Neuron

A Dynamic Bayesian Observer Model Reveals Origins of Bias in Visual Path Integration

Highlights

- Humans navigated in a virtual environment by integrating sparse optic flow
- They consistently overshot goal location—traveled too far and turned too much
- Behavior better explained by a slow-velocity prior than leaky evidence integration
- Growing uncertainty leads to a distance-dependent reversal in the direction of bias

Authors

Kaushik J. Lakshminarasimhan, Marina Petsalis, Hyeshin Park, Gregory C. DeAngelis, Xaq Pitkow, Dora E. Angelaki

Correspondence

jklakshm@bcm.edu

In Brief

Humans are typically biased while navigating by integrating their selfmotion. Using virtual reality and probabilistic modeling, Lakshminarasimhan et al. demonstrate that systematic errors in visual path integration originate from the combined influence of a slow-velocity prior and growing position uncertainty.



A Dynamic Bayesian Observer Model Reveals Origins of Bias in Visual Path Integration

Kaushik J. Lakshminarasimhan,^{1,5,6,*} Marina Petsalis,^{2,5} Hyeshin Park,¹ Gregory C. DeAngelis,³ Xaq Pitkow,^{1,4} and Dora E. Angelaki^{1,4}

¹Department of Neuroscience, Baylor College of Medicine, Houston, TX 77030, USA

²McGovern Medical School, Houston, TX 77030, USA

³Department of Brain and Cognitive Sciences, University of Rochester, Rochester, NY 14627, USA

⁴Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005, USA

⁵These authors contributed equally

⁶Lead Contact

*Correspondence: jklakshm@bcm.edu

https://doi.org/10.1016/j.neuron.2018.05.040

SUMMARY

Path integration is a strategy by which animals track their position by integrating their self-motion velocity. To identify the computational origins of bias in visual path integration, we asked human subjects to navigate in a virtual environment using optic flow and found that they generally traveled beyond the goal location. Such a behavior could stem from leaky integration of unbiased self-motion velocity estimates or from a prior expectation favoring slower speeds that causes velocity underestimation. Testing both alternatives using a probabilistic framework that maximizes expected reward, we found that subjects' biases were better explained by a slowspeed prior than imperfect integration. When subjects integrate paths over long periods, this framework intriguingly predicts a distance-dependent bias reversal due to buildup of uncertainty, which we also confirmed experimentally. These results suggest that visual path integration in noisy environments is limited largely by biases in processing optic flow rather than by leaky integration.

INTRODUCTION

The world is inherently noisy and dynamic. To act successfully, we must continuously monitor our sensory inputs, gather evidence in favor of potential actions, and make subjectively good decisions in the face of uncertain evidence. Traditional binary decision tasks lack the temporal richness to shed light on continuous behaviors in demanding environments (Lee et al., 2014; Pitkow and Angelaki, 2017). Newer "continuous psychophysics" methods (Bonnen et al., 2015; Mulligan et al., 2013) can constrain models much more rapidly and can be elaborated to simulate real world challenges. Here, we develop a naturalistic visuomotor virtual navigation task with controllable sensory uncertainty and provide a unified framework to under-

stand how dynamic perceptual information is combined over time. We then use this framework to understand the origins of bias in path integration—a natural computation that involves sensory perception, evidence accumulation, and spatial cognition.

Path integration is a navigation strategy used to maintain a sense of position solely by integrating self-motion information. Humans and animals are capable of path integrating (Alyan and Jander, 1994; Benhamou, 1997; Collett and Collett, 2000; Etienne, 1992; Loomis et al., 1999; Seguinot et al., 1998), albeit often with systematic errors (or biases). Biases have been observed in many species under a variety of experimental conditions involving visual (Bakker et al., 2001; Frenz and Lappe, 2005; Lappe and Frenz, 2009; Redlick et al., 2001) and/or body-based (Jürgens et al., 1999; Klatzky et al., 1990; Loomis et al., 1993; e.g., vestibular and proprioceptive) self-motion cues, yet their origins are not fully understood.

Broadly speaking, path integration entails two stages-estimating one's self-motion and integrating that estimate over time. Most previous accounts of behavioral biases in path integration implicate the latter, arguing for suboptimal integration of movement velocity that produces errors that increase with time (Mittelstaedt and Glasauer, 1991; Mittelstaedt and Mittelstaedt, 2001; Vickerstaff and Di Paolo, 2005) or distance (Bergmann et al., 2011; Brossard et al., 2016; Lappe et al., 2007, 2011). However, past modeling approaches were dominated by attempts to fit empirical functions using only subjects' final states at the end of the integration process, without considering the performance constraints imposed by noise in the sensory inputs. This has led to the view that bias in path integration is due to leaky integration—a severely suboptimal strategy that is inconsistent with studies in other domains demonstrating statistically optimal behavior in static and dynamic binary tasks (Bogacz, 2007; Brunton et al., 2013; Ernst et al., 2012; Fiser et al., 2010; Issen et al., 2015; Jogan and Stocker, 2013).

An alternative explanation is that the bias stems from errors in sensory estimates—e.g., from bias in velocity estimation or from incorrectly accumulating perceptual uncertainty over time. For example, human judgement of retinal speed is known to be biased, and this is well explained by a Bayesian observer model with a slow-speed prior (Hürlimann et al., 2002; Sotiropoulos

et al., 2014; Stocker and Simoncelli, 2006; Weiss et al., 2002). If a similar prior biases our judgement of self-motion velocity, this will naturally lead to path integration biases, even if the integration itself is perfect. This alternative explanation has been invoked previously to explain bias in integrating passive angular rotations (Jürgens and Becker, 2006) and proposed to be important for path integration in general (Petzschner and Glasauer, 2011). However, this proposal has never been rigorously tested against explanations based on leaky integration. These two alternatives have very different implications for neural models of evidence accumulation under uncertainty, so it is important to distinguish them. Finally, to our knowledge, past studies of path integration considered integration of either angles or distance, but not both. Therefore, it is unclear whether those findings would generalize to navigation in the real world.

To more directly determine whether bias in path integration under naturalistic conditions stems mainly from a slow-speed prior or suboptimal integration, we tested human subjects on a task in which they navigated within a horizontal plane using sparse optic flow. We found that subjects underestimated both linear and angular displacements when navigating short distances. We analyzed these data using a mathematical theory that includes components for sensory processing, integration dynamics, and decision making. Our analysis revealed that the behavioral errors can be explained by a model in which subjects maximized their expected reward under the influence of a slowspeed prior rather than by leaky integration of unbiased velocity estimates. This result was confirmed in a separate experiment, in which we tested the predictions of both models by manipulating the reliability and the range of optic flow. In addition, when extended to longer distance scales, the model predicts a potential reversal in the pattern of bias from overshooting to undershooting due to buildup of uncertainty, and we also confirmed this prediction experimentally. These findings suggest that human subjects can maintain a dynamic probabilistic representation of their location while navigating, and their ability to path integrate is limited largely by brain structures that process selfmotion rather than by circuits that integrate velocity estimates based on optic flow.

RESULTS

We asked human subjects to perform a visual navigation task in which they used a joystick to steer to a cued target location in a virtual environment devoid of allocentric reference cues (Figure 1A; Videos S1 and S2; STAR Methods). At the beginning of each trial, a circular target blinked briefly (\sim 1 s) at a random location on the ground plane, after which it disappeared and the joystick controller was activated. The joystick had two degrees of freedom that controlled forward and angular velocities, allowing the subject to steer freely in two dimensions (Figure 1B). Subjects were instructed to stop steering when they believed their position fell within the target, but did not receive any performance-related feedback. Target locations were varied randomly across trials and were uniformly distributed over the ground plane area within the subject's field of view (Figure 1C, top). The subject's movement trajectory was recorded throughout each trial (Figure 1C, bottom).

Behavioral Data

The subject's "response location" was given by their stopping position at the end of each trial. We quantified behavioral error on each trial by comparing the response location against the target location. Figure 1D (left) shows an aerial view of the target location and one subject's trajectory during a representative trial. On this trial, the error vector points radially outward and away from straight ahead, implying that the subject overshot the target both in terms of the net distance moved as well as the net angle rotated. The vector field of errors across all trials revealed a qualitatively similar pattern of errors throughout the experiment (Figure 1D, right).

To quantify these errors, we separately compared the radial distance and angular eccentricity of the target to those of the subject's response location in each trial. We found a systematic bias underlying the response errors in both quantities: this subject consistently traveled a greater distance and rotated through a greater angle than necessary (Figure 1E). We observed similar biases across subjects (Figure S1A), and these biases were well described by a simple linear model with multiplicative gain scaling, rather than an additive shift, of the subjects' estimates of their net displacement and rotation (mean coefficient of determination, R^2 across subjects-distance, 0.70±0.12; angle, 0.92 ± 0.11). Moreover, this multiplicative model was able to capture \sim 95% of the variance explained by more complicated nonlinear models (see legend of Figure S1A). Therefore, instead of quantifying bias in units of distance and angle, we used the slopes of the corresponding linear regressions (i.e., the gain) as a measure of bias in radial distance and angle for each subject. In this scheme, slopes greater than and less than unity correspond to overshooting and undershooting, respectively, whereas unity (not zero) slope corresponds to unbiased performance. Both radial and angular biases were significantly greater than unity across subjects (Figure 1F; mean distance bias [±SE], $\Gamma_r = 1.19 \pm 0.07$, $p = 4.1 \times 10^{-2}$; t test; mean angle bias, $\Gamma_{\theta} = 1.78 \pm 0.16$, $p = 2.8 \times 10^{-3}$).

We varied target locations across trials to preclude the use of strategies based only on movement duration. Nevertheless, subjects may have been encouraged to use such a strategy due to the inherent relationship between distance and time. To test this, we randomly interleaved trials in which we removed all ground plane elements, thereby eliminating optic flow. The correlation between target and response locations dropped substantially in these trials (Figure S1B), implying that subjects relied heavily on optic flow cues, rather than a mental clock, to perform the task.

