Neuron

Tracking the Mind's Eye: Primate Gaze Behavior during Virtual Visuomotor Navigation Reflects Belief Dynamics

Highlights

- Humans and monkeys virtually navigated to a memorized goal by integrating optic flow
- Eye movements tracked the goal location as subjects steered toward it
- Better goal tracking was associated with better navigation
 performance
- Forcing humans to fixate impaired navigation precision

Authors

Kaushik J. Lakshminarasimhan, Eric Avila, Erin Neyhart, Gregory C. DeAngelis, Xaq Pitkow, Dora E. Angelaki

Correspondence jkl9@nyu.edu

In Brief

Humans and animals constantly maintain and update beliefs about latent world states, but those belief dynamics are notoriously difficult to measure experimentally. Using a naturalistic task in which subjects navigated to a hidden goal in virtual reality, Lakshminarasimhan et al. demonstrate that this could be achieved by tracking eye movements.



Article

Tracking the Mind's Eye: Primate Gaze Behavior during Virtual Visuomotor Navigation Reflects Belief Dynamics

Kaushik J. Lakshminarasimhan,^{1,2,7,8,*} Eric Avila,^{1,7} Erin Neyhart,³ Gregory C. DeAngelis,⁴ Xaq Pitkow,^{3,5} and Dora E. Angelaki^{1,6}

¹Center for Neural Science, New York University, New York, NY, USA

²Center for Theoretical Neuroscience, Columbia University, New York, NY, USA

³Department of Neuroscience, Baylor College of Medicine, Houston, TX, USA

⁴Brain and Cognitive Sciences, University of Rochester, Rochester, NY, USA

⁵Department of Electrical and Computer Engineering, Rice University, Houston, TX, USA

⁶Tandon School of Engineering, New York University, New York, NY, USA

⁷These authors contributed equally

⁸Lead Contact

*Correspondence: jkl9@nyu.edu

https://doi.org/10.1016/j.neuron.2020.02.023

SUMMARY

To take the best actions, we often need to maintain and update beliefs about variables that cannot be directly observed. To understand the principles underlying such belief updates, we need tools to uncover subjects' belief dynamics from natural behavior. We tested whether eye movements could be used to infer subjects' beliefs about latent variables using a naturalistic navigation task. Humans and monkeys navigated to a remembered goal location in a virtual environment that provided optic flow but lacked explicit position cues. We observed eye movements that appeared to continuously track the goal location even when no visible target was present there. Accurate goal tracking was associated with improved task performance, and inhibiting eye movements in humans impaired navigation precision. These results suggest that gaze dynamics play a key role in action selection during challenging visuomotor behaviors and may possibly serve as a window into the subject's dynamically evolving internal beliefs.

INTRODUCTION

Rational behavior often requires predicting latent states from sensory observations. Since latent variables cannot be directly observed, and since the utility of actions depends on the status of latent variables in the future, we must use statistical regularities in space and in time to predict them. There is a large body of studies that not only demonstrate that humans exploit regularities in feature space (Knill and Pouget, 2004) but also show how to infer the associated subjective priors from data (Houlsby et al., 2013; Paninski, 2006; Smith et al., 2012; Stocker and Simoncelli,

2006; Turnham et al., 2011). In contrast, we know relatively little about how physical laws that govern the temporal dynamics of inputs are internalized and used to guide time-evolving beliefs in the absence of reliable observations (Lee et al., 2014). We refer to these subjective beliefs about time-varying latent states as belief dynamics.

The reasons for limited progress in understanding belief dynamics are twofold. First, psychophysics continues to be dominated by experimental paradigms in which actions are discrete (e.g., binary choice) and sporadic (e.g., at the end of the trial). In contrast, continuous tasks (Bonnen et al., 2015; Huk et al., 2018; Knöll et al., 2018; Pitkow and Angelaki, 2017) provide subjects the opportunity to reveal more information about their beliefs and predictions as they unfold in time. Second, although theoretical techniques to infer latent beliefs from actions are slowly becoming available (Kumar et al., 2019; Reddy et al., 2018; Wu et al., 2019), they have yet to be successfully applied to settings in which state and action spaces are both continuous. Consequently, principled ways to reliably uncover subjects' belief dynamics from natural behavior are still lacking. Meanwhile, a practical way to overcome this hurdle would be by covertly "measuring" those beliefs. One candidate tool to accomplish this is eye tracking (Spivey, 2007). Saccadic eye movements have previously been used to understand mental processes underlying a wide variety of abstract tasks such as language comprehension (Tanenhaus et al., 1995), reading (Rayner, 1998), mental imagery (Spivey and Geng, 2001), evidence accumulation (Gold and Shadlen, 2000), visual search (Zhang et al., 2018), and even random number generation (Loetscher et al., 2010). Furthermore, it has recently been argued that smooth-pursuit eye movements may be influenced by short-term memory (Deravet et al., 2018; Orban de Xivry et al., 2013). By formulating oculomotor pursuit to transiently occluded moving targets as an active inference process, these eye movements have been used to infer subjects' internal beliefs (Adams et al., 2015). We wanted to know whether eye movements also reflect belief dynamics for extended periods of time under more naturalistic conditions.

To test this, we created a virtual environment with an unstructured ground plane on which subjects steered to a transiently cued target location. To successfully perform the task, subjects had to first infer their own movements based on the sparse optic flow cues generated while steering and then integrate them over time to estimate the relative target location. Although the target appears briefly at the beginning of the trial, the location of the target relative to the subject becomes latent as soon as they start steering, because thereafter, the relative location is not directly observed, only inferred by discounting one's own displacements. This task differs from traditional paradigms used to study latent state inference in two important ways that make it both challenging and better suited to understanding belief dynamics in the real world. First, in contrast to tasks in which latent states remain unchanged throughout the trial, such as right or left in a random-dot motion (Britten et al., 1992) or heading-discrimination task (Britten, 2008), here, the latent state dynamically varies over the course of each trial, under the subject's control. Second, unlike tasks that use pulsatile evidence and discrete numbers of latent states, such as an auditory-clicks (Brunton et al., 2013) or accumulating-towers task (Pinto et al., 2018), both the sensory input (self-motion) and latent states (relative target position) are continuous valued. To test whether eye movements were informative regarding subjective beliefs about those time-varying, continuous-valued latent states, we recorded the gaze behavior of humans and rhesus macaques while they performed this task. Parallel experiments in the two species allowed us to test whether the eye movements were evolutionarily conserved. We found that both humans and monkeys tend to follow the location of the unseen target with their gaze until they reach it, and their success in tracking the target over time predicted their final behavioral accuracy. These findings suggest that gaze dynamics reflect internal beliefs and could help shed light on the computations that transform visual perception to action in naturalistic settings.

RESULTS

Monkeys and humans performed a visual navigation task in which they used a joystick to steer to a transiently cued target location in a three-dimensional virtual reality (VR) environment without allocentric reference cues (i.e., stable landmarks) (Figure 1A; STAR Methods). Individual visual elements comprising the ground plane were transient and could not be used as landmarks. At the beginning of each trial, a circular target blinked briefly at a random location within the field of view on the ground plane and then disappeared. The joystick controlled forward and angular velocities, allowing subjects to steer freely in two dimensions (Figure 1B). The subjects' goal was to steer toward the target and stop when they believed their position fell within a circular reward zone centered on the target. They received feedback about their performance at the end of each trial.

Monkeys were first trained extensively using a staircase procedure (see STAR Methods) until their performance stopped improving. Here, we will focus only on their post-training behavior. At this point, the radius of the reward zone was fixed across trials (see STAR Methods) and they received juice reward at the end of the trial for correctly stopping within this zone (Figure 1C). In contrast, human subjects received no prior training on this task. Instead, we used an adaptive feedback scheme in which the radius of the reward zone was dynamically scaled using a staircase procedure to match individual subjects' abilities (Figure S1A; see STAR Methods). In practice, it took less than 50 trials for the performance of humans to stabilize (Figure S1B). Therefore, we ignored the first 50 trials collected from human subjects and focused our analyses on the remaining data.

Target locations were uniformly distributed at random over the ground plane area within the subject's field of view (Figure 1D, left). The stimulus was nearly identical for both species except for minor details such as the range of target distances and the duration for which the target was visible (see STAR Methods). All subjects were head-fixed, and we recorded each subject's movement trajectory (Figure 1D, middle) as well as eye position (Figure 1D, right) throughout each trial.

Behavioral Performance

Figure 1E shows the performance of the monkeys in this task. Both radial distance (Figure 1E, left) and angular eccentricity (Figure 1E, right) of the monkeys' responses (stopping location) were highly correlated with the target location across trials (n = 3 monkeys, Pearson's $r \pm$ standard deviation, radial distance: 0.72 ± 0.1 , angle: 0.84 ± 0.1), suggesting that their behavior was appropriate for the task. To test whether their performance was accurate, we regressed their responses against target locations. The slope of the regression was close to unity both for radial distance (mean \pm standard deviation = 0.92 ± 0.06) and angle (0.98 ± 0.1), suggesting that the monkeys were nearly unbiased (Figure 1F, green). We did notice modest undershooting for distant targets, an effect that is likely due to growing position uncertainty described in previous work (Lakshminarasimhan et al., 2018a).

We showed previously that humans are systematically biased when performing this task without feedback (Lakshminarasimhan et al., 2018a). Consistent with those findings, human subjects overshot the target in an initial block of trials in which no feedback was provided (Figure S1C; n = 5, mean slope \pm standard deviation, radial distance: 1.21 ± 0.2 , angle: 1.78 ± 0.3) to a degree that was proportional to target distance. With feedback, however, the same subjects quickly adapted their responses to produce nearly unbiased performance (Figure 1F, purple; see Figure S1D for individual trials; mean slope \pm standard deviation, radial distance: 0.95 \pm 0.1, angle: 1.15 ± 0.2). Notably, this improvement in performance was maintained in a final block of trials in which feedback was withheld (Figures S1E and S1F; radial distance: 1.03 ± 0.15 , angle: 1.2 ± 0.2), suggesting that learning of this task was stable. To be consistent with monkey data, we only consider human subjects' data collected during the block of trials with feedback in the remainder of this work.

