

Belief embodiment through eye movements facilitates memory-guided navigation

Received: 4 September 2023

Accepted: 29 October 2025

Published online: 24 November 2025

 Check for updates

Akis Stavropoulos ^{1,5}✉, Kaushik J. Lakshminarasimhan ^{2,3,5} & Dora E. Angelaki^{1,4}

The brain evolved to navigate a dynamic and uncertain world, but the mechanisms underlying ethologically-relevant behavioral strategies remain unclear. In the real-world, such strategies are shaped both by task demands and by the cognitive resources available to the animal. We hypothesized that eye movements constitute a vital cognitive resource to support neural computations for memory-guided navigation. We tested this using a naturalistic task in which humans use a joystick to steer and catch flashing targets in a virtual environment lacking explicit position cues. While navigating to the goal, participants physically track the latent target position with their gaze even in the absence of optic flow, demonstrating that these task-relevant eye movements reflect an embodiment of the subjects' dynamic internal beliefs about the goal location. We developed a neural network model with tuned recurrent connectivity between oculomotor and evidence-integrating fronto-parietal circuits to account for this behavioral strategy. We show that this model better explained neural data from male monkeys' posterior parietal cortex compared to models optimized solely for task performance and unconstrained by such an oculomotor-based strategy. These results highlight the importance of eye movements in working memory computations and establish a functional significance of oculomotor signals for evidence-integration and navigation computations via embodied cognition.

The brain evolved complex recurrent networks to interpret and act upon a dynamic and uncertain world, but its computational powers and mechanisms generating natural behavior remain cryptic^{1–3}. Most of our insights into neural computation are based on binary tasks with highly constrained actions that are artificially segregated from perception^{4–7}. Tightly controlling laboratory behavior by preventing natural, continuous movements has simplified interpretability but also hindered our ability to gain insights from natural behavioral strategies. Artificially keeping the eyes fixed, for example, has been standard in monkey studies of working memory and decision-making^{8–12}. In contrast, natural behavior involves continuous eye movements^{13–15}. Thus, there is concern that traditional experimental paradigms, such as

those requiring fixation or tightly constrained actions, deprive subjects of natural affordances¹⁶, and thus may hide fundamental neural mechanisms that are only expressed in a dynamic, closed-loop context. Can more naturalistic laboratory behaviors with free eye movements shed new light onto neural mechanisms of ethological behaviors?

A major gap in understanding is epitomized by an emerging tool to probe neural mechanisms: neural network models optimized to perform neuroscience tasks. The representations learned by the networks often resemble the response properties of brain areas that drive behavior in those tasks^{17,18}. However, such models are typically grounded in generic neural architectures—feedforward or recurrent,

¹Center for Neural Science, New York University, New York, NY, USA. ²Department of Neuroscience, University of Texas, Dallas, TX, USA. ³Center for Theoretical Neuroscience, Zuckerman Institute, Columbia University, New York, NY, USA. ⁴Tandon School of Engineering, New York University, New York, NY, USA. ⁵These authors contributed equally: Akis Stavropoulos, Kaushik J. Lakshminarasimhan. ✉e-mail: ges6@nyu.edu

depending on the task—and cannot explain why neural computations are distributed across functionally distinct brain areas. A jarring example of distributed brain computation is the prevalence of motor signals in sensory and association areas^{19,20}, and sensory signals in the motor and frontal areas^{21,22}. Existing models of distributed neural representations appeal to the brain's recurrent architecture to capture multi-area data but fall short of providing a normative account of such representations^{23,24}. There is a growing realization that building task-optimized neural network models with brain-inspired modular architectures provides limited insights beyond what is already determined by the task goal^{25,26}. To gain new insights from adopting brain-like architectures, we need to additionally incorporate the specific strategy used by animals to solve the task²⁷. Traditional neuroscience tasks like binary decision-making are too simple to admit interesting cognitive strategies, especially when participants are mechanically restrained in fixation-based paradigms.

To unravel the neural mechanisms of natural behavior and to interrogate alternatives to the traditional approach, we have developed a naturalistic navigation task featuring action/perception loops in virtual reality (VR) with unconstrained eye movements. This continuous foraging task requires participants to steer towards remembered target locations by using sensory evidence, working memory, and continuous actions constituting a naturalistic visual perception-action loop^{28–30}. Participants observe a briefly flashed target in the distance (like the blinking of a firefly) and steer to the remembered target location using optic flow feedback from a virtual environment comprising an unstructured ground plane with no landmarks. Importantly, in contrast to traditional tasks such as evidence accumulation or delayed discrimination in which the latent world states and/or contents of working memory remain unchanged throughout the trial, the latent state (i.e., egocentric target location) dynamically varies over the course of each trial, under the participant's control, and must be mentally tracked in order to know precisely when to stop steering.

In principle, this task can be performed without physically tracking the believed goal location with one's eyes. Yet, Lakshminarasimhan et al.³⁰ found that both humans and monkeys tend to follow the location of the invisible target with their gaze until they reach it, and noticed a significant decline in steering performance when eye movements were suppressed. Given the visually guided nature of the steering task, such eye movements may reflect a strategy to gather information about self-motion: since subjects must integrate optic flow to dynamically update their beliefs about the relative goal location, directing gaze to specific regions of the environment such as the focus of expansion, might help acquire more information about their movement velocity (*active sensing* hypothesis). Alternatively, these task-relevant eye movements may reflect an embodiment of subjects' dynamically evolving internal beliefs about the goal: by allowing dynamic beliefs about the relative target location to continuously modulate eye movements in this task, the computational burden on circuits involved in working memory is reduced by recruiting the oculomotor (OC) circuit in belief updating, despite its primary function being unrelated to working memory computations (*cognitive embodiment* hypothesis). The latter hypothesis predicts that these eye movements should also govern other types of navigation, e.g., inertially guided steering, where the joystick controls inertial accelerations in the absence of visual cues.

Here, we first provide strong support for the embodiment hypothesis by analyzing eye movements under both inertially and visually guided versions of the task. We found that under both sensory conditions, eye movements reflect the evolving belief dynamics about the relative target location. We then used this behavioral strategy as an additional constraint for training a distributed neural network model and found that it recapitulated the behavioral and neural data more accurately with fewer tunable parameters than purely task-optimized models. These results lend support to the notion that ethologically

valid paradigms can help constrain modeling to provide insights into the neural mechanisms and the emergence of distributed neural representations.

Results

Participants sitting on a motion platform performed a VR navigation task using a joystick to steer freely and catch targets that pop up transiently (like fireflies) on the ground plane, one at a time (Fig. 1a). Participants' steering was coupled either to a visual environment that provided optic flow but devoid of landmarks ("Visual" condition) or to the platform's motion in complete darkness ("Inertial" condition³¹). At the beginning of each trial, a target ("firefly") appears for 1 s at a random location on the ground plane within the field of view (Fig. 1b). When it disappears, an analog joystick controlling linear and angular motion is activated, allowing participants to navigate toward the remembered target location. Participants could steer freely on the ground plane and integrate momentary sensory evidence about their movements based on either visual (optic flow) or inertial (vestibular with somatosensory/tactile) sensory cues. This task has a crucial time-varying latent variable: position of the target relative to oneself, which must be computed by integrating noisy angular/linear sensory cues (visual or inertial), which in turn are controlled dynamically by the participant's joystick actions. A time constant governed the control dynamics (CD): in trials with a *small* time constant, joystick position mainly controlled velocity (Velocity Control; VC); when the time constant was *large*, joystick position mainly controlled acceleration (Acceleration Control; AC), mimicking inertia under viscous damping (Fig. 1c). Across trials, visual and inertial sensory conditions were randomly interleaved while manipulation of the time constant followed a bounded random walk³¹.

Subjects did not receive any performance-related feedback. By varying motion dynamics across trials and by eliminating performance feedback, we aimed to induce variable behavioral performance, which ensures greater statistical power in the analyses needed to decouple subjective beliefs from the true latent states. Indeed, subjects showed strong biases, especially in the inertial condition (Fig. 1d, e). Biases, defined as the regression slope between target and stopping positions (a value of 1 indicates no bias), are strongly correlated with the CD (Fig. 1f), a pattern that has been described in detail previously³¹. Here, we compare the eye movements generated in the inertial, compared to the visual version of the task, to distinguish between the *active sensing* and *cognitive embodiment* hypotheses.

Eye movements track beliefs about the latent goal location

Participants received no instruction about their gaze behavior, yet eye movements tracked the memorized location of the goal, and this was true not only in the visual condition (as previously shown by Lakshminarasimhan et al.³⁰), but also during inertially guided steering. As the target is on the ground plane below the subject's eye level (Fig. 1a, b), the relative position of the (invisible) target approaches the midline and moves downward in the visual field as the participant steers towards it. Eye movements mirrored this same pattern: horizontal eye position converged toward zero (midline) and vertical eye position descended over time, even under the Inertial condition, where no visual cues were provided after the target disappeared (Fig. 2a). As shown in a typical example, there was an initial saccade towards the target immediately after target onset (Fig. 2b—gray region, 0–1 s), followed by a mostly smooth tracking until the end of the trial.