We allowed subjects to freely control their velocity at all times and found modest variability in average velocity across trials. This trial-by-trial variability in velocity was uncorrelated with trial-by-trial variability in subjects' radial and angular position biases (Figure S1C), suggesting that movement velocity within the range we observed does not influence subjects' path integration errors during self-generated movement. Such a velocity-independent bias can arise if bias in velocity estimation is multiplicative, and we quantitatively model and test this phenomenon later. Velocities varied across time differently for different subjects as well: four of the seven subjects used a serial strategy, first rotating and then moving straight ahead to reach the target



Figure 1. Task Structure and Behavioral Response

(A) Subjects use a joystick to navigate to a cued target (yellow disc) using optic flow cues generated by ground-plane elements (orange triangles). The ground-plane elements appeared transiently at random orientations to ensure that they cannot serve as spatial or angular landmarks (STAR Methods).

(B) Left: the time course of linear (top) and angular (bottom) speeds during one example trial. Time is also encoded by line color. Right: aerial view of the subject's spatial trajectory during the same trial is shown.

(C) Top: aerial view of the spatial distribution of target positions across trials. Bottom: subject's movement trajectories during a representative subset of trials are shown.

(D) Left: target location (solid black) and subject's steering response (colored as in B) during a representative trial. Red arrow represents the error vector. Right: vector field denoting the direction of errors across trials is shown. The tail of each vector is fixed at the target location, and vectors were normalized to a fixed length for better visibility. The grayscale background shows the spatial profile of the error magnitude.

(E) Top right: comparison of the radial distance \bar{r} of the subject's response (final position) against radial distance r of the target across all trials for one subject. Bottom right: angular eccentricity of the response $\bar{\theta}$ versus target angle θ is shown. Black dashed lines have unity slope, and the red solid lines represent slopes of the regression fits. Left: geometric meaning of the quantities in the scatterplots is shown.

(F) Radial and angular biases were quantified as the slopes of the corresponding regressions and plotted for individual subjects. Error bars denote 95% confidence intervals of the respective slopes. Dashed lines indicate unbiased radial or angular position responses. Solid diagonal line has unit slope.

See also Figures S1 and S2 and Videos S1 and S2.

(Figures S2A and S2B), while the remaining subjects traveled along curvilinear trajectories. Subjects with both strategies had comparable radial and angular biases (Figure S2C), suggesting that they do not benefit from integrating the angular and linear components separately. This finding also shows that overshooting is not restricted to cases in which subjects make curvilinear trajectories.

Finally, we introduced angular landmarks in the environment by displaying a distant mountainous background (Figure S2D). This manipulation did not alter the radial bias but eliminated angular bias almost completely ($\Gamma_r = 1.29 \pm 0.08$; $\Gamma_{\theta} = 1.1 \pm 0.04$; Figure S2E). This suggests that biases measured in the absence of landmarks reflect errors in spatial perception rather than problems associated with motor control. To further validate this, we conducted an additional experiment in which we passively transported subjects over trajectories that passed through the targets at a constant velocity, thereby eliminating motor control (STAR Methods; Figure S2F). Subjects simply pressed a button to indicate when they believed they had reached the target. Again, we observed overshooting that scaled linearly with the radial distance of the target ($\Gamma_r = 1.38 \pm 0.1$; Figure S2G). Note that a delay in pressing the button would produce an identical bias at all distances (i.e., additive bias) and thus cannot explain the above result.

Together, these data suggest that subjects overshoot when using optic flow to navigate modest distances regardless of the precise speed or curvature of the trajectory, and this bias is due to a systematic error in the subject's perception, not action.

Dynamic Bayesian Observer Model

Past studies have attributed biases in path integration to leaky integration (Lappe et al., 2007, 2011; Mittelstaedt and Glasauer, 1991; Mittelstaedt and Mittelstaedt, 2001; Vickerstaff and Di Paolo, 2005). According to those behavioral models, subjects forget part of their movement history, leading to sub-additive accumulation of self-motion information. Consequently, they underestimate their distance moved and end up traveling further than necessary, overshooting the target. We asked whether the overshooting could instead result from accurate integration



Figure 2. Dynamic Bayesian Observer Model Subjects combine noisy sensory evidence from optic flow with prior expectations about self-motion speed to perform probabilistic inference over their movement velocity. The resulting noisy velocity estimates are integrated to generate beliefs about one's position. Bias in position estimation might come about from two extreme scenarios – a velocity prior that favors slower speeds coupled with perfect integration (green) or a uniform prior over velocity coupled with leaky integration (purple). For simplicity, this schematic shows the onedimensional case. For general planar motion, both linear and angular velocity must be inferred and integrated to update position in two dimensions.

of inaccurate, biased velocity estimates. Specifically, if subjects were to underestimate their linear and/or angular movement velocities, accurate integration might yet lead to overshooting. In fact, human subjects are known to underestimate retinal velocities, and those effects have been successfully attributed to a slow-speed prior using Bayesian theories (Hürlimann et al., 2002; Sotiropoulos et al., 2014; Stocker and Simoncelli, 2006; Weiss et al., 2002).

We hypothesized that such a slow-speed prior might also underlie the biases observed in our experiments. We tested this possibility against the alternative of leaky temporal integration using the framework of a dynamic Bayesian observer model. In this framework, we explicitly model the subject's belief, i.e., the subjective posterior distribution, which is the posterior over position given its model assumptions. This is computed across two stages: combining noisy optic flow input with a prior belief to compute the posterior over self-motion velocity (inference step) and integrating the resulting posterior with a constant leak rate (integration step). Because the position estimate is uncertain, we used this framework to identify model parameters that maximized the expected reward, a quantity that takes both the mean and uncertainty in position into account. Although we will shortly show that the above behavioral results can be understood purely in terms of a bias in subjects' mean position estimates, we will also show in a later section that uncertainty plays a pivotal role in influencing responses when navigating larger distances.

Because position is computed by integrating velocity, bias in position estimates can originate either from bias in velocity estimation or from imperfect integration. We modeled the distinction between these two hypotheses within the proposed framework by manipulating the shape of the prior to be exponential or uniform, and the nature of integration to be perfect or leaky (Figure 2). At one extreme, the combination of an exponential prior and perfect integrator would attribute path integration bias entirely to underestimation of self-motion velocity. At the other extreme, a uniform prior would yield unbiased velocity estimates, which, if integrated with leak, could also lead to a bias as proposed by other studies. We will refer to the above two instantiations as the "slow-speed prior" and the "leaky integrator" models, respectively. We assumed a Gaussian velocity likelihood whose variance scales linearly with the magnitude of measurement, as it yields a convenient mathematical form for the mean and variance of velocity estimates (STAR Methods; Equations 1.1 and 1.2). This choice of likelihood function does not conform to Weber's law. However, because the same parameterization was used for both models, this assumption does not intrinsically favor one model over another. We show later that our main conclusions also hold for alternative parametrizations and in particular also for log-normal likelihoods, which do follow Weber's law. Furthermore, we assumed that the noise in optic flow measurements is temporally uncorrelated so that the mean and variance of the integrated position estimates change at the same rate in both models (STAR Methods). Later, we relax this assumption to examine path integration bias for a more general class of integrated noise models.

Although both the slow-speed and the leaky integration models can lead to overshooting, they attribute the bias to two very different sources—velocity underestimation or leaky dynamics. For uniform motion in one dimension, this difference can be readily detected by observing how the subject's bias scales with distance: the bias due to a slow-speed prior will scale linearly, whereas leaky integration produces a sub-linear scaling, ultimately leading to saturating estimates of position. However, when velocity changes over time, distinguishing the models requires analyzing the subject's entire movement history rather than just comparing the pattern of bias in the stopping position. This framework allows us to incorporate measurements of the subject's time-varying velocities to fit and distinguish the models.

Because the task was performed on a two-dimensional ground plane, subjects had to infer and integrate two components of their velocity (linear and angular). We assumed the two velocity components were integrated by separate integrators with possibly different time constants (STAR Methods; Equations 2.1 and 2.2). Consequently, both models had four free parameters (STAR Methods): two likelihood widths to represent uncertainties in linear and angular velocity and either two exponents to represent priors for those same components (for the slow-speed prior model) or two time constants to represent rates of leak in integrating them (for the leaky integrator model). Additionally, we fit a two-parameter "null" model that attributed subjects' movements entirely to random variability, as well as a "full" model with six parameters that featured both exponential priors and leaky integrators.



Figure 3. Model Comparison and Validation

(A) Posterior probability distribution over position implied by the best-fit slowspeed prior (left, green) and leaky integrator (right, purple) models, swept over time during an example trial for the subject with the largest bias. The distributions at different time points were rescaled to the same height, so these plots reflect this subject's relative beliefs about his location across the duration of the trial. Target location (yellow dot) and the actual trajectory (black line) have been overlaid. Yellow ellipses depict an isoprobability contour (68% confidence interval) of the model posteriors over position at the end of the trial.

(B) Vector field of errors in the mean estimate of final position across trials for the two models. Error vectors of both models were rescaled to minimize overlap. The spatial profiles of the error magnitude (distance between target

Model Fitting and Comparison

For each subject, we fit the models using the sequences of velocities along each trajectory. The models infer and integrate these velocity inputs and, depending on their parameters, generate specific trajectory estimates. Trajectories of different models correspond to the subject's believed (rather than actual) positions during the trial. Our probabilistic framework assumes that subjects maintain estimates of both the mean and the uncertainty about their location and steer to the target to achieve the greatest possible reward. We therefore fit the models to maximize the subject's expected reward, defined as the overlap between the posterior distribution over their position and the target region at the end of each trial (STAR Methods; Equation 3).

We found that the slow-speed prior model was about 1.35 times more likely per trial than the leaky integrator model for each individual subject. Multiplying this ratio over all trials, this means that speed misperception from a slow-speed prior is an overwhelmingly more likely explanation of subjects' path integration biases than leaky integration (Figure S3A; mean [\pm SE] log-likelihood ratio across subjects, 66.6 \pm 18.2). Both models had substantially greater likelihoods than the null model, with larger improvements when biases were larger because the null model could not explain any bias (Figure S3B).

Because the evidence supporting both the slow-speed prior and leaky integration models was correlated, we asked whether subjects' behavior may have been influenced by both. To test this, we fit a model that incorporated both exponential prior and leaky integration. This full model was not much better at explaining subjects' responses than the slow-speed prior model (Figure S3C). Moreover, for all subjects, the best-fit time constants of integration in the full model were much greater than the average trial duration (Figure S3D), implying that integration was nearly perfect in this model. Therefore, leaky integration could not explain any appreciable variability in the data in excess of what was already explained by the slow-speed prior.