We wanted to know whether humans and monkeys had comparable accuracies. Because we used a slightly larger range of target distances for humans (see STAR Methods), travel durations were longer (median travel time \pm interquartile range [IQR]: monkeys: 1.9 ± 0.8 s, humans: 2.7 ± 0.6 s). Consequently, we could not directly compare the mean error magnitude of the subjects as it ignores differences in task difficulty. Instead,

Cell²ress



Figure 1. Primates Can Navigate by Integrating Optic Flow

(A) Monkeys and human subjects use a joystick to navigate to a cued target (yellow disc) using optic flow cues generated by ground plane elements (brown triangles).

(B) The time course of linear (top) and angular (bottom) velocities during one example trial. Yellow shaded region corresponds to the time period when the target was visible on the screen. Time is also coded by color.

(C) Example trials showing incorrect (left) and correct (right) responses of a monkey.

(D) Left: overhead view of the spatial distribution of target positions across trials. Middle: movement trajectories of one monkey during a representative subset of trials. Blue dot denotes starting location. Right: first-person view of the trajectories of eye movements during the same trials. Abscissa and ordinate show horizontal version and elevation of the eyes. Blue dots represent the initial eye position (when the target was turned OFF) on each trial.

(E) Left: comparison of the radial distance of the monkey's response (stopping location) against radial distance of the target across trials. Right: angular eccentricity of the response versus target angle. Black dashed lines have unity slope. The subject's starting location was taken as the origin.

(F) Subjects' accuracy in radial distance (top) and angular eccentricity (bottom) were quantified as the slopes of the corresponding linear regressions and plotted for individual monkeys and human subjects. Horizontal dashed lines denote the value of the slope that corresponds to unbiased behavior. Error bars denote ±1 SEM across trials.

(G) Left: the proportion of correct trials of one monkey for various values of hypothetical reward window size (black). Shuffled estimates are shown in gray. Right: ROC curves for all subjects, obtained by plotting their true proportion of correct trials (from unshuffled data) against the corresponding chance-level proportions (from shuffled data) for a range of reward windows. Shaded area denotes standard deviation across subjects. Inset shows the average area under the ROC curve (AUC) for monkeys and human subjects.

See also Figure S1.

we used an approach that is conceptually similar to receiver operating characteristic (ROC) analysis to objectively compare the performance of monkeys and human subjects on a common scale. For each subject, we constructed a "psychometric function" by computing reward probability as a function of hypothetical reward window size (Figure 1G; see STAR Methods). By plotting the true psychometric function against one obtained by shuffling target locations across trials, we obtain the subject's ROC curve. Chance-level performance would correspond to an area under the ROC curve (AUC) of 0.5, while perfectly accurate responses (zero error) will yield an AUC of one. The AUCs for both monkey and human subjects were quite large and

statistically indistinguishable (mean \pm standard deviation, monkeys: 0.85 \pm 0.03, humans: 0.84 \pm 0.05; t test: p = 0.41), suggesting that they performed comparably. Although it is possible, in principle, to avoid integrating optic flow by learning the precise sensorimotor transformation implemented by the joystick controller, we previously showed that the variability of human subjects is greatly affected by removing optic flow cues (Lakshminarasimhan et al., 2018a) (Figure S1G). Likewise, monkeys rapidly adapt their actions in response to gain changes of the joystick controller (Figure S1H). This suggests that both monkeys and humans use optic flow to perform this task.

Pattern of Eye Movements

To understand the role of eye movements, we recorded the position of the subjects' eyes while they performed the task. Figure 2A shows the vertical and horizontal eye positions of one monkey during an example trial. On this trial, we noticed saccades (eye movements exceeding 200°/s) before the target was turned off (henceforth called start of the trial) and around the time when the monkey stopped moving (end of steering), but not in between. This pattern was evident across trials, as seen in the trial-averaged density of saccades (Figure 2B). Across all datasets from monkeys, the average frequency of saccades during the trial was significantly smaller than that during the inter-trial interval (mean saccade rate ± standard deviation, during trials: 0.5 ± 0.3 Hz, between trials: 0.9 ± 0.5 Hz; paired t test: p = 0.02). We noticed a similar tendency among human subjects, although the comparison was not statistically significant (Figure S2A; during trials: 0.8 ± 0.5 Hz, between trials: 1.4 ± 1 Hz; p = 0.11). Moreover, the velocity of eye movements during steering was generally low, with magnitudes well below 20°/s both in monkeys (Figure 2C; mean ± standard deviation: 16.2 ± 2.1°/s) and in humans (Figure S2B; 11.4 \pm 3.2°/s).

Because saccades were mostly confined to periods when the animal was not actively steering and subjects appeared to make slowly varying eye movements while steering, we asked whether they may be continuously "tracking" the (invisible) target with their eyes while they navigated to it. Note that as one steers toward the target location, the target becomes progressively less eccentric and moves downward in the visual field. Therefore, if subjects' eyes were to track the target, the magnitude of lateral version would tend toward zero and the eye elevation would become more negative with time (Figure S2C). To quantitatively test whether subjects tracked the target, we first generated ground truth theoretical predictions for the binocular position of their eyes during each trial, assuming that they maintained fixation at the center of the target throughout the trial (Figure S2D; STAR Methods, Equation 1). Note that at each moment in time, the predicted eye position depends only the relative target position at that moment regardless of whether the subject accurately stopped on target, but we will examine the relationship to the latter in the next section. We then compared this prediction against the observed eye position of the subject by expressing both quantities in terms of three standard components: lateral version, elevation, and vergence (Figure S2E; see STAR Methods).

We expect subjects' eyes to be drawn to the target when it appears on the screen. Indeed, the model predictions were highly correlated with the measured values of lateral version (Figure 2D, left; and Figure S3A, left; Pearson's r ± standard deviation, monkeys: 0.91 ± 0.1 , humans: 0.85 ± 0.1) as well as elevation (Figure 2D, right; and Figure S3A, right; monkeys: 0.60 ± 0.2 , humans: 0.42 ± 0.2) at the beginning of the trial. The somewhat lower correlations for the latter are understandable, because it is difficult to precisely fixate at the elevations for distant targets since they subtend a smaller visual angle. Next, we examined the time course of eye movements during the trial and found a striking qualitative correspondence to the predicted dynamics (Figure 2E and S3B); as the trial progressed, lateral version became increasingly more concentrated around zero (Figure S3C, left) while eye elevation was significantly lower (Figure S3C, right). The correlation between predicted and observed values remained significantly greater than zero throughout the trial for both components (Figure S3D). This is guite remarkable, because the target appeared only transiently at the beginning of the trial.

On the other hand, the correspondence between predicted and observed vergence was less clear. Performing this comparison for our task was challenging, because ~90% of the full range of vergence angles is known to occur within gaze distances <1 m (Howard, 2012), and the predicted change in vergence is negligible for gaze distances >2 m (Figure S2E, bottom right). Only two of the monkeys exhibited vergence values that weakly correlated with the predictions at trial onset (Figure S3A) and a tendency to make convergent eye movements as they approached the target (Figure S3B), an effect that was also absent in human subjects (Figures S3B-S3D). It is possible that this inconsistency is due to the previously documented difficulty in executing voluntary vergence movements to imagined moving targets (Erkelens et al., 1989). This difficulty is likely exacerbated in VR, where vergence eye movements must be executed without changing accommodation to maintain a clear retinal image of onscreen objects (Hoffman et al., 2008; Shibata et al., 2011). Therefore, we did not consider the vergence component for further analyses.

To quantify how well subject's eyes tracked the target, we expressed the eye position as a two-dimensional vector comprised of lateral version and elevation and computed a target-tracking index. Specifically, this quantity was given by the square root of the fraction of variance in the observed eye position that was explained by the prediction (STAR Methods, Equation 2). An index of one implies that the subject consistently looked at the center of the (invisible) target while steering toward it, while zero denotes lack of correspondence between target and gaze locations. The target-tracking index was quite high at trial onset (during the first 500 ms) when the target had just disappeared (Figure 2F; mean \pm standard deviation, monkeys: 0.73 ± 0.05 , humans: 0.71 ± 0.05). Although this slowly dropped during the trial, the index at the end of the trial (during the last 500 ms) remained well above zero (Figure 2G; mean ± standard deviation, monkeys: 0.35 ± 0.1 , humans: 0.18 ± 0.05), implying that subjects tend to maintain gaze at the target location while they steer toward it. To estimate this timescale of the correlation between gaze and target location, we analyzed the cross-correlogram between them and found that subjects' eye positions did not systematically lead or lag the predictions based on the relative

CellPress



Figure 2. Eye Movement Dynamics during the Task

(A) Time course of vertical and horizontal (bottom) positions of the left and right eyes of a monkey during one example trial. Yellow region shows the period when a target was visible on the screen. Red dashed line corresponds to the end of steering in this trial.

(B) The time course of the rate of saccades during the trial, averaged across all trials separately for each monkey. Trial averaging was done by aligning trials relative to target onset (yellow region, before the break on the x axis) and end of steering (red dashed line, following the break). Gray line denotes mean saccade rate across monkeys during the period between trials.

(C) Joint probability distribution over horizontal and vertical eye velocities, averaged across monkeys, while they steered toward the target. Marginals are shown in black.

(D) Comparison of the predicted and true eye positions in a subset of trials for all monkeys at the moment when the target was just turned OFF.

(E) Time-course of the eye position during a random subset of trials taken from one monkey. Blue and red dots denote the times at which the target was turned OFF and the end of steering, respectively.

(F) Target-tracking index when the target turned OFF for individual monkeys and humans. Error bars denote ±1 SEM obtained either by averaging across recording sessions (for monkeys) or bootstrapping (for humans).

(G) Time course of the target-tracking index, averaged across monkeys and humans. Gray arrow denotes the chance-level tracking index verified by shuffling procedure. Shaded region denotes ±1 SEM across datasets.

See also Figures S2–S4A.

target location (Figure S4A). This suggests that eye movements reflect the current relative target position rather than predict its future value (although the computations used to estimate it could still be predictive).

The tracking index quantifies how subjects' dynamical state (relative target position) is encoded in their continuous-valued eye position while they navigate toward the target. However, recent work has highlighted the importance of discrete saccadic

CellPress



Figure 3. Saccadic Eye Movements Aid Target Tracking

(A) Time course of observed (black) and predicted (gray) vertical position of the eyes of a monkey. Black arrows indicate saccades made during three different task epochs (inter-trial, target presentation, and steering periods). Yellow region shows the period when a target was visible on the screen. (B) Empirical cumulative distribution function of saccade amplitudes conditioned on the task epoch, averaged across monkeys. Inset shows amplitudes of individual saccades as a function of their timing. (C) Average saccade-triggered target tracking error during a time window around saccades made during steering. Shaded region denotes ±1 SEM obtained by bootstrapping.

