in several cases, after both leftward and rightward learning, learned decreases and increases in eye acceleration resulted in opposite directions of movement (e.g., monkey K).

Nonspecific effects of learning on responses to two-target stimuli

In Figure 5, we analyzed the prelearning and postlearning responses to a given two-target stimulus in one monkey that had relatively little nonspecific effect. This analysis provided perfectly controlled comparisons based on data obtained within single experiments. In Figures 6–8, we finessed nonspecific effects and presented data from each of four monkeys by comparing the responses to a given two-target stimulus after increase and decrease learning experiments conducted on different days. Now we document nonspecific effects by comparing the responses to a

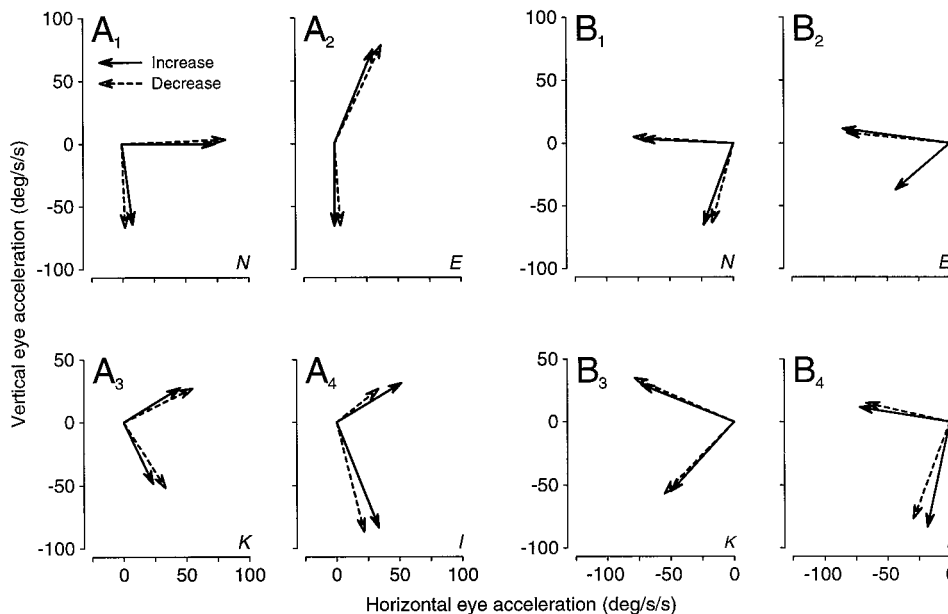


Figure 7. Summary showing the absence of specific effects of learning on responses to two-target stimuli that combined motion in the control direction with vertical target motion. In *A* and *B*, the four separate plots summarize data from four animals: monkeys N, E, K, and I. *A*₁–*A*₄, Responses to two-target stimuli that interacted rightward target motion with vertical target motions after experiments that caused learning for leftward target motion. *B*₁–*B*₄, Responses to two-target stimuli that interacted leftward target motion with vertical target motions after experiments that caused learning for rightward target motion. *Solid arrows* plot the vectors of average postlearning eye acceleration after experiments that increased eye acceleration. *Dashed arrows* show the vectors of average postlearning eye acceleration after experiments that decreased eye acceleration.

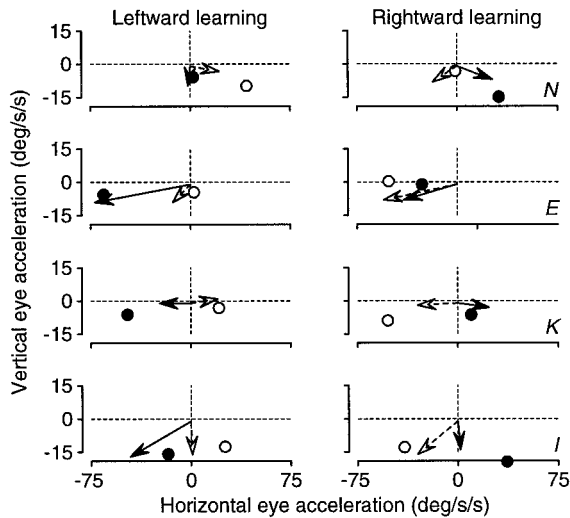


Figure 8. Effects of learning on responses to two-target stimuli that combined leftward and rightward target motions. *Left* and *right* columns show results from experiments with learning for leftward and rightward target motion. From *top* to *bottom*, the four graphs in each column show data from monkeys N, E, K, and I, respectively. *Solid arrowheads* show the vectors describing average postlearning eye acceleration for two-target stimuli after experiments that caused increases in eye acceleration for single targets. *Open arrowheads* show the vectors describing average postlearning eye acceleration for two-target stimuli after experiments that caused decreases in eye acceleration for single targets. *Circles* display the predictions of model 1. *Open* and *filled circles* show predictions that decreased or increased eye acceleration, respectively. *Horizontal* and *vertical dashed lines* show zero vertical and horizontal eye acceleration, respectively.

given two-target stimulus for all four combinations of learning direction and learned increases versus decreases in eye acceleration. Although less well controlled than the earlier comparisons in the sense that we are now comparing responses obtained in four different groups of four daily experiments, this approach revealed consistent nonspecific effects of learning that we needed to analyze to be able to interpret the specific effects.

The four vector plots in Figure 9*A–D* show average eye acceleration for the most compelling example we found of a nonspecific effect of learning. Each plot shows the responses before and after learning for two-target stimuli that delivered rightward and upward target motions. The four plots summarize groups of experiments that used different learning conditions. For example, Figure 9*A* shows that right-decrease learning caused the response to this two-target stimulus to show a large change in direction. After learning (*bold solid arrow*), eye acceleration was nearly upward, whereas before learning (*fine solid arrow*), eye acceleration was more rightward than upward. Figure 9*B* shows the seemingly paradoxical finding that right-increase learning also caused the response to this two-target stimulus to be rotated toward upward eye acceleration. This apparent paradox is consistent with Figure 6*B*₂, however, because the upward rotation of the responses was much greater after right-decrease learning (Fig. 9*A*) than after right-increase learning (Fig. 9*B*). The explanation for the apparent paradox appears in Figure 9, *C* and *D*. Left-decrease and left-increase learning both caused a large upward rotation of the responses to two-target stimuli consisting of upward and rightward motions. Thus, every learning condition caused the response to rightward and upward targets to rotate upward. We conclude that the inescapable upward rotation in Figure 9*A–D* represents a nonspecific effect of learning.

