

Task-Dependent Modulation of the Sensorimotor Transformation for Smooth Pursuit Eye Movements

VINCENT P. FERRERA

Department of Psychiatry and Center for Neurobiology and Behavior, Columbia University, New York, New York 10032

Received 19 January 1999; accepted in final form 18 August 2000

Ferrera, Vincent P. Task-dependent modulation of the sensorimotor transformation for smooth pursuit eye movements. *J Neurophysiol* 84: 2725–2738, 2000. To investigate the transformation of retinal image velocity into smooth pursuit eye velocity, eye movements were measured in the presence of two moving targets. In the first experiment, the targets were identical in all respects except for direction of motion, and the monkey was not cued to attend to either target. In this experiment, smooth pursuit eye velocity elicited by two targets was the vector average of the response evoked by each target alone. In subsequent experiments, we examined the effects of stimulus and task parameters on the selectivity of pursuit. When the targets were made different colors and monkeys were cued for the color of the rewarded target, their pursuit eye movements were biased in the direction of the rewarded target, but still showed a substantial influence of the nonrewarded target (distractor). It did not matter whether the same target color was used for an entire session or whether the color was randomized from trial to trial. Reducing uncertainty about the axis of motion of the rewarded target also had little effect. However, the pattern of image motion appeared to have a substantial effect; radial image motion favored averaging, and winner-take-all pursuit was found only with nonradial image motion. We conclude that the sensorimotor interface for pursuit uses a flexible decision rule that can vary continuously from vector averaging to winner-take-all. We present a simple recurrent network model that reflects this range of behavior. The model has allowed us to identify three computational elements (selection bias, competitive inhibition, and response normalization) that should be taken into consideration in future models of smooth pursuit.

INTRODUCTION

To execute voluntary goal-directed movements in complex environments, sensorimotor systems must be able to select individual targets and to filter out distracting stimuli. How well these systems are able to suppress irrelevant stimuli limits the accuracy with which movements can be made. Several lines of evidence have suggested that sensorimotor systems use either a vector averaging or vector summation strategy to compute movement trajectories given a topographically distributed pattern of neuronal activity (Georgopoulos et al. 1993; Glimcher and Sparks 1993; Groh et al. 1997; Lee et al. 1988; Lisberger and Ferrera 1997; Robinson and Fuchs 1969). However, any computation that averages or sums information over multiple stimuli could impair the ability to make precise movements directed toward a single stimulus. Presumably, one of the functions of selective attention is to overcome averaging by restricting the set of sensory signals that are used in program-

ming movements to those that arise from attended objects or locations. Little is known about how effectively focused attention is able to suppress signals from nonselected stimuli, although it is often assumed that attention is highly selective. In the current study, we describe experiments on rhesus monkeys in which we quantified the selectivity of voluntary smooth pursuit eye movements made in the presence of two moving stimuli. We show that the selectivity of these movements can vary from completely nonselective (vector averaging) to highly selective (winner-take-all), depending on stimulus and task-specific factors. Surprisingly, prior knowledge of the correct target does not always result in accurate pursuit, but only biases the direction of the initial eye movement.

The results of this study indicate that the sensorimotor transformation for smooth pursuit, i.e., the transformation of image velocity into eye velocity, is not characterized by a single fixed computation, but rather that it is flexible and subject to attentional modulation. We show that the results can be accounted for by a simple recurrent network model (Koch and Ullman 1985) in which direction-selective neurons compete with one another through mutual inhibition to drive the motor response. In fact, the entire range of observed behavior, including vector-averaging and winner-take-all, can be simulated by varying parameters that determine the strength of attentional bias and inhibitory interactions in the network. This suggests that competitive inhibition may play a role in the attentional control of eye movements and is consistent with physiological observations in the frontal eye field (Burman and Bruce 1997), but not visual areas MT or MST (Ferrera and Lisberger 1997b; Groh et al. 1997; Recanzone et al. 1997). The current study fills a gap between two previous reports that demonstrated vector-averaging for pursuit in the absence of attentional bias (Lisberger and Ferrera 1997), and winner-take-all pursuit when attentional bias was present (Ferrera and Lisberger 1997a). The current report extends these findings by showing that there is a graded transition between these two extremes of behavior, and by providing a mechanism to account for the transition.

METHODS

Experiments were conducted on four juvenile male rhesus monkeys (*Macaca mulatta*). *Monkey K* was trained and tested at University of California at San Francisco, while *monkeys A, C, and D* were trained

Address for reprint requests: Center for Neurobiology and Behavior, PI Annex Rm. 729, Columbia University, 722 W. 168th St., New York, NY 10032 (E-mail: vpf3@columbia.edu).

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

and tested at Columbia University. Our methods were approved by the Institutional Animal Care and Use Committees at UCSF and Columbia University. Monkeys were trained to move voluntarily from their home cage to a primate chair. A method modified from Wurtz (1969) was used to train each monkey to attend a stationary target. Surgery was then performed under sterile conditions to implant a coil of wire on one eye (Judge et al. 1980) and to secure a post to the skull for head restraint (Miles and Eighmy 1980). For all subsequent training and experiments, the monkey's head was secured to the primate chair, and a set of field coils was lowered over the chair so that a magnetic search-coil could be used to monitor horizontal and vertical eye position. The eye coil was calibrated by having the monkey fixate targets at different positions, and the monkey was subsequently required to keep the direction of gaze within 2–3° of target position. Correct performance of the task was rewarded with drops of fruit juice or water.

Behavioral tasks

Monkeys were trained to track moving targets presented on a color cathode ray tube (CRT) monitor. We used a step-ramp target motion paradigm to minimize the occurrence of saccades during pursuit initiation (Rashbass 1961). Trials were initiated by requiring the monkey to look at a stationary central fixation light. After a short interval, one or two perifoveal moving targets appeared. At the same time, the central fixation light was turned off, and the monkey was required to initiate a smooth pursuit eye movement. The target always moved at 15°/s. The initial target position was set so that targets moving toward the initial fixation position would cross it after 200 ms. The monkey was given a liquid reward provided that he kept his gaze directed toward the correct target for the duration of the trial. The monkey's performance was monitored by tracking his eye position relative to a $\pm 3^\circ$ fixation window centered around the target. Several steps were taken to ensure that the monkeys did not make anticipatory eye movements. First, the initial position and direction of the target were randomized trial-to-trial so that the monkey could not anticipate correctly the direction of the required eye movement. Second, the time, relative to the start of the trial, at which the target was presented or started to move was randomized. Finally, trials were aborted if the monkey initiated an eye movement before the fixation light went off. Trials were randomized within blocks, and the monkey was required to correctly complete exactly one trial of each type before proceeding to the next block of trials.

Visual stimulation

Visual stimuli were generated by a Univision Piranha or CRS VSG2/3F video framebuffer with an on-board microprocessor (Texas Instruments TMS 34020). The output from the video board was displayed on a calibrated 20-in. (Barco) or 27-in. (Mitsubishi) color monitor with a 60-Hz noninterlaced refresh rate. The monitors stood at a viewing distance of 30 in. so that the display area subtended roughly 30° horizontally by 20° vertically (Barco) or 40 × 30° (Mitsubishi). The spatial resolution of both displays was 1,280 pixels by 1,024 lines, and the depth was 8 bits/pixel. Pursuit targets were small (1.0°) colored squares presented on a uniform gray background. The target luminance was 15.0 cd/m², while the background was 5.0 cd/m². The framebuffer was programmed to send out digital pulses (frame sync) for timing purposes at the beginning of each frame in which a stimulus first appeared or started to move. These pulses were sampled by the computer and stored along with the eye movement data.

Eye movement recording

Eye position was monitored using a monocular scleral search coil system (CNC Engineering). Separate horizontal and vertical eye po-

sition signals were fed through an analog differentiator (low-pass, -3 dB at 25 Hz) to yield horizontal and vertical eye velocity. The eye position and eye velocity signals were then digitally sampled by computer with 12-bit resolution (0.025° for position, 0.092°/s for velocity) at 1 kHz/channel and stored on disk for further analysis. Eye acceleration was computed off-line by digital differentiation of eye velocity.

To avoid contaminating estimates of smooth pursuit eye velocity with the much higher eye velocities that accompany saccades, we used an automatic algorithm to detect and remove saccades that occurred during a 400-ms interval starting at the onset of target motion. The algorithm used an acceleration criterion of 500°/s² to detect the beginning and end of each saccade. Eye velocity was linearly interpolated between the beginning and end of the saccade. Saccade cutting was applied separately to both the horizontal and vertical eye velocity records. We excluded trials when the monkey blinked, failed to attend to the task, or otherwise failed to produce clean eye velocity data. Out of a total of 43,456 trials, 571 (1.3%) were excluded for these reasons.

To estimate smooth pursuit latency, we developed a method based on signal detection theory (SDT) (Green and Swets 1966). For each trial, the horizontal and vertical eye velocity traces $\dot{H}(t)$ and $\dot{V}(t)$ were first combined into a single radial eye velocity trace $[\dot{R}(t)]$. SDT was used to determine the earliest time, T , such that the distribution of eye velocity samples within a fixed time window following T was significantly different from the distribution within a similar window preceding T . For a given time, T , relative to the onset of target motion, we set windows of 100 ms duration before and after T . We then made frequency distributions of the eye velocity samples within the pre- T and post- T time windows. We determined the degree of overlap between the two distributions by constructing a "receiver operating characteristic" (ROC) curve based on the two distributions. To make the ROC curve, we chose a criterion velocity and then computed the proportion of each distribution (relative to its total area) above the criterion value. The number derived from the post- T distribution is called the "hit-rate," and that derived from the pre- T distribution is called the "false-alarm-rate." By plotting the hit-rate versus the false-alarm-rate, we obtained a single point on the ROC curve. We generated the entire ROC curve by varying the velocity criterion in small steps over the entire range of the two distributions. Next, we calculated the area under the ROC curve. If the two distributions overlapped completely, then the hit-rate and false-alarm-rate were equal and the area under the ROC curve was 0.5. If the two distributions were nonoverlapping, then as the velocity criterion was reduced, the hit-rate approached 1.0 while the false-alarm rate remained zero, and therefore the area under the ROC was 1.0. This entire procedure was repeated for each point in time following the onset of target motion so that the eye velocity trace, $\dot{R}(t)$, was converted into a probability function, $p(t)$, which varied between 0.5 and 1.0. The onset of pursuit was taken to be the time at which the probability function first reached a criterion of 0.95 and remained at or above that level for at least 50 ms. Using these parameters, the algorithm performed in a manner that agreed well with the experimenter's determination of pursuit onset by visual inspection. No other means were used to optimize the performance of the algorithm.