We wanted to know why the slow-speed prior model performed better. A good behavioral model will believe that subjects *should* stop moving where they *do* stop. This means that the model's beliefs about its position should be concentrated near the true target, even when the actual position has overshot. To evaluate this, we used the best-fit model parameters to reconstruct the subjects' beliefs, given by the posterior distribution over their position throughout each trial as they steered toward the target. Belief trajectories implied by the two models during an example trial are shown in Figure 3A. Because the model

and mean estimated final position) for the two models are shown by the grayscale background.

⁽C) Subject's internal estimates of the radial distance (\hat{r} , left) and angle ($\hat{\theta}$, right), not to be confused with the subject's actual response \bar{r} and $\bar{\theta}$ in Figure 1E, are plotted against target distances and angles for the subject in (A) and (B). Internal estimates implied by the slow-speed prior and leaky integrator models are shown in green and purple, respectively. Model estimates for each trial are shown as vertical bars centered on the mean and ±1 SD in length.

⁽D) Bias in model estimates (termed "residual bias") of radial distance and angle for the two models, obtained by cross-validation (STAR Methods). Error bars denote ±1 SEM obtained via bootstrapping. Dashed lines indicate unbiased radial or angular position estimates. Solid diagonal line has unit slope. See also Figure S3.

has a cloud of uncertainty over position, the plots actually show this cloud of beliefs swept out over time. This is overlaid with the subject's actual trajectory and the target position. On this trial, the beliefs implied by the slow-speed prior model (Figure 3A, left) terminated near the target, indicating that the subject strongly (and wrongly) believed he steered to the target location. On the other hand, the leaky integrator model believes it completely missed the target (Figure 3A, right), contradicting the basic premise that the subject is making a subjectively good decision.

This difference between the models' estimates of the final position was consistent across trials, as revealed by the much greater estimation error magnitudes for the leaky integrator model (Figure 3B, gray level). Moreover, unlike the slow-speed prior model, the vector field of errors in the estimates generated by the leaky integrator model was non-random (Figure 3B, arrows), betraying this model's inability to fully account for the subject's systematic errors. Note that the above model estimates are meant to reflect subjects' internal beliefs about their position (which should be nearly unbiased) rather than their actual positions (which we know are biased). So the near-zero estimation error magnitude of the slow-speed prior model (Figure 3B) suggests that this model is almost entirely able to explain away the subject's bias.

To assess the difference in the quality of fits of the two models, we compared the final position estimates generated by each of the two models against the target position (Figure 3C). This comparison is similar to the one used to evaluate subjects' responses (Figure 1E), except that we now replace the subject's actual position with the model estimates. We emphasize that we are modeling the subject's internal estimates, not predicting their external response. Therefore, the quality of the model estimates should be assessed by comparing against the unit slope, not by comparing with Figure 1E. For the example subject shown in Figure 3C, it can be readily seen that the estimates of the slow-speed prior model were in reasonably good agreement with target distances and angles. However, estimates generated by the leaky integrator model were still biased, and those biases were particularly large for nearby targets. Intuitively, this is because nearby targets only require short integrations, so the leak does not have time to take effect. Consequently, the leaky integrator model is objectively accurate at short times and thus cannot account for the subjective biases in those trials, leading to a relatively poor fit. On the other hand, the slow-speed prior model attributes path integration bias to velocity underestimation, a bias that persists at all times and thus generalizes well across trials.

We quantified the goodness of fit of the models by computing residual biases in the model estimates of radial distance and angle using a 4-fold cross-validation procedure (STAR Methods). These residual biases were not significantly different from unity across subjects for the slow-speed prior model (mean \pm SEM residual radial bias = 1.03 ± 0.04 , p = 0.27, t test; residual angular bias = 1.01 ± 0.1 , p = 0.36). On the other hand, residual biases of the leaky integrator model were significantly greater than unity (residual radial bias = 1.09 ± 0.06 , p = 3.2×10^{-2} ; residual angular bias = 1.31 ± 0.14 , p = 3.4×10^{-3}). Therefore, the slow-speed prior model provided a much better account of subjects' biases (Figure 3D).

Recent studies on path integration have modeled leak using space constants instead of time constants, so that the integration dynamics are only active during movement. We formulated a variation of the model based on space constants (STAR Methods) and found that it still performed worse in predicting subjects' responses than the slow-speed prior model (Figure S3E). This is not surprising because spatial leak suffers from the same problem responsible for the relatively poor performance of the model with temporal leak. Finally, to verify consistency with the family of Bayesian models previously used in psychophysics, we considered other parameterizations of the slow-speed prior and leaky integrator models that incorporate a Gaussian prior (Hürlimann et al., 2002) or a log-normal likelihood (Jürgens and Becker, 2006; Stocker and Simoncelli, 2006), and we fit each alternative (STAR Methods). On average, the slow-speed prior model continued to outperform the leaky integrator model in all cases (Figure S3F), indicating that our conclusions are robust to the specific choices of shapes of the prior and likelihood.

Test of Model Predictions

The likelihood comparison above clearly favors attributing the path integration bias to a slow-speed prior over leaky integration of velocity estimates. This makes new predictions, which we tested experimentally by manipulating parameters of the task. One manipulation involved changing the reliability of optic flow by varying the density of the ground plane elements between two possible values ("sparse" and "dense"). A hallmark of Bayesian inference is that, for unimodal non-uniform priors and symmetric likelihood functions, the bias increases for less reliable observations. Therefore, if subjects had a slow-speed prior, sparse optic flow would increase how much they underestimate their velocity, leading to a larger bias (Figure 4A). However, if the prior is uniform, the density of optic flow would merely affect subjects' uncertainty about their speeds while the instantaneous optic flow estimates themselves would still be unbiased under both conditions. The leaky integrator model thus predicts that changing the texture density would leave position bias unaffected.

The performance of an example subject is shown in Figure 4B. Sparsifying optic flow had a detrimental effect on behavior as indicated by a steeper relationship between true and perceived distance moved as well as angle rotated. As before, we quantified the bias as the slope of this regression and found similar effects across subjects (Figures 4C and S4A). Consistent with the prediction of the slow-speed prior model, decreasing the density led to a significantly greater bias both in distance moved (mean ± SEM radial bias, Γ_r , high density, 1.27 ± 0.1 ; low density, 1.46 ± 0.1 ; p = 2.5×10^{-2} ; paired t test) and in angle rotated (mean angular bias, Γ_{θ} , high density, 1.58 ± 0.1 ; low density, 2.13 ± 0.1 ; p = 9.1×10^{-4}).

In a second manipulation, we imposed two different speed limits ("slow" and "fast") on different trials, which we implemented by randomly switching the gain by which the joystick controlled velocity. To avoid inducing different effects on biases in distance and angle, both linear and angular velocities were scaled by the same gain factor (STAR Methods). Because the leaky integrator model incorporates a uniform prior, subjects' estimates of speeds will always be unbiased in this model.



Figure 4. Test of Model Predictions

(A) Reliability of optic flow was manipulated by altering the density of ground-plane elements. Decreasing the density will increase subjects' bias only if they have a slow-speed prior.

(B) Scatterplots showing the effect of density on radial and angular bias of one subject. Dashed line represents unity slope (unbiased performance), and solid lines represent slopes of regression fits. Trials are colored according to density: red, high-density trials; blue, low-density trials.

(C) Effect of density manipulation on radial (left) and angular (right) biases of individual subjects. Asterisks denote significant difference between means (see text). (D) Subjects' speed limit was manipulated by altering the gain of the joystick. The leaky integrator model predicts that subjects' biases will be reduced in the high-speed condition.

(E) Scatterplots showing the effect of speed on distance and angle bias of one subject. Trials are colored according to speed: red, high-speed trials; blue, low-speed trials.

(F) Speed manipulation does not affect subjects' biases in a systematic way. n.s., not significant (see text). See also Figure S4.

However, a fundamental feature of this model is that the integration error accumulates over time, so lowering the speed limit should lead to a larger positional bias due to increased travel time (Figure 4D). On the other hand, for a Gaussian likelihood whose variance scales linearly with speed, an exponential slow-speed prior predicts that the velocity would be underestimated by the same multiplicative factor at all velocities. Therefore, the slow-speed prior model predicts that subjects will accurately perceive the relative change in their speeds and thus be biased to the same extent under both conditions. Note that this latter prediction strictly holds only for model parametrizations that produce a multiplicative bias in velocity estimates and may not be applicable to alternative formulations of the model. However, the prediction of the leaky integrator model's speed dependence does not depend on the form of the velocity likelihood and can therefore be unambiguously tested.

We analyzed subjects' biases and found that their performance was, on average, unaffected by the speed manipulation (Figures 4E, 4F, and S4B) both for distance (Γ_r , high speed, 1.33 ± 0.1; low speed, 1.38 ± 0.1; p = 0.59, paired t test) and angle (Γ_{θ} , high speed, 1.92 ± 0.1; low speed, 1.72 ± 0.1; p = 0.15). This result is consistent with our preliminary observation of a velocity-independent distance bias (Figure S1C). This once again argues against the leaky-integrator model and favors, in particular, a slow-speed prior model that can induce multiplicative biases in velocity estimates.

Distance-Dependent Bias Reversal

Because subjects compute their position by integrating noisy velocity estimates, their position estimates are uncertain. When traveling modest distances, such as those tested in the above experiments, the integrated uncertainty in position is relatively small. Although we took this uncertainty into account, we could qualitatively explain overshooting solely in terms of a bias in the subject's mean position estimates resulting from integrating biased velocity estimates. In this section, we show that, when path integrating over larger distances, the influence of position uncertainty can produce a reversal in the pattern of bias—from overshooting to undershooting—and we provide experimental evidence for this phenomenon.