(D) The time course of coefficients obtained by linearly regressing the amplitudes of the two components of saccades (blue, vertical; red, horizontal) against the corresponding components of the target tracking error (STAR Methods).

(E) Similar to (D), but showing coefficients for regression done separately for the first, second, and third saccades made during steering. See also Figures S4B and S4C.

eye movements in mediating flow-tracking behavior of humans and primates (Knöll et al., 2018). Therefore, we wanted to know whether saccades aided the target-tracking behavior of our subjects. To test this, we first compared the distribution of saccade amplitudes during three non-overlapping epochs of the experiment: inter-trial periods when saccades tend to be exploratory, target-presentation phase when saccades are auided by the external stimulus, and the ensuing task phase when subjects steered using optic flow (Figure 3A). We found that across monkeys, the amplitude of saccades was much lower during the task phase than during other epochs (Figure 3B; mean \pm SE: inter-trial, 10.2 \pm 1.6°; target-presentation, 14.4 \pm 2.2°; task phase, 7.1 \pm 1.2°), suggesting that saccades made while steering were qualitatively different from saccades made at other times. To directly test whether those saccades served to correct errors in target tracking, we computed a saccade-triggered average of the target-tracking error and found that this error dropped significantly (peak decrease of $2.4 \pm 0.4^{\circ}$; p < 10^{-10} , t test) shortly after saccade onset (Figure 3C). Following Knöll et al. (2018), we used lagged regression analysis to determine the precise relationship between saccade amplitude and the dynamics of target-tracking error (STAR Methods). The amplitude of both vertical and horizontal components of the saccade were influenced by tracking error during the previous 200 ms, suggesting that these saccades were indeed made toward the target (Figure 3D). Moreover, the regression kernels were biphasic, implying that the saccades overcompensated for the tracking errors. Finally, if these saccades were corrective, they should depend on the subjects' internal estimate of the target location, making them increasingly unreliable over time due to the buildup of uncertainty. Indeed, the strength of the regression kernel was weaker for later saccades (Figure 3E; peak-to-peak difference in weights for vertical component: first saccade, 0.44 ± 0.1 ; third, 0.20 ± 0.15 ; horizontal component: first saccade, 1.5 ± 0.4 ; third, 0.15 ± 0.2), thereby signaling a drop in

saccadic precision over time. This suggests that these saccades were not stereotyped reflexive responses but were dynamically modulated by ongoing cognitive computations analogous to "catch-up" saccades observed during smooth pursuit of visible targets (Daye et al., 2014; Orban de Xivry et al., 2008). Although the amplitudes of saccades made by human subjects were not significantly smaller while steering (Figure S4B), regression analysis revealed a strong association between tracking error and saccade amplitude but with slightly shorter integration windows. As observed for monkeys, the strength of this association was lower for saccades that happened later (Figure S4C), reflecting a potential influence of noisy integration.

Eyes Convey Internal Beliefs about Target

Subjects could not have possibly been tracking the observed target location, because the target disappeared at the beginning of the trial. A plausible explanation for their pattern of eye movements is that they tracked the location at which they believed the target was present. As they integrate their movements, subjects need to continuously update their internal estimate of the relative goal location, and perhaps their eye movements reveal those estimates. If this is the case, then we should be able to better predict their eye position when their beliefs are more accurate. We tested this both across subjects and across trials within each subject.

To test this across subjects, we used the variability in subjects' stopping positions to first quantify the level of uncertainty in their position estimates (STAR Methods). Due to the low trial count of individual human subjects, we pooled trials from all humans into a single dataset. Because uncertainty in knowing one's location should limit one's ability to visually track the target, we used the estimated uncertainties to calculate an approximate upper bound on the target-tracking index for each dataset (Figure 4A; STAR Methods, Equation 3). This upper bound serves to capture the heterogeneity in the spatial profile of uncertainty both across subjects (Figure 4B, left) and across sessions within each



Figure 4. Accurate Target Tracking Is Associated with Increased Task Performance

(A) Time course of the target-tracking index for one session computed using a monkey's actual eye movements (black solid) and its theoretical upper bound (black dashed) determined using variability in stopping positions (STAR Methods, Equations 3 and 4).

(B) Left: overhead view of the spatial maps showing the standard deviation of stopping positions as a function of target location for individual monkeys and the average human subject. Each wedge corresponds to the map of one subject, calculated by binning target locations (see Figure 1D) and smoothed using a Gaussian filter. The maps of monkeys S and Q, and of the humans, have been rotated for compactness. Right: Comparison of the observed target-tracking index against the theoretical upper bound (averaged over the last 500 ms of the trials) across all individual datasets. Dashed line has unity slope and error bars denote ± 1 SEM obtained by bootstrapping.

(C) Top: time course of the target-tracking index for one example monkey shown separately for trials in which he stopped within the reward zone (blue) or stopped outside it (red). Shaded regions denote ± 1 standard error estimated by bootstrapping. Bottom: the difference between tracking coefficients the two sets of trials for all subjects. For human subjects, trials in which the subject's final position was within 0.6 m of the center of the target were considered "rewarded." (D) Top: we divided trials into five groups based on the magnitude of behavioral error. Time courses of the



monkey (Figure S4D). Across all datasets (subjects × sessions), the target-tracking index observed toward the end of the trial (during the last 500 ms) was weakly but significantly correlated with the theoretical upper bounds (Figure 4B, right; Pearson's r = 0.26, p = 0.029). This suggests that differences in the ability to track the target with the eyes is due, at least in part, to differences in the magnitudes of positional uncertainty between subjects.

We also tested whether eye movements reflect fluctuations in the subject's belief about their location across trials. Because subjects were more precise during rewarded (Figure S4E, left) than during unrewarded trials (Figure S4E, middle), we expect them to track the target more accurately during rewarded trials (Figure S4E, right). We computed the target-tracking index separately for the two groups of trials and found that it was indeed higher during rewarded trials (Figure 4C, top). The difference between the target-tracking indices during the two sets of trials grew as the trial progressed and was significantly greater than zero at the end of the trial (Figure 4C, bottom; mean difference standard deviation during the period shaded in gray: ± monkeys: 0.19 ± 0.05 , $p = 4.8 \times 10^{-3}$; humans: 0.13 ± 0.05 , p =3.1×10⁻²; bootstrap test, 10,000 bootstrap samples). In fact, when trials were stratified based on behavioral accuracy, we found that the tracking index increased with behavioral accuracy (Figure 4D). To more directly test for a fine-grained relationship between eye movements and task performance, we estimated the correlation between the behavioral error (distance between the stopping location and the target) and the target-tracking error (mean absolute difference between the actual eye position and the theoretical prediction, see STAR Methods) across trials (Figure 4E, top). To control for possible spurious effects of trial difficulty, we computed a shuffled estimate by subdividing the trials into groups based on initial target distance and then shuffling the trials within each group (see STAR Methods). We found that the behavioral and target-tracking errors were significantly correlated across trials (Figure 4E, bottom; Pearson's $r \pm$ standard deviation across all datasets: true: 0.14 ± 0.04; controlled shuffle: 0.04 ± 0.02 ; $p = 9.1 \times 10^{-3}$, paired t test), further reinforcing the view that subjects track their internally estimated goal location with their eyes.

Purely Reflexive Eye Movements Do Not Explain Target-Tracking Behavior

In principle, the above results could also be produced by purely reflexive eye movements, driven solely by optic flow (ocular following response [OFR]). For instance, if subjects' eye velocity



is perfectly correlated with their perceived movement velocity, then oculomotor errors would be proportional to perceptual errors, explaining the relatively poor target-tracking in erroneous trials. However, past studies have shown that errors in reflexive eye movements are uncorrelated with perceptual errors (Blum and Price, 2014; Boström and Warzecha, 2010; Glasser and Tadin, 2014; Price and Blum, 2014), suggesting that the observed eye movements are not reflexive. Two further pieces of evidence in our own monkey data support this.

First, in a subset of sessions, we recorded the stimulus movie of the complete block of trials and replayed them back to the animal at the end of the session, but with the joystick withheld (see STAR Methods). All aspects of the task structure during this replay block were identical to the initial block of trials (e.g., the monkey still received juice reward at the end of the corresponding trials), except the animal only viewed a movie of the stimulus rather than actively performing the task. Importantly, monkeys were still free to move their eyes. Eye movements were weaker during passive viewing than during active task (Figures S5A and S5B), and the magnitude of eye velocity was much smaller during passive block even though both blocks had identical visual stimuli (Figure S5C). We analyzed the target tracking behavior by computing the target-tracking index separately for the two blocks of trials. Figure 5A (top panel) shows the time course of the target-tracking index of one monkey during the both blocks of trials. In this monkey, the tracking index was much lower during passive viewing (red versus blue). Because OFR is, by definition, involuntary and difficult to suppress, this suggests that eye movements contributing to the high targettracking index during active steering must be voluntary. Note,

Figure 5. Steering-Induced Eye Movements Are Not Reflexive

(A) Top: time course of target-tracking index for one monkey during trials in which he performed the task (blue) or passively viewed the stimulus identical to the one generated when performing the task (red). Black trace shows the tracking index of the OFR model. Tracking indices at time points with negative variance explained were clipped to zero. Shaded region denotes ±1 SEM obtained by bootstrapping. Bottom: time course of the cumulative difference between the target-tracking index on active trials and the OFR model for individual monkeys.

(B) Top: time course of the tracking index of one monkey during trials in which the ground plane density was either high (blue) or low (red). Bottom: difference between target-tracking index under high- and low-density conditions for individual monkeys. Brown shaded regions in the bottom panels correspond to the time window considered for statistical testing.