It will be noted that this is a theoretically unorthodox application as SDT is normally used to compare independent observations drawn from stationary distributions, whereas eye velocity samples are serially correlated and the mean of the underlying distribution changes continuously during pursuit initiation. Empirically, however, the method satisfied our goal of providing accurate and reproducible estimates of pursuit latency. The performance of the algorithm is demonstrated in Fig. 1, which shows eye velocity traces for a single trial with one upward-moving target (A) together with the derived probability (area under the ROC curve) trace. The minimum latency allowed was 40 ms, so the probability trace is initially flat. Figure 1C shows all trials with one upward-moving target for one animal (K)

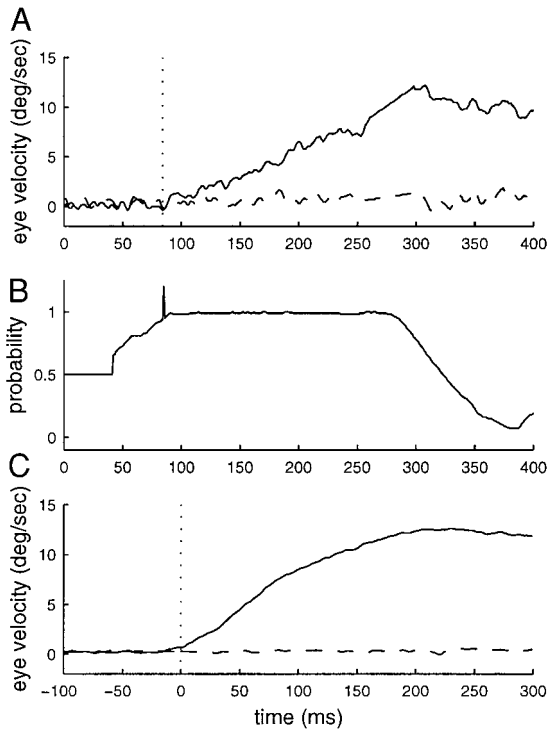


FIG. 1. Method of determining smooth pursuit latency. *A*: horizontal (dashed) and vertical (solid) eye velocity traces for one trial with a single upward-moving target. *B*: probability [area under receiver operating characteristic (ROC) curve] trace derived from the data in *A*. A "spike" was inserted in the probability trace to indicate the derived latency. This corresponds to the dotted line in *A*. *C*: eye velocity data aligned on pursuit onset and averaged over 36 trials.

aligned on pursuit onset ($t = 0$). Ideally, the average eye velocity should be indistinguishable from zero prior to $t = 0$.

RESULTS

Experiment 1

We measured smooth pursuit eye movements evoked by two small targets moving in different directions (Fig. 2, *A* and *B*). In *experiment 1*, the two targets were identical in size (1.0°), shape (square), color (yellow), luminance (15 cd/m^2), and speed ($15^\circ/\text{s}$; Fig. 2, *A* and *B*). The targets could move in any of eight directions spaced 45° apart (4 cardinal directions plus 4 obliques). There were 64 different trial types that exhausted all possible combinations of two targets and eight directions. In cases where the two target directions would have been identical, a single target (luminance 15 cd/m^2) was presented instead. The monkey was trained to initiate smooth pursuit as soon as any target appeared, but was not trained to selectively track or attend to a particular target. On two-target trials, both targets were present only for the first 200 ms of target motion. Thereafter, only one target was present, and the monkey was required to keep his eye position within 3° of this target to get his reward. It was not possible for the monkey to predict in advance which target would remain at the end of the trial.

Figure 2*C* shows the eye velocity response for *monkey A* recorded during four trial types. For two of the trial types, two moving targets were presented (Fig. 2, *A* and *B*), and for the other two trial types each target was presented alone. The horizontal component of target motion was the same for both

targets, and hence the horizontal eye velocities in each of the four trial types were indistinguishable and are not shown. Figure 2*C* shows vertical eye velocity. Each eye velocity trace is the average of approximately 20 individual trials. The solid traces correspond to single target trials, and the dashed traces correspond to two-target trials. For the first 200 ms of pursuit, the vertical eye velocity traces for the two-target trials are not distinguishable from each other or from zero velocity, indicating that the monkey was tracking in a purely horizontal direction during this interval. This is borne out in Fig. 2*D*, in which the horizontal and vertical eye velocities during the first 200 ms of pursuit have been combined to show the direction of pursuit at each point in time.

We performed a two-dimensional vector analysis of the eye movements evoked by two moving targets. In previous work (Lisberger and Ferrera 1997), this analysis had been based on eye acceleration during an interval 156–206 ms after the onset of target motion. This analysis was prone to the following drawbacks. 1) Eye acceleration is noisier than eye velocity due to an additional numerical differentiation. 2) The constant time interval relative to target motion onset ignores variations in the latency of pursuit, which can introduce additional variability into the acceleration measurements. 3) The use of a single time interval does not give a sense of the dynamics of the response. To address these issues, we chose to analyze average eye

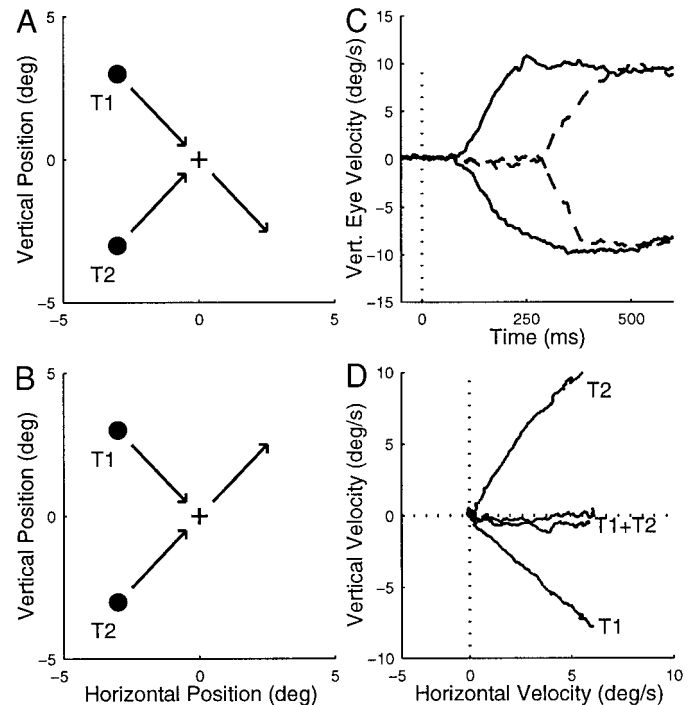


FIG. 2. *A* and *B*: the stimuli for 2 trials conditions of *experiment 1*. The monkey first fixates a small target (+) in the center of the cathode ray tube (CRT) display. Two identical targets appear simultaneously 3° away and move back toward the fixation point at $15^\circ/\text{s}$. *C*: vertical eye velocity records averaged over 20+ trials. Solid lines are single target trials; dashed lines are paired target trials. *D*: eye velocity samples for *monkey A* during the interval 150–300 ms after target motion onset to show the direction and speed of eye movement. There are 4 conditions: T1 alone, T2 alone, and 2 conditions with T1 + T2 as shown in *A* and *B*. Each of the 4 trajectories represents horizontal and vertical eye velocity parameterized as a function of time. Each trajectory was constructed by 1st averaging over 20+ trials the separate horizontal [$H'(t)$] and vertical [$V'(t)$] eye velocity traces. Then, for each point in time $V'(t)$ is plotted against $H'(t)$.

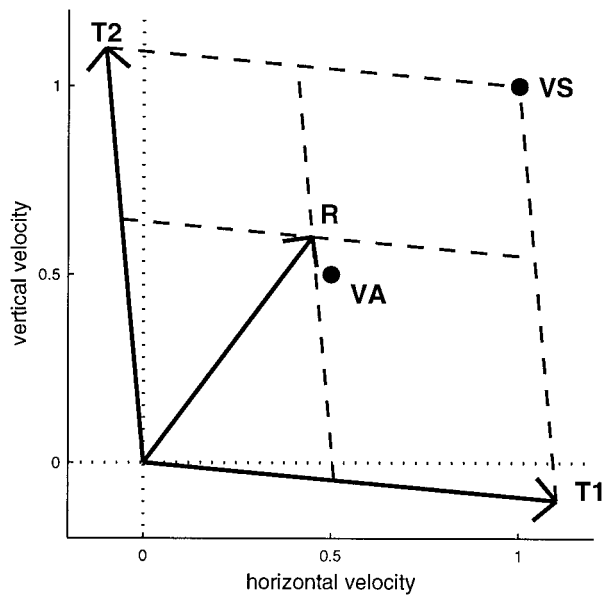


FIG. 3. Vector analysis of eye velocity. The response to 2 targets presented together (R) is decomposed into a weighted sum of the responses obtained with each target individually (T_1 and T_2). VA is the 2-target response expected under vector averaging of T_1 and T_2 . VS is the vector sum of T_1 and T_2 . Winner-take-all would correspond to $R = T_1$ or $R = T_2$.

velocity measured during six 25-ms time intervals starting from the onset of pursuit. One theoretical objection to the choice of eye velocity (and the reason previous studies have favored the use of eye acceleration) is that eye velocity is not constant during the initiation of pursuit. The lack of stationarity is ameliorated by the use of relatively short time intervals and by the fact that the mean velocity over a given interval is a good approximation of the eye velocity sampled at the interval midpoint provided that acceleration is constant.

The first step in the vector analysis was to compute eight basis vectors that represent the average eye velocity evoked by single targets moving in any of eight directions. It was then possible to express the eye velocity (R , Fig. 3) for each two-target trial as a weighted sum of the component basis vectors (T_1, T_2)

$$\vec{R} = w_1 \vec{T}_1 + w_2 \vec{T}_2 \quad (1)$$

Cases where the target directions were 180° apart were excluded from this analysis because the basis vectors tended to be co-linear and this made it difficult to recover the orthogonal dimension. For example, if both basis vectors were close to horizontal, then extremely large weights might be required to represent responses that had a small vertical component, and these large weights would distort the results. We also excluded trials where the eye movement latency was outside the range 50–300 ms or where either weight was outside the range -1.0 to 2.0 . This resulted in the loss of a further 3.7% of the data in addition to the 1.3% lost to blinks, etc. We did not exclude any data based on whether we judged the animal's response to be correct or incorrect, as there was no objective basis for doing so in the first experiment, and we wanted to treat the data in the same manner for all experiments. Vector decomposition allowed us to identify several interesting outcomes according to the values of the weights w_1 and w_2 . Pure vector averaging (VA) corresponds to the case where $w_1 = w_2 = 0.5$. Vector

summation (VS) corresponds to $w_1 = w_2 = 1.0$. If the response to two targets is identical to the response to a single target ($w_1 = 1.0$ and $w_2 = 0.0$, or $w_1 = 0.0$ and $w_2 = 1.0$), then the outcome is said to be “winner-take-all” (WTA) for the target with the nonzero weight. Winner-take-all tracking implies that the pursuit system is operating in a highly selective mode, whereas anything short of WTA implies that the output of the system reflects influences from both target motions.

An example of data for one condition (*monkey K*, interval 6: 125–150 ms after pursuit onset, $T_1 = 0^\circ$, $T_2 = 90^\circ$) are shown in Fig. 4A. The results of the vector analysis for this condition are shown in Fig. 4B. The weights computed for the entire experiment are shown in Fig. 5. Each small dot in Fig. 5A represents the weights calculated for a single two-target trial for *monkey K*. Despite the variability in the data, the centroid of the weight distribution (Fig. 5A, open circle) is very close to pure vector averaging. The individual trial weights for *monkeys A, C, and D* are not shown, but the centroids of their weight distributions are indicated by the square, plus sign, and x , respectively, in Fig. 5A.