Recall that the proposed framework assumes that subjects incorporate their knowledge of position uncertainty by tracking the expected reward of stopping at a given location. When this expected reward reaches its maximum, they stop moving. At any given moment during the trial, the expected reward is given by the overlap of the probability distribution over their position with the target. Let us examine how it should change as a function of their position by considering uniform motion in one dimension for clarity. Subjects integrate both the mean (signal) and random fluctuations ("noise") in their velocity estimates. If integration is leak free, then their uncertainty in position would gradually keep building up over time (STAR Methods; Equation 4). The rate at which position uncertainty builds up depends on the nature of sensory noise (independent or temporally correlated) as well as ability to represent and integrate large uncertainties. In principle, drift-diffusion models can be used to model noisy integration (Simen et al., 2011; Thurley, 2016). However, constructing such mechanistic models would entail making uninformed assumptions about the noise process. Here, rather than positing a particular mechanism, we choose a phenomenological model for position uncertainty, assuming that the standard deviation σ of the position distribution grows as a power-law function of time t, as $\sigma(t) \propto t^{\lambda}$. For uniform motion, this can also be expressed as a distance-dependent scaling of the width with the same power-law exponent so that $\sigma(r) \propto r^{\lambda}$ for distance r (Figure 5A). A scaling exponent of $\lambda = 0.5$ (Wiener process) would result from integrating velocity estimates with independent Gaussian noise. Other types of noise may yield smaller (sub-diffusion) or larger (super-diffusion) exponents, depending on whether variance in the position estimate (σ^2) scales faster or slower than the mean.

We analyzed how the expected reward should qualitatively depend on distance for a range of exponents. Intuitively, one would expect it to be greatest when the probability distribution over position is centered on the target. However, this is not always true. Figure 5B shows how the expected reward evolves with distance for near and far targets for one example case ($\lambda = 1.5$). When steering to nearby targets, the built up uncertainty is relatively small, so the expected reward is indeed greatest when the mean of the distribution over distance moved roughly matches the target distance. For faraway targets, however, the expected reward actually peaks *before* reaching the target. This happens because, if the subject moves beyond that optimal distance, the probability distribution over their position becomes so wide that its overlap with the target begins to

decrease. Therefore, when steering toward sufficiently distant targets, an ideal observer should stop short of the target (Figure 5C).

The precise extent of undershooting depends on the noise process, with larger exponents producing greater undershooting due to a faster buildup in uncertainty (Figure 5D, top left). Furthermore, for exponents larger than one, the tendency to undershoot grows stronger with distance. Note that the degree of undershooting does not depend on the uncertainty in estimating the target position but rather on the change in uncertainty over distance. Any uncertainty about the target would reduce the magnitude of expected reward, but not the optimal distance at which maximum reward is achieved. Thus, potentially, large biases in path integration can stem solely from a subject hedging their bets against increasingly uncertain position estimateseven when those estimates are unbiased. We have already demonstrated that velocity, and consequently the distance moved, is likely underestimated due to a slow-speed prior (Figure 5D, bottom left). The two factors will have opposing effects on bias, with potentially different spatial dependences: whereas the slow-speed prior causes overshooting through a perceptual bias that scales linearly with distance, growing uncertainty does not alter the perceptual bias but generates an increasing tendency for responses to undershoot. This undershooting can increase linearly or supra-linearly depending on whether uncertainty scales slower ($\lambda < 1$) or faster ($\lambda > 1$) than Weber's law. The combined effect of the two factors is shown in Figure 5D (right). For sub-Weber-law scaling in uncertainty, bias will increase linearly with distance, consistently producing either overshooting or undershooting depending on the relative strength of the two effects. For scaling exponents larger than one, the different spatial scaling from the slow-speed prior and from growing positional uncertainty leads to a rather surprising prediction: when position uncertainty grows faster than the mean, bias in the subjects' responses should gradually reverse from overshooting to undershooting when navigating to increasingly distant targets.

The above prediction also holds for motion in two dimensions. In this case, both linear and angular components of motion are subject to the effects of growing uncertainty and may eventually lead to undershooting both in radial as well as angular responses. To test whether there is such a bias reversal, we conducted an additional experiment in which we asked subjects to steer to targets that were much further away. Target locations were discretized, and their distances were varied on a logarithmic scale (STAR Methods). Because the limited viewing angle in our setup restricted the angular eccentricity of the targets, we did not test for bias reversal in the angular domain. Similar to our original experiment, subjects continued to exhibit a significant angular bias ($\Gamma_{\theta} = 2.32 \pm 0.6$; $p = 1.6 \times 10^{-4}$; t test; Figure S5A), turning much more than required.

On the other hand, the pattern of radial bias was strikingly consistent with our prediction. Figure 5E shows how the radial distance of an example subject scaled with target distance in this task. The subject exhibited overshooting in trials with nearby targets (2, 4, 8, and 16 m), as was observed in the original task, but this pattern of bias was replaced by significant undershooting for the farthest targets (32 m). Note that, when steering



Figure 5. Model Explains Bias Reversal with Distance

(A) The width of the subjects' probability distribution over their position (black) is modeled as a power law with exponent λ . The overlap (gray shade) of the probability distribution with the target (orange) corresponds to the subjects' expected reward.

(B) Evolution of subjects' expected reward when steering to a nearby (left) and distant (right) target for and proportionality constant equal to one. Insets show probability distributions over position at three different locations indicated by solid circles of the corresponding color on the reward curve. The peaks of the reward curves correspond to the optimal response distance. Orange bars denote the width of the target and dashed vertical lines the target center. (C) The optimal response distance as a function of target distance for the above case.

(D) The effect of position uncertainty (top) and the effect of slow-speed prior (bottom) combine to determine the model prediction for path integration bias, shown for various values of the power-law exponent (right). The interaction scales the optimal response distance by the slope Γ of the relation between actual and perceived distance moved.

(E) Mean net distance moved by one subject in response to targets at five different distances. Gray solid line corresponds to the best-fit model.

(F) Gray circles denote mean responses of individual subjects. Black line corresponds to the subject-averaged response.

(G) Mean response of one subject under conditions of low-density (blue) and high-density (red) optic flow. Asterisks denote a significant difference between mean responses under the two conditions (2 m, p = 0.029; 4 m, p = 0.007; 32 m, $p = 4.1 \times 10^{-4}$; paired t test).

(H) Mean responses of individual subjects under the two conditions. Asterisks denote a significant difference between mean responses (across subjects) under the two conditions (2 m, p = 0.035; 32 m, p = 0.013; paired t test). Solid lines correspond to subject-averaged response.

In (C)–(H), black dashed lines have unit slope; in (E) and (G), error bars denote SEM across trials. See also Figure S5.

to distant targets, the effect of the slow-speed prior would still persist, but its effect is outweighed by that of increasing positional uncertainty. To quantify the relative strength of the two effects, we simultaneously fit a multiplicative constant Γ and exponent λ to the subject's data (STAR Methods; Equation 5). The multiplicative constant captures the linear effect of velocity underestimation that causes overshooting, whereas the exponent reveals the rate of scaling of uncertainty with distance that causes undershooting to faraway targets. Both parameters must be greater than unity in order to produce a reversal from overshooting to undershooting. This was indeed the case for this subject (Figure 5E, gray curve; $\Gamma = 2.2$; $\lambda = 2.4$). A similar

pattern of bias reversal was observed across subjects (Figure S5B; $\Gamma = 1.5 \pm 0.2$, $p = 3.6 \times 10^{-5}$, t test; $\lambda = 1.8 \pm 0.4$, $p = 8 \times 10^{-5}$) and can be noticed in the subject-averaged responses (Figure 5F). The supra-Weber-law scaling in subjects' uncertainty can also be explicitly seen in the manner by which SDs in response scaled with the mean distance moved (Figure S5C).

The undershooting observed for distant targets could simply have been due to motor fatigue. To test whether the bias was influenced by sensory uncertainty, we re-analyzed our data by dividing the trials into two groups based on the density of optic flow cues. If sensory uncertainty contributes to undershooting, decreasing the reliability of sensory cues should cause greater undershooting. The behavior of an example subject shown in Figure 5G confirms this assertion. When steering to the farthest target, this subject covered significantly less distance when the density of optic flow was reduced. Note that, for nearby targets, the effect is reversed because the influence of the slow-speed prior is stronger than that of position uncertainty. These effects of density manipulation were observed across subjects (Figures 5H and S5D), suggesting that subjects' bias is likely not due to fatigue but rather governed by sensory uncertainty.

The bias-reversal reported here should not be confused with an experience-based regression effect in which subjects' responses are biased toward the mean stimulus (Jazayeri and Shadlen, 2010; Petzschner and Glasauer, 2011). Specifically, due to noise in the measurement of the target location, subjects' perceptual estimates of that location may be biased toward the mean of the distribution of targets across trials; in turn, this could manifest in path integration behavior that is similarly biased. However, the pattern of bias shown in Figure 5 was present from the beginning of the experiment (Figure S5E), suggesting that this bias is not likely due to effects of adaptation. Furthermore, to ensure that manipulating ground-plane density does not influence judgment of target location, we conducted a separate experiment in which we manipulated ground-plane density only after the target was turned off so that it selectively affects reliability of optic flow. Once again, we observed that bias reversal was sensitive to this manipulation (Figures S5F and S5G), implying that it is induced by the process of path integration rather than a bias in target perception. While regression to the mean may still play a role in this task, our results suggest that prior expectations about self-motion velocity and uncertainty in position due to accumulated uncertainty about optic flow have a dominant influence.

DISCUSSION

We have presented a unified framework that combines Bayesian inference, evidence integration, and the principle of utility maximization to explain human behavior in a naturalistic navigation task. This framework yields a parsimonious account of bias in visually guided path integration, in which bias stems from prior expectations and sensory noise associated with self-motion. Our claim is based on four primary findings. First, when navigating modest distances using optic flow, humans overshoot the goal location, implying that they underestimated both their net translation and rotation. Second, analysis of subjects' movement trajectories using a dynamic observer model revealed that their bias was more likely to originate from a slow-speed prior rather than forgetful integration of self-motion. Third, experimental outcomes of manipulating the reliability of self-motion cues and speed confirmed the predictions of the slow-speed prior model. Finally, when navigating long distances, the model predicts a possible reversal in the direction of bias due to the growing influence of uncertainty on the expected reward, a phenomenon that was confirmed experimentally.

In order to study visual path integration, we used virtual reality so that vestibular and proprioceptive inputs were no longer relevant for the task. Specifically, subjects used a joystick to steer to a cued target location based solely on optic flow. To perform accurately on this task, participants had to determine the location of the target, remember that location, and integrate their own movements until they reached that location. Each of those steps is a potential source of behavioral errors.