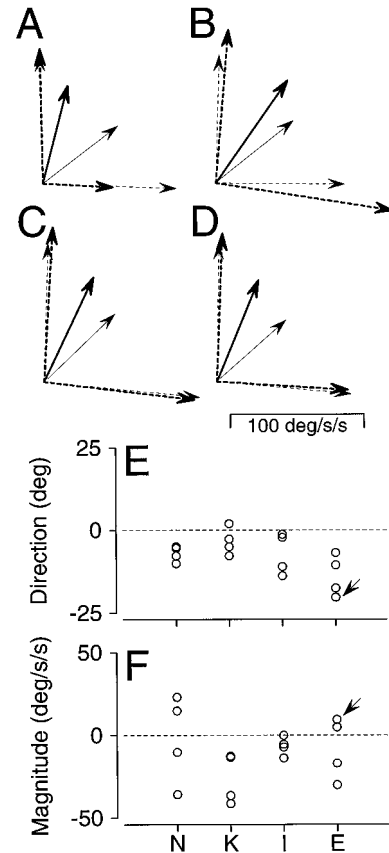


Figure 9. Summary of nonspecific effects of learning on responses to two-target stimuli. *A–D*, Data from four different experiments on one monkey to document the largest nonspecific shift we observed in the direction of eye acceleration. Each vector plot shows the effect of one learning condition on the responses to upward or rightward motion of single targets and two-target stimuli consisting of rightward and upward target motion. Learning conditions are right-decrease (*A*), right-increase (*B*), left-decrease (*C*), and left-increase (*D*). *Fine* and *bold arrows* show responses before and after learning, respectively. *Dashed arrows* show responses to single target stimuli. *Solid arrows* show responses to two-target stimuli. *E*, Summary of nonspecific effects of learning on the direction of responses to two-target stimuli. Positive changes in direction describe shifts toward the horizontal axis. Negative changes in direction describe shifts toward the vertical axis. *F*, Summary of nonspecific effects of learning on the magnitude of responses to two-target stimuli. Each graph plots responses for the different monkeys at different locations along the *x*-axis. For each monkey, the four points quantify nonspecific effects for each of the four combinations of orthogonal two-target motions. The *oblique arrows* in *E* and *F* indicate the measurements taken from the examples in *A–D*.

Figure 9, *E* and *F*, estimates the nonspecific changes separately for each two-target stimulus that paired horizontal and vertical target motion and each of our four monkeys. To isolate the nonspecific changes, we analyzed the effects of learning on the responses to two-target stimuli in which horizontal target motion was in the control direction for the learning condition. For each two-target pair, we averaged the magnitude and direction of the nonspecific changes from learned increases and decreases in eye acceleration across all repetitions of the relevant learning experiments. For the experiments summarized by Figure 9*A–D*, for example, the nonspecific effect of learning on the response to right and up targets was calculated as the mean of the magnitude and direction changes measured from left-decrease and left-increase experiments (Fig. 9*C,D*). The analysis of nonspecific

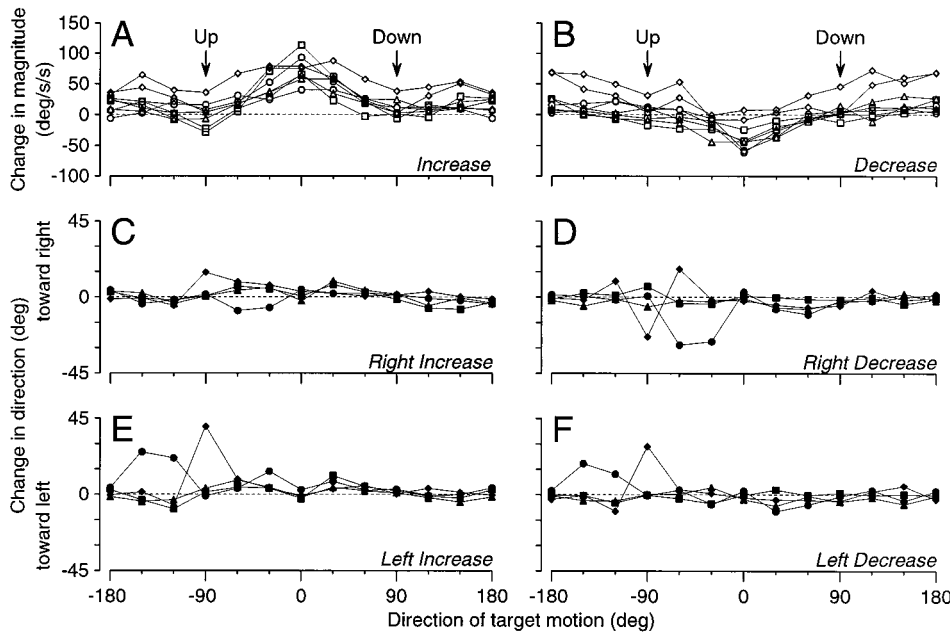


Figure 10. Generalization of learning to pursuit evoked by 12 directions of single target motion. *A, B*, Changes in the magnitude of eye acceleration are plotted as a function to the direction of target motion, separately for experiments that increased (*A*) or decreased (*B*) eye acceleration. *C–E*, Changes in the direction of eye acceleration are plotted as a function of the direction of target motion, separately for right-increase (*C*), right-decrease (*D*), left-increase (*E*), and left-decrease (*F*) experiments. Changes in direction that rotated the vector toward rightward and leftward are plotted as positive and negative values on the y-axis. In all six graphs, 0° on the x-axis represents the learning direction, -90° represents upward target motion, and 90° represents downward target motion. Therefore, for left-increase and left-decrease experiments, responses to single targets moving at 30° and 60° left and down plot at $+30^\circ$ and $+60^\circ$ on the x-axis. For right-increase and -decrease experiments, responses to single targets moving

at 30° and 60° right and up plot at -30° and -60° , and responses to single target moving at 30° and 60° right and down plot at $+30^\circ$ and $+60^\circ$ on the x-axis. Different symbols show data from different animals: circles, monkey E; squares, monkey K; triangles, monkey I; diamonds, monkey N.

effects for the four combinations of orthogonal two-target stimuli yielded the four observations plotted in Figure 9, *E* and *F*, for each monkey. Inspection of Figure 9*E* reveals that monkeys N and K showed only small nonspecific effects on the direction of the response to two targets, monkey I showed slightly larger effects, and monkey E had the largest nonspecific effects on response direction. Figure 9*F* reveals quite a few examples of nonspecific decreases in magnitude of responses to two-target stimuli. For reference, the arrows in Figure 9, *E* and *F*, indicate the results of analyzing the vector plots in Figure 9*A–D*.