Figure 5, *B* and *C*, shows the individual weight distributions

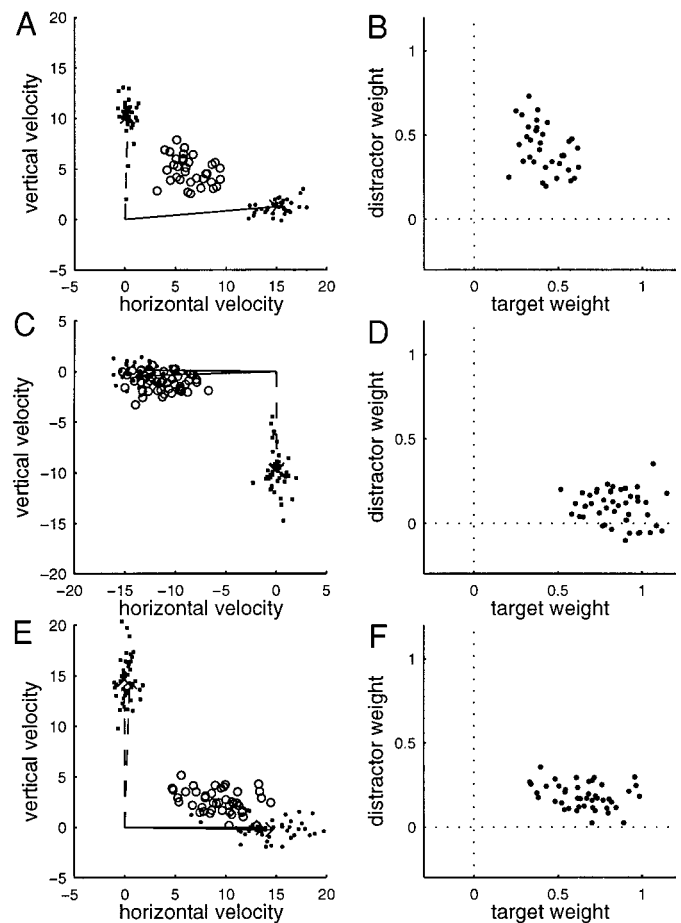


FIG. 4. Raw eye velocity data for selected conditions (*A, C, and E*) and derived weights (*B, D, and F*). *A*: experiment 1, *monkey K*, interval 6. Small circles and squares are eye velocity measured on individual trials with single targets. Solid and dashed arrows are the mean single target velocity vectors. Open circles are the eye velocity response to the 2 targets presented simultaneously. *B*: solid circles are the target vs. distractor weight for each 2 target trial in *A*. *C* and *D*: experiment 2, *monkey A*, interval 6. Same conventions as *A* and *B*. *E* and *F*: experiment 2, *monkey C*, interval 6. Same conventions as *A* and *B*.

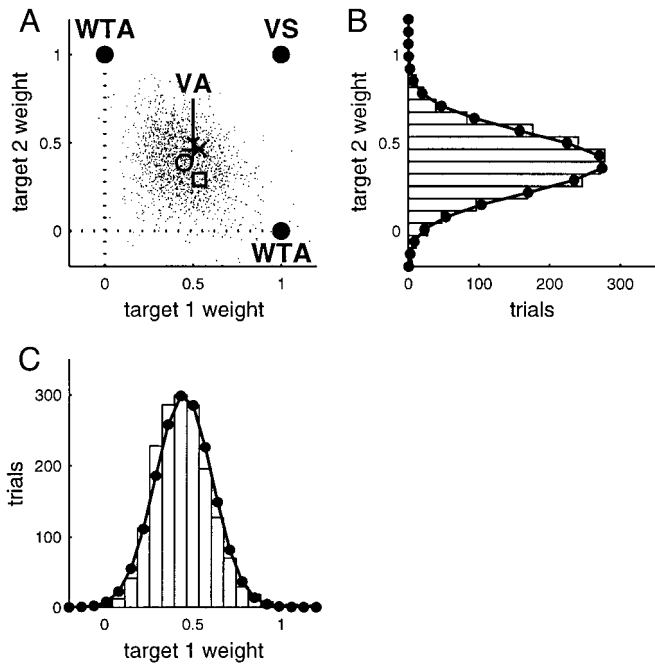


FIG. 5. Weight vectors for *experiment 1*. *A*: the dots represent weight vectors calculated for each 2-target trial for *monkey K*. The open circle, cross, square, and plus sign are the centroids of the weight distributions for *monkeys K, A, C, and D*, respectively. *B* and *C*: frequency distributions of the 2 target weights for *monkey K*. Solid circles and lines are the best-fitting Gaussians. WTA, winner-take-all.

(30 bins spanning the range -1 to 2) and their best-fitting Gaussians. Gaussian parameters were optimized using the Nelder-Mead algorithm in Matlab 5.3. The optimization was successful in 288/288 conditions (4 monkeys \times 6 experiments \times 6 time intervals \times 2 wt distributions), indicating that all of the weight distributions were well approximated by Gaussians. This is supported by Fig. 6, showing the correlation between the fitted functions and the data for the complete set of experiments in this study. The distributions (h) and fitted functions (y) were normalized to $\max(h, y)$ before performing the correlation. The clustering of points along the lines $x = 1$ and $y = 1$ is an expected consequence of this normalization.

To test for bimodality in the weight distributions, we did the following. First, we fit each weight distribution with a sum of two Gaussians, where the mean, standard deviation, and height of each Gaussian component was independently varied to provide the optimal fit using the Nelder-Mead algorithm. We then identified all local minima and maxima of the fitted double Gaussian. We considered the fitted function to be bimodal if there were two peaks and the smaller peak was at least 3% the size of the larger peak. Seven of 288 weight distributions were bimodal according to this criterion. There was no single condition (i.e., experiment, monkey, and time interval) for which both weight distributions were bimodal. To quantify the degree of bimodality, we calculated the average height of the two peaks (P_1, P_2) relative to the valley (V) separating them in proportion to the average height of the peaks overall according to the following formula $BMI = (P_1 + P_2 - V)/(P_1 + P_2)$. The bimodality index (BMI) can range from 0.0 [no valley, $V = (P_1 + P_2)/2$] to 1.0 (full valley, $V = 0$). The mean bimodality index was 0.08 ± 0.02 (mean \pm SE; range 0.03–0.17).

Vector analysis was performed to derive the target weights

for each two-target trial. The means of these weight distributions for all four monkeys are plotted as a function of time interval relative to pursuit onset in Fig. 7A. The dotted line is the expected value for vector averaging, and the dashed lines are the values expected for WTA. A paired t -test was performed for each monkey and time interval to test whether the mean weights were significantly different. The small open symbols plotted slightly to the left of each interval show the resulting P values (2-tailed).

Figure 7, *B* and *C*, shows the direction and amplitude of the weight vectors. Again, the dotted and dashed lines are the VA and WTA predictions, respectively. Note that the expected amplitude for vector summation ($\sqrt{2}$) is off scale. Two-way analyses of variance (independent variables: monkey, interval; dependent variables: weight vector direction and amplitude) showed that monkey and interval were significant factors for both the amplitude ($p_{\text{monk}} = 0.000$; $p_{\text{int}} = 0.000$) and direction ($p_{\text{monk}} = 0.000$; $p_{\text{int}} = 0.011$) of the weight vector.

Experiments 2 and 3

The results of *experiment 1* show the monkey's behavior when he does not know in advance which target will lead to a reward. These results provide a baseline for evaluating the selectivity of smooth pursuit when the animal is given prior information that might bias his response toward one target or the other. In the next two experiments, we provided this information by making the targets different colors (red or green, 15 cd/m² luminance) and making the reward contingent on tracking a particular color. The nonrewarded target is referred to as the distractor. During one recording session the monkey was rewarded only when he tracked green targets and during an-