There are several reasons why systematic errors seen in our data cannot be attributed to biased perception of the initial target location. First, we used stereoscopic stimuli to generate an immersive virtual environment with depth cues that facilitate judgement of target distances, at least for short distances. Second, judging target angles is more straightforward and does not require depth cues, yet subjects exhibited a large angular bias in the task. Notably, introducing angular landmarks in the virtual environment abolished this angular bias. The landmarks themselves were uninformative about target angles but helped obviate the need to integrate angular velocity by providing a direct estimate of the subject's orientation in the virtual environment. Thus the large angular biases seen in the absence of landmarks must be related to the perception of optic flow cues. Finally, and perhaps most importantly, manipulating the density of the ground plane significantly altered subjects' biases at all distance scales. This effect persisted even when we confined the above manipulation to the period of movement, leaving target distance estimation relatively intact. Therefore, the behavioral bias seen in our task likely reflects error in estimating one's own position rather than difficulties associated with estimating the target location.

Relation to Past Experiments

Past studies on visual path integration employed visually simulated motion along a straight line or along predetermined curvilinear trajectories. In contrast, our experimental task allowed subjects to actively steer using two degrees of freedom, allowing for precise control of their self-motion velocity at all times, as would be the case during natural foraging. This design was motivated by the need to engage neural mechanisms and computations that likely underlie path integration in the real world. Yet our results regarding biases are qualitatively similar to those of previous studies, even though those studies tested motion along a one-dimensional hallway. Specifically, studies that tested visual path integration over short distances found that subjects overshoot the target (Frenz and Lappe, 2005; McManus et al., 2017), whereas studies that used long-range targets found the opposite (Brossard et al., 2016; Lappe et al., 2007; Redlick et al., 2001). A noteworthy departure from past studies was observed in the manner by which variability of our subjects' responses scaled with distance. Like sensory perception, odometric ability has been shown to follow Weber's law in humans (Durgin et al., 2009; Petzschner and Glasauer, 2011) and invertebrates (Cheng et al., 1999, 2006). In contrast, we observed supra-Weber-law scaling. Given that the overall magnitude of response variability was also larger in our study, we suspect that the use of sparse optic flow in our experiments-needed to probe the internal model of subjects' self-motion velocityunderlies the above discrepancy. If this is the case, then it suggests that Weber's law might break down in the regime of high uncertainty.

Relation to Past Models

To explain our subjects' behavior, we tested two different instantiations of a dynamic Bayesian observer model and found that bias in path integration appears to stem mainly from a slowspeed prior that causes subjects to underestimate their velocity. Unlike a prior over retinal speed (Hürlimann et al., 2002; Sotiropoulos et al., 2014; Stocker and Simoncelli, 2006; Weiss et al., 2002), the prior in our Bayesian model corresponds to subjects' prior expectation of their self-motion velocity. Nonetheless, the latter might be inherited from low-level visual sensory priors that govern human perception of local image velocities. Alternatively, the prior over self-motion velocity could reflect the statistics of sensory inputs experienced during natural self-motion, which is known to be biased toward slower velocities (Carriot et al., 2014). In this case, one would expect the prior to be shared across all modalities. Our data qualitatively agree with this view for the prior over angular velocities, because angular bias persisted when estimates were based on efference copy alone, although the variability was prohibitively large to allow for a proper quantitative assessment (Figure S1B). Regardless of its specific origins, this work demonstrates that sensory priors can have tangible consequences for complex dynamic behaviors, such as path integration, well beyond the realm of traditional binary decision-making tasks. Although we focused on visual self-motion, this model is also applicable to other modalities. Availability of additional modalities should diminish the effect of the prior, leading to reduced bias. Such a reduction has in fact been observed when path integrating using multimodal cues (Bakker et al., 1999; Becker et al., 2002; Sun et al., 2004).

While the slow-speed prior can explain why subjects would travel beyond the goal, it cannot account for undershooting reported in previous studies that used distant goals (Brossard et al., 2016; Lappe et al., 2007; Redlick et al., 2001). However, analysis of our model revealed that, when path integrating over longer distances, the effect of growing uncertainty can eventually override the effect of perceptual bias induced by prior expectations and cause undershooting in subjects' responses. This is a spatial analog of a model that explains early abandonment on a waiting task as a rational response to increasing uncertainty about the next reward (McGuire and Kable, 2013). We tested this prediction and found that the pattern of bias changed from overshooting to undershooting when navigating to increasingly distant targets. This phenomenon of bias-reversal is also discernable in the results of previous visual (Brossard et al., 2016; Lappe et al., 2007; Redlick et al., 2001) and non-visual (Bergmann et al., 2011; Lappe and Frenz, 2009) path integration studies. Traditional leaky integration models cannot explain why subjects would undershoot. To account for undershooting, such models have had to be modified to update distance to target rather than distance moved (Frenz and Lappe, 2005). However, such a change of variable neither explains why subjects overshoot to relatively nearby goals nor why the degree of undershooting is sensitive to the reliability of optic flow. Here, we show that a distance-dependent reversal in the response bias naturally emerges when performing probabilistic inference over position under the influence of a slow-speed prior to maximize expected reward. Although velocity estimates are more reliable at smaller speeds, moving slowly would entail integrating over longer times, and consequently, subjects would not benefit from such a strategy (Figure 4F). In fact, longer integration times would result in an even greater accumulated uncertainty due to the supra-Weber-law scaling effect (Figure S5C). Perhaps this explains why subjects tended to travel close to the speed limit in our task (Figure S1C).

Recent models based on iterative Bayesian estimation suggest that subjects may exploit trial history to update an explicit prior over net distances and angles turned (Petzschner and Glasauer, 2011; Prsa et al., 2015). While such models can explain responses that exhibit a regression toward the mean of previously experienced stimuli, they cannot account for the unidirectional response bias observed in many studies, including our own (Figure 1E). Moreover, such models have limited potential for explaining systematic biases seen in the first few trials of the experiment when the subjects have not yet accurately learned the environmental statistics.

Other models attribute bias in path integration primarily to either a path-dependent (Bergmann et al., 2011; Brossard et al., 2016; Lappe et al., 2007, 2011) or temporal (Mittelstaedt and Glasauer, 1991; Mittelstaedt and Mittelstaedt, 2001; Vickerstaff and Di Paolo, 2005) decay in integrating self-motion. However, leaky integration cannot explain the effect of reliability of optic flow cues reported here. Moreover, it is worth noting that, in addition to the leak factor, a recent model of leaky-integration incorporated a gain factor that rescaled subjects' displacement in each step. The best-fit gain factors were generally less than unity (Bergmann et al., 2011; Brossard et al., 2016; Lappe et al., 2007, 2011), which essentially amounts to velocity underestimation. Coupled with small leak rates (~0.01–0.02 m⁻¹) found in those studies, it is clear that the performance of that model for short-range targets is in fact dominated by velocity underestimation.

Although the precise neural circuit underlying path integration has not been worked out, there is physiological evidence for near-perfect integration of visual motion cues by neurons in macaques (Huk and Shadlen, 2005; Kiani et al., 2008), suggesting that our model is neurally plausible. Our work is also supported by recent studies showing lossless evidence accumulation of temporally disjoint sensory inputs in rats (Brunton et al., 2013; Raposo et al., 2014), humans (Brunton et al., 2013; Kiani et al., 2013), and monkeys (Huk and Shadlen, 2005; Kiani et al., 2008) performing binary decision tasks. Subjects may benefit from imperfect integration when the statistical structure of sensory inputs is unpredictable (Carland et al., 2016; Glaze et al., 2015) or when signal strength fluctuates wildly (Ossmy et al., 2013; Veliz-Cuba et al., 2016). However, when sensory dynamics are known a priori or are predictable from physical laws, it makes sense that behavior is limited by sensory inputs rather than leaky integration. One limitation of this work is that it is based solely on the principle of probabilistic perceptual inference and ignores the costs incurred in performing actions. Because navigation is effortful, future extensions should test whether the subjects also optimize their actions at finer timescales to minimize the total cost during goal-oriented navigation.

Toward Understanding the Neural Basis

Disentangling the neural representation of sensory inputs from internal beliefs is notoriously difficult. Binary decision tasks and other paradigms measuring discrete responses are ill suited for this purpose, as behavioral readouts are often sparse in time and may conflate belief states and actions. The task and the modeling approach used in this study offer an ideal setting to overcome such limitations to investigate "dynamic belief state representations" (Lee et al., 2014) in a widely studied neural circuit. Specifically, theoretical models suggest that grid cells may be involved in path integration (Burak and Fiete, 2009; Giocomo et al., 2011), yet experimental evidence remains indirect (Chen et al., 2015; Gil et al., 2018) and controversial (Bjerknes et al., 2018). By properly modeling the beliefs of animals trained to path integrate, one can distinguish the spatiotemporal dynamics of their belief states from actual states (e.g., Figure 3A, left) and then look for the neural correlates of those states in the grid cell network.

Beyond characterizing neural representations, the approach outlined here would also allow one to directly test the optimality of the underlying neural computations. For instance, brain areas encoding self-motion velocity have recently been identified in rodents (Kropff et al., 2015; Whitlock et al., 2012; Wilber et al., 2017) and primates (Avila et al., 2017). By simultaneously recording the activity of areas representing velocity and internal estimates of position, one can test whether brain-wide neural circuits can optimally integrate dynamic, non-stationary sensory inputs, which would help unravel how the brain works in the real world.

STAR***METHODS**

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

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT AND RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
 - Behavioral Task
 - Stimulus and Data acquisition
 - Dynamic Bayesian observer model
 - Model fitting
 - Modeling position uncertainty
 - Fitting distance-dependent bias reversal
- QUANTIFICATION AND STATISTICAL ANALYSIS
 - Estimation of Bias
 - Model comparison and validation
 - Test of model predictions
 - Quantifying bias-reversal
- DATA AND SOFTWARE AVAILABILITY

SUPPLEMENTAL INFORMATION

Supplemental Information includes five figures and two videos and can be found with this article online at https://doi.org/10.1016/j.neuron.2018.05.040.