We computed the Euclidean distance between *eye* and *target* position as the "target-tracking error (TTE)" (Fig. 2b, bottom, see "Methods"). TTE at target offset was low across subjects (mean TTE at trial onset $\pm$ standard deviation (SD): $5.60 \pm 0.38^\circ$) and increased as the trial progressed (Fig. 2c). Despite the very long trial durations (across subjects trial duration mean $\pm$ SD—inertial: 14.1 ± 5.1 s, visual: 13.3 ± 4 s), TTE remained significantly below chance level obtained by shuffling




Fig. 1 | Experimental design and task performance. **a** Experimental set up. **b Left:** Illustration of a virtual environment. Participants steer towards a briefly cued target (yellow disc) using optic flow cues available on the ground plane (visual condition only; platform motion is the only available cue in the inertial condition). During steering, the target becomes less eccentric over time (towards the participant's midline), while it lowers in the participant's field of view (color-coded arrow). **b Right:** Overhead view of the spatial distribution of target positions across trials and the corresponding trial trajectories. Red dot shows the starting position of the participant. **c** Simulated maximum pulse joystick input and the corresponding velocity output under Velocity Control (VC; beige) and Acceleration Control (AC; brown). The input is low-pass filtered to mimic the existence of inertia. The time constant of the filter varies across trials (time constant τ), along with maximum velocity, to ensure comparable travel times across trials. Gray zone: brief cueing period of the target at the beginning of the trial. **d** Target vs. Response. **Left:** Comparison of the radial distance of a typical subject (stopping location) against the radial distance of

the target across all trials. **Right:** Comparison of the angular eccentricity of the stopping location against the angular eccentricity of the target (both with respect to the starting position) across all trials. Black dashed lines have a unity slope (unbiased performance). Solid lines: linear regression. Data colored according to the sensory condition (red: inertial, cyan: visual). Radial and angular response biases were defined as the slope of the corresponding regressions. **e** Scatter plot of radial and angular biases in each sensory condition plotted for each individual participant. Ellipses show 68% confidence intervals of the distribution of data points for the corresponding sensory condition. Diamonds (centers of the ellipses) represent the mean radial and angular response biases across participants. Dashed lines indicate unbiased radial or angular position responses. Solid diagonal line has a unit slope. **f** Participant average of radial and angular response biases in each condition, with trials grouped into tertiles of increasing time constant τ . Error bars denote ± 1 SEM across participants ($n = 8$).

("Methods"; Fig. 2c, gray line) for $68.8 \pm 4.9\%$ (visual) and $51.9 \pm 4.4\%$ (inertial) of the trajectory (mean $\pm$ SEM of percentage distance traveled until TTE crosses chance level); for data from individual subjects see (Supplementary Fig. S1a). These results are consistent and build upon findings from a purely velocity-controlled visual steering task of much shorter trial durations (~ 2 s) with performance feedback³⁰. Notably, these results hold true also for inertial navigation in the absence of optic flow, suggesting that the pattern of eye movements reflects a strategy of embodiment and is not linked solely to active sensing of optic flow patterns, if at all. While TTE was larger for the Inertial condition compared to the Visual condition (Fig. 2c, red vs. blue), this is due to increased behavioral variability in the former condition (Supplementary Fig. S1b) rather than an inability to track the memorized goal.

A trivial explanation for the increase in TTE over time is that eye movements become progressively more random with time. Alternatively, the increase in TTE could arise if eye movements track the participant's *belief* about the goal location rather than the true goal location. In this case, TTE should correlate with steering error (SE, distance of stopping from actual target position). This is because steering decisions are based on beliefs, and consequently, any error in belief must manifest as an error in the eventual stopping position. Indeed, the correlation between TTE and SE increased as the trial progressed, reaching a peak at about 70% (visual) and 50% (inertial) into the trial, and decreased sharply thereafter (Fig. 2d). Peak correlations were statistically significant ($p < 0.05$) in 8/8 and 7/8

participants in the inertial and visual conditions, respectively. This result supports the hypothesis that the eyes track the *believed* location of the target in the virtual environment. In fact, when regressed against both the initial target and stop positions, eye movements were driven mostly by target position at trial onset and mostly by stopping position at the end of the trial (Fig. 2e), revealing how the believed target location drifts gradually from the target to the stop location over the course of a trial.

Small saccades aid belief tracking

During steering, saccade frequency (Supplementary Fig. 2a, top) and amplitude (Supplementary Fig. 2a, bottom) were both suppressed. Nevertheless, the infrequent small saccades contributed to goal tracking, as there was a drop in the correlation between steering and tracking error when saccades were removed from the data (correlation drop mean $\pm$ SEM, inertial: 0.091 ± 0.058 , visual: 0.033 ± 0.023 ; Fig. 3a). This was particularly notable in the inertial condition, where the horizontal slow eye movements are strongly affected by the yaw vestibulo-ocular reflex (VOR; horizontal component correlation drop mean $\pm$ SEM, inertial: 0.186 ± 0.052 , visual: 0.080 ± 0.022 ; t -test p value, inertial: 0.012, visual: 0.011; Supplementary Fig. 2b). Indeed, there were significant correlations between the cumulative horizontal saccade amplitude and the angular steering errors, revealing that saccades made during the VOR had a major contribution in "undoing" the effects of the VOR such that the eye position could still reflect the internal belief of goal location (Fig. 3b). We previously have shown that




Fig. 2 | Target-tracking with eye movements. **a** Time-course of horizontal and vertical eye position during a random subset of trials from one participant; time at 0 denotes disappearance of target (target offset); Black dots illustrate end of trial (clipped at 11 s). **b** Time-course of horizontal (top) and vertical (middle) eye position (black solid lines), the respective target position (gray dashed lines), as well as target-tracking error (TTE, orange), during a representative trial. Gray region denotes the period when a target was visible. Red dashed line corresponds to the end of the trial. **c** Normalized TTE over trial progression (percentage of total distance traveled). TTE was normalized by the chance-level TTE obtained by shuffling

(gray line). The area above the gray line corresponds to TTE worse than chance. Shaded regions: ± 1 SEM across subjects. **d** Correlation coefficient of TTE and steering error (SE). Shaded regions: ± 1 SEM across participants. **e** Multilinear regression of eye position against initial target and stop positions over trial progression, for both horizontal and vertical components. Light and dark shades represent the target and stop position coefficients, respectively. Notice how the modulation of eye position by the target and stop positions gradually reverses as the trial progresses. Shaded regions: ± 1 SEM across participants.




Fig. 3 | Saccades contribute to evolving beliefs about goal location. **a** Top: actual eye position (black line) and the corresponding saccade-free eye position (red line) of an example trial. Bottom: Comparison of correlation coefficients of target-tracking error (TTE) and steering error (SE) for the actual eye position and saccade-free eye position (horizontal component). Peak correlations within 50% of the distance traveled were selected for each participant. **b** Correlation between cumulative horizontal saccade amplitude and angular steering error in the visual (top, blue) and inertial (middle, red) conditions. Saccades (within 50% of distance traveled) were pooled across participants for better visualization. Bottom: average

correlation coefficients across participants. Asterisks denote the level of statistical significance of the correlation difference within each condition using a paired *t*-test ($p < 0.05$). Error bars: ± 1 SEM across participants ($n = 8$). **c** Time-course (kernel) of coefficients obtained by linearly regressing the amplitude of the horizontal/vertical component of saccades evoked within 50% of the distance traveled against the corresponding target-tracking error (light blue/red), stop position-tracking error (dark blue/red), and believed target-tracking position (gray). Shaded regions denote ± 1 SEM obtained by bootstrapping.

forcing participants to fixate (in the visual condition) substantially affects task performance³⁰. Combined with the persistent target-tracking in darkness while under the influence of the VOR, these findings strongly suggest that this embodiment has a computational role.