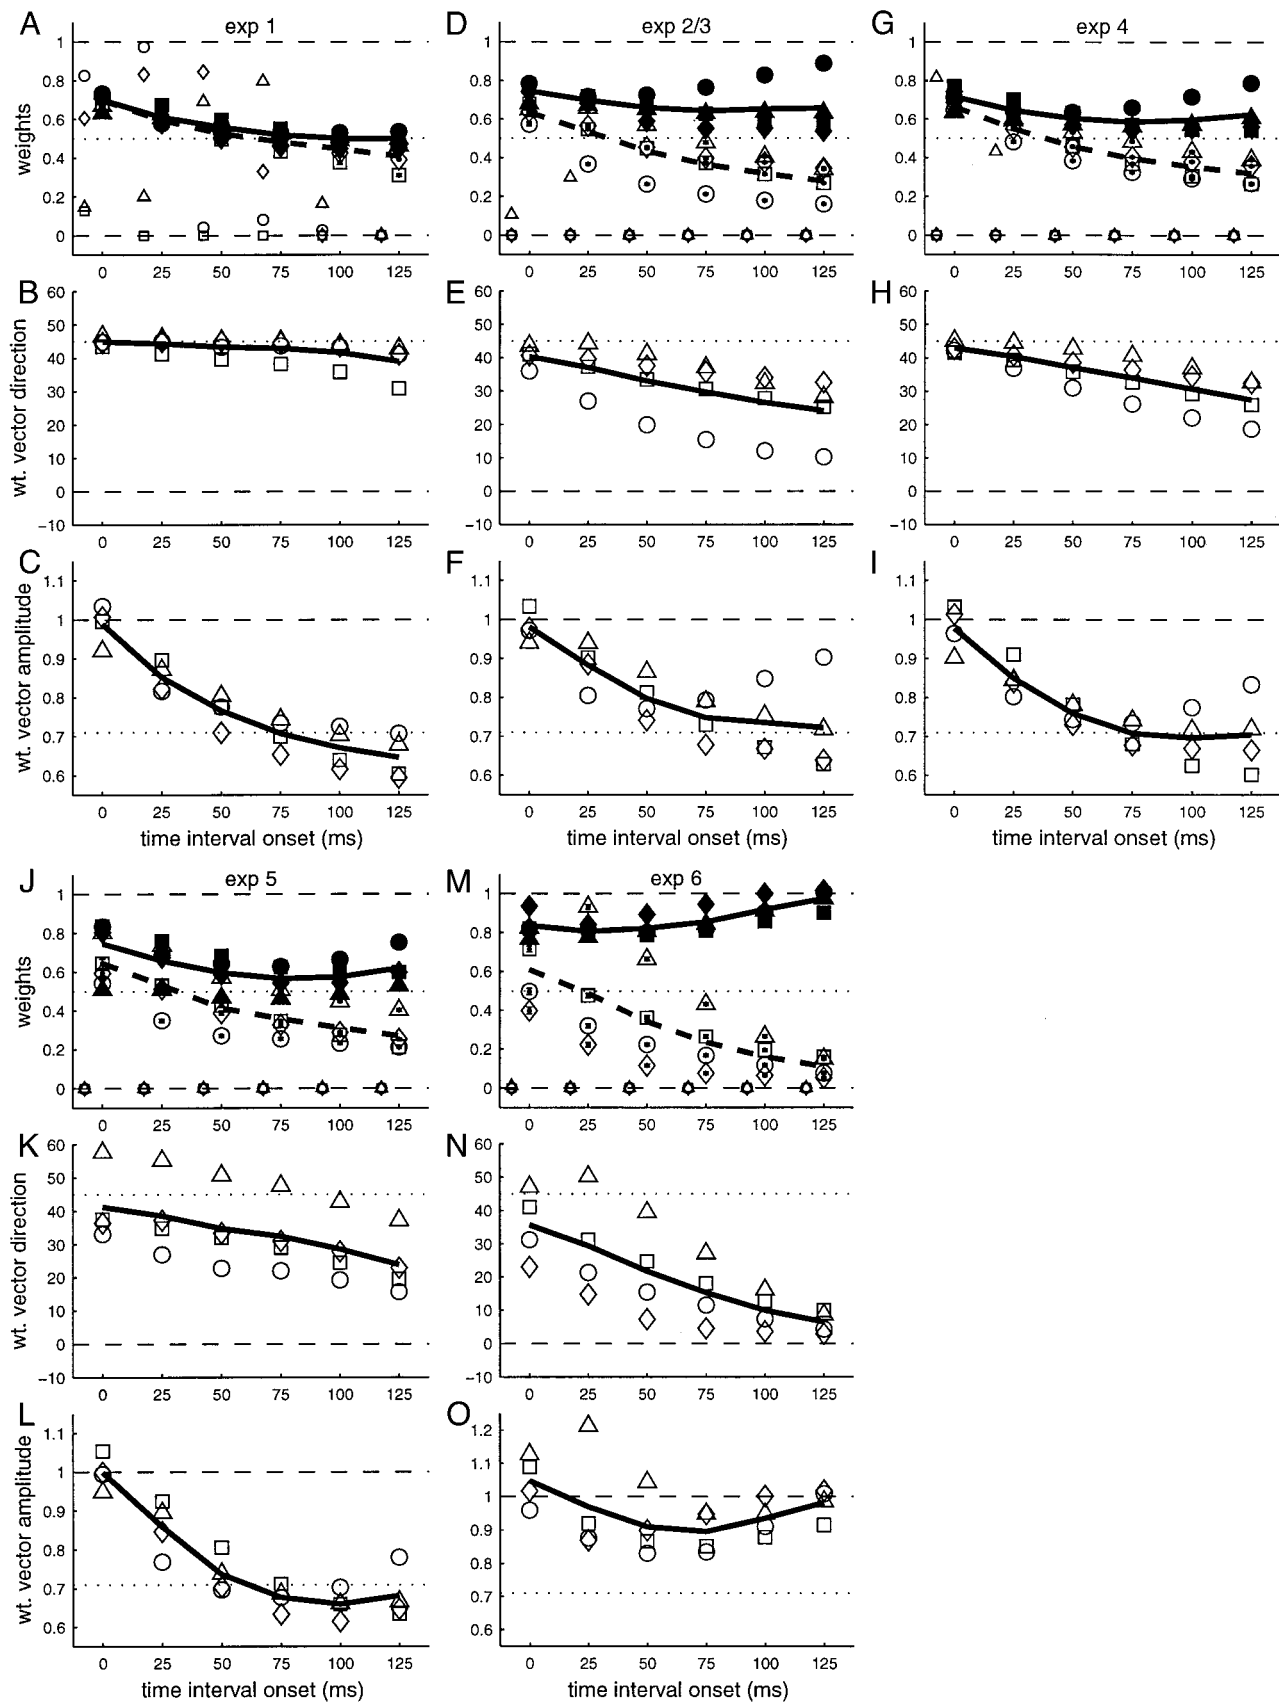


FIG. 6. Correlation between weight distributions and Gaussian fits. Each point is the fitted vs. observed value for a single histogram bin. The data were grouped according to target vs. distractor (w_1 or w_2) and subject to linear regression; $m_i, b_i,$ and r_i ($i = 1, 2$) are the slopes, intercepts, and correlation coefficients for the regressions.



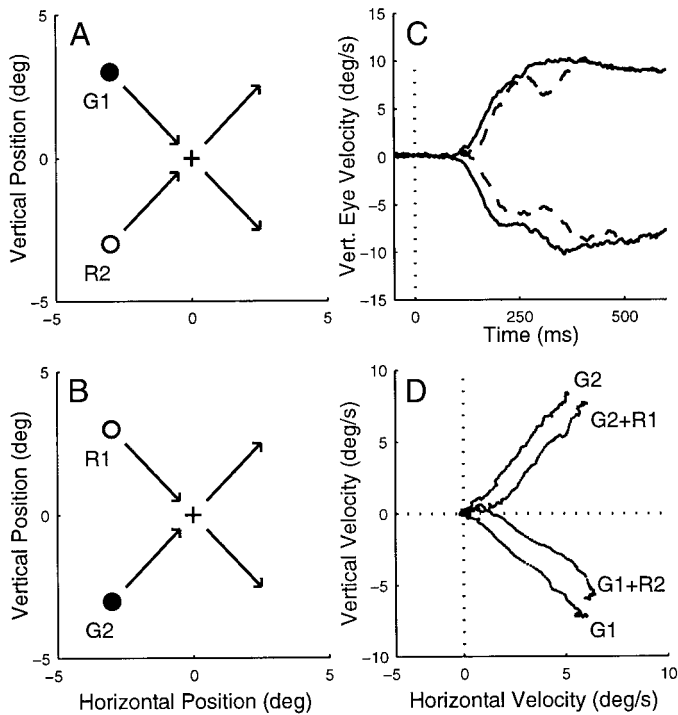


FIG. 8. *A* and *B*: the stimuli for 2 trials conditions of *experiment 2*. Two targets of different colors appear simultaneously 3° away and move toward the fixation point at $15^\circ/\text{s}$. *C*: vertical eye velocity records averaged over 20+ trials. Solid lines are single target trials; dashed lines are paired target trials. *D*: eye velocity samples for *monkey A* during the interval 150–300 ms after target motion onset are plotted to show the direction and speed of eye movement. There are 4 conditions: G1 alone, G2 alone, G1 + R2, and G2 + R1 as shown in *A* and *B*. Each of the 4 trajectories represents horizontal and vertical eye velocity parameterized as a function of time.

other, he was rewarded only for tracking red targets. The order was not the same for all monkeys, but for expository purposes we refer to the experiment with green targets as *experiment 2* and the one with red targets as *experiment 3*. Either target or distractor could move in any of eight directions. Sessions were spaced 2 or 3 days apart, and the monkey had one day of practice with each version of the task before data were recorded. For this and all following experiments, both targets were present and moving from the time they appeared until the end of the trial.

Figure 8 shows responses elicited by four trial types: two in which there was only a single green target (G1 or G2 in Fig. 8, *A* and *B*), and two in which there was a green target with a red distractor (G1 + R2 or G2 + R1 in Fig. 8, *A* and *B*). Figure 8*C* shows vertical eye velocity averaged over roughly 20 trials of each type. The two solid lines correspond to G1 or G2 alone, while the dashed lines correspond to G1 + R2 and G2 + R1. There appears to be some averaging during the initial 20 ms of pursuit, which is why the onset of pursuit occurs later for the two target trials. Thereafter, the distractor appears to have very little effect on the direction of pursuit. [The oscillations in eye velocity that occur roughly 300 ms after target motion onset

were not a consistent feature of the data. Such oscillations are not uncommon during pursuit initiation and have been described and modeled by Goldreich et al. (1992).] This is corroborated by Fig. 8*D*, where the trajectories for single and paired targets are nearly parallel. Comparing Fig. 8, *C* and *D*, to Fig. 2, *C* and *D*, shows that prior knowledge of the target color has a striking effect on the initial direction of pursuit.

The data of *experiments 2* and *3* were subjected to vector decomposition according to Eq. 1 on a trial-by-trial basis. Figure 4, *C–F*, shows examples of the raw data and derived weights for two conditions (Fig. 4, *C* and *D*: *monkey A*, interval 6, Target = 180° , Distractor = 270° ; Fig. 4, *E* and *F*: *monkey C*, interval 6, $T = 0^\circ$, $D = 90^\circ$). The means of the target and distractor weight distributions for all conditions are shown in Fig. 7*D*. The weight vector directions and amplitudes are shown in Fig. 7, *E* and *F*. A three-way ANOVA (independent factors: experiment, monkey, interval; dependent variables: weight vector amplitude and direction) showed that there was no difference between experiments for weight vector direction ($P = 0.11$) and that the effect on amplitude, while significant ($P = 0.009$), was small (*experiment 2* mean = 0.87 vs. 0.90 for *experiment 3*). There was no significant effect of monkey (amplitude: $P = 0.745$; direction: $P = 0.562$), but the effect of time interval was highly significant ($P = 0.000$ and 0.002 for amplitude and direction, respectively). Because the results for the two experiments were similar, they were combined. The data shown in Fig. 4, *C–F*, and Fig. 7, *D–F*, are for the combined experiments.

The difference between animals probably is not related to their individual training histories or the behavioral requirements of the task, which were similar if not identical for all animals. There was a grace period of 300 ms starting at the onset of target motion during which there was no fixation requirement. Only after 300 ms was the monkey required to have his eye position within 3.0° of the rewarded target. The reason for this was to avoid artificially biasing the animal's behavior in favor of any particular outcome. The same grace period and fixation window were used in *experiment 1*, thus showing that vector averaging was not ruled out by the behavioral requirements. Although winner-take-all tracking of the rewarded target might have been a somewhat better strategy, the monkeys were free to track in any direction for the first 300 ms and then switch to the rewarded target without substantially affecting their reward rate. Using a smaller grace period and fixation window might have altered the animal's behavior and would certainly have caused a larger proportion of trials to be aborted. This would have the effect of artificially screening the data in a way that favored winner-take-all pursuit.

Experiment 4

In *experiment 4*, the color of the rewarded target was randomized from trial to trial, and the monkeys were trained to use an instructional cue that indicated the color of the rewarded target on each trial. The cue was presented centrally by chang-

FIG. 7. Weight distribution means for all experiments. *A*: *experiment 1*: w_1 (filled symbols, thick solid line) and w_2 (open symbols, thick dashed line) plotted as a function of time interval for 4 monkeys (*monkey A*, circle; *monkey C*, square; *monkey D*, triangle; *monkey K*, diamond). Error bars are ± 1 SE but are generally smaller than plotting symbols. Thick lines are the mean weights averaged over monkeys. Small symbols are paired *t*-test probabilities that the mean weights for each monkey are different. *B*: weight vector directions calculated from w_1 and w_2 . *C*: weight vector amplitudes. In all plots, dashed lines are the WTA expectations, dotted lines are the VA expectation. *D–O*: results of *experiments 2–6* in same format as *A–C*.

ing the color of the white fixation target to red or green for 300 ms before the tracking targets appeared. This experiment comprised 128 trial conditions: all combinations of eight directions for two targets (64 conditions, as in the previous experiment) \times two instructions (track red or track green). The mean weights for this experiment are plotted in Fig. 7G. The small symbols show the outcomes (P values, 2-tailed) of paired t -tests on the mean weights for each monkey and time interval. Weight vector directions and amplitudes are plotted in Fig. 7, *H* and *I*. Two-way ANOVAs (independent factors: monkey and interval; dependent variables: weight vector direction and amplitude) showed highly significant effects of monkey and interval on weight vector direction and amplitude (all P values < 0.000). A three-way ANOVA was used to determine whether there were significant differences between *experiment 4* and *experiments 2/3* (grouped together) when monkey and time interval were factored in as independent variables. The results were mildly significant (direction: $P = 0.039$; amplitude $P = 0.011$). It seems reasonable to conclude that the results for blocked colors (*experiment 2/3*) are very similar to those for randomly interleaved colors (*experiment 4*).