Generalization of learning to different directions of single-target motion

Because learning caused consistent changes in the initial pursuit to two-target stimuli, it seemed important to ascertain whether similar changes in direction or magnitude of initial eye acceleration were observed in the responses to single targets moving in the directions of the vector-averaged responses to two targets. We tested the effect of learning on the initiation of pursuit for single targets moving in 12 directions at 30° intervals. For each learning condition and each monkey, we averaged the changes in the magnitude and direction of eye acceleration across three repeats of each of the four learning conditions. Changes in the magnitude of eye acceleration were tuned around the learning direction, which is plotted at 0° on the x-axis (Fig. 10*A, B*). These graphs show eight curves each for learned increases and decreases in eye acceleration: one each for leftward and rightward learning directions in each of four monkeys. The generalization bandwidth at half-height for pursuit learning was $\sim 60^\circ$, and pursuit learning rarely generalized from the learning direction to orthogonal directions, except for monkey N (Fig. 10, diamonds). There was a tendency for eye accelerations in the opposite, control direction (plotted at $\pm 180^\circ$) to increase slightly regardless of the learning protocol. Again, monkey N (Fig. 10*B*, diamonds) provided the only exception: left-decrease and right-decrease learning experiments caused little or no decrease in eye acceleration in the learning direction but still caused large increases in eye accel-

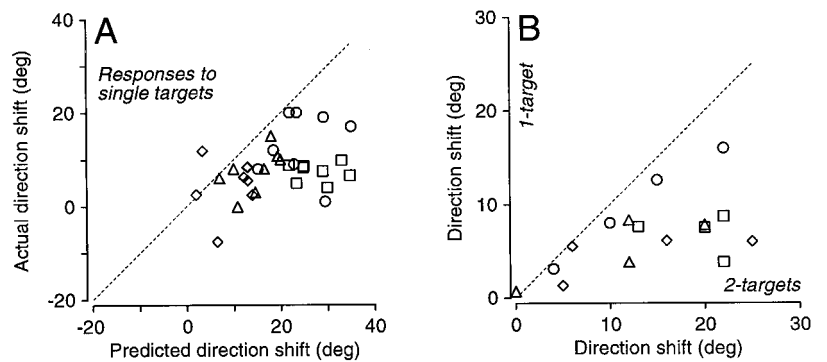
ation in the opposite, control direction. Separate experiments in which animals tracked only test trials in both directions suggested that some of the small increases in the opposite direction were general effects of pursuing targets for 1000 trials (data not shown).

Figure 10*C–F* shows that learning for target motion along the horizontal axis usually had only small effects on the direction of pursuit for targets moving in other directions. Inspection of each of these graphs reveals that there was essentially no change in the direction of eye movement for single targets moving in the learning direction, plotted at direction of target movement of 0° . However, there was a tendency for organized changes in the direction of the eye movement evoked by single targets in directions within 60° of the learning direction. When the learning direction was rightward (Fig. 10*C, D*), learned increases in eye acceleration caused eye movement to rotate toward the right in most cases (Fig. 10*C*), and learned decreases in eye acceleration caused the eye movement to rotate away from the right (Fig. 10*D*). The opposite tendencies were present when the learning direction was leftward. Learned increases in eye acceleration caused the responses to single targets to rotate toward the left (Fig. 10*E*), and learned decreases caused the responses to rotate away from the left (Fig. 10*F*). Except for a few points, these changes are quite subtle. Furthermore, not all of the small changes were present in all four monkeys.

Lack of generalization of learning to the same eye movement evoked by different stimuli

We showed above that the responses to single targets tended to deviate in the same direction as the specific effects of learning on the responses to two-target stimuli. In the present section, we evaluate the possibility that learning-induced changes in the direction of responses to two-target stimuli might generalize to all smooth eye movements in a given direction. For example, one might see the same change in the direction of a response to a single target moving rightward and upward as to a two-target

Figure 11. Quantitative rejection of hypotheses predicting learning-induced changes in the direction of the initial eye acceleration that depend only on the direction of the prelearning smooth eye movement and not on whether the stimulus consisted of one or two targets. **A**, Comparison of the actual changes in direction of responses to oblique motion of single targets with the shift predicted if learning altered only the horizontal component of pursuit. Eight points are plotted for each monkey: one point for each of four single-target stimuli and each of two learning directions. The four single-target stimuli moved obliquely 30 and 60° up and down relative to the learning direction. **B**, Comparison of the shifts in direction of responses to two-target stimuli (*x*-axis) with those for single-target stimuli (*y*-axis) that evoked responses in the same direction. Each data point shows the average for the upper and lower quadrants in the learning direction, so four data points are plotted for each monkey: one each for right-up, right-down, left-up, and left-down two-target motions. Different symbols show data from different animals: circles, monkey E; squares, monkey K; triangles, monkey I; diamonds, monkey N.



stimulus consisting of rightward and upward motions. If true, this possibility would make it difficult to interpret our data.

Figure 11A evaluates the most extreme example of this class of explanation for our data, which is the motor hypothesis we defined earlier. According to the motor hypothesis, the site of learning would be after the pursuit commands have been divided into the horizontal and vertical components of eye movements. If the motor hypothesis were true, then learning should generalize to the initial pursuit evoked by any target motion with a horizontal component, as it does to target motion at different speeds in the learning direction (Kahlon and Lisberger, 1996). According to the motor hypothesis, it should be possible to predict the postlearning responses to the oblique motion of single targets by simply adjusting the prelearning horizontal component by the same gain factor obtained for single-target motion in the learning direction. To make this prediction, we calculated the horizontal components of the prelearning response for all eight oblique single-target motions in each of the four monkeys, scaled each horizontal component by the gain change induced in the appropriate learning direction, and predicted the direction of the postlearning responses. Figure 11A plots the actual learning-induced shift in direction of the single-target responses versus the shift in direction predicted by the motor hypothesis. All but two of the points plot below the line with slope of 1, showing that the actual changes in direction were considerably smaller than the predictions. Linear regression with errors in both coordinates (Press et al., 1992) under the assumption of equal variances along the *x*- and *y*-axes revealed a regression slope of 0.37. Of course, there need not be a single site of learning; sites could be distributed across different levels of the pursuit system, with one site in the motor pathways. This slope places an upper limit of 37% on the fraction of learning that could occur in the motor pathways, after separate commands have been formed for the horizontal and vertical components of eye movement.