To explore more directly what drives saccadic eye movements during steering, we ran a regression analysis that shows how saccade amplitude is modulated by errors in tracking the actual target position, the stop position, or the participant's reconstructed dynamic *belief* about goal location—the latter computed as the weighted average of the actual target and stop positions over time obtained for each participant (from Fig. 2e). Saccades are indeed modulated by beliefs, as illustrated by the fact that, before saccade onset, the kernel (time-course of the coefficients obtained by linear regression) is larger for the tracking error corresponding to the believed goal rather than the actual target or stop position (Fig. 3c; pre-saccadic peak of regression kernel mean $\pm$ SEM, horizontal component—inertial: [target: -0.05 ± 0.02 , stop: 0.03 ± 0.07 , belief: 0.43 ± 0.21], visual: [target: -0.02 ± 0.06 , stop: -0.002 ± 0.05 , belief: 0.34 ± 0.09]; vertical component—inertial: [target: -0.16 ± 0.16 , stop: 0.05 ± 0.03 , belief: 0.42 ± 0.19], visual: [target: -0.10 ± 0.06 , stop: 0.03 ± 0.02 , belief: 0.32 ± 0.10]. These results support the hypothesis that the small saccades generated during steering stabilize the gaze towards the believed goal location.

In summary, we conclude that participants integrate movement velocity to track their position relative to the goal using an OC-based cognitive strategy: the evolving belief about goal location relative to their current position is embodied in eye position—and this cognitive embodiment has a computational role. We now turn to modeling to understand how such a strategy of embodiment can inform the underlying neural mechanisms. Specifically, we train different neural models optimized to do this task both with and without this cognitive

strategy and evaluate how well each model predicts behavioral and neural data recorded in monkeys.

A frontoparietal network model constrained by behavioral strategy

We previously demonstrated that both posterior parietal and dorso-lateral prefrontal cortices represent latent beliefs in this task^{32,33}. For simplicity, here we consider a combined frontoparietal recurrent neural network (FPN) as a stand-in for computations across both cortical areas. To investigate the mechanistic contribution of eye movements to this network, we simulated four models that are architecturally identical but differ in which connections are tuned (Fig. 4a—green and crimson). They feature a frontoparietal (FPN) module (see “Discussion”) comprising 100 recurrently connected nonlinear (“sigmoidal”) units that receive two-dimensional sensory inputs (linear and angular velocity) and a two-dimensional pulse whose amplitude encodes the target position (x - y coordinates) at the beginning of each trial. The FPN module sends projections to motor units that drive two-dimensional joystick actions (linear and angular acceleration) and has bidirectional connections with OC units that drive two-dimensional eye movements (horizontal and vertical). Two of the models are optimized solely for task performance, i.e., to minimize the discrepancy between the stopping position and the target position, by tuning either just the readout weights onto the motor units (*Model 1*) or both the readout and recurrent weights within FPN (*Model 2*). The remaining models are optimized for task performance by tuning the readout weights, while also being constrained by the strategy used by humans and monkeys. Specifically, we tune the weights from FPN to OC to minimize an auxiliary loss such that OC could dynamically decode position relative to the target from FPN activity (*Models 3 and 4*). We additionally tune the feedback projection from OC to FPN to optimize for task performance (*Model 4*). For




Fig. 4 | Model comparison. **a** Model Schematic. Models comprise a posterior parietal cortex (FPN, $n = 100$ units) module, which continuously receives 2D self-motion velocity and transiently receives 2D target position as inputs. PPC is recurrently connected with the oculomotor (OC, $n = 2$) module. Linear readout of FPN activity drives motor output, which encodes 2D acceleration (linear/angular). For simplicity, the traces show only one dimension. **b** In all models, connectivity weights shown in gray are fixed, while weights in green are optimized for task performance by minimizing the squared error between stopping position (z_t) and target position (z^*), averaged across trials. In models 3 and 4, weights shown in crimson are estimated by linear regression to minimize the squared error between

the activity of the OC units (r_t) and the relative target location ($\tilde{z}_t = z^* - z_t$), averaged across trials. Number of tuned parameters—200, 10,200, 400, and 600 for models 1, 2, 3, and 4, respectively. **c** *Left*: Correlation between gaze position (activity of the OC units) and the relative target position across test trials. *Middle*: Model performance in the test trials, quantified as variance in target position explained by stopping position (R -squared). *Right*: Correlation between target-tracking error (TTE) and steering error (SE). In all panels in **c**, violin plots represent distributions of statistics across 50 realizations of each model type, and cyan horizontal lines denote average statistics from data.

simplicity, we ignore recurrence within modules other than FPN. We trained the network weights (green weights) to reach the target in a time-bound manner via backpropagation-through-time. OC neurons in *Models 3 and 4* are constrained to encode the relative target location via linear regression (crimson weights) (Fig. 4b). In all models, observation and process noise were added to the sensory and motor units, respectively to prevent the models from developing a purely feed-forward control strategy (see “Methods”).

By construction, only the models constrained by the behavioral strategy explain participants’ eye movements (Fig. 4c–left; Pearson’s r : Model 1, 0.04 ± 0.1 ; Model 2, -0.02 ± 0.2 ; Model 3, 0.72 ± 0.1 ; Model 4, 0.7 ± 0.2). Of these two models, only the one with tuned feedback from OC to FPN had good task performance (Fig. 4c–middle; Supplementary Fig. 3a, B; R^2 : Model 1, 0.32 ± 0.2 ; Model 2, 0.69 ± 0.2 ; Model 3, 0.56 ± 0.1 ; Model 4, 0.67 ± 0.2). Notably, the performance of this model (*Model 4*) was almost as good as the task-optimized model in which all recurrent weights are tuned (*Model 2*) despite having substantially fewer tunable parameters (see figure caption). This suggests that recurrent frontoparietal-oculomotor (FPN-OC) interactions serve as a useful anatomical motif to support task performance. By selectively funneling the subspace of FPN activity that encodes the latent state (relative target position) into OC, *Model 4* enables learning efficiently (i.e., with fewer tunable parameters), thereby highlighting the computational significance of embodied cognition. Unlike other models, *Model 4* also recapitulates the trial-by-trial correlations that arise between TTE and steering error (SE) (Fig. 4c–right; Pearson’s r : Model 1, -0.06 ± 0.1 ; Model 2, 0.13 ± 0.2 ; Model 3, 0.25 ± 0.2 ; Model 4, 0.45 ± 0.2). This is because, in this model, error in estimating position results in poor target-tracking by the OC units that control eye

movements, and this error propagates to joystick movements via tuned feedback connections from OC to FPN.

To test whether *Model 4* recapitulates more granular aspects of the data, we tested this model under two different levels of observation noise that differed by an order of magnitude (to simulate visual and inertial conditions) and found that the performance generalized to settings with greater observation noise albeit at a lower precision (Fig. 5a; compare with Fig. 1d). At the same time, the activity of the OC units in the model recapitulated the dynamics of eye movements seen in experiments (Fig. 5b; compare with Fig. 2a; Supplementary Fig. 3c). Specifically, the two OC units appeared to track the x - y components of the relative target position, regardless of the magnitude of observation noise. However, it is impossible to precisely discern the actual target position at any given time due to the accumulation of noise. Consequently, the influence of target position on the OC network activity decreased during navigation. On the other hand, the influence of the stopping position increased, consistent with our experimental findings (Fig. 5c; compare with Fig. 4b). This suggests that the OC network encodes an internal estimate of the target location (i.e., belief), which is then used by FPN to control steering. Consistent with this interpretation, navigation performance deteriorated substantially when we prevented “eye movements” in the model by clamping the activity of OC units to zero (Fig. 5d, left; increase in error, ε : low noise, 95 ± 10 cm; high noise, 112 ± 16 cm), which agrees with the performance drop in humans instructed to avoid eye movements (Lakshminarasimhan et al.³⁰; Fig. 5d, right; $\varepsilon = 62 \pm 27$ cm).

Finally, we tested whether distributing the computation across the frontoparietal network in this manner enables the network to learn representations that resemble the brain. We previously showed that it




Fig. 5 | Strategy-constrained model (*Model 4*) recapitulates behavioral data.

a Comparison of the radial distance (from origin) of the model’s stopping position against the radial distance of the target (left), as well as the angular eccentricity of the model’s stopping position versus target angle (right) across all trials. Black lines have unit slope. Visual (cyan) and vestibular (red) conditions are simulated by changing the variance of observation noise during testing (see “Methods”).

b Activity dynamics of the two OC units that were constrained to encode the x and y components of the relative target position in a subset of test trials. Horizontal line denotes zero. **c** Coefficients corresponding to the target and stopping position

from the linear regression model that best explained the OC unit activity (averaged across the two units). Error bars denote standard errors estimated by bootstrapping. **d** Error in stopping position, on typical trials (control) and trials in which the model OC units are inhibited by clamping their activity at zero (*inh*). For comparison, the steering error in free-gazing (control) and inhibited gaze (*inh*) trials in humans is shown on the right. Error bars denote standard errors of the mean estimated by bootstrapping across 50 realizations of the model or across participants, as applicable.