See also Figure S5.

however, that the tracking index during passive viewing is poor right from trial onset, perhaps because the monkey did not consistently look at the target initially when it appeared on the screen. We wanted to know whether OFR dynamics,

coupled with the appropriate boundary condition (looking at the target when it initially appears), might be sufficient to give the impression that the animal is tracking the target. We simulated this model by shifting the initial eye position on each trial of the passive block to match the corresponding trial in the active block, a procedure that left the eye movement dynamics unaltered (Figure 5A, black). The tracking index of this simulated model was substantially lower than that observed during the active block of trials, suggesting that the target-tracking behavior is voluntary. In all monkeys, the target tracking during the active task was significantly stronger than during either the passive viewing condition or the OFR model (Figure 5A, bottom; mean difference ± standard deviation during the period shaded in brown, active: 0.27±0.1, passive: 0.08±0.1, OFR model: 0.07 ± 0.1 ; p < 0.01, bootstrap test). The difference between conditions was small in one monkey (labeled "Q" in Figure 5A, bottom; Figure S5, rightmost), possibly because this animal was mentally performing the task even during passive viewing.

Second, OFR is known to be sensitive to signal strength (Barthelemy et al., 2009; Quaia et al., 2012). To test whether target tracking depends on signal strength, we manipulated stimulus reliability by randomly interleaving trials with two different densities of ground plane elements by more than an order of magnitude (see STAR Methods). We analyzed the two sets of trials separately but found no significant difference between the target-tracking index (Figure 5B; mean \pm standard deviation across subjects, low density: 0.28 ± 0.1 , high density: 0.31 ± 0.1). Therefore, the pattern of eye movements observed during this task likely represents volitional movements rather than reflexive ones.



Figure 6. Fixation Affects Task Performance

(A) Trial-averaged temporal variability of subjects' eye position, quantified by standard deviation (see STAR Methods) during "eyes-moving" (blue) and "eyes-fixed" (red) trials. Error bars denote standard deviation across subjects (* $*p = 1.2 \times 10^{-3}$, paired t test).

(B) ROC curves averaged across subjects, for trials in the eyes-moving (blue) and the eyes-fixed condition (red). Inset shows the area under the two curves. Error bars denote standard deviation across subjects ($*p = 2.5 \times 10^{-3}$, paired t test).

(C) Top: comparison of the radial distances of the response and the target on trials under the two conditions. Different symbols denote different human subjects. Bottom: comparison of the (absolute) angular eccentricity of the response and target.

(D) Top: Pearson's correlation coefficient between the radial distance of subjects' response and the target for all individual subjects. Bottom: similar comparison for the absolute angular eccentricity of target and response under the two conditions. See also Figure S6.

Inhibiting Eye Movements Worsens Task Performance

Since eye movements were predictive of subjects' navigational performance, we wanted to know if they were essential for performing the task. To test this, we asked five human subjects to perform a variation of the task in which we overlaid a cross on top of the target location and instructed them to fixate on this cross for as long as it appeared on the screen. In half the trials ("eyes-moving" condition), the fixation cross disappeared along with the target so that subjects were free to produce eye movements as before. In the remaining trials ("eyes-fixed' condition), the cross remained at the same location on the screen throughout the trial, and subjects had to perform the task without moving their eyes (see STAR Methods). Although we did not penalize subjects for breaking fixation, we verified offline that they maintained fixation as instructed (Figure 6A and S6). We assessed their behavioral performance by comparing the AUC and found that performance was significantly impaired in the eyes-fixed condition (Figure 6B; n = 5 humans, mean AUC \pm standard deviation; eyes moving: 0.85 ± 0.07 , eyes fixed: 0.77 ± 0.07 , $p = 2.5 \times$ 10⁻³, paired t test). Figure 6C shows the responses of individual subjects. Although subjects were nearly unbiased under both conditions, the correlation between target and response locations was significantly lower in the absence of eye movements (Figure 6D; mean \pm standard deviation; $Corr(r, r^*)$, eyes moving: 0.71 ± 0.1 , eyes fixed: 0.49 ± 0.2 , p = 0.011, paired t test; $Corr(|\theta|, |\theta^*|)$, eyes moving: 0.92 ± 0.03 , eyes fixed: 0.82 ± 0.1 , p = 0.035). These results suggest that subjects benefit when their eyes can track the internally estimated goal location in this task.

DISCUSSION

Although the tracking index remained significantly above chance throughout the trial in these experiments, it nonetheless decreased over time. This is expected, because the target disappears, so subjects cannot directly measure its true position but must instead rely on an internal estimate computed by integrating optic flow with knowledge of the controller dynamics. We have previously shown that human subjects perform nearperfect integration in this task (Lakshminarasimhan et al., 2018a). Nevertheless, due to noise in the integration process, the error in the internal estimate of target location on any given trial should grow over time. Consequently, even if those estimates are unbiased, their precision worsens, leading to a decrease in the target-tracking index (Figure 4A, dashed line). Consistent with this, the precision of error-correcting saccades gradually deteriorated as the trial progressed (Figure 3E). Therefore, the observed decrease in target tracking is an inevitable consequence of noisy observations and noisy integration and in fact serves to expose the growing uncertainty in subjective beliefs.

Inferring Belief Dynamics

The task design used in this study was motivated by the need to ultimately understand neural computations governing belief dynamics that transform sensory inputs to motor output. In the real world, these belief dynamics correspond to subjective estimates of latent-variable dynamics, thus making them difficult to measure and relate to neural activity. Classic experimental paradigms used with primates attempt to measure this by training animals to provide discrete responses to simple stimuli, which reduces the dimensionality of state and action spaces, limiting their potential to shed light on natural computations. On the other hand, paradigms that let rodents loose in open arenas mimic natural behavior at the expense of sacrificing control over the nature of computations they perform. One exception is a recent paper (Knöll et al., 2018) in which the authors used stimuli with rich spatiotemporal dynamics to elicit continuous oculomotor behavior from primates with minimal reinforcement. While the Knöll et al. (2018) task is similar to ours in that it leverages unconstrained eye movements to remedy several shortcomings of historical approaches, the computation performed by the animals in their task was somewhat straightforward: inferring the current focus of expansion from noisy optic flow. In contrast, we asked animals to use optic flow to navigate to a goal location without providing explicit position cues. This requires the animal to first infer their self-motion from optic flow and then temporally integrate the resulting time-varying estimates to track the dynamics of a latent variable (position of the target relative to them)-a more challenging set of computations. Moreover, we trained our animals only by rewarding them for reaching the goal location. Eye movements were not explicitly reinforced, yet our post hoc analysis revealed that the continuous-valued, time-varying eye position encoded subjective beliefs about the time-varying latent variable. We hope that this approach of covertly measuring belief dynamics will serve as a useful template for future studies.

The Nature of Eye Movements

To understand the nature of slow eye movements made while steering toward the target, we analyzed individual components of eye position and found that both lateral version and elevation were largely smooth and consistent with the predicted dynamics for pursuing the invisible target. By analyzing eye movements during stimulus playback, we ruled out the possibility that the smooth dynamics correspond to pure ocular following response (OFR) induced by optic flow. Because these eye movements were always preceded by fixating a visible target and occurred in parallel with computations for mentally tracking that same target, they are functionally more similar to smooth-pursuit eye movements. Despite ample evidence for smooth-pursuit eye movements in the absence of foveal stimulation in humans (Becker and Fuchs, 1985; Missal and Heinen, 2017; Wyatt et al., 1994) and rhesus macaques (Ilg and Thier, 1999), smoothly tracking a purely imaginary object is thought to be difficult (Spering and Montagnini, 2011). This is because in the absence of dynamic information about target motion, the pursuit velocity gradually decays to zero (Barnes, 2008; Missal and Heinen, 2017). However, when the underlying model for target motion is known, subjects can use their dynamic internal representation of the target to make predictive smooth pursuit during target blanking (Adams et al., 2012; Orban de Xivry et al., 2008, 2013). In our task, the dynamics of optic flow completely determine the (relative) motion of the target and can subsequently drive eye movements. Furthermore, the flow fields were self-generated rather than simulated, a condition that has previously been shown to improve pursuit of occluded targets (Danion et al., 2017; Gauthier et al., 1988; Vercher and Gauthier, 1992). Nevertheless, a moderate contribution of OFR induced by optic flow cannot be completely excluded, so it is possible that the eye movements reported here are composed of a mixture of reflexive signals that encode velocity of self-motion and predictive signals that encode the latent state.

Finally, saccadic eye movements, although infrequent, contributed to tracking the target. The amplitude of these saccades was largely influenced by target-tracking error during the previous ~200 ms. These results suggest that the mechanism responsible for generating saccades in this paradigm are similar to the ones at play in flow tracking (Knöll et al., 2018) and smooth pursuit of visible objects (Daye et al., 2014; Orban de Xivry et al., 2008). One reason for the relatively low frequency of saccades in this study could be that motion in our task was self-generated and predominantly smooth, whereas saccades in smooth-pursuit experiments are primarily due to unexpected jumps in target velocity (de Brouwer et al., 2002).

Computational Role of Tracking Eye Movements

The experimental task was specifically designed to ensure that subjects would attempt to mentally track the goal location by integrating momentary sensory evidence about movement provided by optic flow. In principle, this can be accomplished without physically tracking the believed goal location with one's eyes. Yet, we noticed a significant decline in task performance when eye movements were suppressed. This is consistent with previous results that demonstrated that real-world driving performance is impaired when eye movements are constrained (Wilson et al., 2008). Although this does not demonstrate a need to make tracking eye movements, it suggests that eye movements play an important role in neural computations for navigation. Indirect evidence of a role for slow eye movements in visually guided navigation comes from a recent study of path integration in which subjects used a joystick to reproduce previously experienced self-motion (Churan et al., 2018). Eye movements during the reproduction phase were similar to those during initial exposure even when optic flow was removed. This suggests that eye movements constitute a form of mental imagery that, if suppressed, hamper memory retrieval (Johansson and Johansson, 2014; Johansson et al., 2012). Our findings extend this to naturalistic settings and argue that eye movements have a more dynamic role in path integration. The precise computational advantage of the specific eye movement dynamics observed in our task is unclear. Below, we propose two potential theories.

One possibility is that eye movements directed toward the intended goal location stabilize the mental image of the goal and could reduce the computational complexity of estimating selfmotion from optic flow similar to the effect of foveal image stabilization (Lappe et al., 1999; Longuet-Higgins and Prazdny, 1980;

Perrone and Stone, 1994; Sandini and Tistarelli, 1990). Normative mathematical theories posit that maintaining gaze at a point on the intended path can greatly simplify the problem of exploiting optic flow (Glennerster et al., 2001; Kim and Turvey, 1999; Wann and Swapp, 2000). Therefore, the eye movements reported here may constitute a closed-loop visuomotor process in which subjects integrate sense data (optic flow) to dynamically update their beliefs about the relative goal location and in turn use them to guide future eye movements in order to acquire new sense data in a computationally useful format. In this view, eye movements primarily aid optic flow processing.