We next tested a more general formulation of the motor hypothesis in which changes in the direction of the initial eye acceleration depend only on the direction of the prelearning smooth eye movement and not on whether the stimulus consisted of one or two targets. Figure 11B analyzes whether learning-induced changes in the direction of the responses to two-target stimuli were the same size as the changes in the direction of same-direction responses to single-target stimuli. For two-target responses, we calculated the difference between the directions of responses to a given pair of targets after increase and decrease learning in a given direction. For single-target responses, we computed the same difference but did so only after interpolating

along the curves in Figure 10C–F to estimate the effect of learning on the responses to single targets in the direction of the prelearning response to each combination of two targets. In Figure 11B, a plot of the direction shift for single-target responses as a function of that for two-target responses reveals that almost all of the data plot below the line of slope 1. Thus, as a general rule the changes in direction of responses to two-target stimuli were larger than those to single-target stimuli in the same direction. Monkey E (circles) comes closest to being an exception to the general rule, because his data plot only slightly below the line of slope 1. Linear regression with errors in both coordinates (Press et al., 1992) under the assumption of equal variances along the *x*- and *y*-axes yielded a regression slope of 0.33. Again, if learning is distributed across multiple neural sites, then this slope places an upper limit of 33% on the amount of the learning-induced change in the direction of smooth pursuit that depends on the direction of the evoked eye movement and not on the exact stimulus.

Monkey N provided a final, serendipitous opportunity to test directly whether learning generalized equally to eye movements evoked by single- and two-target stimuli when the resulting eye movements had approximately the same magnitude and direction. This monkey emitted almost zero smooth eye acceleration for upward motion of single targets and showed pure horizontal eye acceleration for two-target stimuli consisting of horizontal and upward target motion at 20°/sec. The magnitude of the response to two targets was slightly smaller than that for single horizontal target motion at 20°/sec. Before taking advantage of this opportunity, we attempted to match the sizes of the prelearning responses to horizontal single targets and horizontal and upward two-target stimuli by using single-target motion at 15°/sec. We then evaluated the generalization of learning to eye movements evoked by single- and two-target stimuli when both evoked the same eye movements before learning. Different effects on the same eye movements would argue strongly that the site(s) of learning are upstream from the conversion of pursuit commands to motor coordinates.

The poor generalization of learning between the two matched eye movements is shown in Figure 12 for two-target stimuli that consisted of rightward and upward or leftward and upward target motion. The *left column* summarizes responses to single rightward or leftward moving targets. For each graph in the *left column*, the companion graph in the *right column* shows the responses to two-target stimuli that paired upward with rightward or leftward targets. Each *point* represents the end of a vector that starts at the intersection of the *horizontal* and *vertical dashed lines* in each graph. Consider first the right-increase experiments sum-

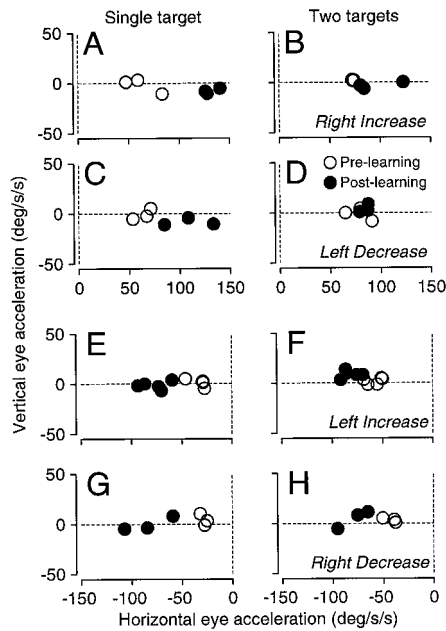


Figure 12. Weak generalization of learning from single-target stimuli to two-target stimuli that evoked eye movements in the same direction in monkey N. Each graph compares the responses before and after a given learning condition for a given single-target or two-target stimulus. The left column of graphs contains data from single-target trials, and the right column contains data from two-target trials. Each row of graphs shows the responses from a given set of experiments. From top to bottom, learning conditions were right-increase, left-decrease, left-increase, and right-decrease. Open and filled symbols show data from prelearning and postlearning trials, respectively. Dashed vertical and horizontal lines show zero eye acceleration on the horizontal and vertical axes.

marized in Figure 12, *A* and *B*. Before learning, rightward target motion at 15°/sec evoked eye accelerations that averaged 65°/sec² (Fig. 12*A*, open circles), and rightward and upward two-target stimuli evoked eye accelerations that averaged 75°/sec² (Fig. 12*B*, open circles). Learning caused increases in the average eye acceleration evoked by the single target to 135°/sec² (Fig. 12*A*, filled circles). In contrast, the eye acceleration evoked by the rightward and upward two-target stimulus was unchanged in two of the three experiments and increased to an average of only 95°/sec² (Fig. 12*B*, filled circles).

The rest of Figure 12 shows that we obtained the same result for leftward target motion after left-increase learning (Fig. 12*E,F*), for rightward target motion after left-decrease learning (Fig. 12*C,D*), and for leftward target motion after right-decrease learning (Fig. 12*G,H*). Even though we were not able to obtain learned decreases in eye acceleration in monkey N, Figure 12, *C*, *D*, *G*, and *H*, takes advantage of the associated increase in eye acceleration in the control direction in this monkey. From top to bottom in Figure 12, the mean changes in the responses to two-target stimuli were 29, 11, 31, and 38% (mean, 27%) of those to single-target stimuli. Thus, the eight examples in Figure 12 show that learning generalizes only partly according to the direction of the eye movement evoked by a stimulus. Instead, learning appears to generalize according to the evoking stimulus configuration (two-target or single target) and the directions of target motion relative to the learning stimulus. Again, if learning is distributed across multiple sites, then this analysis places an upper limit of 27% on the amount of learning that can occur in pathways that are organized according to the direction of the ultimate eye movement.