Experiment 5

In *experiments 1–4*, either target could move in any of eight directions, and the monkey could not predict these directions before the trial started. In *experiment 5*, we reduced directional uncertainty by restricting the motion of the rewarded target to the horizontal axis. In this experiment, the rewarded target moved along the horizontal meridian, either to the left or right, while the nonrewarded target moved in any of the eight directions. As in the previous experiment, the color of the rewarded target was randomized from trial to trial.

The mean weight vectors for both monkeys are plotted in Fig. 7J, with the vector directions and amplitudes in Figs. 7, *K* and *L*. The data for both cue colors were combined. Note that for this experiment, a full set of basis vectors could not be computed because monkeys made only horizontal eye movements. We therefore used the basis vectors from *experiments 2* and *3*. Two-way ANOVAs (independent factors: monkey and interval; dependent variables: weight vector direction and amplitude) showed highly significant effects of monkey and interval on weight vector direction and amplitude (all P values < 0.01). Although the results of *experiments 2–5* appear qualitatively similar, three-way ANOVAs (independent factors; experiment, monkey, and interval) showed significant effects of experiment on both weight vector amplitude ($P = 0.000$) and direction ($P = 0.001$). Interestingly, there was no significant effect of monkey ($P = 0.592$ and 0.095 for amplitude and direction, respectively) even though the effect of monkey was significant within *experiments 4* and *5*. The effect of time interval was highly significant across experiments and monkeys ($P = 0.000$ for both direction and amplitude).

Experiment 6

In the five prior experiments, target motion was along a straight line path that went through the location of the initial fixation target, which was also the center of the display. We refer to this pattern of target movement as radial image motion. We suspected that radial image motion might favor a vector

averaging outcome. We therefore repeated *experiment 5* but changed the trajectory of the rewarded target so that instead of being exactly on the horizontal meridian, its path was 3° above or below the horizontal meridian (HM) and parallel to it. The sign of the vertical offset (above or below the HM) was randomized from trial to trial. The nonrewarded target started from a position that was symmetric about the HM with respect to the rewarded target. The initial horizontal position was 1.5° from the vertical meridian and both targets started from the same hemifield. We refer to this pattern of movement as parallel image motion, meaning that the path of the rewarded target was parallel to the HM. The nonrewarded target moved in one of eight directions as in the previous experiment, so the paths of the two targets were not necessarily parallel.

The results of *experiment 6* appeared qualitatively different from *experiments 2–5* in that the weight vectors for all monkeys more closely approximated the expected pattern for a WTA outcome. The mean weight vectors are plotted in Fig. 7M (note: we again used the basis vectors from *experiments 2* and *3*). The weight vector direction (Fig. 7N) and amplitude (7O) appear to show a stronger tendency toward the WTA expectation (dashed lines) than *experiments 2–5*. However, when the data for *experiments 2–5* were grouped together and tested against *experiment 6* in three-way ANOVAs, the effects of experiment were not significant ($P = 0.052$ for amplitude; $P = 0.154$ for direction). However, when three-way ANOVAs (independent factors: *experiment 2–5*, monkey, and interval) were run to look for variability among *experiments 2–5*, the effect of experiment was highly significant on the weight vector amplitude ($P = 0.000$) and direction ($P = 0.001$). This suggests that variability between *experiments 2* and *5* might have washed out the difference between those experiments and *experiment 6*. To look at this using a different statistical measure, we performed unpaired t -tests to compare directly the mean weight vectors in *experiments 2–5* (grouped) with those in *experiment 6*. For these tests, data were also averaged across monkeys, but not time intervals. We found a significant effect (2-tailed $P = 0.00000$) for the differences in direction and amplitude for each time interval except *interval 2* (25–50 ms), where only the difference in amplitudes was significant.

To look at the effect of radial versus nonradial image motion, we compared *experiments 5* and *6*. We again used an unpaired t -test to compare the mean weight vectors, and combined data across monkeys but not time intervals. We found that for each time interval, the weight vector direction and amplitude differed significantly (2-tailed $P = 0.0000$) between the two experiments.

Summary of experiments 1–6 and weight vector dynamics

To summarize the results for all six experiments, weight vectors are plotted in Fig. 9 for the latter four time intervals. Each data point represents the results from one experiment from one monkey. The earliest two time intervals were left out because the results were less reliable due to the fact that eye velocities were low. (The early intervals clearly showed a tendency toward vector summation that can be seen in Fig. 10.) The data in Fig. 9 show variability between experiments and monkeys. However, the variability is not random, but appears to be constrained along a line that includes both vector averaging and WTA outcomes. Linear regression analysis supports

the impression that there is a strong negative correlation between the two elements of each weight vector. The slope of the regression line becomes progressively shallower at later time intervals.

The same data can be plotted as a function of time-interval relative to pursuit onset to give a sense of the dynamics of the response. Such plots are shown in Fig. 10. Early time intervals actually favor vector summation. For *experiment 1*, the weight vectors shift over time from vector summation to averaging. For *experiments 2–5*, there is first a shift toward vector-averaging and then a later shift toward WTA. For *experiment 6*, the transition goes from vector summation to WTA without passing through vector averaging. For many experiments, the weight vectors follow curved trajectories. In the DISCUSSION, we will consider how these trajectories might be accounted for by a simple recurrent network model.

Smooth pursuit latency

We have previously reported that the latency of smooth pursuit depended on the relative direction between the target and distractor (Ferrera and Lisberger 1997a). The current set of experiments allowed us to test whether this effect depends on whether the outcome of the experiment was vector-averaging or winner-take-all. Figure 11 shows smooth pursuit latency as a function of relative direction for all six experiments sorted by monkey. Relative latency was calculated by subtracting the average latency for single target trials with the same target direction as the two-target trial. We considered only the absolute value of relative direction, hence the data are symmetric about zero. It should be noted that, in contrast to the other experiments, *experiment 6* included a condition where the

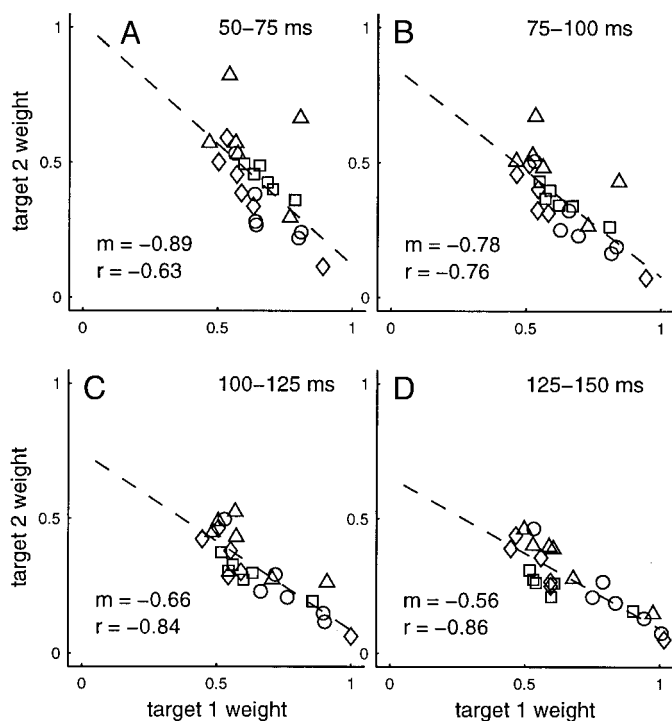


FIG. 9. Summary of results for *experiments 1–6*. Each subpanel is a different time interval (indicated at top of subpanel). Monkeys are represented by same symbols as in Fig. 7. m , linear regression slope; r , correlation coefficient.

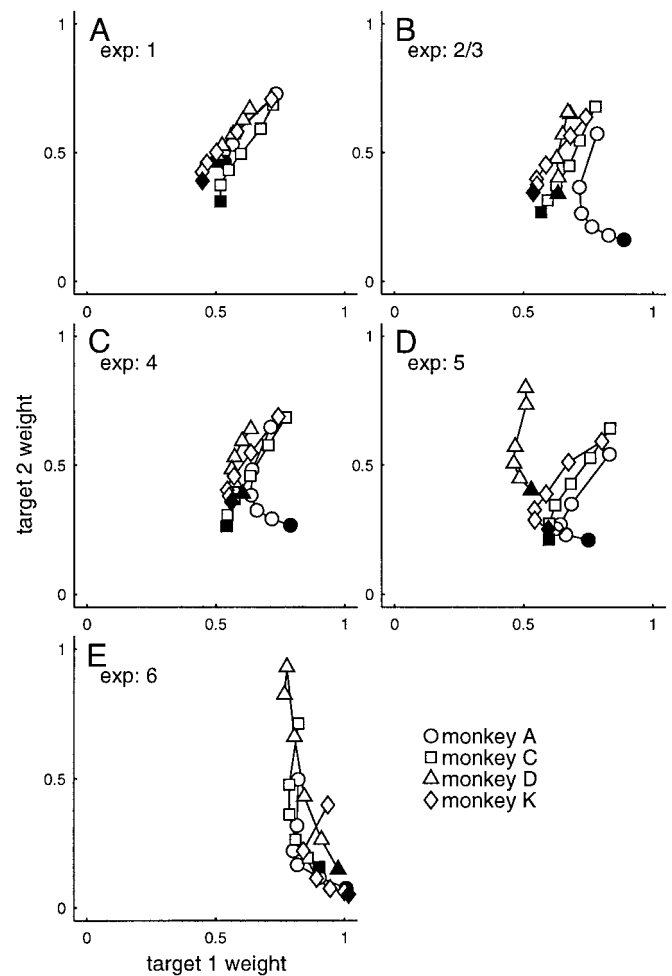


FIG. 10. Weight vector dynamics. Each subplot shows data for 1 experiment, as indicated. The order of the data points in each curve is the same as the order of the time intervals they represent. Filled symbols are the weight vectors for the final time interval (125–150 ms).

target and distractor moved in the same direction, along parallel trajectories. This accounts for the single data point at $x = 0$.