Alternatively, the observed eye movements might simply be an embodiment of subjects' dynamically evolving internal beliefs about the goal. Humans have a well-documented tendency for externalizing their internal representations (Barsalou, 2008; Spivey, 2007), with eye movements sometimes employed as a pointing device to visible as well as invisible objects, much like one's index finger (Ballard et al., 1995, 1997; Spivey and Geng, 2001). By allowing dynamic beliefs about the relative target location to continuously modulate eve movements in this task, the brain could piggyback on the oculomotor circuit and reduce the computational burden on working memory. Consistent with this interpretation, there is overwhelming evidence for decision-related responses in primate oculomotor brain areas (de Lafuente et al., 2015; Shadlen and Newsome, 1996), and such responses are thought to drive eye movements (Joo et al., 2016). Therefore, in this view, primates use gaze as an affordance to efficiently update and store the output of integrating optic flow.

Although the above accounts are not mutually exclusive, simultaneously recording the neural activity from the primate sensory, oculomotor, and decision areas during this task might shed light on the dominant role of eye movements and how they link perception and action. A candidate brain area is the primate posterior parietal cortex, where there is ample evidence for convergence of self-motion (Avila et al., 2019; Britten, 2008; Gu et al., 2012), gaze (Andersen, 1989; Andersen et al., 1987), and decision-related (Gold and Shadlen, 2007; Ibos and Freedman, 2017; Lakshminarasimhan et al., 2018b) signals. In any case, regardless of the mechanistic and computational explanations for these eye movements, the paradigm used here offers a useful approach to directly readout dynamical internal beliefs in real time, simply by tracking subjects' eyes.

STAR*METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- LEAD CONTACT AND MATERIALS AVAILABILITY
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
 - Experimental setup
 - Behavioral Task
 - Feedback
 - Stimulus and Data acquisition
 - Model predicted eye position

• QUANTIFICATION AND STATISTICAL ANALYSIS

- Bias estimation
- Psychometric analysis
- Characterizing eye position
- Saccade detection and pre-processing
- Comparing predicted and observed eye positions
- Correlation between saccade amplitude and targettracking error
- Estimation of position uncertainty
- Deriving an upper bound on the target-tracking index
- Comparing behavioral and target-tracking errors
- Assessing performance in the fixation task
- DATA AND CODE AVAILABILITY

SUPPLEMENTAL INFORMATION

Supplemental Information can be found online at https://doi.org/10.1016/j. neuron.2020.02.023.

ACKNOWLEDGMENTS

We thank Samir Saidi for assisting with the human experiments and Jing Lin and Jian Chen for their help in programming the stimulus. This work was supported by NIH grants EY022538 and R01-DC014678. G.C.D. was supported by NIH grant EY016178. K.J.L. was supported in part by NSF NeuroNex Award DBI-1707398 and the Gatsby Charitable Foundation.

AUTHOR CONTRIBUTIONS

Conceptualization, K.J.L., G.C.D., X.P., and D.E.A.; Methodology, K.J.L., E.A., X.P., and D.E.A.; Investigation, K.J.L., E.A., and E.N.; Data Curation, K.J.L. and E.A.; Formal Analysis, K.J.L.; Writing – Original Draft, K.J.L.; Writing – Review & Editing, K.J.L., E.A., G.C.D., X.P., and D.E.A.; Funding Acquisition, G.C.D., X.P., and D.E.A.; Supervision, X.P. and D.E.A.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: July 3, 2019 Revised: December 24, 2019 Accepted: February 19, 2020 Published: March 13, 2020

REFERENCES

Adams, R.A., Perrinet, L.U., and Friston, K. (2012). Smooth pursuit and visual occlusion: active inference and oculomotor control in schizophrenia. PLoS ONE 7, e47502.

Adams, R.A., Aponte, E., Marshall, L., and Friston, K.J. (2015). Active inference and oculomotor pursuit: the dynamic causal modelling of eye movements. J. Neurosci. Methods *242*, 1–14.

Andersen, R.A. (1989). Visual and eye movement functions of the posterior parietal cortex. Annu. Rev. Neurosci. *12*, 377–403.

Andersen, R.A., Essick, G.K., and Siegel, R.M. (1987). Neurons of area 7 activated by both visual stimuli and oculomotor behavior. Exp. Brain Res. 67, 316–322.

Avila, E., Lakshminarasimhan, K.J., DeAngelis, G.C., and Angelaki, D.E. (2019). Visual and vestibular selectivity for self-motion in macaque posterior parietal area 7a. Cereb. Cortex *29*, 3932–3947.

Ballard, D.H., Hayhoe, M.M., and Pelz, J.B. (1995). Memory representations in natural tasks. J. Cogn. Neurosci. *7*, 66–80.

Ballard, D.H., Hayhoe, M.M., Pook, P.K., and Rao, R.P.N. (1997). Deictic codes for the embodiment of cognition. Behav. Brain Sci. *20*, 723–742, discussion 743–767.

Barnes, G.R. (2008). Cognitive processes involved in smooth pursuit eye movements. Brain Cogn. 68, 309–326.

Barsalou, L.W. (2008). Grounded cognition. Annu. Rev. Psychol. 59, 617–645.

Barthelemy, F.V., Fleuriet, J., and Masson, G.S. (2009). Temporal dynamics of 2D motion integration for ocular following in macaque monkeys. J. Neurophysiol. *103*, 1275–1282.

Becker, W., and Fuchs, A.F. (1985). Prediction in the oculomotor system: smooth pursuit during transient disappearance of a visual target. Exp. Brain Res. 57, 562–575.

Blum, J., and Price, N.S.C. (2014). Reflexive tracking eye movements and motion perception: one or two neural populations? J. Vis. 14, 23.

Bonnen, K., Burge, J., Yates, J., Pillow, J., and Cormack, L.K. (2015). Continuous psychophysics: target-tracking to measure visual sensitivity. J. Vis. *15*, 14.

Boström, K.J., and Warzecha, A.K. (2010). Open-loop speed discrimination performance of ocular following response and perception. Vision Res. *50*, 870–882.

Britten, K.H. (2008). Mechanisms of self-motion perception. Annu. Rev. Neurosci. 31, 389–410.

Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. J. Neurosci. *12*, 4745–4765.

Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. Science *340*, 95–98.

Churan, J., von Hopffgarten, A., and Bremmer, F. (2018). Eye movements during path integration. Physiol. Rep. 6, e13921.

Danion, F., Mathew, J., and Flanagan, J.R. (2017). Eye tracking of occluded self-moved targets: role of haptic feedback and hand-target dynamics. eNeuro 4, ENEURO.0101-17.2017.

Daye, P.M., Blohm, G., and Lefèvre, P. (2014). Catch-up saccades in head-unrestrained conditions reveal that saccade amplitude is corrected using an internal model of target movement. J. Vis. *14*, 14.

de Brouwer, S., Missal, M., Barnes, G., and Lefèvre, P. (2002). Quantitative analysis of catch-up saccades during sustained pursuit. J. Neurophysiol. *87*, 1772–1780.

de Lafuente, V., Jazayeri, M., and Shadlen, M.N. (2015). Representation of accumulating evidence for a decision in two parietal areas. J. Neurosci. *35*, 4306–4318.

Deravet, N., Blohm, G., de Xivry, J.O., and Lefèvre, P. (2018). Weighted integration of short-term memory and sensory signals in the oculomotor system. J. Vis. *18*, 16.

Erkelens, C.J., Van der Steen, J., Steinman, R.M., and Collewijn, H. (1989). Ocular vergence under natural conditions. I. Continuous changes of target distance along the median plane. Proc. R. Soc. Lond. B Biol. Sci. 236, 417–440.

Gauthier, G.M., Vercher, J.L., Mussa Ivaldi, F., and Marchetti, E. (1988). Oculomanual tracking of visual targets: control learning, coordination control and coordination model. Exp. Brain Res. *73*, 127–137.

Glasser, D.M., and Tadin, D. (2014). Modularity in the motion system: independent oculomotor and perceptual processing of brief moving stimuli. J. Vis. *14*, 28.

Glennerster, A., Hansard, M.E., and Fitzgibbon, A.W. (2001). Fixation could simplify, not complicate, the interpretation of retinal flow. Vision Res. *41*, 815–834.

Gold, J.I., and Shadlen, M.N. (2000). Representation of a perceptual decision in developing oculomotor commands. Nature *404*, 390–394.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci. 30, 535–574.

Gu, Y., Deangelis, G.C., and Angelaki, D.E. (2012). Causal links between dorsal medial superior temporal area neurons and multisensory heading perception. J. Neurosci. *32*, 2299–2313.

Hoffman, D.M., Girshick, A.R., Akeley, K., and Banks, M.S. (2008). Vergenceaccommodation conflicts hinder visual performance and cause visual fatigue. J. Vis. 8, 1–30.

Houlsby, N.M.T., Huszár, F., Ghassemi, M.M., Orbán, G., Wolpert, D.M., and Lengyel, M. (2013). Cognitive tomography reveals complex, task-independent mental representations. Curr. Biol. *23*, 2169–2175.

Howard, I.P. (2012). Vergence eye movements. In Perceiving in Depth (Oxford University Press), pp. 475–548.

Huk, A., Bonnen, K., and He, B.J. (2018). Beyond trial-based paradigms: continuous behavior, ongoing neural activity, and natural stimuli. J. Neurosci. *38*, 7551–7558.

lbos, G., and Freedman, D.J. (2017). Sequential sensory and decision processing in posterior parietal cortex. eLife 6, e23743.

Ilg, U.J., and Thier, P. (1999). Eye movements of rhesus monkeys directed towards imaginary targets. Vision Res. 39, 2143–2150.

Johansson, R., and Johansson, M. (2014). Look here, eye movements play a functional role in memory retrieval. Psychol. Sci. 25, 236–242.

Johansson, R., Holsanova, J., Dewhurst, R., and Holmqvist, K. (2012). Eye movements during scene recollection have a functional role, but they are not reinstatements of those produced during encoding. J. Exp. Psychol. Hum. Percept. Perform. *38*, 1289–1314.

Joo, S.J., Katz, L.N., and Huk, A.C. (2016). Decision-related perturbations of decision-irrelevant eye movements. Proc. Natl. Acad. Sci. USA *113*, 1925–1930.