Figure 13. Comparison of the errors between the data and the predictions of models 1–3. Histogram bars labeled, 1–3 show the average prediction error across monkeys for models 1–3. Groups of bars labeled *Decrease* and *Increase* show average prediction error across monkeys for experiments in which learning either decreased or increased eye acceleration. Symbols show average prediction error for each monkey, averaged across the three two-target conditions that interacted the learning direction of target motion with two orthogonal directions or the opposite direction and over experiments that caused learning in pursuit for leftward and rightward target motion. Open symbols show errors associated with the winning model 1. Filled symbols show errors associated with other models. Different symbols display data for different monkeys: circles, monkey E; triangles, monkey I; squares, monkey K; diamonds, monkey N.

Modeling the effects of learning on vector averaging

The effects of learning on the responses to two-target stimuli provide qualitative support for the hypothesis outlined in Figure 2*B*: learning is upstream from vector averaging. To provide a quantitative analysis, we now test our data against linear models that formalize the two hypotheses described in Figure 2 and those of a third hypothesis that is a variant on the first.

Model 1 places learning upstream from vector averaging. Model 2 places learning downstream from vector averaging. Model 3 implements learning as changes in the weights used for vector averaging.

Using equations and methods described in Appendix, we compared the ability of each model to predict the effects of learning on the responses to two-target stimuli. We wish to emphasize that this was not a fitting procedure. Rather, we predicted the responses to two-target stimuli after learning using a deterministic procedure based on the weights afforded each target motion for vector averaging before learning, the eye accelerations induced by single targets before and after learning, and the direction generalization data. Each prediction was compared with the actual postlearning response by measuring the distance from the prediction to the data point in an *x*–*y* coordinate framework. Prediction errors were computed as the mean error across the three two-target stimuli that combined motion in the learning direction with upward, downward, or control direction target motion. Errors were averaged across learning directions, but those associated with learned increases and decreases in eye acceleration were computed separately. Figure 13 shows that the prediction error was almost always smallest for model 1, which placed learning upstream of averaging. The only exception was for learned decreases in eye acceleration in monkey N (filled diamond), whose postlearning responses after increases in eye acceleration were predicted best by model 3. The predictions of model 1 are plotted as circles in Figures 6 and 8 to allow direct comparison with the data. The predictions fit the data well for two-target stimuli that paired horizontal target and vertical target

motion (Fig. 6). They fit the data less well (albeit better than other models) when the two-target stimuli consisted of oppositely directed target motions (Fig. 8). There are two possible reasons for this. First, we were unable to estimate any nonspecific offset that may be associated with interactions of leftward and rightward moving targets and therefore could not compensate for these nonspecific effects in our model predictions. Second, two-target stimuli using opposite directions of motion generally yielded more variable data than did two-target stimuli using orthogonal target motions.

It proved difficult to distinguish the three models statistically, because the differences between the predictions of the different models were themselves small. However a number of our findings disagreed qualitatively with the predictions of models 2 and 3. (1) Model 2 fails because it predicts that learning will cause very little or no change in the direction of vector-averaged responses to two-target stimuli that pair orthogonal target motion. (2) Model 2 also fails to reproduce some of the data for two-target stimuli that pair opposite direction target motion. Because the gain factor (g_{ab}) is outside the averaging expression, model 2 cannot reproduce changes in the left–right direction of the responses to these stimuli and therefore cannot account for the responses of monkeys N and K in Figure 6. (3) Model 3 fails because it predicts that learning will cause changes in the magnitude of responses to two-target stimuli that contradict the data in some instances. It predicts no change in the magnitude of the responses to two targets if the magnitude of the response to the vertical component of a two-target stimulus is equal to the horizontal component. It also predicts changes in the magnitude of the responses to two targets that are opposite in direction to the learning if the magnitude of the response to the vertical target motion singly is greater than that for the horizontal target motion. In contrast to these predictions, we always observed changes in the magnitude of the responses to two-target stimuli that were in the same direction as learning: for each two-target pair of horizontal and vertical target motion, we recorded increases (decreases) in magnitude when learning caused increases (decreases) in eye acceleration, whether the response to the horizontal target motion singly was larger or smaller than that to the vertical target motion singly. Thus, model 1 had both quantitative and qualitative response properties that provided the best prediction of the full range of results we obtained when we tested two-target responses after learning in single target responses.

DISCUSSION

Relative location of vector averaging and learning

We have analyzed the effects of motor learning in horizontal pursuit eye movements on the initiation of pursuit for stimuli consisting of two identical targets that moved either in orthogonal directions or in opposite directions along the horizontal axis. Our results revealed “specific” effects that were modulated in a consistent way by the learning condition and “nonspecific” effects that were the same across all learning conditions. We were able to devise data analysis procedures to segregate nonspecific effects from specific effects of learning, but we did not attempt to determine the site or mechanism of nonspecific effects.

Specific effects of learning on the responses to two-target stimuli were consistent across monkeys, and we have evaluated them in relation to three hypotheses for the sites of learning and vector averaging: (1) learning occurs before averaging; (2) learning occurs after averaging but before the creation of separate commands for horizontal and vertical smooth pursuit; and (3) learn-

ing is mediated by changes in the weights used for averaging. Both quantitative and qualitative observations revealed that the specific effects were predicted most closely by a model that implemented hypothesis 1: learning is upstream of vector averaging. We did not attempt to localize the nonspecific effects. Perhaps they are an example in monkeys of the compelling effects of the history of target motion on pursuit in humans (Kowler, 1990).

The conclusion that learning is upstream of vector averaging is based heavily on the finding that learning in the responses to single-target stimuli causes changes in the direction of the vector averaged responses to two-target stimuli. In principle, this finding would be compatible with learning downstream from vector averaging if learning occurred entirely in the motor system, after the creation of separate commands for the horizontal and vertical extraocular muscles. In practice, however, our data are incompatible with the prediction of the motor hypothesis that learning should have the same effect on the direction and amplitude of pursuit in a given direction, whether the stimulus consisted of one or two targets. In almost all monkeys, learning caused small changes in the direction of responses to single targets, but these changes could account for at most 33–37% of the change in the direction of responses to two-target stimuli (from the analyses of Fig. 11*A,B*). In monkey N, we observed different effects of learning on the horizontal pursuit evoked by single-target and two-target stimuli. When the prelearning eye movements were similar for single-target and two-target stimuli, the changes in the responses to two-target stimuli were only ~27% as large as those for single targets (from the analysis of Fig. 12).