The data show effects similar to those reported previously; a slight decrease in latency when the target and distractor move in similar directions and an increase in latency when their directions of motion are $>90^\circ$ apart. At first glance the effect appears to vary little between experiments, particularly if one excludes the 12 obvious outliers (e.g., *monkey C, experiment 1*, angle = -180°) that account for only 6.1% of the data. To confirm this impression, we performed a three-way ANOVA (independent variables: monkey, experiment, and relative angle; dependent variable: latency) and found a significant variation between animals ($P = 0.027$) but not between experiments ($P = 0.169$). The effect of relative direction was highly significant ($P < 0.0001$).

Relationship between latency and accuracy

It is well established in the saccade system that short latencies favor averaging saccades while long latencies favor accurate (i.e., winner-take-all) saccades (Ottes et al. 1985). We found a similar trend in our smooth pursuit data. For example,

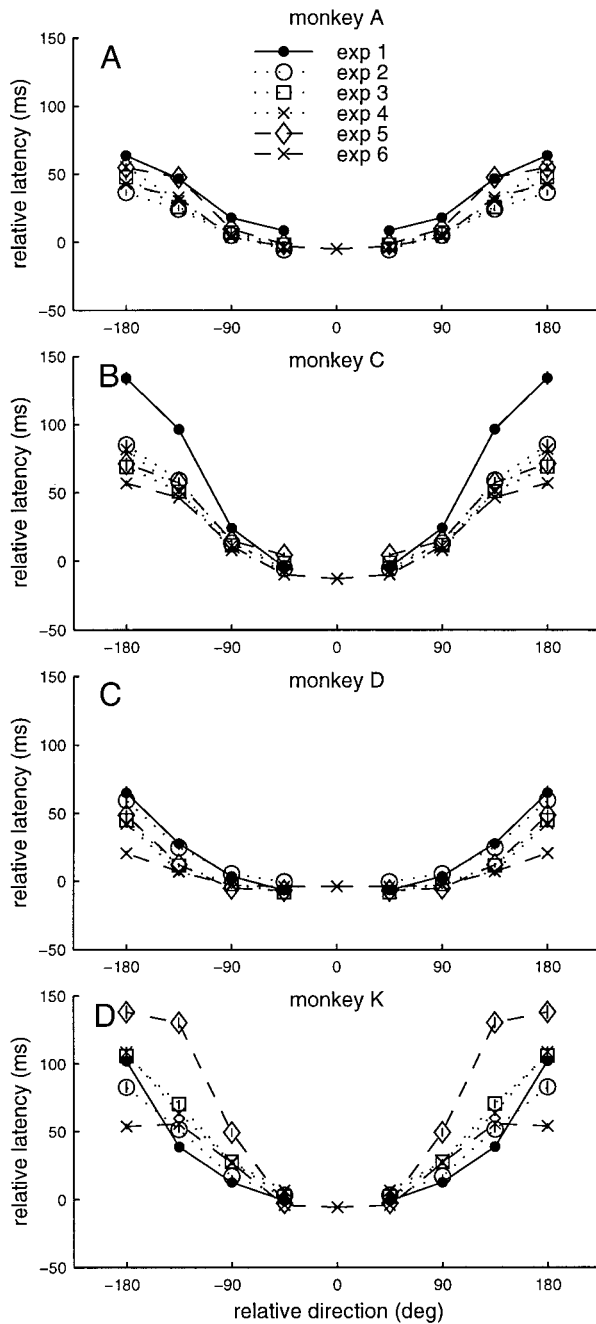


FIG. 11. Smooth pursuit latency as a function of relative direction of motion for target and distractor. Each subpanel shows results for a different monkey (indicated at top of subpanel). Error bars are ± 1 SE.

Fig. 12A shows target and distractor weights on individual trials from *monkey A, experiment 4, time interval 5* (125–150 ms). For short latencies, the target and distractor weight distributions overlap with both means close to 0.5. For longer latencies, the distributions diverge as the target weight approaches a mean near 1.0, and the distractor weight goes to zero.

To explore this effect quantitatively, we grouped the target and distractor weights into 10-ms latency bins and plotted the results for each experiment in Fig. 12, B–F. Data were averaged over monkeys and time intervals. Within each experiment, the mean target and distractor weights were compared

using a paired *t*-test. An asterisk indicates that the corresponding weights were significantly different ($P < 0.001$, 2-tailed). Across experiments, the point at which the weights diverged shifts to progressively shorter latencies. This shift parallels the shift from VA to WTA pursuit. One might question how it is possible for the weights to diverge for latencies shorter than 200 ms in *experiment 1*, given that there was no information about which target would be rewarded during the first 200 ms of image motion. The answer is that, because the data were

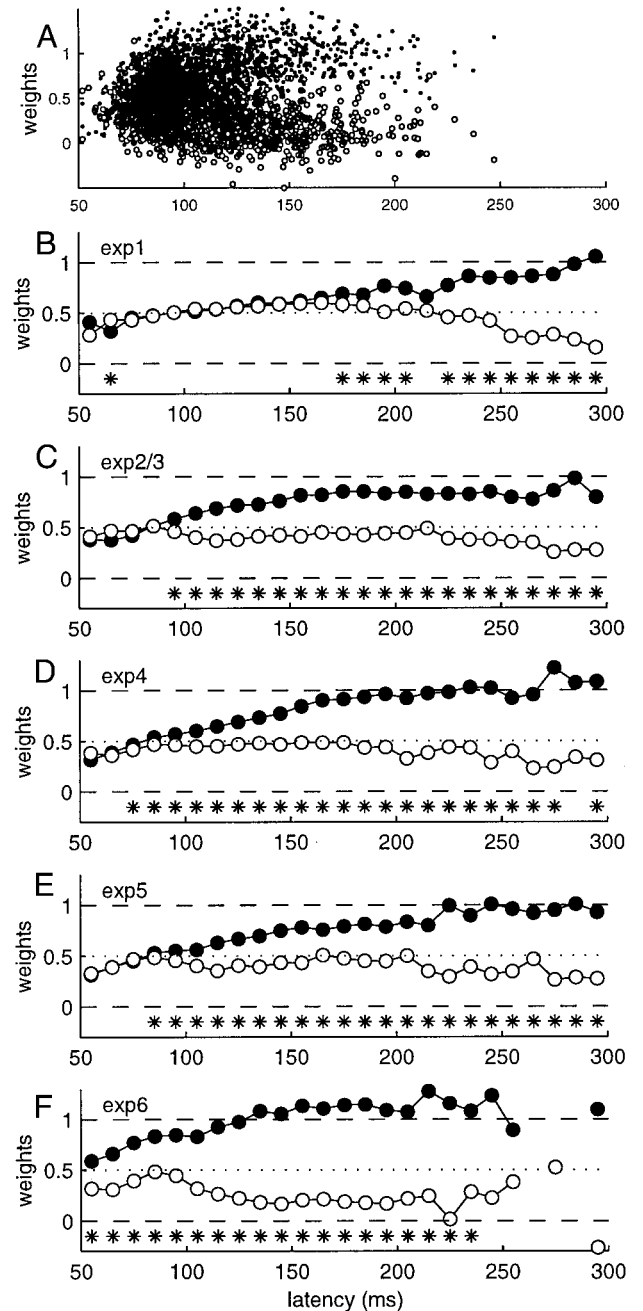


FIG. 12. Target and distractor weights as a function of smooth pursuit latency. A: target (solid circles) and distractor (open circles) weights for *monkey A, experiment 4, interval 5*. Each point is an individual trial. B–F: target and distractor weights averaged over monkeys and time intervals for *experiments 1–6*. Latency binwidth is 10 ms. Asterisks indicate significant differences (paired *t*-test, $P < 0.001$). Dashes and dotted lines represent the WTA and VA expectations, respectively.

averaged over different time intervals, each latency bin reflects information that was available to the monkey up to 150 ms later. For example, the bin labeled “150 ms” reflects data sampled up to 300 ms after the onset of target motion. A delay of 300 ms is long enough for changes in the visual display to affect pursuit, given feedback delays in the range of 80–120 ms. It is reasonable to conclude that the divergence seen in *experiment 1* reflects changes in pursuit subsequent to the disappearance of the distractor. Supporting this interpretation is the finding that one-way ANOVAs (independent factor: latency; dependent variable: weight difference, $w_1 - w_2$) run on each time interval for *experiment 1* showed significant effects ($P, 0.05$) of latency only in *intervals 4, 5, and 6* (i.e., those starting 75 ms or later after pursuit onset). A similar interpretation is not possible for the other experiments first because the distractor was always present, and second because significant divergence occurred for latencies much shorter than 150 ms. One-way ANOVAs (independent factor: latency; dependent variable: $w_1 - w_2$) for *experiments 2/3, 4, 5, and 6* showed highly significant effects of latency for each time interval (average P value = 0.0054, 24 conditions: 4 experiments \times 6 time intervals) except for *experiment 5, interval 1* ($P = 0.072$).

DISCUSSION

In agreement with an earlier report (Lisberger and Ferrera 1997), we found that in the presence of two identical targets, smooth eye velocity during the initial 150 ms of the eye movement was very close to a pure vector average of the eye velocity evoked by either target presented alone. The direction and speed of smooth pursuit were not generally consistent with vector summation, nor did the monkey arbitrarily choose one target or the other on any given trial. This result implies that, on average, smooth pursuit eye movements are driven by an equally weighted combination of the signals provided by the different neuronal pools responding to each target. In addition, the fact that the data show averaging rather than summation implies that there is normalization of the response to target motion somewhere along the sensorimotor pathway. It should be noted that there was a tendency toward vector summation during the first 50 ms of pursuit, which suggests that the onset of normalization is not simultaneous with the onset of pursuit. These observations are consistent with a model developed below which qualitatively accounts for the transitions from vector summation to vector averaging and from vector averaging to WTA.

In subsequent experiments, we asked what happens when the task is structured so that the monkey is able to 1) distinguish between the targets and 2) form an a priori notion of which would be the rewarded target. We investigated this by training the monkeys to select a target based on its color (Ferrera and Lisberger 1995, 1997a). We found one condition (*experiment 6*) where the initial eye movement showed WTA behavior for the direction of the correct target, which agrees with earlier results using a similar task (Ferrera and Lisberger 1997a). In four other experiments, the behavior was intermediate between vector-averaging and WTA.

The observation that monkeys did not perform winner-take-all pursuit in *experiments 2–5* suggests that the vector-averaging obtained in *experiment 1* is not simply an optimal cognitive

strategy that monkeys use when they do not know in advance which target will be rewarded. If monkeys always used an optimal strategy, then one would expect pursuit to be WTA in all cases where monkeys know in advance the identity of the correct target, because WTA is the optimal strategy in these cases. The observation that pursuit is often intermediate between vector averaging and WTA suggests that monkeys cannot easily switch from one strategy to the other, but rather that some effort is required to overcome an inherent tendency toward vector averaging. This supports the idea that vector averaging is a default computation that is performed automatically at a stage of visual motion processing that contributes to pursuit initiation.