Kim, N.G., and Turvey, M.T. (1999). Eye movements and a rule for perceiving direction of heading. Ecol. Psychol. *11*, 233–248.

Knill, D.C., and Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. Trends Neurosci. 27, 712–719.

Knöll, J., Pillow, J.W., and Huk, A.C. (2018). Lawful tracking of visual motion in humans, macaques, and marmosets in a naturalistic, continuous, and untrained behavioral context. Proc. Natl. Acad. Sci. U S A. *115*, E10486–E10494.

Kumar, A., Wu, Z., Pitkow, X., and Schrater, P. (2019). Belief dynamics extraction. arXiv, 1902.00673 https://arxiv.org/abs/1902.00673.

Lakshminarasimhan, K.J., Petsalis, M., Park, H., DeAngelis, G.C., Pitkow, X., and Angelaki, D.E. (2018a). A dynamic bayesian observer model reveals origins of bias in visual path integration. Neuron *99*, 194–206.e5.

Lakshminarasimhan, K.J., Pouget, A., DeAngelis, G.C., Angelaki, D.E., and Pitkow, X. (2018b). Inferring decoding strategies for multiple correlated neural populations. PLoS Comput. Biol. *14*, e1006371.

Lappe, M., Bremmer, F., and Van Den Berg, A.V. (1999). Perception of self-motion from visual flow. Trends Cogn. Sci. *3*, 329–336.

Lee, D.D., Ortega, P.A., and Stocker, A.A. (2014). Dynamic belief state representations. Curr. Opin. Neurobiol. 25, 221–227.

Loetscher, T., Bockisch, C.J., Nicholls, M.E.R., and Brugger, P. (2010). Eye position predicts what number you have in mind. Curr. Biol. *20*, R264–R265.

Longuet-Higgins, H.C., and Prazdny, K. (1980). The interpretation of a moving retinal image. Proc. R. Soc. Lond. B. Biol. Sci. 208, 385–397.

Missal, M., and Heinen, S.J. (2017). Stopping smooth pursuit. Philos. Trans. R. Soc. Lond. B Biol. Sci. 372, 20160200.

Orban de Xivry, J.J., Missal, M., and Lefèvre, P. (2008). A dynamic representation of target motion drives predictive smooth pursuit during target blanking. J. Vis. 8, 1–13.

Orban de Xivry, J.-J., Coppe, S., Blohm, G., and Lefèvre, P. (2013). Kalman filtering naturally accounts for visually guided and predictive smooth pursuit dynamics. J. Neurosci. *33*, 17301–17313.

Paninski, L. (2006). Nonparametric inference of prior probabilities from Bayesoptimal behavior. Advances in Neural Information Processing Systems *18*, 1067–1074.

Perrone, J.A., and Stone, L.S. (1994). A model of self-motion estimation within primate extrastriate visual cortex. Vision Res. *34*, 2917–2938.

Pinto, L., Koay, S.A., Engelhard, B., Yoon, A.M., Deverett, B., Thiberge, S.Y., Witten, I.B., Tank, D.W., and Brody, C.D. (2018). An accumulation-of-evidence task using visual pulses for mice navigating in virtual reality. Front. Behav. Neurosci. *12*, 36.

Pitkow, X., and Angelaki, D.E. (2017). Inference in the brain: statistics flowing in redundant population codes. Neuron *94*, 943–953.

Price, N.S.C., and Blum, J. (2014). Motion perception correlates with volitional but not reflexive eye movements. Neuroscience 277, 435–445.

Quaia, C., Sheliga, B.M., Fitzgibbon, E.J., and Optican, L.M. (2012). Ocular following in humans: spatial properties. J. Vis. *12*, 13.

Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. Psychol. Bull. *124*, 372–422.

Reddy, S., Dragan, A.D., and Levine, S. (2018). Where do you think you're going? Inferring beliefs about dynamics from behavior. arXiv, 1805.08010 https://arxiv.org/abs/1805.08010.

Sandini, G., and Tistarelli, M. (1990). Active tracking strategy for monocular depth inference over multiple frames. IEEE Trans. Pattern Anal. Mach. Intell. *12*, 13–27.

Shadlen, M.N., and Newsome, W.T. (1996). Motion perception: seeing and deciding. Proc. Natl. Acad. Sci. USA 93, 628–633.

Shibata, T., Kim, J., Hoffman, D.M., and Banks, M.S. (2011). The zone of comfort: predicting visual discomfort with stereo displays. J. Vis. *11*, 11.

Smith, M.L., Gosselin, F., and Schyns, P.G. (2012). Measuring internal representations from behavioral and brain data. Curr. Biol. 22, 191–196.

Spering, M., and Montagnini, A. (2011). Do we track what we see? Common versus independent processing for motion perception and smooth pursuit eye movements: a review. Vision Res. *51*, 836–852.

Spivey, M. (2007). The Continuity of Mind (Oxford University Press).

Spivey, M.J., and Geng, J.J. (2001). Oculomotor mechanisms activated by imagery and memory: eye movements to absent objects. Psychol. Res. *65*, 235–241.

Stocker, A.A., and Simoncelli, E.P. (2006). Noise characteristics and prior expectations in human visual speed perception. Nat. Neurosci. *9*, 578–585.

Tanenhaus, M.K., Spivey-Knowlton, M.J., Eberhard, K.M., and Sedivy, J.C. (1995). Integration of visual and linguistic information in spoken language comprehension. Science *268*, 1632–1634.

Turnham, E.J., Braun, D.A., and Wolpert, D.M. (2011). Inferring visuomotor priors for sensorimotor learning. PLoS Comput. Biol. 7, e1001112.

Vercher, J.L., and Gauthier, G.M. (1992). Oculo-manual coordination control: ocular and manual tracking of visual targets with delayed visual feedback of the hand motion. Exp. Brain Res. *90*, 599–609.

Wann, J.P., and Swapp, D.K. (2000). Why you should look where you are going. Nat. Neurosci. 3, 647–648.

Wilson, M., Chattington, M., and Marple-Horvat, D.E. (2008). Eye movements drive steering: reduced eye movement distribution impairs steering and driving performance. J. Mot. Behav. 40, 190–202.

Wu, Z., Kwon, M., Daptardar, S., Schrater, P., and Pitkow, X. (2019). Rational thoughts in neural codes. bioRxiv. https://doi.org/10.1101/765867.

Wyatt, H.J., Pola, J., Fortune, B., and Posner, M. (1994). Smooth pursuit eye movements with imaginary targets defined by extrafoveal cues. Vision Res. *34*, 803–820.

Zhang, M., Feng, J., Lim, J.H., Zhao, Q., and Kreiman, G. (2018). What am I searching for? arXiv, 1807.11926 https://arXiv.org/abs/1807.11926.

STAR***METHODS**

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Equipment, Software and Algorithms		
M20U9T-N82 USB joystick	CTI Electronics	http://www.ctielectronics.com/
Power 1401 MkII data acquisition system and Spike2 software	Cambridge Electronic Design Ltd.	http://ced.co.uk/
Open Graphics Library (OpenGL)	Khronos Group	https://www.opengl.org/
Custom-built analysis code	MATLAB	https://github.com/kaushik-l/monkey-dj

LEAD CONTACT AND MATERIALS AVAILABILITY

Further information and requests for resources should be directed to and will be fulfilled by the Lead Contact, Kaushik Lakshminarasimhan (jkl9@nyu.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Three rhesus macaques (all male, 7-8 years. old) and ten human subjects (six males, all adults in the age group 18-32 years.) participated in the experiments. All but one subject were unaware of the purpose of the study. All surgeries and experimental procedures were approved by the Institutional Review Board at Baylor College of Medicine, and were in accordance with National Institutes of Health guidelines. All human subjects signed an approved consent form. In the following sections, the term subject is used to denote both monkey and human subjects, unless specified otherwise or implied by the context.

METHOD DETAILS

Experimental setup

Monkeys were chronically implanted with a lightweight polyacetal ring for head restraint, and scleral coils for monitoring eye movements (CNC Engineering, Seattle WA, USA). At the beginning of each experimental session, monkeys were head-fixed and secured in a primate chair placed on top of a platform (Kollmorgen, Radford, VA, USA). A 3-chip DLP projector (Christie Digital Mirage 2000, Cypress, CA, USA) was mounted on top of the platform and rear-projected images onto a 60 × 60 cm tangent screen that was attached to the front of the field coil frame, ~30cm in front of the monkey. The projector was capable of rendering stereoscopic images generated by an OpenGL accelerator board (Nvidia Quadro FX 3000G).

Human subjects wore a custom-fit thermoplastic mask (CIVCO Medical Solutions) that was screwed to the back of the chair to restrain their head. The mask was mounted with a binocular eye tracker (ISCAN Inc.) to record the position of the subjects' pupils at 60Hz. All other aspects of the setup were similar to the one used for monkeys, but with subjects seated 67.5cm in front of a 149 \times 127 cm² (width \times height) rectangular screen. Although humans and monkeys were head-fixed, they were both free to move their eyes when performing the task, except under one experimental manipulation in humans (noted toward the end of the section below).

Behavioral Task

Subjects used an analog joystick (M20U9T-N82, CTI electronics) with two degrees of freedom and a circular displacement boundary to control their linear and angular speeds in a virtual environment. This virtual world comprised a ground plane whose textural elements had limited lifetime (~ 250ms) to avoid serving as landmarks. The ground plane was circular with a radius of 70 m (near and far clipping planes at 5cm and 4000cm respectively), with the subject positioned at its center at the beginning of each trial. Each texture element was an isosceles triangle (base × height: $8.5 \times 18.5 \text{ cm}^2$) that was randomly repositioned and reoriented anywhere in the arena at the end of its lifetime, making it impossible to use as a landmark. The maximum linear and angular speeds were fixed to $v_{\text{max}} = 2\text{ms}^{-1}$ and $\omega_{\text{max}} = 90^{\circ}$ /s respectively, and the density of the ground plane was either held fixed at $\rho = 2.5$ elements/m² or varied randomly between two values ($\rho = 2.5$ elements/m² and $\rho = 0.1$ elements/m²) in a subset of recording sessions (see below). The stimulus was rendered as a red-green anaglyph and projected onto the screen in front of the subject's eyes. Subjects wore goggles fitted with Kodak Wratten filters (red #29 and green #61) to view the stimulus. The binocular crosstalk for the green and red channels was 1.7% and 2.3% respectively.