Our data and modeling do not exclude the possibility that there are multiple sites of learning. For example, ~35% of the effect of learning on responses to two-target stimuli could be attributed to changes in the direction of responses to single-target stimuli. This implies that it would be possible to explain the effects of learning on responses to two-target stimuli by placing up to 35% of learning in the final motor pathways, after the creation of separate commands for the horizontal and vertical extraocular muscles. Even if some of the learning occurs in the motor final pathways, our data imply that the remaining 65% occurs before the pursuit signals are converted into commands for the horizontal and vertical extraocular muscles and before vector averaging. Furthermore, our data do not support the idea that learning and vector averaging are widely distributed. Model 3 would be one way to formalize the hypothesis that learning and vector averaging are codistributed over many sites, and it is consistently less able than model 1 to reproduce our data.

Functional organization of the pursuit system

Observations of the pursuit response to a single moving target suggest models of pursuit in which simple, serial computations convert visual inputs into commands for smooth eye velocity (Krauzlis and Lisberger, 1994). Recent observations from a number of laboratories, however, imply that pursuit results from a much richer and more complex series of neural computations that create properties such as learning (Kahlon and Lisberger, 1996), on-line control of pursuit gain (Schwartz and Lisberger, 1994), and vector averaging (Lisberger and Ferrera, 1997). The results of the present paper provide some insight into the organization of these neural computations and how they map onto the anatomy and physiology of the pursuit system. The general flow of signals for pursuit is diagrammed in Figure 14, *center panel*, *green arrows*. Anatomical studies (Tusa and Ungerleider, 1988) have shown that visual signals flow from the primary visual cortex (V1)

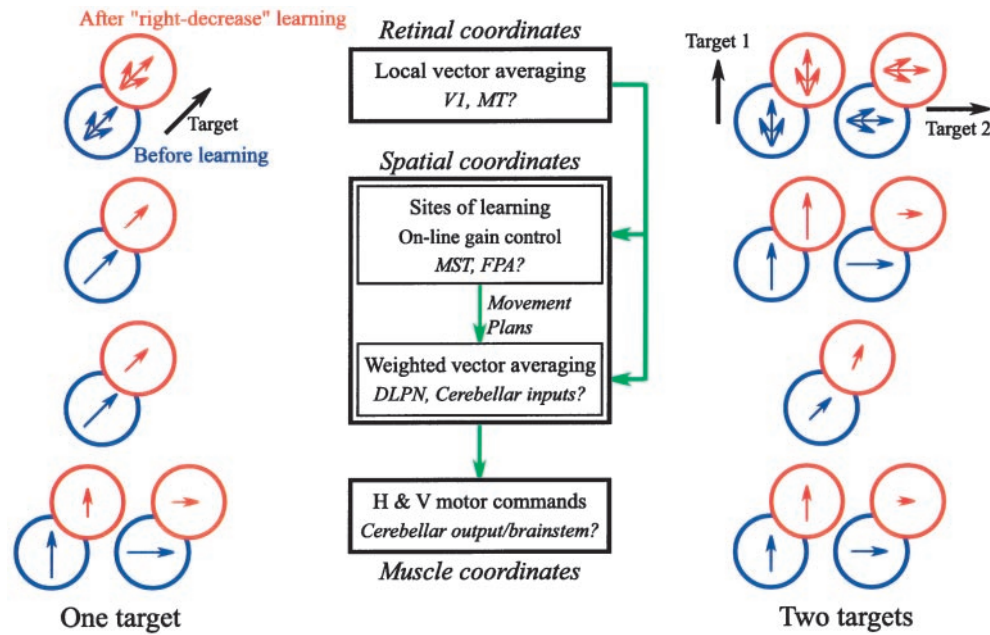


Figure 14. Schematic diagrams showing a hypothesis for signal processing operations and coordinate transformations underlying the generation of pursuit eye movements. The *center panel* uses *green arrows* to summarize the postulated flow of signals through the pursuit system. *VI*, Primary visual cortex; *MT*, middle temporal visual area; *MST*, medial superior temporal area; *FPA*, frontal pursuit area; *DLPN*, dorsolateral pontine nucleus. The *center panel* uses *boxes* to indicate signal processing operations and coordinate transformations. The *three main boxes* suggest a progression from retinal to spatial to muscle coordinates. The *four boxes* containing *names* of signal processing operations and anatomical areas indicate the postulated order of signal processing in the pursuit system. The *left and right panels* show how single-target and two-target stimuli would be processed according to the hypothesis in the *center panel*. The *left panel* shows processing for a single target moving obliquely at 45° right and up. The *right panel*

shows processing for two targets moving purely rightward and upward. At each stage, the *blue and red icons* represent prelearning and postlearning representations, respectively, for right-decrease learning experiments. The *top pairs of icons* are meant to show population responses in *V1* and *MT*, whereas the *icons in the second through fourth groups* indicate the direction and magnitude of representations at downstream levels.

through the extrastriate middle temporal visual area (MT) to the medial superior temporal area (MST) and the part of the frontal cortex that we call the “frontal pursuit area” (FPA). MT, MST, and FPA then project in parallel through the dorsolateral pontine nucleus (DLPN) to the cerebellum and on to the final ocular motor pathways in the brainstem.

Figure 14 outlines a hypothesis for the localization of functions such as learning, on-line gain control, and vector averaging. It suggests that learning occurs downstream from area MT and upstream from Purkinje cells in the cerebellum, perhaps in MST, FPA, or DLPN. Our previous paper (Kahlon and Lisberger, 1996) showed that learning alters the response to a brief perturbation of target velocity only if the perturbation and the ongoing target motion are both in the learning direction. The gating of learning according to the direction of eye and target motion suggested that learning and on-line gain control (Schwartz and Lisberger, 1994) might occur at the same site. The gating also implies that the site of pursuit learning must relay signals related to eye or target movement. This probably excludes MT, where the firing of most neurons encodes only image motion and does not reflect extraretinal signals related to eye or target motion (Newsome et al., 1988; Ferrera and Lisberger, 1997). In contrast, neurons in MST (Newsome et al., 1988) and FPA (Tanaka and Fukushima, 1998) seem to encode both retinal and extraretinal events and may represent smooth pursuit in a spatial, rather than retinal, coordinate system. One common behavioral finding is explained most easily if some of the cortical signals for pursuit are in a spatial frame of reference. In two cases in which pursuit eye velocity was maintained well below target velocity for a single direction of motion (Grasse and Lisberger, 1992; Kiorpes et al., 1996), the deficit was related to the direction of *target* motion, not the direction of *image* motion. Pursuit became entirely normal if image motion in the direction of the deficit was presented during pursuit of target motion in the opposite direction.