One of the goals of these experiments was to determine how stimulus and task-related factors affect the selectivity of smooth pursuit. The results of *experiments 2–5* suggest that attention to color biases the direction of pursuit toward the attended target. It appeared to make little difference if the monkey tracked one color for an entire session or if the color of the rewarded target varied randomly from trial to trial (*experiment 2/3 vs. 4*). Uncertainty about the direction of target motion also appears to play only a small role (*experiments 2–4 vs. 5*).

The pattern of image motion appears to be an important factor; winner-take-all pursuit was found only when the rewarded target moved parallel to the horizontal meridian (*experiment 6*). Radial image motion favored vector-averaging, perhaps by evoking a larger contribution from cells in MST that respond to radial flow patterns (Graziano et al. 1994; Saito et al. 1986). It should be noted, however, that MST cells tend to prefer flow patterns with elements moving away from one another (expansion) rather than the contracting patterns formed by the stimuli used in the current experiments. Large field expanding optic flow patterns have also been shown to induce short-latency vergence eye movements (Busettini et al. 1997). The contribution of vergence to the eye movements measured here is unknown as all eye movements were measured monocularly. Busettini et al. recorded peak vergence velocities around 2°/s, an order of magnitude smaller than the pursuit velocities measured here, and other considerations make it unlikely that the small amount of vergence that may have been evoked by our impoverished stimuli would have contaminated our results significantly.

The differences between monkeys suggest that factors such as motivation and individual preference also affect the selectivity of pursuit. It is difficult to control for these factors without designing the experiment in a way that predetermines the outcome. Our main conclusions therefore are simply that the selectivity of smooth pursuit can vary from pure vector averaging to winner-take-all, that the selectivity is influenced by prior knowledge of the color of the rewarded target, which may be construed as feature-selective attention, and that there is an influence of the pattern of image motion (radial vs. parallel). One of the more surprising conclusions is that attention appears not to act as an all-or-none gate, but rather exhibits a graded, modulatory influence on the sensorimotor transformation. The modulation of pursuit selectivity is not uniformly distributed over the entire two-dimensional weight space, but appears to be constrained along a one-dimensional subspace (see Fig. 9). Intermediate outcomes along this dimension appear to be stable, as are the extremes of VA and WTA. This

suggest the possibility that the behavior of the pursuit system might be constrained by a “line-attractor” in weight-space, i.e., a one-dimensional locus of asymptotically stable fixed points. The concept of a line-attractor has been used to model eye position memory in the oculomotor “neural integrator” (Seung 1996).

The transition from vector-averaging to WTA pursuit depended not only on task and stimulus conditions, but also on the latency of pursuit, as shown in Fig. 12. When prior information was available (*experiments 2–6*), pursuit that was initiated after long latencies tended to be more WTA-like for all time intervals following pursuit initiation. This result suggests that target selection and pursuit initiation are governed by independent processes, each with its own time course. Hence, pursuit can be initiated at a relatively early stage of target selection, resulting in VA pursuit, or at a later stage, which results in WTA pursuit. Similar latency-accuracy trade-offs have been observed in the saccade and limb movement systems (Hening et al. 1988; Ottes et al. 1985).

These experiments raise the issue of how attention affects the way that motor outputs are coded based on the pattern of distributed activity in a neural map. The issue of distributed coding of movement has arisen previously with regard to the control of saccades by the superior colliculus (reviewed by McIlwain 1991). Vector averaging or summation and winner-take-all computations have been suggested as alternative means of decoding neuronal population activity (Georgopoulos et al. 1993; Groh et al. 1997). Our results indicate that smooth pursuit eye movements in monkeys exhibit a continuous range of behavior from pure vector averaging to winner-take-all for target direction. We were interested to determine whether this range of behavior could be reproduced in a simple recurrent neural network (Ferrera and Lisberger 1997a; Koch and Ullman 1985; Williams et al. 1986; Wilson and Cowan 1972; Yuille and Grzywacz 1989). We tested a network model similar to that used in our previous work (Ferrera and Lisberger 1997a), the only difference being the number of neurons and the distribution of recurrent inhibition across neurons. In the previous model, we used directionally tuned inhibition, whereas the current model used uniform inhibition. However, we have found that tuned inhibition was not needed to model our results and the untuned model is preferable in that it requires fewer parameters. The network comprised 16 “neurons,” each of which signaled a different direction of motion. The output of the network was coded as a population vector (Georgopoulos et al. 1993). Further details are provided in the APPENDIX.

We tested the behavior of the model by running simulations with either a single input (target) or two inputs (target plus distractor). The direction of the distractor varied across runs (8 directions, 22.5–180°). The target direction was always 0° (horizontal). For each set of runs, two-dimensional weights were computed by vector decomposition of the population vectors. Two parameters were varied over different sets of runs: the strength of the recurrent inhibition (w_i) and a selection bias that was added to the input of the neuron tuned to the target direction. The dynamics of the network response are illustrated in Fig. 13A, which shows the weight vectors computed from every 25th iteration. The squares correspond to a set of runs where $w_i = -0.2$ and bias = 0. Initially, the response is close to vector summation, but it quickly evolves to

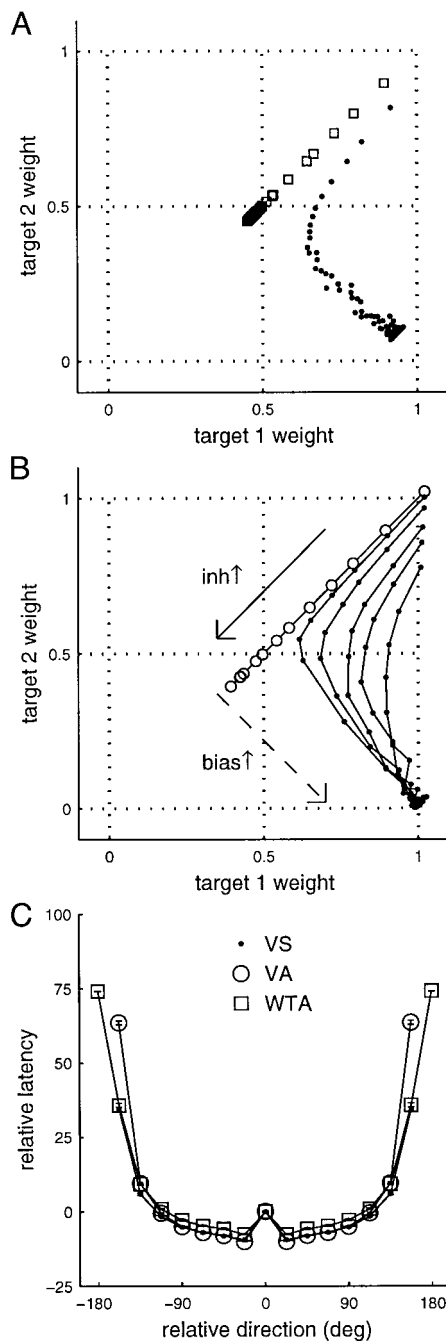


FIG. 13. Behavior of a competitive network model of target selection. *A*: population vector dynamics for 2 sets of runs with the same strength of inhibition ($w_i = -0.2$) but different levels of bias (open squares, bias = 0; closed circles, bias = 0.5, or 5% as strong as the sensory input). *B*: steady-state behavior. Open circles are runs with no bias, only increasing inhibition (range 0 to -0.2). Closed symbols are runs with nonzero, positive bias and the same range of inhibition. Lines connect sets of runs with equal bias. The large arrows labeled “inh” and “bias” indicate the general direction of the shift in the steady-state equilibrium caused by increasing recurrent inhibition and selection bias, respectively. *C*: response latency based on population vector amplitude. Error bars are ± 1 SE.

a pure vector average. This behavior is comparable to the weight vector dynamics shown in Fig. 10A. The filled circles correspond to a set of runs with $w_i = -0.2$ and bias = 0.5. In this case the network follows a curved trajectory that ends up near WTA. The curved trajectory shows some similarity to the

trajectories of the data in Fig. 10, *B–E*. With stronger biases, the trajectory shifts to the right and flattens out, i.e., it makes more of a beeline between VS and WTA, reminiscent of the data in Fig. 10*E*. It should be noted that the shape of the trajectory does not reflect the dynamics of either the selection bias or recurrent inhibition as these are fixed at the beginning of the simulation. The trajectory shape is simply a consequence of the internal dynamics of the network. Similarly, it would be unwise to conclude that a shift in the data from VA to WTA reflects the dynamics of attentional or decision processes that are accumulating information over time (see Leon and Shadlen 1998). Such shifts may simply reflect the response dynamics of cortical networks.

The steady-state behavior of the network is shown in Fig. 13*B*. Each point represents the weight vector calculated from one set of runs, and the selection bias and recurrent inhibition were varied over different sets of runs. When there was no selection bias, increasing inhibition ($w_i = 0.0$ to -0.3) resulted in a transition from vector summation to vector averaging (Fig. 13*B*, open circles). Increasing the selection bias (range: 0.0–2.0) resulted in a shift toward WTA. The maximum bias was 20% as strong as the direct sensory input (I). The lines in Fig. 13*B* connect sets of runs with the same inhibition level, but increasing bias. If the inhibition is strong enough for vector averaging with bias = 0, then varying the bias results in outcomes that are constrained along a line that runs from VA to WTA. This suggests that there is a direct relationship between the strength of attentional bias and the locus of the weight vector along the VA-WTA dimension, and therefore that a single parameter can account for the variability in the summary data shown in Fig. 9.

It was possible to assign a latency to the response of the network model by looking at the evolution of the amplitude of the population vector. Figure 13*C* shows the “relative latency,” or number of iterations before the population vector amplitude reached a fixed arbitrary threshold, as a function of the relative direction of target and distractor. Runs were grouped according to whether the steady-state weight vector was within a radius of 0.05 around the VS, VA, or WTA outcomes. (Each data point in Fig. 13*C* is the mean latency over several runs; the SEs are also plotted but are generally smaller than the plotting symbols.) For VS and VA outcomes, the population vector never reached criterion when the target and distractor moved in opposite directions, so there is no data for $x = -180, 180$. For all three outcomes, the latency depends on relative direction in a manner similar to that shown in Fig. 11. This result suggests that the directional tuning of the behavioral latency functions does not depend on directionally tuned recurrent inhibition. Indeed, there is directional tuning even when the outcome is VS, which only happens when there is negligibly weak inhibition. More significantly, there are only small differences in the latency effects for different classes of outcome, which is in general agreement with the data in Fig. 11. Our interpretation of this is that the latency effect reflects fixed properties of the underlying neural architecture, i.e., the pattern of input and recurrent connections, while the VA to WTA transition reflects the action of cognitive signals.