is possible to dynamically decode target distance from population activity in monkey PFC³² and PPC³³ (Fig. 6a, b–top; R^2 : 0.67 ± 0.2). We first verified that target distance can also be decoded from the model FPN activity by training a linear decoder (Fig. 6a–bottom). Decoders trained on the activity of each of the four models explained a substantial variance in target distance (Fig. 6b–bottom; Model 1, $R^2 = 0.78 \pm 0.13$; Model 2, 0.88 ± 0.06 ; Model 3, 0.82 ± 0.08 ; Model 4, 0.8 ± 0.1). We wanted to know whether models differed in their ability to capture the fine-grained structure of neural activity in monkeys. Therefore, we reanalyzed the monkey data to determine the subspace of activity that is most informative about target distance. To do this, we first denoised the data by reducing it to the top 16 principal components. In this denoised subspace, we found that more than 90% of explainable variance in target distance was concentrated within the top five principal components of PPC activity (Fig. 6c–top). We estimated the targeted participation ratio (TPR; see “Methods”) to quantify the extent to which target distance was concentrated within the top few principal components. A TPR of 1 indicates that information about the variable of interest (target distance) is uniformly distributed across all PCs while lower values indicate that the variable can be decoded from activity in the leading PCs. We estimated the TPR and found it to be low (TPR: 0.22 ± 0.1), suggesting that few leading PCs of PPC activity are sufficient to decode target distance. Strikingly, target distance information was also largely contained within the top few principal components of the FPN activity in models with the cognitive constraint but not in purely task-optimized models (Model 1, TPR = 0.46 ± 0.1 ; Model 2, 0.48 ± 0.05 ; Model 3, 0.3 ± 0.05 ; Model 4, 0.28 ± 0.03 ; Fig. 6c–bottom). This is because, by explicitly projecting the belief about the relative target position into a low-dimensional OC activity, the strategy-constrained model allows this signal to undergo recurrent amplification which increases its variance. In contrast, such amplification does not take place in the model lacking cognitive constraints where belief signals remain buried in low-variance modes (bottom principal components) of the population activity.

Model predictions

To study the implications of the embodied strategy, we simulated the effect of stimulating OC units during the trial by systematically injecting a brief (0.2 s) external input pulse of fixed amplitude that was

either *positive* or *negative* into the model OC unit that encoded either the *horizontal* or *vertical* component of the believed target position, yielding four types of perturbations (right/left/up/down). The stimulation was delivered at various intervals following the removal of target position information (0, 0.4, 0.8, 1.2, 1.6 s). Ideally, these perturbations should produce stereotyped rightward, leftward, upward, or downward saccades. In contrast, we noticed substantial variability in the evoked saccades regardless of the timing of the stimulation (Fig. 7a and Supplementary Fig. 4a). This variability could not be simply attributed to OC units participating in recurrent dynamics with FPN: Variability in saccades evoked was substantially lower when simulating the same stimulation protocol on the architecturally identical model lacking cognitive constraints (Model 2, Fig. 7b; Supplementary Fig. 4b).

Furthermore, the variability of the evoked saccade gradually decreases as the stimulation occurred later in the trial (Fig. 7c). Since the variability in the belief (about the relative target position) also decreases as the trial progresses (see Fig. 5b), we asked whether the high variability in evoked saccade in the model was due to the trial-by-trial variability in beliefs. We found that the saccade magnitude was indeed strongly anti-correlated with the belief at the time of stimulation (mean Pearson's ρ corr. coefficient $\pm$ SD across all conditions: *x*-component: -0.39 ± 0.2 , *y*-component: -0.60 ± 0.1 ; Fig. 7d), suggesting that the belief substantially influences the properties of the evoked saccade. The negative sign in the correlation can be understood by recognizing that, for accurate navigation, the belief about the relative target position always approaches zero (recall Fig. 5b). Thus, saccades that are congruent with belief updates should be negative or positive depending on whether the belief is above or below zero, respectively. Furthermore, because the model encodes the *x*-*y* components of the beliefs in the two-dimensional eye position, memory about one component can persist even if the other component is perturbed by stimulation. For example, a stimulation intended to evoke a downward saccade might produce a saccade that is biased rightward or leftward depending on the horizontal component of the belief at the time of stimulation. Consequently, stimulation does not completely disrupt belief updates when the beliefs are embodied in eye movements. Indeed, the effect of stimulation on navigation performance (quantified by steering errors) is relatively small in this model compared to the model that does not rely on the embodied




Fig. 6 | Neural representation in the strategy-constrained model matches monkey data. a Top: Anatomical location of the multi-electrode arrays (Red—monkey B, Green—monkey Q, Blue—monkey S) superimposed on the 3D reconstructed brain of monkey S (IPS—intraparietal sulcus, STS—superior temporal sulcus, LF—lateral fissure). Example trials showing the performance of a linear decoder trained to estimate target distance from a population of simultaneously recorded neurons in left monkey posterior parietal cortex. Bottom: Like the top panel but decoded from the FPN activity of the strategy-constrained model (Model 4).

b Fraction of variance in actual target distance explained by a linear decoder trained on the activity of monkeys' PPC neurons (top) and the models' FPN neurons (bottom). **c** Fraction of cumulative variance explained by the decoder as a function of the number of principal components used for decoding in the monkey data (top) and the models (bottom). Data for the average monkey is overlaid in the bottom panel for easy comparison. Error bars in (b) denote standard errors in the mean estimated by bootstrapping.




Fig. 7 | Model predictions for stimulation experiments. a Saccades evoked by stimulating the OC units at different delays with respect to the time at which the target position cue disappeared, while Model 4 performed the navigation task. Each line corresponds to a different trial, and color denotes the type of stimulation (see text). **b** Similar to A but showing the response of Model 2. **c** Variability in the final position of the saccade across trials as a function of the timing of the stimulation. Error bars denote ± 1 standard error of the mean estimated by bootstrapping. **d** Comparison of the magnitude of saccade against the model's belief about the

relative target position at the time of stimulation for 4 different stimulation sites (generating rightward, leftward, downward, or upward saccades). *Thin and thick circles* denote the *x* (horizontal) and *y* (vertical) eye movement components, respectively. Data plotted are from all stimulation times. **e** Movement trajectories of the model under baseline trials (left) and during stimulated trials (middle). Right: Trajectory of Model 2 during stimulation. **f** Steering errors of Model 4 (solid) and Model 2 (dashed) in stimulated trials. *Gray line* denotes the error during baseline trials. Error bars denote ± 1 standard error of the mean estimated by bootstrapping.

strategy, where stimulation can disrupt both components of the belief (Model 2, Fig. 7e, f). Therefore, the model predicts that the embodied strategy should lead to a paradoxical effect wherein stimulating the OC areas should evoke highly variable saccades, yet only modestly affect task performance.

Discussion

Using a naturalistic behavioral paradigm defined by dynamic action/perception loops and unconstrained eye movements, we show that the dynamic belief about goal location is reflected in the subjects' OC behavior. By demonstrating that goal tracking is also observed in a purely inertial navigation version of the task in the absence of optic flow, we showed that this behavioral strategy is not driven by an active sensing strategy and instead provides strong support for the embodiment hypothesis. Specifically, we also show that these task-relevant eye movements reflect an embodiment of the *subjects' dynamically evolving internal beliefs about the goal*, and not just the initial location of the target. Furthermore, we found that a neural model constrained by the cognitive strategy adopted by animals explains behavioral and neural data better than purely task-optimized models. Thus, we believe that computations needed for steering could be distributed across multiple brain networks, including frontoparietal and OC networks, where the frontoparietal network would be involved in temporally

integrating self-motion but outsources belief state representation to the OC network, resulting in eye movements that dynamically track beliefs. We propose that mixing of signals between association and (oculo-)motor areas results from a distributed brain architecture that evolved to implement computations by grounding subjective beliefs about latent world states in states of the body.

We show for the first time that humans persistently use their eyes to track latent goal locations, even in the absence of visual navigational cues. This was made possible by using a naturalistic behavioral paradigm, as opposed to highly controlled tasks that restrict motor behavior and hinder the ability of the brain to use the algorithms that generate natural behaviors. In our previous study³⁰, we showed that the eyes follow the latent target in a visual-only condition where trial durations were much smaller (~ 2 s), while inhibiting these eye movements worsened performance significantly, highlighting their computational importance. Here, we show that target-tracking can happen for much longer trial durations (> 8 s), even in the absence of visual stimuli (inertial condition), and despite the presence of reflexive oculomotor processes (i.e., VOR). Specifically, we showed that TTE was kept low for most of the trial in both visual and inertial conditions. Although the error increased faster in the inertial condition, the ability to execute smooth-pursuit-like eye movements lasting several seconds in complete darkness is nonetheless unexpected and surprising. The

ability to execute smooth eye movements in the absence of visual stimulation suggests that embodiment arises from flexible rerouting of signals within the brain to achieve computational efficiency, rather than a byproduct of activating the sensory pathways engaged by the task. We were able to map tracking errors to steering errors, where the correlation between these two quantities was higher later within a trial. Overall, these findings show that the eyes follow the believed goal location, which shifts over time, from the actual location of the target when first presented, to the final stopping location. Eye movements have been found to facilitate working memory computations in non-navigation settings^{34–37}, foveal processing of optic flow³⁸, and other discrete domains^{39–41}. Our findings complementarily expand this body of work on embodied cognition to a naturalistic sequential decision task like navigation.