At the other end of the pursuit system, current evidence suggests that learning is upstream of Purkinje cells, which are the

output neurons of the cerebellar cortex. The outputs from at least one of the relevant parts of the cerebellar cortex appear to be organized as separate commands for horizontal and vertical eye motion (Miles et al., 1980; Krauzlis and Lisberger, 1996). Vector averaging describes the mechanisms used to create these separate commands and must, therefore, reside either in or before the cerebellar circuits that create Purkinje cell simple spike discharge related to horizontal or vertical eye movements. If vector averaging occurs at the level of or upstream from Purkinje cell output of the cerebellum, then learning occurs even further upstream and may reside before the cerebellum. This conclusion is consistent with the results of single-unit recordings from cerebellar Purkinje cells during pursuit learning (Kahlon and Lisberger, 1997; Kahlon, 1998). Figure 14 implies that the DLPN is a site of vector averaging that is downstream from learning, but available evidence is equally compatible with DLPN as a site of learning that is upstream from vector averaging.

Our data place the site of vector averaging surprisingly far downstream in the pursuit system. If learning is downstream of area MT, and vector averaging is even further downstream, then vector averaging occurs well beyond the immediate outputs from area MT. This does not negate the possibility of vector averaging as a mechanism of readout from cells with similar receptive fields in MT (Groh et al., 1997; Recanzone et al., 1997) but instead raises the possibility that vector averaging happens at multiple levels of the pursuit system. For the stimulus configuration we used, in which the two targets stimulated different parts of the visual field and activated different groups of MT cells, the site responsible for behavioral vector averaging seems to be quite far downstream in the system. Perhaps a more upstream site would be responsible for “local” vector averaging if the two targets stimulated the same area of visual field.

Coordinate transformations in the pursuit system

The coordinate system of area MT appears to be retinal, and its output therefore represents the direction and speed of image

motion as a population code (Maunsell and Van Essen, 1983; Ferrera and Lisberger, 1997). As signals pass from MT through MST and FPA, they are converted to a representation of target motion in space by the addition of extraretinal signals presumably related to eye motion in space (Newsome et al., 1988; Gottlieb et al., 1994; Tanaka and Fukushima, 1998). The diagrams in Figure 14, *left* and *right*, show how the signals needed to drive pursuit responses to single- and two-target stimuli would be represented at each level in our hypothetical pursuit system. At the level of V1 and MT, the direction and speed of each target are represented by a population code. Within the spatial coordinate frame at the site of learning, the visual inputs from the two targets are processed separately to create plans for movements to each of the two targets. The gains of the planned movements are modified according to previous experience and subjected to weighted vector averaging to create a command for a single movement in spatial coordinates. The cerebellum then divides this unified command into control signals for the horizontal and vertical extraocular muscles. At first blush, it seems cumbersome to maintain separate representations of two potential tracking targets far into the pursuit system and to combine those representations by vector averaging only to separate them later into control signals for the horizontal and vertical extraocular muscles. However, the complexity of the signal transformations suggested in Figure 14 is commensurate with the diversity of the basic properties of pursuit behavior and may reflect the necessity of guiding a complex, voluntary movement during natural visual stimuli.

APPENDIX

We derived three linear models in which the response to a given two-target stimulus (\mathbf{AB}_{pre} and $\mathbf{AB}_{\text{post}}$) is described in terms of the weights for vector averaging before learning (w_a and w_b), the responses to the motion of single targets before learning (\mathbf{A} and \mathbf{B}), and the changes in pursuit gain caused by learning (g_a , g_b , and g_{ab}). For the weights and the gains, the subscript indicates the direction of eye and target motion that applies: subscript a means direction \mathbf{A} , subscript b means direction \mathbf{B} , and subscript ab means the direction of the eye movement that resulted from a two-target stimulus consisting of directions \mathbf{A} and \mathbf{B} .

Model 1 places learning upstream from vector averaging:

$$\mathbf{AB}_{\text{post}} = \frac{w_a(g_a\mathbf{A}) + w_b(g_b\mathbf{B})}{w_a + w_b}. \quad (1)$$

Model 2 places learning downstream from vector averaging:

$$\mathbf{AB}_{\text{post}} = g_{ab} \left[\frac{w_a\mathbf{A} + w_b\mathbf{B}}{w_a + w_b} \right]. \quad (2)$$

Model 3 implements learning as changes in the weights used for vector averaging:

$$\mathbf{AB}_{\text{post}} = \frac{(w_a g_a)\mathbf{A} + (w_b g_b)\mathbf{B}}{g_a w_a + g_b w_b}. \quad (3)$$

We assessed the performance of each model against the averaged performance of each monkey separately in each learning condition. For each set of data, we first derived the values of w_a and w_b that provided the best fit weighted vector averaging to the prelearning responses to each two-target stimulus:

$$\mathbf{AB}_{\text{pre}} = \frac{w_a\mathbf{A} + w_b\mathbf{B}}{w_a + w_b}. \quad (4)$$

The weights obtained for any one direction were very similar for all two-target stimuli that included that direction, even though we did not include this constraint in the fitting procedure. Next, we used the three models to predict the eye accelerations evoked by two-target stimuli after learning ($\mathbf{AB}_{\text{post}}$). Note that the procedures used to make these predictions do not use any additional fitting. Instead, we derived the parameters on the right side of each equation from the data, evaluated the equations, and asked which model performed better. The weights (w_a and w_b) assigned each target in two-target stimuli before learning were obtained from Equation 4, and the prelearning eye accelerations (A and B) evoked by single targets were taken directly from the data. The gains representing the effects of learning on the responses to single targets in a given direction (g_a , g_b , and g_{ab}) were derived using approaches described below.