It is possible to speculate on how three features of the network discussed above, namely competitive inhibition, selection bias, and response normalization, might correspond to various stages of the neural pathway for smooth pursuit. There

is evidence that response normalization is performed by neurons in visual areas MT and MST, and this might underly the computation of a vector average (Groh et al. 1997; Qian and Andersen 1994; Recanzone et al. 1997; Simoncelli and Heeger 1998; Snowden et al. 1991). However, it is not clear that the spatial scale of interactions in MT/MST can account for interactions between widely spaced pursuit targets that have been observed (Ferrera and Lisberger 1997a). There is some evidence against the idea that competitive inhibition, as implemented in the network, is present in MT or MST (Ferrera and Lisberger 1997b). However, it has been found that electrical stimulation of the frontal eye field (FEF) with currents that are subthreshold for eliciting saccades is effective in inhibiting the production of voluntary saccades to locations outside the movement field of the stimulation site (Burman and Bruce 1997). This result suggests that competitive interactions that could mediate target selection are a feature of FEF circuitry, which may include the part of the FEF that is specialized for smooth pursuit (frontal pursuit area) (Gottlieb et al. 1994; MacAvoy et al. 1991). The selection bias used in the network might correspond to attentional modulation of sensory responses seen in visual areas MT and V4 (Ferrera and Lisberger 1997b; Moran and Desimone 1985; Motter 1993; Treue and Maunsell 1996). A bias of 2.0 (i.e., 20% of the sensory input) is commensurate with the strength of attentional effects in those areas. Alternatively, the attentional signals observed in extrastriate cortex may reflect top-down feedback from FEF (Schall et al. 1995b), which has been shown to play a role in color-based target selection for saccadic eye movements (Ferrera et al. 1999; Schall and Hanes 1993; Schall et al. 1995a).

In summary, this study supports the notion that target selection for smooth pursuit eye movements is mediated by an attentional bias amplified by a competitive network. One possible scenario is that the attentional bias is introduced in prefrontal and/or extrastriate cortex and that competitive interactions in prefrontal cortex then amplify the selection bias, transforming it into a motor command. The degree to which the motor command represents a vector average or winner-take-all response to the visual input depends on the strength of the selection bias. The selection bias and competitive inhibition, as well as response normalization, are three computational elements that should be considered in future models of the smooth pursuit system.

APPENDIX

This appendix provides details of a recurrent network model of smooth pursuit target selection consisting of an array of 16 direction selective units with Gaussian tuning ($\sigma = 15.0^\circ$). The activity of each unit (E) was a nonlinear function (g) of the input to that unit (I) plus the weighted activity of all other units

$$\frac{dE_i}{dt} = -E_i + g\left(k \sum_j W_{ij} E_j + I_i\right) \quad (A1)$$

A weak bias was added to the target signal so that $I_{\text{target}} = 10.0 + \text{bias}$, and $I_{\text{distractor}} = 10.0$. The weights were specified by a 16-by-16 matrix where each element $W_{ij} = w_e$ (the strength of the excitatory weight) if $i = j$, and w_i (the strength of the inhibitory weight) otherwise. This means that each unit excited itself and no others and uniformly inhibited and was inhibited by all other units. This model differs from those we have used previously (Ferrera and Lisberger

1997a) in that the inhibition was uniform rather than directionally tuned. In all simulated runs of the model, w_e was fixed at 0.0, while w_i varied between runs. The constant k determined the relative strength of recurrent versus feed-forward connections in the network. This was set to a value of 50.0. Smaller values of this constant did not substantially affect the qualitative behavior of the network, but mainly slowed the time course with which the network reached its steady-state output.

The sigmoidal activation function, g , used was the logistic function

$$g(x) = \{1 + \exp[-\eta(x - \delta)]\}^{-1} - [1 + \exp(-\eta\delta)]^{-1} \quad (\text{A2})$$

with $\eta = 0.5$, and $\delta = 9.0$.

The output of the network was coded as a population vector (\vec{P})

$$\vec{P} = \sum_i^n E_i \vec{V}_i \quad (\text{A3})$$

where \vec{V}_i is the unit vector oriented toward the preferred direction of each unit. It should be noted that, although the population vector was coded as a vector sum of the weighted unit activations, this does not imply that the network always performed vector summation. The population vector is simply a characterization of the pattern of activity during a single run. One must compare population vectors from different runs with single and paired targets to ascertain what computation the network is performing.

This network was simulated in Matlab v. 5.2 on a DEC workstation with an Alpha 21164 microprocessor (500 MHz) running Digital Unix 4.0C. The number of iterations per simulation was 3,000 to allow the network to reach steady state.

I am thankful to S. Lisberger in whose laboratory some of this work was performed and who provided helpful comments on the manuscript along with N. Qian and P. Glimcher, and A. Barborica. Excellent technical assistance was provided by J. Cohen, S. Tokiyama, A. Rocca, and J. Willi.

This research was supported by the McDonnell-Pew Program in Cognitive Neuroscience (JSMF 92-38), the Alfred P. Sloan Foundation (BR-3630), the Whitehall Foundation (S97-11), and the Klingenstein Foundation (CU512694).

REFERENCES

- BURMAN DD AND BRUCE CJ. Suppression of task-related saccades by electrical stimulation in the primate's frontal eye field. *J Neurophysiol* 77: 2252–2267, 1997.
- BUSETTINI C, MASSON GS, AND MILES FA. Radial optic flow induces vergence eye movements with ultra-short latencies. *Nature* 390: 512–515, 1997.
- FERRERA VP, COHEN JK, AND LEE BB. Activity of prefrontal neurons during location and color delayed matching tasks. *NeuroReport* 10: 1315–1322, 1999.
- FERRERA VP AND LISBERGER SG. Attention and target selection for smooth pursuit eye movements. *J Neurosci* 15: 7472–7484, 1995.
- FERRERA VP AND LISBERGER SG. The effect of a moving distractor on the initiation of smooth pursuit eye movements. *Vis Neurosci* 14: 323–337, 1997a.
- FERRERA VP AND LISBERGER SG. Neuronal responses in visual areas MT and MST during smooth pursuit target selection. *J Neurophysiol* 78: 1433–1446, 1997b.
- GEORGOPOULOS AP, TAIRA M, AND LUKASHIN A. Cognitive neurophysiology of the motor cortex. *Science* 260: 47–52, 1993.
- GLIMCHER PW AND SPARKS DL. Representation of averaging saccades in the superior colliculus of the monkey. *Exp Brain Res* 95: 429–435, 1993.
- GOLDREICH D, KRAUZLIS RJ, AND LISBERGER SG. Effect of changing feedback delay on spontaneous oscillations in smooth pursuit eye movements of monkeys. *J Neurophysiol* 67: 625–638, 1992.
- GOTTLIEB JP, MACAVOY MG, AND BRUCE CJ. 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, 1994.
- GRAZIANO MSA, ANDERSEN RA, AND SNOWDEN RJ. Tuning of MST neurons to spiral motions. *J Neurosci* 14: 54–67, 1994.
- GREEN DM AND SWETS JA. *Signal Detection Theory and Psychophysics*. New York: Wiley, 1966.
- GROH JM, BORN RT, AND NEWSOME WT. 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, 1997.
- HENING W, FAVILLA M, AND GHEZ C. Trajectory control in targeted force impulses. V. Gradual specification of response amplitude. *Exp Brain Res* 7: 116–128, 1988.
- JUDGE SJ, RICHMOND BJ, AND CHU FC. Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res* 20: 535–538, 1980.
- KOCH C AND ULLMAN S. Shifts in selective attention: towards the underlying neural circuitry. *Hum Neurobiol* 4: 219–227, 1985.
- LEE C, ROHRER WH, AND SPARKS DL. Population coding of saccadic eye movements by neurons in the superior colliculus. *Nature* 332: 357–360, 1988.
- LEON MI AND SHADLEN MN. Exploring the neurophysiology of decisions. *Neuron* 21: 669–672, 1998.
- LISBERGER SG AND FERRERA VP. Vector averaging for smooth pursuit eye movements initiated by two moving targets in monkeys. *J Neurosci* 17: 7490–7502, 1997.
- MACAVOY MG, GOTTLIEB JP, AND BRUCE CJ. Smooth pursuit eye movement representation in the primate frontal eye field. *Cereb Cortex* 1: 95–102, 1991.
- MCILWAIN JT. Distributed spatial coding in the superior colliculus: a review. *Vis Neurosci* 6: 3–13, 1991.
- MILES FA AND EIGHMY BB. Long-term adaptive changes in the primate vestibulo-ocular reflex. I. Behavioral observations. *J Neurophysiol* 43: 1406–1425, 1980.
- MORAN J AND DESIMONE R. Selective attention gates visual processing in the extrastriate cortex. *Science* 229: 782–784, 1985.
- MOTTER BC. Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *J Neurophysiol* 70: 909–919, 1993.
- OTTES FP, VAN GISBERGEN JA, AND EGGERMONT JJ. Latency dependence of colour-based target vs. nontarget discrimination by the saccade system. *Vision Res* 25: 849–862, 1985.
- QIAN N AND ANDERSEN RA. Transparent motion perception as detection of unbalanced motion signals. II. Physiology. *J Neurosci* 14: 7367–7380, 1994.
- RASHBASS C. The relationship between saccadic and smooth tracking eye movements. *J Physiol (Lond)* 159: 326–338, 1961.
- RECANZONE GH, WURTZ RH, AND SCHWARZ U. Responses of MT and MST neurons to one and two moving objects in the receptive field. *J Neurophysiol* 78: 2904–2915, 1997.
- ROBINSON DA AND FUCHS AF. Eye movements evoked by stimulation of frontal eye fields. *J Neurophysiol* 32: 637–649, 1969.
- SAITO HA, YUKIE M, TANAKA K, HIKOSAKA K, FUKADA Y, AND IWAI E. Integration of direction signals of image motion in the superior temporal sulcus of the macaque monkey. *J Neurosci* 6: 145–157, 1986.
- SCHALL JD AND HANES DP. Neural basis of saccade target selection in frontal eye field during visual search. *Nature* 366: 467–469, 1993.
- SCHALL JD, HANES DP, THOMPSON KG, AND KING DJ. Saccade target selection in frontal eye field of macaque. I. Visual and premovement activation. *J Neurosci* 15: 6905–6918, 1995a.
- SCHALL JD, MOREL A, KING DJ, AND BULLIER J. Topography of visual cortex connections with frontal eye field in macaque: convergence and segregation of processing streams. *J Neurosci* 15: 4464–4487, 1995b.
- SEUNG HS. How the brain keeps the eyes still. *Proc Natl Acad Sci USA* 93: 13339–13344, 1996.
- SIMONCELLI EP AND HEEGER DJ. Model of neuronal responses in visual area MT. *Vision Res* 38: 743–761, 1998.
- SNOWDEN RJ, TREUE S, ERICKSON RG, AND ANDERSEN RA. The response of area MT and V1 neurons to transparent motion. *J Neurosci* 11: 2768–2785, 1991.
- TREUE S AND MAUNSELL JHR. Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* 382: 539–541, 1996.
- WILLIAMS D, PHILLIPS G, AND SEKULER R. Hysteresis in the perception of motion direction as evidence for neural cooperativity. *Nature* 324: 253–255, 1986.
- WILSON HR AND COWAN JD. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys J* 12: 1–24, 1972.
- WURTZ RH. Visual receptive fields of striate cortex neurons in awake monkeys. *J Neurophysiol* 32: 727–742, 1969.
- YUILLE AL AND GRZYWACZ NM. A winner-take-all mechanism based on presynaptic inhibition feedback. *Neural Comp* 1: 334–347, 1989.