We generated a belief estimate as a dynamic weighted sum of the relative target and stopping positions, whose weights exhibit an almost perfect reversal between start and end of trial. This reconstructed belief modulates the saccades' amplitude and direction, which proved crucial in the inertial condition, as they allowed the eyes to successfully counter the VOR and track the believed goal location. VOR cancellation has been previously studied using targets that participants were required to fixate during passive yaw rotations^{42,43}. Here, we present evidence of volitional target-tracking eye movements countering the VOR in a naturalistic navigation setting. Importantly, these eye movements were driven dynamically by the belief about the relative goal location as participants actively steered toward it.

These experimental results provide strong support for the *cognitive embodiment* hypothesis, whereby allowing dynamic beliefs about the relative target location to continuously modulate eye movements, the OC circuit reduces the computational burden on circuits involved in working memory. This perspective is related but possibly somewhat distinct from prevailing views on embodiment. Traditional accounts emphasize the idea that cognitive representations are not symbolic but rather instantiated in sensory and motor pathways⁴⁴. In contrast, we propose a computational perspective, where embodiment arises from cognitive computations that repurpose specialized systems honed through evolution. Whether the embodiment in this computation requires the actual movement of the eyes or merely the OC neural circuit dynamics remains to be explored.

Motivated by the support that our findings offer for the embodiment hypothesis as a strategy for navigational control, we propose a recurrent neural network (RNN) model of the underlying computation that the brain uses to exploit eye movements: a circuit model in which the believed target location is encoded in OC neurons that have tuned bidirectional connections with the frontoparietal cortex (FPN) that integrates self-motion signals. This model, with substantially fewer tuned connections, was able to perform similarly to a model in which learning was accomplished by tuning all recurrent connections within the FPN. Notably, in addition to performing the steering task accurately, this model recapitulated human eye movements, thereby providing a normative explanation for why subjective beliefs are externalized in eye movements. In contrast to purely task-optimized models, this strategy-constrained model also correctly predicted that the leading principal components of the monkey posterior parietal cortex activity should encode their position relative to the goal.

The ability to predict neural responses accurately has made task-optimized neural network models an increasingly common tool for probing neural mechanisms underlying a wide range of computations, including image recognition, speech perception, working memory, and motor control^{45–48}. However, such an approach neither explains why computations are distributed across functionally distinct modules nor allows modularity to emerge on its own. Our findings directly address this dual challenge by providing a possible computational benefit: both can be explained by augmenting task-optimized models with constraints obtained by analyzing the strategy used by animals to

solve the task. Since naturalistic tasks increase the likelihood of engaging strategies that the brain evolved to use in the real world, we believe combining such task designs with strategy-constrained computational modeling can shed further light on distributed neural computations in other domains.

Multiple brain areas, such as the hippocampus, entorhinal cortex, retrosplenial cortex, posterior parietal cortex, and prefrontal cortex, contribute to navigation computations^{49,50}, but parietal circuits are considered to be of greater importance in egocentric navigation⁵¹. Additionally, neurons in the monkey posterior parietal cortex and dorsolateral prefrontal cortex have been identified as candidate regions involved in computing beliefs during this task^{32,33}. Anatomical studies in monkeys have also found extensive reciprocal connectivity between frontoparietal brain regions and neural circuits involved in eye movements, including frontal eye fields (FEFs), supplemental eye fields, and area 8ar^{52–55}; thus, we focused here on FPN and its connections to the OC circuit.

For ease of interpretability, we have considered a minimal model of the OC module with only two units. However, similar results could also be obtained by modeling the OC module as another RNN with a 2-dimensional output that controls horizontal and vertical eye position, following previous work^{56,57}. Such an expanded model would still account for the amplification of belief signals seen in the monkey neural data as long as the neural activity in the OC module is low-dimensional. Furthermore, the computational benefit of learning FPN-OC interactions (over recurrent weights within FPN) will also hold, provided the OC module has fewer units than FPN.

The model makes two concrete predictions to be tested in future experiments. First, the communication subspace between the FPN and OC regions should represent the subjective beliefs about the relative position of the target. Second, stimulation of the OC regions that provide feedback to FPN should have a modest yet clear effect on navigation performance. Regions with bidirectional connectivity with the posterior parietal and dorsolateral prefrontal cortex, such as area 8ar, FEF, supplementary eye fields^{52–55,58} are all excellent candidates for testing these predictions. More broadly, the proposed circuit model suggests that embodied cognition might be the reflection of a strategy by which the brain exploits distributed neural circuits and sensorimotor pathways structured through evolution in order to learn efficiently.

Navigation is a complex sensorimotor process in which multiple sensory modalities take part. Previous studies have shown the contributions of proprioceptive⁵⁹, tactile⁶⁰, and auditory⁶¹ stimuli in navigation performance, in which eye movements remained task-relevant (i.e., tracking self-motion) despite the lack of visual cues. In our non-visual condition, vestibular cues were the dominant sensory modality. Nevertheless, tactile/somatosensory cues from the air or the seat on the participants' skin due to platform motion were inevitable. Since such cues existed and all sensory cues provided under this condition are associated with inertial motion, we chose to name this condition "Inertial" instead of "Vestibular," despite tactile contributions in navigation being small^{61–63}.

Although our paradigm is less restrictive compared to traditional neuroscience tasks, head fixation was necessary to render controlled vestibular stimuli due to limitations of the motion platform (see "Methods"). Previous studies have shown that head movements facilitate gaze shifts towards targets in space, but are necessary only when the required gaze shift exceeds 40°, which is more than the target angles presented in our task⁶⁴. Head movements have also been associated with navigation performance through the control of gaze, in synergy with eye movements^{65–67}. Recent studies showed that gaze location in space (i.e., sum of eye plus head movement) is similar between head-free and head-fixed navigation, which was not the case when restricting eye movements, suggesting that gaze control and its effect on navigation depend primarily on the OC system^{68,69}.

Therefore, although head fixation deprives participants of head movements, the task-related variable, which is gaze, should be unaffected as it is sufficiently subserved by eye movements in our task.

Our emphasis on the role of eye movements in dynamically tracking latent beliefs complements previous studies that highlight the information-gathering role of temporally structured eye movements^{70–74} and contextualize findings from controlled studies that report an influence of short-term memory on smooth pursuit eye movements^{75–77}.

The proposed model builds on recent efforts that take advantage of well-characterized behavioral strategies to gain mechanistic insights via neural network models. For example, one study demonstrated a need to incorporate structural priors into RNNs (via pre-training) for recapitulating suboptimal choice by rats that fail to account for serial correlations in stimulus statistics across trials⁷⁸. Likewise, another study varied interaction strengths in an RNN model to account for a stress-induced switch from active to passive coping strategy in zebrafish⁷⁹. Another recent study endowed RNNs with an auxiliary loss function to mimic human error patterns in an intuitive physics task⁸⁰. However, to our knowledge, no study has harnessed a dynamic, within-trial behavioral strategy to inform the design of such models, nor shown the need to use modular architectures to replicate animal behavior. The present study achieves both by using a naturalistic task to tap into an innate, evolutionarily conserved behavioral strategy for tracking one's beliefs over time.

Multiple brain areas contribute to navigation computations, such as the hippocampus, entorhinal cortex, retrosplenial cortex, posterior parietal cortex, and prefrontal cortex.

Recent work has contributed statistical tools to infer latent beliefs from behavior^{75,81–84}; our findings and proposed model could facilitate the development and application of these tools in sequential decision behaviors. Additionally, our model agrees with recent work showing the benefit of modular architectures in computing latent beliefs⁸⁵. Although it is simplistic, it can guide future studies that probe neural mechanisms underlying the involvement of the OC system in cognition. Also, the learning efficiency of the distributed architecture has important implications for realizing biologically inspired artificial intelligence in embodied agents, especially robotics.

Embodiment and its computational role in cognition have been largely overlooked by the neuroscience community, and yet their importance on artificial agents is the subject of an ongoing debate with decades-long roots^{86–92}. Our study underlines embodiment as a cornerstone of human intelligence that any attempt to human-like computations and representations in machines should seriously consider.