For models 1 and 3, it was possible to compute the gains that represent the effects of learning on single targets (g_a and g_b) from the responses to single-target stimuli that were embedded in the two-target experiments. For each cardinal direction, gain was computed as the magnitude of postlearning eye acceleration for single-target motion in that direction divided by the magnitude of the prelearning response. For model 2, it was not possible to compute the gain (g_{ab}) directly from the responses to single targets in two-target experiments, because the vector-averaged response was almost always in a noncardinal direction. Instead, we interpolated based on the results of the separate experiments on the direction generalization of learning for single target (Fig. 10) to estimate the prelearning and postlearning eye accelerations for the direction of the vector-averaged response. We then computed g_{ab} as the postlearning response divided by the prelearning response. As mentioned previously, one monkey (monkey N) completed both direction generalization and two-target tests in the same experiments. Thus the values of g used for the data of this monkey were taken from the same experiments for all three models.

For each model, appropriate values of w , g , and single-target accelerations (\mathbf{A} and \mathbf{B}) were plugged in to the equations to predict postlearning responses to two-target stimuli. However, these procedures incorporated only the specific effects of learning. To make a valid comparison with the data, it was also necessary to estimate and include nonspecific effects. We estimated the nonspecific effects exactly as we had from the data: we computed the average difference between the output of each equation and the actual data for orthogonal two-target stimuli that included target motion in the control direction. We then added the estimated nonspecific effects to the predictions for all orthogonal two-target stimuli. This yielded perfect fits for control quadrants and enabled a valid test of how well the model fitted data from two-target stimuli that included the learning direction. We chose to compute the nonspecific effects from the predictions of the models because they included estimates of both the actual shifts and any additional nonspecific error that may have arisen in the estimation of weights from prelearning data. We obtained essentially the same final model predictions by adding the nonspecific effects measured from the data, although the results were noisier for two-target pairs that were relatively ill fit in prelearning data.

REFERENCES

- Andersen RA, Snyder LH, Li CS, Striccanne B (1993) Coordinate transformations in the representation of spatial information. *Curr Opin Neurobiol* 3:171–176.

- 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.
- Goldreich D, Krauzlis RJ, Lisberger SG (1992) Effect of changing feedback delay on spontaneous oscillations in smooth pursuit eye movements of monkeys. *J Neurophysiol* 67:625–638.
- 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.
- Grasse KL, Lisberger SG (1992) Analysis of a naturally occurring asymmetry in vertical smooth pursuit eye movements in a monkey. *J Neurophysiol* 67:164–179.
- 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.
- Judge SJ, Richmond BJ, Chu FC (1980) Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res* 20:535–538.
- Kahlon M (1998) Learning in a cortical motor system. PhD Thesis, University of California, San Francisco.
- Kahlon M, Lisberger SG (1996) Coordinate system for learning in the smooth pursuit eye movements of monkeys. *J Neurosci* 16:7270–7283.
- Kahlon M, Lisberger SG (1997) Neuronal correlate of pursuit learning in the floccular complex of the cerebellum. *Soc Neurosci Abstr* 23:1298.
- Kiorpes L, Walton PJ, O'Keefe LP, Movshon JA, Lisberger SG (1996) Effects of early-onset artificial strabismus on pursuit eye movements and on neuronal responses in area MT of macaque monkeys. *J Neurosci* 16:6537–6553.
- Kowler E (1990) The role of visual and cognitive processes in the control of eye movement. *Rev Oculomot Res* 4:1–70.
- Krauzlis RJ, Lisberger SG (1994) Simple spike responses of gaze velocity Purkinje cells in the floccular lobe of the monkey during the onset and offset of pursuit eye movements. *J Neurophysiol* 72:2045–2050.
- Krauzlis RJ, Lisberger SG (1996) Directional organization of eye movement and visual signals in the floccular lobe of the monkey cerebellum. *Exp Brain Res* 109:289–302.
- Lisberger SG, Ferrera VP (1997) Vector averaging for smooth pursuit eye movements initiated by two moving targets in monkeys. *J Neurosci* 17:7490–7502.
- Lisberger SG, Movshon JA (1999) Visual motion analysis for pursuit eye movements in area MT of macaque monkeys. *J Neurosci* 19:2224–2246.
- Lisberger SG, Westbrook LE (1985) Properties of visual inputs that initiate smooth pursuit eye movements in monkeys. *J Neurosci* 5:1662–1673.
- Maunsell JHR, Van Essen DC (1983) Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. *J Neurophysiol* 49:1127–1147.
- Newsome WT, Wurtz RH, Komatsu H (1988) Relation of cortical areas MT and MST to pursuit eye movements. II. Differentiation of retinal from extraretinal inputs. *J Neurophysiol* 60:604–620.
- Ogawa T, Fujita M (1997) Adaptive modifications of human postsaccadic pursuit eye movements induced by a step-ramp-ramp paradigm. *Exp Brain Res* 116:83–96.
- Optican LM, Zee DS, Chu FC (1985) Adaptive responses to ocular muscle weakness in human pursuit and saccadic eye movements. *J Neurophysiol* 54:110–122.
- Press WH, Teukolsky SA, Vetterling WT, Flannery BP (1992) Numerical recipes in C: the art of scientific computing, Ed 2, pp 666–670. Cambridge, UK: Cambridge UP.
- Recanzone GH, Wurtz RH, Schwarz U (1997) Responses of MT and MST neurons to one and two moving objects in the receptive field. *J Neurophysiol* 78:2904–2915.
- Schwartz JD, Lisberger SG (1994) Initial tracking conditions modulate the gain of visuo-motor transmission for smooth pursuit eye movements in monkeys. *Vis Neurosci* 11:411–424.
- Shidara M, Kawano K, Gomi H, Kawato M (1993) Inverse-dynamics model of eye movement control by Purkinje cells in the cerebellum. *Nature* 365:50–52.
- Skavenski AA, Robinson DA (1973) Role of abducens neurons in vestibuloocular reflex. *J Neurophysiol* 36:724–738.
- 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.
- Tusa RJ, Ungerleider LG (1988) Fiber pathways of cortical areas mediating smooth pursuit eye movements in monkeys. *Ann Neurol* 23:174–183.
- Van der Steen J, Simpson JJ, Tan J (1994) Functional and anatomic organization of three-dimensional eye movements in rabbit cerebellar flocculus. *J Neurophysiol* 72:31–46.
- van Donkelaar P, Fisher C, Lee RG (1994) Adaptive modification of oculomotor pursuit influences manual tracking responses. *NeuroReport* 5:2233–2236.