Human subjects pressed a button on the joystick to initiate each trial, and the task was to steer to a random target location that was cued briefly at the beginning of the trial (Figure 1A). Monkeys performed the same task, but each trial was programmed to start after a

variable random delay (0.5 - 1.1 s) following the end of the previous trial. The target was a circular disc of radius 20cm whose luminance was matched to the texture elements. It appeared at a random location between $\theta = \pm 40^{\circ}$ of visual angle at a distance of r = 0.7 - 4m (up to 6 m for human subjects) relative the subject at the beginning of the trial. For human subjects, the target disappeared after one second, which was a cue for the subject to start steering, and the joystick controller was activated. In the case of monkeys, the target only appeared on the screen for 300ms, and the joystick was always active.

Monkeys typically performed two blocks of ~750 trials in each experimental session, and received feedback at the end of each trial. Monkeys performed a total of ~6,000 trials (4 sessions) each. Eye tracking was performed either using scleral coils (monkey Q & B) or a head-mounted eye tracker (monkey S). In one of the above recording sessions in each monkey, we saved the stimulus movie and replayed them to the animal at the end of the block. Both the visual stimulus and the schedule of rewards during this replay block were identical to the active navigation block, with the only difference being that the joystick was withheld and monkeys passively viewed the stimulus. Furthermore, a subset of the recording sessions (two sessions in each monkey) contained two randomly interleaved sets of trials that differed in terms of the density of optic flow (ρ = 0.1 elements/m² and ρ = 2.5 elements/m²).

Of the ten human subjects, five subjects performed a total of 600 trials spread equally across three blocks. The blocks were identical in all respects, except no feedback was provided at the end of the trials in the first and third blocks. The purpose of using this block structure was to study how feedback affected learning in humans. Although data collected in the absence of feedback (first and last blocks) are briefly described in Figure S1, the key results of the paper are based only on data collected during the intermittent block with feedback. Furthermore, during the block with feedback, the performance of human subjects typically stabilized within fifty trials (Figure S1B). Because we wanted to ensure that the performance was stable during the course of testing, we ignored the first fifty trials of this block for all our analysis (Figures 1, 2, and4). The remaining five human subjects participated in a version of the experiment that was designed to study the effect of inhibiting eye movements on task performance (Figure 6). These subjects first performed a block of fifty trials with feedback to allow their performance to stabilize. Following this pre-training block, they performed a test block comprising 400 trials of a version of this task in which a fixation cross was overlaid on top of the target in each trial, again with feedback. In a random subset of trials (50%), this fixation cross remained on the screen even after the target disappeared and subjects were instructed to maintain fixation on the cross while steering to the target. The location of the cross remained fixed in screen coordinates and thus carried no dynamic information about stimulus location.

Feedback

Monkeys received binary feedback following a variable waiting period after stopping (range: 0.1–0.6 s, mean waiting period: 0.25 s). They received a drop of juice if their stopping position was within 0.6 m away from the center of the target. No juice was provided otherwise. The fixed reward boundary of 0.6 m was determined using a staircase procedure prior to the experiment to ensure that monkeys received reward in approximately two-thirds of the trials.

Human subjects received a somewhat richer, adaptive feedback in the form of a bullseye pattern that appeared on the ground at the end of steering upon pushing a button. The bullseye was centered on the target, with the innermost region having the highest luminance. The pattern comprised of five zones (Figure S1A), and the radii of the rings were continuously scaled (up or down by 5%) during the experiment using a 1-up 2-down staircase procedure. Additionally, an arrowhead pointing to the target also appeared on the ground in front of the subjects, colored green or red depending on whether the subject's stopping position was inside or outside the reward boundary. The adaptive feedback procedure ensured that human subjects, like monkeys, stopped within the reward boundary in roughly two-thirds of the trials. Unlike monkeys, human subjects did not receive juice at the end of each successful trial, but instead received monetary compensation that was commensurate with their performance.

Stimulus and Data acquisition

All stimuli were generated and rendered using C++ Open Graphics Library (OpenGL) by continuously repositioning the camera based on joystick inputs to update the visual scene at 60 Hz. The camera was positioned at a height of 1 m above the ground plane (10cm for monkeys). Spike2 software (Power 1401 MkII data acquisition system from Cambridge Electronic Design Ltd.) was used to record and store the target location (r^* , θ^*), subject's position (r, θ), horizontal positions of left and right eyes (α_l and α_r), vertical eye positions (β_l and β_r) and all event markers for offline analysis at a sampling rate of 833(1/3) Hz.

Model predicted eye position

To test whether subjects' eyes tracked the location of the (invisible) target, we generated predictions for subjects' instantaneous eye positions by assuming that they maintained fixation at the center of the target. (x_t, y_t, z_t) denotes the location of the target relative to the mid-point of the subject's eyes at time *t*. The *mean* predicted lateral displacement (relative to fixating at the point ($0, \infty, 0$)) of the left and right eyes ($\overline{\alpha}_t$ and $\overline{\alpha}_r$) are geometrically related to the target location and the inter-ocular distance (2 Δ) as (Figure S2D):

$$\overline{\widehat{\alpha}}_{I}(t) = \tan^{-1}\left(\frac{x_{t}+\Delta}{\sqrt{y_{t}^{2}+z_{t}^{2}}}\right); \overline{\widehat{\alpha}}_{r}(t) = \tan^{-1}\left(\frac{x_{t}-\Delta}{\sqrt{y_{t}^{2}+z_{t}^{2}}}\right)$$
(1.1)

Likewise, the vertical displacement of the two eyes $(\overline{\hat{\beta}}_l \text{ and } \overline{\hat{\beta}}_r)$ should be:

$$\overline{\widehat{\beta}}_{I}(t) = \tan^{-1}\left(\frac{z_{t}}{\sqrt{y_{t}^{2} + (x_{t} + \Delta)^{2}}}\right); \overline{\widehat{\beta}}_{I}(t) = \tan^{-1}\left(\frac{z_{t}}{\sqrt{y_{t}^{2} + (x_{t} - \Delta)^{2}}}\right)$$
(1.2)

Note that z_t is determined entirely by the camera height and hence time-invariant. In contrast, x_t and y_t change continuously as the subject steers to the target, and are both equal to zero in the special case when the subject's location coincides with the center of the target. The predicted eye positions also have *variances* associated with them, which we derive in a later section (Equation 4).

QUANTIFICATION AND STATISTICAL ANALYSIS

Customised MATLAB code was written to analyze data and to fit models. Depending on the quantity estimated, we report statistical dispersions either using 95% confidence interval, standard deviation, or standard error in the mean. The specific dispersion measure is identified in the portion of the text accompanying the estimates. For error bars in figures, we provide this information in the caption of the corresponding figure. We report exact *p*-values for all statistical tests, and describe the outcome as significant if p < 0.05.

Bias estimation

We regressed (with an intercept term) each subject's response positions (r, θ) against target positions (r^*, θ^*) separately for the radial $(r \text{ versus } r^*)$ and angular $(\theta \text{ versus } \theta^*)$ co-ordinates, and the radial and angular multiplicative biases were quantified as the slope of the respective regressions (Figure 1F). The intercept terms of the regression models denote additive bias. For each subject, we estimated the 95% confidence intervals for the biases by bootstrapping.

Psychometric analysis

As described in the section on feedback, reward boundaries were chosen to ensure that all subjects correctly stopped within the reward zone in about two-thirds of the trials. However, the precise radius of these boundaries varied across human subjects, as well as between humans and monkeys. To objectively compare the performance of different subjects on a common scale, we performed ROC analysis as follows. For each subject, we first constructed a psychometric function by calculating the proportion of correct trials as a function of (hypothetical) reward boundary (Figure 1G). In keeping with the range of target distances used for the two species, we varied the reward boundary between 0–4 m for monkeys and 0–6 m for human subjects. Whereas an infinites-imally small boundary will result in all trials being classified as incorrect, a large enough reward boundary will yield near-perfect accuracy. To define a chance-level psychometric function, we repeated the above procedure but now by shuffling the target locations across trials, thereby destroying the relationship between target and response locations. Finally, we obtained the ROC curve by plotting the proportion of correct trials in the original dataset (true positives) against the shuffled dataset (false positives) for each value of hypothetical reward boundary. We used the area under this ROC curve to obtain an accuracy measure that was independent of the reward boundary used for various subject.

Characterizing eye position

For convenience, we express the subject's *actual* eye position using the following three standard degrees of freedom: (*i*) Conjunctive horizontal movement of the two eyes or 'lateral version' quantified here as the mean lateral position of the two eyes, $\alpha = (\alpha_l + \alpha_r)/2$, (*ii*) Conjunctive vertical movement of the two eyes or 'elevation' quantified here as $\beta = (\beta_l + \beta_r)/2$, (*iii*) Disjunctive horizontal eye movements or 'vergence' quantified here as $\gamma = (\alpha_l - \alpha_r)/2$. Disjunctive eye movements along the vertical direction (vertical vergence) were an order of magnitude smaller than the precision of our measurements, and therefore we ignored them in all our analyses. We also transformed the *predicted* eye positions given by Equation 1 into the above three degrees of freedom using analogous definitions to obtain $\overline{\alpha}$, $\overline{\beta}$, and $\overline{\gamma}$.

Saccade detection and pre-processing

We estimated the instantaneous speed of eye movements as $(\dot{\alpha}^2 + \dot{\beta}^2)^{1/2}$ where α and β denote lateral version and elevation respectively (as defined above), and a dot denotes a time derivative. Saccades were detected by identifying the time points at which the speed of eye movements crossed a threshold of 200°/s from below (a threshold of 50°/s yielded similar results). Although saccades were mostly confined to periods immediately following target onset and end of steering (Figure 2B), we removed a period of 100ms immediately following the onset of saccades for visualizing the time-course of eye movements during the trial (Figure 2E) and for all subsequent temporal analyses described below. We verified that this procedure had minimal effect on the results. In approximately 10% of the trials in monkeys and ~30% in human subjects, the subject traveled beyond the target. The predicted eye positions toward the end of these trials were outside the range that was physically possible. Therefore, we removed time points at which any of the four predicted components of eye movements in Equation 1 exceeded 60° before further analysis. Such time points constituted less than 3% of the dataset, and including them did not qualitatively alter the results.