Methods

Experimental model and subject details

Eight subjects (6 male, 2 female; all adults in the age group 18–32) participated in the eye-tracking experiments. Apart from two subjects, all subjects were unaware of the purpose of the study. Experiments were first performed on the above two subjects before testing others. All experimental procedures were approved by the Institutional Review Board at the authors' former institution, Baylor College of Medicine, and all subjects signed an approved consent form. Participants were compensated at the rate of \$20/h, regardless of their task performance.

Method details

Behavioral task—visual, inertial, and multisensory motion cues. The task required subjects to navigate to a remembered location on a horizontal virtual plane using a joystick, rendered in 3D from a forward-facing vantage point above the plane. Visual and/or vestibular sensory feedback was provided. Visual feedback was stereoscopic, composed of flashing triangles to provide self-motion information, but

no landmarks. Vestibular feedback was generated by a moving platform approximating the properties of their virtual self-motion.

Participants pressed a button on the joystick to initiate each trial and were tasked with steering to a randomly placed target that was cued briefly at the beginning of the trial. A short tone at every button push indicated the beginning of the trial and the appearance of the target. After one second, the target disappeared, which was a cue for the subject to start steering. Participants were instructed to stop at the remembered target location, and then push the button to register their final position and start the next trial. Participants did not receive any feedback about their performance. Prior to the first session, all participants performed about ten practice trials to familiarize themselves with joystick movements and the task structure.

Participants performed the task under three sensory conditions, which were interleaved randomly across trials. In the visual condition, participants had to navigate towards the remembered target position given only visual information (optic flow); no vestibular sensory feedback was provided during motion. In the multisensory (combined) condition, subjects were provided with both visual and inertial (vestibular/somatosensory) information during their movement. In the inertial condition, after the target disappeared, the entire visual stimulus was shut off too, leaving the subjects to navigate in complete darkness using only inertial cues.

Independent of the manipulation of the sensory information, the properties of the motion controller also varied from trial to trial. Participants experienced different time constants in each trial, which affected the type and amount of control that was required to complete the task. In trials with short time constants, joystick position mainly controlled velocity, whereas in trials with long time constants, joystick position approximately controlled the acceleration (explained in detail in the *Control Dynamics* Methods section in Stavropoulos et al.³¹).

Each participant performed a total of about 1450 trials (mean $\pm$ SD: 1450 ± 224), split equally among the three sensory conditions (mean $\pm$ SD—vestibular: 476 ± 71 , visual: 487 ± 77 , multisensory: 487 ± 77).

Performance under the multisensory condition was comparable to the visual condition, in both steering and target-tracking using eye movements (see Supplementary Fig. S5; there was no statistically significant advantage in performance of the multisensory over the visual condition), and was thus omitted from the analysis.

Visual stimulus. The virtual world comprised a ground plane whose textural elements whose lifetimes were limited (~ 250 ms) to avoid serving as landmarks. The ground plane was circular with a radius of 37.5 m (near and far clipping planes at 5 cm and 3750 cm, respectively), with the subject positioned at its center at the beginning of each trial. Each texture element was an isosceles triangle (base $\times$ height $\times 5.95 \times 12.95$ cm) that was randomly repositioned and reoriented at the end of its lifetime. The floor density was held constant across trials at $\rho = 2.5$ elements/m². The target, a circle of radius 25 cm whose luminance was matched to the texture elements, flickered at 5 Hz and appeared at a random location between $\theta = \pm 38^\circ$ of visual angle at a distance of $r = 2.5 - 5.5$ m (average distance $\bar{r} = 4$ m) relative to where the participant was stationed at the beginning of the trial. The stereoscopic visual stimulus was rendered in an alternate frame sequencing format, and subjects wore active-shutter 3D goggles to view the stimulus.

Experimental setup. The participants sat comfortably on a chair mounted on an electric motor allowing unrestricted yaw rotation (Kollmorgen motor DH142M-13-1320), itself mounted on a six-degree-of-freedom motion platform (comprised of MOOG 6DOF2000E). Subjects used an analog joystick (M20U9T-N82, CTI electronics) with two degrees of freedom and a circular displacement boundary to control their linear and angular speed in a virtual environment based

on visual and inertial stimuli. The visual stimulus was projected (Canon LV-8235 UST Multimedia Projector) onto a large rectangular screen (width $\times$ height: 158 $\times$ 94 cm or 136° $\times$ 110° in visual angle) positioned in front of the subject (77 cm from the rear of the head) and centered such that there is 65° of visual angle available on the screen below the participants' viewing height (i.e., eye level). Participants wore crosstalk-free ferroelectric active-shutter 3D goggles (RealD CE4s) to view the stimulus. Participants wore headphones generating white noise to mask the auditory motion cues. The participant's head was fixed on the chair using an adjustable CIVCO FirmFit Thermo-plastic face mask. Eye movements were monitored at 120 Hz using ISCAN 06-604-0302 binocular eye-tracking and ISCAN ETL 500 software.

Joystick control. Participants navigated in the virtual environment using a joystick placed in front of the participant's midline, in a holder mounted on the bottom of the screen. This ensured that the joystick was parallel to the participant's vertical axis, and its horizontal orientation aligned with the forward movement axis. The joystick had two degrees of freedom that controlled linear and angular motion. Joystick displacements were physically bounded to lie within a disk, and digitally bounded to lie within a square. Displacement of the joystick over the anterior-posterior (AP) axis resulted in forward or backward translational motion, whereas displacement in the left-right (LR) axis resulted in rotational motion. The joystick was enabled after the disappearance of the target. To avoid skipping trials and abrupt stops, the button used to initiate trials was activated only when the participant's velocity dropped below 1 cm/s.

The joystick controlled both the visual and inertial stimuli through an algorithm that involved two processes. The first varied the CD, producing velocities given by a leaky integration of the joystick input, mimicking an inertial body under viscous damping. The time constant of the leak (leak constant) was varied from trial to trial, according to a random walk. The maximum linear and angular velocities are scaled together with the leak constant across trials, such that a target at a given distance can be reached in the same amount of time under different leak constants (assuming an ideal bang-bang controller).

The second process was a motion cueing (MC) algorithm applied to the output of the CD process, which defined physical motion that approximated the accelerations an observer would feel under the desired CD, while avoiding the hardwired constraints of the motion platform. This MC algorithm trades translation for tilt, allowing extended acceleration without hitting the displacement limits of the platform.

Each motion trajectory consisted of a linear displacement in the 2D virtual space combined with a rotation in the horizontal plane. While the motion platform could reproduce the rotational movement using the yaw motor (which was unconstrained in movement range and powerful enough to render any angular acceleration or speed in this study), its ability to reproduce linear movement was limited by the platform's maximum range in total displacement and maximum velocity (but not acceleration). The MC algorithm takes advantage of the gravito-inertial ambiguity⁹³ inherent to the vestibular organs^{94–96} to circumvent this limitation. As the otolith organs in the inner ear sense both linear acceleration (A) and gravity (G), i.e., they sense the gravito-inertial acceleration (GIA): $F = G + A$, a forward acceleration of the head (a_x , expressed in g , with $1g = 9.81 \text{ m/s}^2$) and a backward pitch (by an angle θ , in radians) will generate a total gravito-inertial acceleration $F_x = \theta + a_x$. The MC took advantage of this ambiguity to replace linear acceleration with tilt. Specifically, it controlled the motion platform to produce a total GIA that matched the linear acceleration of the simulated motion in the virtual environment.

Even though this method is generally sufficient to ensure that platform motion remains within its envelope, it does not guarantee it. Thus, the platform's position, velocity, and acceleration commands

were fed through a sigmoid function f . This function was equal to the identity function ($f(x) = x$) as long as motion commands were within 75% of the platform's limits, so these motion commands were unaffected. When motion commands exceed this range, the function bends smoothly to saturate at a value set slightly below the limit, thus preventing the platform from reaching its mechanical range (in position, velocity, or acceleration) while ensuring a smooth trajectory. Thus, if the desired motion exceeds 75% of the platform's performance envelope, the actual motion of the platform is diminished, such that the total GIA actually experienced by the participant ("Actual Platform GIA") may not match the desired GIA. If left uncorrected, these GIA errors would result in a mismatch between inertial motion and the visual VR stimulus. To prevent these mismatches, we designed a loop that estimates GIA error and updates the simulated motion in the visual environment. For instance, if the joystick input commands a large forward acceleration and the platform is unable to reproduce this acceleration, then the visual motion is updated to represent a slower acceleration that matches the platform's motion.

Altogether, the CD and MC algorithms are applied sequentially as follows: (1) The velocity signal produced by the CD process controls the participant's attempted motion in the virtual environment. (2) The participant acceleration in the VR environment is calculated and inputted to the MC algorithm ("Desired Platform GIA"). (3) The MC cueing computes the platform's motion commands and the actual platform GIA is computed. (4) The difference between the Desired GIA motion actual GIA (GIA error) is computed and used to update the motion in the virtual environment. (5) The updated position is sent to the visual display.