ACKNOWLEDGMENTS

We thank Paul Schrater for sharing useful insights on behavioral modeling, Erin Neyhart for assisting with the experiments, and Jing Lin and Jian Chen for their help in programming the stimulus. This work was supported by the Simons Collaboration on the Global Brain, grant no. 324143, and NIH DC007620. G.C.D. was supported by NIH EY016178.

AUTHOR CONTRIBUTIONS

Conceptualization, K.J.L., G.C.D., X.P., and D.E.A.; Methodology, K.J.L., X.P., and D.E.A.; Investigation, K.J.L., M.P., and H.P.; Formal Analysis, K.J.L. and M.P.; Writing – Original Draft, K.J.L.; Writing – Review & Editing, K.J.L., M.P., 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: October 30, 2017 Revised: March 23, 2018 Accepted: May 30, 2018 Published: June 21, 2018

REFERENCES

Alyan, S., and Jander, R. (1994). Short range homing in the house mouse, *Mus Musculus*: stages in the learning of directions. Anim. Behav. 48, 285–298.

Avila, E., Lakshminarasimhan, K.J., DeAngelis, G.C., and Angelaki, D.E. (2017). Speed tuning in the macaque posterior parietal cortex. bioRxiv. https://doi. org/10.1101/204966.

Bakker, N.H., Werkhoven, P., and Passenier, P.O. (1999). The effects of proprioceptive and visual feedback on geographical orientation in virtual environments. Presence *8*, 36–53.

Bakker, N.H., Werkhoven, P.J., and Passenier, P.O. (2001). Calibrating visual path integration in VEs. Presence: Teleoperators Virtual Environ. 10, 216–224.

Becker, W., Nasios, G., Raab, S., and Jürgens, R. (2002). Fusion of vestibular and podokinesthetic information during self-turning towards instructed targets. Exp. Brain Res. *144*, 458–474.

Benhamou, S. (1997). Path integration by swimming rats. Anim. Behav. 54, 321–327.

Bergmann, J., Krauss, E., Münch, A., Jungmann, R., Oberfeld, D., and Hecht, H. (2011). Locomotor and verbal distance judgments in action and vista space. Exp. Brain Res. *210*, 13–23.

Bjerknes, T.L., Dagslott, N.C., Moser, E.I., and Moser, M.-B. (2018). Path integration in place cells of developing rats. Proc. Natl. Acad. Sci. USA *115*, E1637–E1646.

Bogacz, R. (2007). Optimal decision-making theories: linking neurobiology with behaviour. Trends Cogn. Sci. *11*, 118–125.

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.

Brossard, M., Lepecq, J.-C., and Mestre, D. (2016). Oscillatory optical flows improve the perception of travelled distance in static observers. J. Vis. *16*, 768.

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.

Burak, Y., and Fiete, I.R. (2009). Accurate path integration in continuous attractor network models of grid cells. PLoS Comput. Biol. *5*, e1000291.

Carland, M.A., Marcos, E., Thura, D., and Cisek, P. (2016). Evidence against perfect integration of sensory information during perceptual decision making. J. Neurophysiol. *115*, 915–930.

Carriot, J., Jamali, M., Chacron, M.J., and Cullen, K.E. (2014). Statistics of the vestibular input experienced during natural self-motion: implications for neural processing. J. Neurosci. *34*, 8347–8357.

Chen, X., He, Q., Kelly, J.W., Fiete, I.R., and McNamara, T.P. (2015). Bias in human path integration is predicted by properties of grid cells. Curr. Biol. *25*, 1771–1776.

Cheng, K., Srinivasan, M.V., and Zhang, S.W. (1999). Error is proportional to distance measured by honeybees: Weber's law in the odometer. Anim. Cogn. *2*, 11–16.

Cheng, K., Narendra, A., and Wehner, R. (2006). Behavioral ecology of odometric memories in desert ants: acquisition, retention, and integration. Behav. Ecol. *17*, 227–235.

Collett, T.S., and Collett, M. (2000). Path integration in insects. Curr. Opin. Neurobiol. *10*, 757–762.

Durgin, F.H., Akagi, M., Gallistel, C.R., and Haiken, W. (2009). The precision of locomotor odometry in humans. Exp. Brain Res. *193*, 429–436.

Ernst, U.A., Mandon, S., Schinkel-Bielefeld, N., Neitzel, S.D., Kreiter, A.K., and Pawelzik, K.R. (2012). Optimality of human contour integration. PLoS Comput. Biol. 8, e1002520.

Etienne, A.S. (1992). Navigation of a small mammal by dead reckoning and local cues. Curr. Dir. Psychol. Sci. 1, 48–52.

Fiser, J., Berkes, P., Orbán, G., and Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. Trends Cogn. Sci. *14*, 119–130.

Frenz, H., and Lappe, M. (2005). Absolute travel distance from optic flow. Vision Res. 45, 1679–1692.

Gil, M., Ancau, M., Schlesiger, M.I., Neitz, A., Allen, K., De Marco, R.J., and Monyer, H. (2018). Impaired path integration in mice with disrupted grid cell firing. Nat. Neurosci. *21*, 81–91.

Giocomo, L.M., Moser, M.B., and Moser, E.I. (2011). Computational models of grid cells. Neuron 71, 589–603.

Glaze, C.M., Kable, J.W., and Gold, J.I. (2015). Normative evidence accumulation in unpredictable environments. eLife 4, e08825.

Loomis, J.M., Klatzky, R.L., Golledge, R.G., and Philbeck, J.W. (1999). Human navigation by path integration. In Wayfinding Behavior: Cognitive Mapping and Other Spatial Processes, R.G. Golledge, ed. (JHU Press), pp. 125–151.

Huk, A.C., and Shadlen, M.N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J. Neurosci. 25, 10420–10436.

Hürlimann, F., Kiper, D.C., and Carandini, M. (2002). Testing the Bayesian model of perceived speed. Vision Res. 42, 2253–2257.

Issen, L., Huxlin, K.R., and Knill, D. (2015). Spatial integration of optic flow information in direction of heading judgments. J. Vis. 15, 14.

Jazayeri, M., and Shadlen, M.N. (2010). Temporal context calibrates interval timing. Nat. Neurosci. *13*, 1020–1026.

Jogan, M., and Stocker, A.A. (2013). Optimal integration of visual speed across different spatiotemporal frequency channels. Adv. Neural Inf. Process. Syst. *26*, 3201–3209.

Jürgens, R., and Becker, W. (2006). Perception of angular displacement without landmarks: evidence for Bayesian fusion of vestibular, optokinetic, podokinesthetic, and cognitive information. Exp. Brain Res. *174*, 528–543.

Jürgens, R., Boss, T., and Becker, W. (1999). Estimation of self-turning in the dark: comparison between active and passive rotation. Exp. Brain Res. *128*, 491–504.

Kiani, R., Hanks, T.D., and Shadlen, M.N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. J. Neurosci. 28, 3017–3029.

Kiani, R., Churchland, A.K., and Shadlen, M.N. (2013). Integration of direction cues is invariant to the temporal gap between them. J. Neurosci. *33*, 16483–16489.

Klatzky, R.L., Loomis, J.M., Golledge, R.G., Cicinelli, J.G., Doherty, S., and Pellegrino, J.W. (1990). Acquisition of route and survey knowledge in the absence of vision. J. Mot. Behav. *22*, 19–43.

Kropff, E., Carmichael, J.E., Moser, M.B., and Moser, E.I. (2015). Speed cells in the medial entorhinal cortex. Nature 523, 419–424.

Lappe, M., and Frenz, H. (2009). Visual estimation of travel distance during walking. Exp. Brain Res. *199*, 369–375.

Lappe, M., Jenkin, M., and Harris, L.R. (2007). Travel distance estimation from visual motion by leaky path integration. Exp. Brain Res. *180*, 35–48.

Lappe, M., Stiels, M., Frenz, H., and Loomis, J.M. (2011). Keeping track of the distance from home by leaky integration along veering paths. Exp. Brain Res. *212*, 81–89.

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

Loomis, J.M., Klatzky, R.L., Golledge, R.G., Cicinelli, J.G., Pellegrino, J.W., and Fry, P.A. (1993). Nonvisual navigation by blind and sighted: assessment of path integration ability. J. Exp. Psychol. Gen. *122*, 73–91.

McGuire, J.T., and Kable, J.W. (2013). Rational temporal predictions can underlie apparent failures to delay gratification. Psychol. Rev. *120*, 395–410.

McManus, M., D'Amour, S., and Harris, L.R. (2017). Using optic flow in the far peripheral field. J. Vis. *17*, 3.

Mittelstaedt, M.L., and Glasauer, S. (1991). Idiothetic navigation in gerbils and humans. Zool. Jahrbucher-Abteilung Fur Allg. Zool. und Physiol. der Tiere *95*, 427–435.

Mittelstaedt, M.L., and Mittelstaedt, H. (2001). Idiothetic navigation in humans: estimation of path length. Exp. Brain Res. *139*, 318–332.

Mulligan, J.B., Stevenson, S.B., and Cormack, L.K. (2013). Reflexive and voluntary control of smooth eye movements. Proc. SPIE Human Vision and Electronic Imaging XVIII *8651*, 86510Z.

Ossmy, O., Moran, R., Pfeffer, T., Tsetsos, K., Usher, M., and Donner, T.H. (2013). The timescale of perceptual evidence integration can be adapted to the environment. Curr. Biol. *23*, 981–986.

Petzschner, F.H., and Glasauer, S. (2011). Iterative Bayesian estimation as an explanation for range and regression effects: a study on human path integration. J. Neurosci. *31*, 17220–17229.

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

Prsa, M., Jimenez-Rezende, D., and Blanke, O. (2015). Inference of perceptual priors from path dynamics of passive self-motion. J. Neurophysiol. *113*, 1400–1413.

Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. Nat. Neurosci. *17*, 1784–1792.

Redlick, F.P., Jenkin, M., and Harris, L.R. (2001). Humans can use optic flow to estimate distance of travel. Vision Res. 41, 213–219.

Seguinot, V., Cattet, J., and Benhamou, S. (1998). Path integration in dogs. Anim. Behav. 55, 787–797.

Simen, P., Balci, F., de Souza, L., Cohen, J.D., and Holmes, P. (2011). A model of interval timing by neural integration. J. Neurosci. *31*, 9238–9253.

Sotiropoulos, G., Seitz, A.R., and Seriès, P. (2014). Contrast dependency and prior expectations in human speed perception. Vision Res. *97*, 16–23.