Comparing predicted and observed eye positions

Let $\varphi_t = (\alpha_t, \beta_t)$ and $\overline{\varphi}_t = (\overline{\alpha}_t, \overline{\beta}_t)$ denote the observed and mean predicted eye positions respectively at time *t*. For each subject, we computed the square root of the fraction of variance in their eye movements explained by the predictions:

$$\rho_{t} = \sqrt{1 - \frac{\langle \left\| \boldsymbol{\varphi}_{t} - \overline{\boldsymbol{\varphi}}_{t} \right\|_{2}^{2} \rangle}{\langle \left\| \boldsymbol{\varphi}_{t} - \overline{\boldsymbol{\varphi}}_{t} \right\|_{2}^{2} \rangle}}$$
(2)

where $\|\cdot\|_2$ denotes the L_2 norm, $\langle \cdot \rangle$ denotes expectation across trials, and $\overline{\varphi}_t$ denotes the mean observed eye position across trials at time *t*. Because the predictions are based on a model that assumes subjects' eyes track the center of the target, we call ρ the '*target-tracking index*', or simply '*tracking-index*'. A value of 1 corresponds to perfect prediction while zero implies that the predictions were no better than the mean observation. In principle, the deviation from the predictions can be larger than the intrinsic variability of the data. We clipped the target-tracking index to zero whenever this happened. Since trial durations were variable, we aligned all trials relative to the time at which the target was turned off (*t* = 0) to estimate the time course of tracking coefficient $\rho_t^{\text{start}} \forall t \in [0, 1.8s]$. ρ_0^{start} corresponds to the similarity between observed and predicted eye position at the moment when the target was turned off (Figure 2F). We also computed the tracking coefficient by aligning trials with respect to the end of steering (*t* = *T*) to estimate $\rho_t^{\text{storp}} \forall t \in [-1.2s, 0]$. To visualize the time-course of the tracking coefficient, we plot both ρ_t^{start} and ρ_t^{storp} with a break in the x axis (Figures 2G, 4, and 5). To assess standard errors and statistical significance of differences between tracking coefficients from pairs of conditions (e.g., rewarded versus unrewarded trials), we used a bootstrap test with 10,000 bootstrap samples.

Correlation between saccade amplitude and target-tracking error

Because saccades were not very frequent while steering, we pooled data from all subjects (separately for monkeys and humans) for analyzing saccades. The amplitude of saccades was taken to be the average displacement of the position of the two eyes from saccade onset to 100ms later $(\Delta \varphi = (\Delta \alpha^2 + \Delta \beta^2)^{1/2}$; Figure 3B). We quantified the effect of saccadic eye movements on target-tracking error by computing the saccade-triggered average (STA) of the tracking error within a 400ms window centered around time t_s of saccade onset (*i.e.*, $\langle \| \varphi_{t-t_s} - \overline{\hat{\varphi}_{t-t_s}} \|_2 \rangle_{t_s} \forall t \in [-0.2s, 0.2s]$). To quantify the precise relationship between saccade amplitude and tracking error, we simultaneously regressed horizontal and vertical amplitudes of the saccade ($\Delta \alpha$ and $\Delta \beta$) on horizontal and vertical tracking errors ($\alpha - \overline{\hat{\alpha}}$ and $\beta - \overline{\hat{\beta}}$), respectively, at various lags between ± 0.5s with l^2 regularization.

Estimation of position uncertainty

We estimated subjects' position uncertainty by binning the 2D space into $10 \times 10 \text{ cm}^2$ bins. For each bin, we computed the variance in the subject's stopping position across trials in which targets fell in that bin. The resulting spatial map of variability was then convolved with a two-dimensional isotropic Gaussian kernel of width 40cm (equal to the diameter of the target) to yield a smooth estimate of variability as a function of space (Figure 4B, left). Because subjects aimed to stop on the target, variability in their stopping position can be interpreted as the uncertainty in subjects' posterior estimate about their own position.

Deriving an upper bound on the target-tracking index

Once the target disappears, subjects no longer get to directly observe it. To reach the target location, they update their beliefs about the relative location of the target by integrating their self-motion, which in turn must be estimated from the observed optic flow. Even if those beliefs are accurate on average, the uncertainty in believed target location will grow over time on any given trial due to noise both in the observations and in the integration process. Consequently, the degree to which subjects' eyes can track the target (quantified by the tracking index, ρ) should decrease over time. Using the variability in subjects' stopping positions to model their uncertainty in their believed location (see section above), we derived an approximate upper-bound on the temporal dynamics of the tracking-index ρ_t at time *t* assuming inter-ocular distance $\Delta \approx 0$:

$$\rho_t \leq \sqrt{1 - \frac{\langle \|\widehat{\boldsymbol{\varphi}}_t - \overline{\widehat{\boldsymbol{\varphi}}}_t\|_2^2 \rangle}{\langle \|\boldsymbol{\varphi}_t - \overline{\boldsymbol{\varphi}}_t\|_2^2 \rangle}}$$
(3)

where $\overline{\hat{\phi}_t} = \langle \widehat{\phi}_t \rangle$ denotes the mean predicted eye position at time *t*. Note that this represents an upper-bound insofar as the variability in subject's stopping positions stems entirely from uncertainty in their believed location. To derive this approximate bound, we first used the first-order Taylor series approximation of Equation (1) to express the variance of the predicted eye position $(\widehat{\alpha}_t, \widehat{\beta}_t)$ in terms of the variance of the relative target position (x_t, y_t, z_t) as: $Var(\widehat{\alpha}_t) = (\partial f/\partial x)^2 Var(x_t) + (\partial f/\partial y)^2 Var(y_t)$ and $Var(\widehat{\beta}_t) = (\partial g/\partial x)^2 Var(x_t) + (\partial g/\partial y)^2 Var(y_t)$, where $f(x_t, y_t, z_t) = \tan^{-1}(\sqrt{y_t^2 + z_t^2}/x_t)$ and $g(x_t, y_t, z_t) = \tan^{-1}(Z_t/\sqrt{y_t^2 + x_t^2})$ from Equation (1), and we have used the fact that $Var(z_t) = 0$ because there is no motion component perpendicular to the ground plane. Substituting the derivatives, we get:

$$\operatorname{Var}(\widehat{\alpha}_{t}) = \frac{\left(y_{t}^{2} + z_{t}^{2}\right)}{\left(x_{t}^{2} + y_{t}^{2} + z_{t}^{2}\right)^{2}} \operatorname{Var}(x_{t}) + \frac{x_{t}^{2}y_{t}^{2}}{\left(x_{t}^{2} + y_{t}^{2} + z_{t}^{2}\right)^{2}\left(y_{t}^{2} + z_{t}^{2}\right)} \operatorname{Var}(y_{t})$$

$$(4.1)$$

$$\operatorname{Var}(\widehat{\beta}_{t}) = \frac{z_{t}^{2}}{\left(x_{t}^{2} + y_{t}^{2} + z_{t}^{2}\right)^{2} \left(x_{t}^{2} + y_{t}^{2}\right)} \left(x_{t}^{2} \operatorname{Var}(x_{t}) + y_{t}^{2} \operatorname{Var}(y_{t})\right)$$
(4.2)

The above equations are based on first-order Taylor series approximation and hold as long as the higher-order terms are relatively small. Although we cannot not directly measure $Var(x_t)$ and $Var(y_t)$, we could estimate them from the data (see previous section) and use it to determine the variability in predicted eye positions given by Equation 4. Variability in the predictions then implies a lower

bound in the mean squared error achievable by any observation φ_t : $\|\varphi_t - \widehat{\varphi}_t\|_2^2 \ge \|\overline{\varphi}_t - \overline{\widehat{\varphi}}_t\|_2^2$. Substituting this in (2), we obtain an upper bound on the tracking-index given by Equation 3. Note that, in deriving this approximate upper-bound, we ignored the noise in generating an eye movement to an intended location (process noise). So in principle, it is possible to derive a tighter bound by incorporating it. Note that as subjects approach the target, x_t and y_t approach zero, whereas the uncertainty grows so both $Var(x_t)$ and $Var(y_t)$ increase. Together, this leads to an increase in the variance of the predicted eye positions (Equation 4) and consequently, a gradual decrease in the fraction of explainable variance over time (Equation 3).

Comparing behavioral and target-tracking errors

To test whether poor target-tracking was associated with poor behavioral accuracy, we estimated the correlation between behavioral and target-tracking errors across trials of individual recording sessions. Behavioral error was given by the Euclidean distance between the target location and the subject's response (stopping location) on individual trials, while the target-tracking error was given by the Euclidean distance between actual and predicted eye position, averaged over the entire time period of the trial, except for the last 300ms (as the predictions typically broke down when the subject was too close to the target). Because trial difficulty could affect both errors thereby inducing spurious correlations, we estimated the null distribution of correlations using a shuffling procedure where we grouped the trials from each recording session into ten quantiles based on target distance and shuffling only trials within the same group. The results were quite robust to the number of quantiles.

Assessing performance in the fixation task

To assess the behavioral effect of inhibiting eye movements, we compared human subjects' performance across 'eyes-moving' and 'eyes-fixed' trials. Because we did not control for fixation breaks that happened in the 'eyes-fixed' condition during the experiment, we identified and removed such trials offline. Specifically, we removed the trials in which the temporal standard deviation (σ) of subject's eye position during the trial (i.e., from the time when the target disappeared until the end of steering) exceeded 3° (roughly half-width of the fixation cross), from our analysis (~10% of the fixation trials across all subjects). The standard deviation was quantified as $\sigma = \sqrt{\sigma^2(\alpha) + \sigma^2(\beta)}$ where $\sigma(\alpha)$ and $\sigma(\beta)$ denote the temporal standard deviation of lateral version and elevation respectively. To evaluate the role of eye movements, we compared subjects' performance in the fixation trials ('eyes-fixed') with trials that did not require fixation ('eyes-moving'). For both sets of trials, we computed ROC curves for distinguishing 'rewarded' and 'unrewarded' trials (see section 'psychometric analysis' above) and used a paired t test to test whether the mean area under the curves were different. We also computed the correlation between target and response locations and then used a paired t test to test whether there was a significant difference between the correlation coefficients in the two sets of trials across subjects (Figure 6D).

DATA AND CODE AVAILABILITY

MATLAB code implementing all quantitative analyses in this study is available online (see Key Resources Table). Datasets generated by this study are available from the corresponding author upon reasonable request.