These two processes (CD and MC) are explained in more detail in Stavropoulos et al.³¹.

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 70 cm above the ground plane. Spike2 software (Power 1401 MkII data acquisition system from Cambridge Electronic Design Ltd) was used to record and store the target location (r, θ), subject's position ($\tilde{r}, \tilde{\theta}$), horizontal positions of left and right eyes (α_l and α_r), vertical eye positions (β_l and β_r) and all event markers for offline analysis at a sampling rate of 833 $\frac{1}{3}$ Hz.

Quantification and statistical analysis

Statistics and reproducibility. 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, SD, 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 and describe the outcome as significant if $p < 0.05$. No statistical method was used to pre-determine sample size. No data were excluded from the analyses. The three sensory conditions of the experiment were interleaved randomly across trials. Participants were not split into groups, so blinding was not needed. Detailed methods on how different statistics were estimated are below.

Bias estimation. In each sensory condition, we first computed the τ -independent bias for each subject; we regressed (without an intercept term) each subject's response positions ($\tilde{r}, \tilde{\theta}$) against target positions (r, θ) with respect to the starting position, separately for the radial ($\tilde{r}$ vs r) and angular ($\tilde{\theta}$ vs θ) coordinates. The radial and angular multiplicative biases were quantified as the slope of the respective regressions (Fig. 2a). In addition, we followed the same process to calculate bias terms within three τ groups of equal size (Fig. 2c).

Characterizing eye, target and stop position in eye coordinates. For convenience, we express the subject's *actual* eye position using the following two standard degrees of freedom: (i) Conjunctive horizontal movement of the two eyes, quantified here as the mean lateral position of the two eyes, $\alpha = (\alpha_{\text{left}} + \alpha_{\text{right}})/2$, (ii) Conjunctive vertical movement of the two eyes, quantified here as $\beta = (\beta_{\text{left}} + \beta_{\text{right}})/2$. Disjunctive horizontal and vertical eye movements (horizontal and vertical vergence, respectively) were not considered for our analysis, because of the documented difficulty in humans to execute vergence to imagined moving objects^{30,97}.

To test whether participants' eyes tracked the location of the (invisible) target, we need the target and eye positions to be on the same reference frame. Therefore, we transformed the target position from world to eye coordinates. Let s denote the stage of trial evolution, i.e., percentage of total distance traveled from 0 to 100%. We denote the target position in world coordinates as (x_t, y_t, z_t) , relative to the midpoint of the participant's eyes at trial stage s . The target position in eye coordinates—relative to fixating at the point $(0, 0, \infty)$ —relates to its position in world coordinates as (Supplementary Fig. S6):

$$\alpha_t(s) = \tan^{-1} \left(\frac{x_t}{\sqrt{y_t^2 + z_t^2}} \right) \quad (1.1)$$

$$\beta_t(s) = \tan^{-1} \left(\frac{z_t}{\sqrt{y_t^2 + x_t^2}} \right) \quad (1.2)$$

Where $\alpha_t(s)$ and $\beta_t(s)$ are horizontal and vertical target positions at trial stage s , respectively. Note that z_t is determined only by the viewing height, and therefore remains constant. On the contrary, x_t, y_t change continuously as the participants steer in the virtual environment.

In approximately 8% of the trials, the subject traveled beyond the target. The target position in eye coordinates towards the end of these trials was outside the physical range of gaze. Therefore, we removed time points at which any of the two components of the target position in Eq. 1.1 or 1.2 exceeded 60° before further analysis (corresponds to gazing at objects at 40 cm distance, which is within the target radius). Such time points constituted less than 1% of the dataset, and including them did not qualitatively alter the results. Specifically, given that targets appear at distances between 250 and 500 cm, gazing locations between 60° and 90° downward towards the end of a trial correspond to distances of 0–40 cm from the participant's position, which is at most the final 16% of the total distance to the nearest possible target at 250 cm. However, our analysis also demonstrates that the contribution of eye movements to navigation is crucial and predictive of the final steering error until roughly 75% of the distance traveled (Fig. 2d), suggesting that participants themselves stop making use of eye movements when they are about three-fourths of the way to the goal even before the experimental apparatus limits their ability to track the target.

Similarly, we calculated $\alpha_s(s)$ and $\beta_s(s)$ as the horizontal and vertical stopping positions in eye coordinates.

Target-tracking error and belief analysis. We tested how target-tracking performance was associated with steering performance by estimating the correlation between steering and TTEs across trials (Fig. 2c, d). As mentioned above, we scaled trials according to the percentage of total distance traveled and computed this correlation as trials evolved. At every trial stage, *steering error* was given by the Euclidean distance between the target and stop positions in eye coordinates as $\varepsilon_s = \sqrt{(a_t - a_s)^2 + (b_t - b_s)^2}$, while *target-tracking error* (TTE) was given by the Euclidean distance between eye and target

position as $\varepsilon_t = \sqrt{(x_t - \hat{x}_t)^2 + (y_t - \hat{y}_t)^2}$, where ε_t is the TTE, and (x_t, y_t) and $(\hat{x}_t, \hat{y}_t)$ are the horizontal and vertical coordinates of the target and eye position, respectively. Chance-level TTE was estimated as the mean of the null distribution obtained by shuffling target positions across trials. The same method was used to calculate the correlation between steering and tracking errors for the *saccade-free eye movements* (discussed below).

To compute an estimate of the participants' belief about the target location we regressed the participants' eye position against the target and stop positions (multiple regression), obtaining a kernel of weights for each position over trial progression (Fig. 2e). As all trials are scaled equally this way, we regressed the eye positions at trial stage s against the corresponding target and stop positions for the horizontal and vertical components, separately. This provided us with regression weight kernels for the target and stop positions of each component from 0 to 100% of the total distance traveled.

To reconstruct this belief, we simply multiplied the target and stop positions with their respective weights at each trial stage s , for each participant (Fig. 3c).

Saccade analysis. For saccade detection, we estimated the instantaneous speed of eye movements as $(\dot{\alpha}^2 + \dot{\beta}^2)^{1/2}$ where α and β denote horizontal and vertical eye positions, respectively (as defined above), and the dot denotes a time derivative. Saccades were detected by identifying the time points at which the speed of eye movements crossed a threshold of $150^\circ/\text{s}$ (a threshold of $25^\circ/\text{s}$ yielded similar results). Specifically, saccade onset was detected as the time point at which the speed of eye movements crossed the threshold from below, and saccade offset as the time at which the speed dropped below the threshold. The amplitude of saccades was taken to be the average displacement of the position of the two eyes from saccade onset to 150 ms later ($\Delta\varphi = (\Delta\alpha^2 + \Delta\beta^2)^{1/2}$).

To explore the contribution of saccades in target-tracking, we generated saccade-free eye movements by subtracting the displacement of the eye position caused by saccades after target offset ($t \geq 1\text{s}$) (Fig. 3a, b). We removed the periods between saccade onset and offset from the eye velocity signal. The remaining signal was linearly interpolated and then integrated to calculate eye displacement independent of saccades (saccade-free eye displacement). Finally, the eye position at the time of target offset was added to the saccade-free eye displacement. We then computed the correlation between steering and tracking errors for the *saccade-free eye movements* (just as we did for the actual eye position; see *Comparing steering and target-tracking errors*).

To test that the eyes would reflect target beliefs even under the control of the VOR, we explored the relationship between the cumulative saccade amplitude in each trial and the corresponding steering error (Fig. 3b). We only considered the horizontal component of saccades ($\Delta\alpha$), which is aligned to the evoked VOR during rotation. Therefore, we estimated the Pearson's correlation coefficient between angular steering errors and horizontal cumulative saccade amplitudes.

To quantify the precise relationship between saccade amplitude and tracking error (TTE), we obtained a regression weight kernel by regressing horizontal and vertical amplitudes of the saccade ($\Delta\alpha$ and $\Delta\beta$) on horizontal and vertical TTEs ($\alpha - a_t$ and $\beta - b_t$), respectively, at various lags between $\pm 1\text{s}$ with l^2 regularization (Fig. 3c). Similarly, we computed the kernels for the stop position-tracking error (SPTE) and the belief-tracking error (based on reconstructed belief, see "Methods": *Target-tracking error and belief analysis*).

Finally, we computed the gain of the eye position with respect to the target, to evaluate the effect of saccadic eye movements on target-tracking (Supplementary Fig. 2c). Specifically, we regressed (without intercept) the eye positions at time t against the corresponding target

positions for the vertical and horizontal components, separately. We performed this regression for both the actual and the saccade-free eye positions.