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

Sun, H.-J., Campos, J.L., and Chan, G.S.W. (2004). Multisensory integration in the estimation of relative path length. Exp. Brain Res. *154*, 246–254.

Thurley, K. (2016). Magnitude estimation with noisy integrators linked by an adaptive reference. Front. Integr. Neurosci *10*, 6.

Veliz-Cuba, A., Kilpatrick, Z.P., and Josić, K. (2016). Stochastic models of evidence accumulation in changing environments. SIAM Rev. 58, 264–289.

Vickerstaff, R.J., and Di Paolo, E.A. (2005). Evolving neural models of path integration. J. Exp. Biol. 208, 3349–3366.

Weiss, Y., Simoncelli, E.P., and Adelson, E.H. (2002). Motion illusions as optimal percepts. Nat. Neurosci. 5, 598–604.

Whitlock, J.R., Pfuhl, G., Dagslott, N., Moser, M.B., and Moser, E.I. (2012). Functional split between parietal and entorhinal cortices in the rat. Neuron 73, 789–802.

Wilber, A.A., Skelin, I., Wu, W., and McNaughton, B.L. (2017). Laminar organization of encoding and memory reactivation in the parietal cortex. Neuron *95*, 1406–1419.e5.

STAR***METHODS**

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Software and Algorithms		
Open Graphics Library (OpenGL)	Khronos Group	https://www.opengl.org/
Custom-built analysis code	MATLAB	https://www.mathworks.com/
Other		
M20U9T-N82 USB joystick	CTI Electronics	http://www.ctielectronics.com/
Kodak Wratten 2 optical filters – #29 and #61	Kodak	https://www.kodak.com/
Power 1401 MkII data acquisition system and Spike2 software	Cambridge Electronic Design Ltd.	http://ced.co.uk/

CONTACT FOR REAGENT AND RESOURCE SHARING

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

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Nineteen human subjects (Nine Male, Ten Female; all adults in the age group 18-32) participated in the experiments. With the exception of subjects s4 and s5 (Figures S1–S5), all subjects were unaware of the purpose of the study. Experiments were first performed in the above two subjects before testing others. All experimental procedures were approved by the Institutional Review Board at Baylor College of Medicine and all subjects signed an approved consent form.

METHOD DETAILS

Behavioral Task

Subjects used an analog joystick 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 (~ 250 ms) 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: 8.5×18.5 cm) that was randomly repositioned and reoriented at the end of its lifetime, making it impossible to use as a landmark. The stimulus was rendered as a red-green anaglyph and projected onto a large rectangular screen (width × height: 149×127 cm) positioned 67.5cm 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. 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). The target, a circle of radius 20cm whose luminance was matched to the texture elements, blinked at 5Hz and appeared at a random location between $\theta = \pm 42.5^\circ$ of visual angle at a distance of r = 0.7 - 6m relative to where the subject was stationed at the beginning of the trial. After one second, the target disappeared, which was a cue for the subject to start steering, and the joystick controller was activated.

Of the nineteen subjects, seven subjects performed a total of 2000 trials equally spread across eight sessions. Prior to the first session, all subjects were asked to perform around ten practice trials in which they steered to a visible target to familiarize themselves with joystick movements and the task structure. In four of the sessions (two of which contained angular landmarks in the form of a panoramic mountainous background), the maximum linear and angular speeds were fixed to $v_{max} = 2ms^{-1}$ and $\omega_{max} = 90^{\circ}/s$ respectively, with the floor density also held constant at $\rho = 2.5$ elements/m². In the remaining four sessions, trials with two different speed limits ($v_{max} = 2ms^{-1}$ and $\omega_{max} = 90^{\circ}/s$; $v_{max} = 4ms^{-1}$ and $\omega_{max} = 180^{\circ}/s$) and two floor densities ($\rho = 0.1$ elements/m² and $\rho = 2.5$ elements/m²) were randomly interleaved. Six of the seven subjects participated in two additional experimental sessions (250 trials each). The first of these additional experiments was similar to the original experiment, as before, subjects pressed a button on the joystick to initiate each trial. Targets appeared briefly at random locations sampled from a distribution identical to the original experiment. However, rather than actively steering to the target, they were passively transported along trajectories that took them through the target at one of two possible linear speeds ($\nu = 2ms^{-1}$ or $4ms^{-1}$). Since trajectories necessarily passed through the

target and the velocity was held constant throughout the trial, the angular velocity on each trial was constrained by the location of the target. Subjects were instructed to press the button when they believed they had reached the target. Therefore, in this experiment, subjects used the joystick only to initiate trials and register their responses.

Furthermore, six subjects (five of whom did not participate in any of the above studies) were tested on an extended version of the original task wherein the targets were presented at distances of up to 32 m. Due to the longer range of distances used in this version, the target radius was increased from 20cm to 75cm to ensure clear visibility. As before, subjects had to steer to a target location that was cued briefly for a period of 1 s at the beginning of the trial. However in this experiment, target locations were discretized to five possible distances (r = [2, 4, 8, 16, 32] m) and five possible angular eccentricities ($\theta = [0^{\circ}, \pm 15^{\circ}, \pm 30^{\circ}]$) resulting in a total of 25 unique target locations. Subjects performed ten randomized repetitions of each location yielding a total of 250 trials. Trials with two different floor densities ($\rho = 0.1$ elements/m² and $\rho = 2.5$ elements/m²) were randomly interleaved. Finally, seven more subjects (none of whom participated in any of the above studies) performed a slight variation of this extended task in which we manipulated the density only after the target was turned off. This ensured that the manipulation selectively affected the reliability of optic flow, but not the perception of target positions.

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. Spike2 software (Cambridge Electronic Design Ltd.) was used to record and store the subject's linear and angular velocities, target locations, and all event markers for offline analysis at a sampling rate of Hz.

Dynamic Bayesian observer model

To account for the pattern of behavioral results, we considered an observer model comprised of a Bayesian estimator that used noisy measurements m_{ν} and m_{ω} to decode linear and angular self-motion velocities ν and ω , which were then temporally integrated to dynamically update the subject's position. We parameterized the model by making the following three assumptions: First, we chose an exponential function to describe the priors over both linear and angular velocities: $p(\nu) = e^{a_{\nu}|\nu|}$ and $p(\omega) = e^{a_{\omega}|\omega|}$. Second, likelihood functions $p(m_{\nu}|\nu)$ and $p(m_{\omega}|\omega)$ were assumed to be Gaussian, centered around the respective measurements m_{ν} and m_{ω} , with variances proportional to the magnitude of the measurement: $Var(m_{\nu}) = b_{\nu}|m_{\nu}|$ and $Var(m_{\omega}) = b_{\omega}|m_{\omega}|$. Under these conditions, it can be shown that the means and variances of the maximum *a posteriori* estimates $\hat{\nu}$ and $\hat{\omega}$ are given by (Stocker and Simoncelli, 2006):

$$\mathsf{E}\left[\widehat{\nu} \mid m_{\nu}\right] = \beta_{\nu} m_{\nu}; \ \mathsf{E}\left[\widehat{\omega} \mid m_{\omega}\right] = \beta_{\omega} m_{\omega} \tag{Equation 1.1}$$

$$\operatorname{Var}\left[\widehat{\nu} \mid m_{\nu}\right] \approx \beta_{\nu}^{2} \operatorname{Var}(m_{\nu}); \quad \operatorname{Var}\left[\widehat{\omega} \mid m_{\omega}\right] \approx \beta_{\omega}^{2} \operatorname{Var}(m_{\omega}), \quad (\text{Equation 1 .2})$$

where $\beta_v = 1 + a_v b_v$ and $\beta_\omega = 1 + a_\omega b_\omega$ have a straight-forward interpretation in the form of multiplicative biases in the subjects' estimates of their linear and angular speeds respectively. Note that a flat prior corresponds to an exponent of zero yielding an unbiased estimate, while negative/positive values of the exponents would result in under/overestimation of the speeds. The final assumption pertains to the nature of the integrator that computes position from speed. We assume that the integration process is governed by two independent leak time constants τ_d and τ_{φ} that specify the timescales of integration of estimated linear and angular speeds to compute distance *d* and heading φ respectively:

$$d = -d(t)/\tau_d + \hat{\nu}(t); \ \dot{\phi} = -\phi(t)/\tau_{\phi} + \hat{\omega}(t).$$
 (Equation 2.1)

The mean distance and heading at each time point can be determined by convolving the mean velocity estimates with an exponential kernel: $E[\hat{d}(t)] = e^{-t/\tau_d} \circ E[\hat{v}(t)]$ and $E[\hat{\varphi}(t)] = e^{-t/\tau_g} \circ E[\hat{\omega}(t)]$ where the expectations are taken over the corresponding posterior probability distributions. Likewise, if noise in the velocity measurements is temporally uncorrelated, the variance of the distance and heading estimates can be expressed in terms of the variances of the velocity estimates as: $Var[\hat{d}(t)] = e^{-t/\tau_d} \circ Var[\hat{v}(t)]$ and $Var[\hat{\varphi}(t)] = e^{-t/\tau_g} \circ Var[\hat{\omega}(t)]$. Thus, in this case, both mean and variance of the integrated estimates will share the same temporal dynamics. Note that the mean estimates $E[\hat{d}(t)]$ and $E[\hat{\varphi}(t)]$ will be accurate in the limit of large time constants (perfect integration), but are misestimated if the time constants are comparable to travel time, *T*. Since the time course of distance and heading together determine position, it follows that the subjects' mean estimates of their linear and angular coordinates (\hat{r} and $\hat{\theta}$) will also be different from their actual values (*r* and $\hat{\theta}$) when $\tau \approx T$.

We also analyzed a variation of the leaky integration model in which the leak was implemented using space constants π_d and π_{φ} according to:

$$d = -\hat{\nu}(t)d(t)/\pi_d + \hat{\nu}(t); \ \dot{\varphi} = -\hat{\omega}(t)\varphi(t)/\pi_{\varphi} + \hat{\omega}(t).$$
(Equation 2.2)

Note that unlike the temporal leak model in Equation 2.1, this model only integrates when velocity is non-zero. Therefore, position is only updated during movement resulting in estimates that are robust in time. This version of the leaky-integrator model was used for Figure S3E.

Model fitting

In order to determine the key factor underlying subjects' biases, we fit two different variants of the model: (i) A 'slow-speed prior model' (\mathcal{M}_1) in which the integration was assumed to be perfect ($\tau_d = \tau_{\varphi} = \infty$) and (ii) a 'leaky integration model' (\mathcal{M}_2) where the prior was held flat ($a_r = a_\omega = 0$). These models represent the two extreme scenarios in which bias in path integration is attributed exclusively to speed misperception and forgetful evidence accumulation, respectively. The models both had four parameters each: width parameters (b_r and b_ω) of the two likelihood functions to represent how fast the respective widths scale with the magnitude of linear and angular velocity measurements, in addition to either two exponents (a_r and a_ω) to represent priors in \mathcal{M}_1 or two time constants (τ_d and τ_{φ}) to represent the degree of leak in \mathcal{M}_2 . Since subjects' position estimates are probabilistic, we fit model parameters ψ by taking both mean and uncertainty of position into account; this was done by maximizing the expected reward, which is essentially the probability that the subjects believed themselves to be *within* the target at the end of each trial:

$$\mathcal{L}_{\mathcal{M}} = \arg\max_{\psi} \prod_{i=1}^{N} \int_{-\infty}^{\infty} \mathcal{P}\left(\mathbf{x}^{(i)} | \dot{\mathbf{x}}_{1,...,T}^{(i)}\right) \mathbb{1}_{Q}(\mathbf{x}) d\mathbf{x},$$
 (Equation 3)

where **x** is a vector that denotes position on the horizontal plane, $P(\mathbf{x}^{(l)}|\dot{\mathbf{x}}_{1,...,T}^{(l)})$ is the probability distribution over the subject's stopping position on the *i*th trial conditioned on the path taken in that trial, *T* is the duration of that trial, $\mathbb{1}_Q(\mathbf{x})$ is the indicator function which is equal to 1 for all values of **x** that fall within the target Q and zero everywhere else, \mathcal{L}_M denotes the likelihood of model \mathcal{M} with the best-fit parameters, and model parameters $\psi \ni \{a_v, a_\omega, b_v, b_\omega\}$ and $\{b_v, b_\omega, \tau_d, \tau_\varphi\}$ for \mathcal{M}_1 and \mathcal{M}_2 respectively. We fit two more models in addition to the above: a *'null'* model \mathcal{M}_0 that had only two free parameters $-\psi \ni \{b_v, b_\omega\}$ – essentially attributing the biases in subjects' position estimates entirely to random variability in their self-motion speed estimates, and a *'full'* model \mathcal{M}_{12} in which all six model parameters were free such that $\psi \ni \{a_v, a_\omega, b_v, b_\omega, \tau_d, \tau_\varphi\}$.

In addition to the above parametrization, we fit both the slow-speed and leaky-integrator models using three alternative parametrizations constructed by changing either the form of the prior, the form of the likelihood function, or both: (*i*) Gaussian prior and Gaussian likelihood, (*ii*) Exponential prior and log-normal likelihood, and (*iii*) Gaussian prior and log-normal likelihood. Each of the above parametrizations also had four free parameters (Figure S3F). Maximum *a posteriori* estimates $\hat{\nu}$ and $\hat{\omega}$ were computed numerically.

Modeling position uncertainty

Since position is estimated by integrating velocity, uncertainty in velocity estimates will accumulate over time, leading to growing uncertainty in position estimates. If integration is leaky, noise will only accumulate over the time constant of integration, causing position uncertainty to eventually asymptote to a fixed value. However, if the integration is perfect, noise will accumulate perpetually leading to uncertainty that grows with time. Let r(T) denote the subject's one-dimensional position estimate at time T. If $\nu(t)$ denotes subject's instantaneous velocity estimate, we have:

$$r(T) = \int_0^T \nu(t)dt + \int_0^T \eta(t)dt,$$
 (Equation 4)

where $\eta(t)$ represents a noise in the velocity estimate and the integral of this noise corresponds to a random walk. If noise has zero mean, the subject's mean position estimate $\langle r(T) \rangle$ is not affected. However, the noise variance of position estimate $\sigma^2 = \langle (\delta r)^2 \rangle$ will grow with time. For integration of temporally uncorrelated noise, the variance of position uncertainty is proportional to time *T*. We postulate that uncertainty in position, δr , will be proportional to T^{λ} for some exponent λ . Large exponents may occur due to temporal correlations, or computational constraints within the system. For the case of uniform motion, $\nu(t) = \nu$, the mean position estimate is $\langle r(T) \rangle = \langle \int_0^T \nu dt \rangle = \nu T$. Since mean position then scales linearly with time, position uncertainty can be expressed in terms of position as $\sigma(r) = kr^{\lambda}$ for some proportionality constant *k*.

Fitting distance-dependent bias reversal

To simultaneously quantify the effects of position bias (due to velocity underestimation) and position uncertainty leading to a distance-dependent bias-reversal, we modeled the subject's radial distance response as:

$$\widehat{r}_{i} = \Gamma \arg \max_{r_{i}} \int_{-\infty}^{\infty} P(r_{i}|\boldsymbol{k},\gamma) \mathbb{1}_{Q}(r) dr, \qquad (\text{Equation 5})$$

where \hat{r}_i is the model estimated radial distance on the i^{th} trial, $P(r_i|k,\lambda) \sim \mathcal{N}(r_i,kr_i^{\lambda})$ is the modeled probability distribution over the subject's position, $\mathbb{1}_Q(r)$ is the indicator function which is equal to 1 for all values of r within the target Q and zero otherwise, and Γ is a multiplicative constant that captures multiplicative bias in the subject's mean position estimate. If the mean position is underestimated, the multiplicative bias should be greater than unity because the subject would respond by overshooting. The integral on the right-hand side represents the subject's belief that he/she is on target, and captures the effect of position uncertainty, whereas the multiplicative constant captures the bias in mean position induced by prior expectations of self-motion velocity. For each subject, we fit the model parameters Γ , λ , and k by minimizing the squared-error between the radial distance of the model and the subject's actual response across trials according to arg min $\sum_{i=1}^{N} (\hat{r}_i - r_i)^2$ where N is the total number of trials.

QUANTIFICATION AND STATISTICAL ANALYSIS

Customized 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. Model parameters were optimized using MATLAB's *fmincon* function by constraining time constants and likelihood widths to be non-negative.

Estimation of Bias

Behavioral error on each trial was quantified by computing the difference between the subject's response position and the corresponding target position to yield an error vector \vec{e} . Error magnitudes were computed as the Euclidean norm of the error vectors, and were convolved with a two-dimensional isotropic Gaussian kernel g(x, y) with width 50cm to yield smoothed error magnitudes $e_s(x_0, y_0) = \sum_{x,y} g(x - x_0, y - y_0) e(x, y)$ for visualization in Figure 1D. To visualize spatial distribution of errors for the two models (Figure 3C), we replaced subject's response position by the position estimates generated by the models. We regressed each subject's response positions $(\tilde{r}, \tilde{\theta})$ against target positions (r, θ) separately for the radial $(\tilde{r} \text{ versus } r)$ and angular $(\tilde{\theta} \text{ versus } \theta)$ co-ordinates, and the radial and angular multiplicative biases (Γ_r and Γ_{θ}) were quantified as the slope of the respective regressions (Figure 1E). Quantifying the biases in this polar representation of the position sallowed us to qualitatively relate them to perceptual biases in linear and angular speeds — quantities that the subjects controlled using the joystick. For each subject, we estimated the 95% confidence interval in the biases by bootstrapping (Figure 1F). We used a one-sample t test to determine whether the mean multiplicative bias across subjects was significantly greater than unity (n = 7).

Model comparison and validation

For each subject, we estimated the likelihoods of all four models – the slow-speed prior model (\mathcal{L}_1), leaky integration model (\mathcal{L}_2), the null model (\mathcal{L}_0) and the full model (\mathcal{L}_{12}) – by fitting the corresponding model parameters to the subject's response trajectories from all trials as explained in the 'Model fitting' section. We compared relative goodness of models by estimating the ratio of their likelihoods. For each subject, we estimate confidence intervals for these likelihood ratios by bootstrapping (Figure 3A). Additionally, for each trial, we generated the subjects' believed trajectories implied by the best-fit parameters for both models (Figure 3B). We then computed the "residual bias" for both models by regressing the final position estimates corresponding to the resulting trajectories against the target positions both for linear and angular co-ordinates. Prior to doing the regression, we used a 4-fold cross-validation procedure in which we fit both models to 75% of the trials at a time (training set) and generated model estimates for the remaining trials (test set) using the learned model parameters, to avoid overfitting. We repeated this procedure four times so that each trial was allocated to the test set exactly once. We then quantified the residual bias by performing linear regression on the pooled model estimates from all four non-overlapping test datasets (Figure 3E). For each model, we used a one-sample t test to determine whether its residual bias across subjects was greater than unity (n = 7).

Test of model predictions

Fitting and comparison of the two models described above was done using behavioral data collected during the sessions when the ground plane density (ρ) and speed limits (ν_{max} and ω_{max}) were held fixed. The models make distinct predictions for how bias would be affected by the latter quantities, so we manipulated both of those quantities in a separate experiment (as explained in the 'Behavioral task' section) and tested the model predictions. For each subject, we performed linear regression and quantified bias as the regression slope. For estimating the effect of density manipulation, we collapsed trials across both speeds (Figures 4B and 4C). Similarly, the effect of speed manipulation was analyzed by combining trials from the two densities (Figures 4E and 4F). We used paired t tests to determine whether manipulating density (and speed) significantly altered bias across subjects (n = 7).

Quantifying bias-reversal

To quantify reversal of bias with distance (Figure 5F), we used each subject's final response distance to fit the phenomenological model in Equation 5. We then used one-sided t tests to test whether the distribution of the best-fit multiplicative constants (Γ) and power-law exponents (λ) across subjects were, respectively, greater than and less than one. Additionally, we used paired t tests to test whether decreasing the density of optic flow significantly increased (decreased) the mean response distance of the subjects for the closest (farthest) target condition (Figure 5H legend).

DATA AND SOFTWARE AVAILABILITY

The datasets generated in this study and code to analyze them are available from the Lead Contact on request.