Recurrent neural network models. We trained four different RNN models to solve the velocity control version of the task performed by human participants. All models comprise two modules: one recurrently connected population of 100 nonlinear (“sigmoidal”) units that we identify as the frontoparietal circuit (FPN) and an OC module comprising two linear units encoding vertical and horizontal eye position, where, for simplicity, we ignore the biomechanics of eye movement generation. The FPN module contains 4 input channels, two for conveying the 2D target location (z^*) encoded in the amplitude of a transient pulse delivered at the beginning of the trial and two for conveying continuous sensory feedback about the 2D self-motion velocity ($\dot{z}$) throughout the trial. There were 2 output channels from FPN, one each for controlling the velocity of the “hand” along the linear and angular axes of the joystick, i.e., movement acceleration ($\ddot{z}$). To mimic process noise, we added zero-mean additive Gaussian noise to the output channels, and the noisy output is temporally integrated and fed back to the network through the input channels conveying movement velocity, thereby closing the sensorimotor loop. Additive noise is also added to the input channels to simulate observation noise. This feedback mimics the functionality of the VR simulator that uses the joystick output to render real-time sensory feedback in the form of optic flow cues or vestibular cues in our experiments.

The equation governing the network dynamics was:

$$\tau \dot{h} = -h + \varphi(W^{\text{rec}}h + W^{\text{in}}x + W^{\text{PPC} \leftarrow \text{OC}}r) \quad (2.1)$$

$$r = W^{\text{OC} \leftarrow \text{PPC}}h \quad (2.2)$$

$$\ddot{z} = W^{\text{out}}h \quad (2.3)$$

where h and r represent population activity in FPN and OC, respectively. $x = (\tilde{z}, z^*)$ denotes the input to PPC where $\tilde{z} = \dot{z} + \varepsilon$ denotes the velocity corrupted by additive observation noise $\varepsilon \sim \mathcal{N}(0, \sigma_s^2)$. $\tilde{z}$ is the network output representing acceleration such that $\dot{z} = \int dz \tilde{z}$ where $\tilde{z} = \ddot{z} + \eta$ denotes acceleration corrupted by process noise $\eta \sim \mathcal{N}(0, \sigma_p^2)$. τ is the cell-intrinsic time constant, and $\varphi(\bullet) = \tanh(\bullet)$ is the neuronal nonlinearity. Matrices W^{rec} , W^{in} , W^{out} , $W^{\text{PPC} \leftarrow \text{OC}}$, and $W^{\text{OC} \leftarrow \text{PPC}}$ correspond to recurrent, input, output, frontoparietal, and parieto-frontal weights, respectively.

Model training and details. We trained the RNN models defined in Eq. (2) by tuning different sets of model parameters in each model. All models were initialized by drawing parameters from a normal distribution, $\mathcal{N}(0, g^2/M)$, whose variance was scaled down by M to ensure that the total input current to each neuron had unit variance. For example, $M=100$ for the elements of the recurrent weight matrix, W^{rec} , since there were 100 recurrently connected units. On the other hand, $M=4$ for the elements of the input weight matrix, W^{in} , since there were only 4 input channels. To facilitate comparison across models, we used $g=1.2$ uniformly for all models, although varying g ranging from 1 to 1.5 yielded results that were very similar for all models except the reservoir network (Model 1), which performed best for $g=1.2$. This is because, unlike other models, the reservoir network is strongly dependent on chaotic dynamics to generate appropriate outputs.

Models were trained to reach the target location within a certain time t and stay there for 0.6 s. t corresponded to the time taken when traveling along an idealized circular trajectory from the starting location to the target location at maximum speed. The time constant τ was set to 20 ms, and each training trial lasted between 2 and 3 s,

depending on the target location. In all four models, output weights W^{out} were updated at the end of each trial to minimize the loss function $L = \sum_{t > t'} |z(t) - z^*|^2$, using gradient descent. In addition to output weights, we updated the recurrent weights W^{rec} in Model 2 and feedback weights from OC to FPN ($W^{\text{FPN} \leftarrow \text{OC}}$) in Model 4 to minimize L via backpropagation-through-time. In models 3 and 4, behavioral strategy constraint was incorporated by training the weights from FPN to OC ($W^{\text{OC} \leftarrow \text{FPN}}$) were trained to minimize an auxiliary loss function $L_{\text{aux}} = \sum_t |r(t) - \tilde{z}(t)|^2$ by linear regression, where the relative target location $\tilde{z}(t) = z(t) - z^*(t)$. Since abruptly updating $W^{\text{OC} \leftarrow \text{FPN}}$ in conjunction with other weights hampered learning, we updated $W^{\text{OC} \leftarrow \text{PPC}}$ incrementally as $W^{\text{OC} \leftarrow \text{PPC}} \leftarrow \alpha W^{\text{OC} \leftarrow \text{PPC}}_{\text{old}} + (1 - \alpha) W^{\text{OC} \leftarrow \text{PPC}}_{\text{new}}$ where $\alpha = 0.99$. The total number of free parameters were 200, 10 200, 400, and 600 for models 1, 2, 3, and 4, respectively.

Targeted participation ratio. PCA was performed by first concatenating data from all the trials into a single data matrix, so the PCs contain a mixture of task-relevant and intrinsic fluctuations in population activity. PCA was used as a denoising step because fitting decoders in a high-dimensional space can lead to unreliable weights. This enables us to reliably estimate dimensionality and perform decoding on a low-dimensional (16-dim, see below) subspace within which the population activity is largely confined in this task.

We estimated the TPR to quantify the extent to which a target variable of interest is concentrated within the few leading principal components of the population activity as:

$$\text{TPR} = \frac{\left(\sum_{i=1}^N \lambda_i \right)^2}{\sum_{i=1}^N \lambda_i^2} \quad (3)$$

where λ_i denotes the fraction of variance in the target variable that is explained by the i^{th} principal component of the population activity. It can be readily seen that $\text{TPR} = 1$ if $\lambda_i = \lambda \forall i$, i.e., all principal components explain an equal amount of variance, while $\text{TPR} = 1/N$ if the variance in the target variable is exclusively explained by the first principal component. We used $N = 16$ both for our analysis of PPC data and the neural network models.

Neural recordings. Three rhesus macaques (*Macaca mulatta*) (all male, 7–8 years old)—referred to as B, S, and Q for simplicity—participated in the experiments. All surgeries and experimental procedures were approved by the Institutional Review Board at Baylor College of Medicine and were in accordance with National Institutes of Health guidelines.

Monkeys were chronically implanted with a lightweight polyacetal ring for head restraint and scleral coils for monitoring eye movements (CNC Engineering, Seattle, WA, USA). Utah arrays were chronically implanted in area 7a of the Posterior Parietal Cortex in the left hemisphere of all three monkeys using craniotomy. Prior to the surgery, the brain area was identified using structural MRI to guide the location of the craniotomy. After craniotomy, the array was pneumatically inserted after confirming the coordinates of the target area using known anatomical landmarks.

At the beginning of each experimental session, monkeys were head-fixed and secured in a primate chair placed on top of a platform (Kollmorgen, Radford, VA, USA). All methods regarding these recordings have been previously described in Lakshminarasimhan et al.³³.

Recordings were performed using 96-channel multielectrode arrays in monkeys Q and B and a 48-channel array in monkey S. All channels were functional, and spike sorting was performed as explained in Lakshminarasimhan et al.³³. The vast majority of the channels had at least one single-unit. Some channels (~20%) had only multi-unit activity, which was not considered for the analysis. Since the

decoding analysis benefits from datasets with a large number of simultaneously recorded units, we restrict our focus to a subset of six sessions with the highest yield, two from each monkey (with an average yield of 85 units, 101 units, and 41 units from monkeys Q, B, and S, respectively).

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

Data generated in this study are available online at the G-Node database (https://gin.g-node.org/akis_stavropoulos/belief_embodiment_through_eye_movements_facilitates_memory-guided_navigation).

Code availability

MATLAB and Python code^{98,99} implementing all quantitative analyses and modeling, respectively, is available online (https://github.com/AkisStavropoulos/eye_movement_analysis, DOI: 10.5281/zenodo.15646521 and <https://github.com/kaushik-l/firefly-rnn>, DOI: 10.5281/zenodo.7434801).

References

- Cooper, R. P. & Peebles, D. Beyond single-level accounts: the role of cognitive architectures in cognitive scientific explanation. *Top. Cogn. Sci.* **7**, 243–258 (2015).
- Krakauer, J. W., Ghazanfar, A. A., Gomez-Marín, A., MacIver, M. A. & Poeppel, D. Neuroscience needs behavior: correcting a reductionist bias. *Neuron* **93**, 480–490 (2017).
- Miller, C. T. et al. Natural behavior is the language of the brain. *Curr. Biol.* **32**, R482–R493 (2022).
- Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
- Freedman, D. J. & Assad, J. A. A proposed common neural mechanism for categorization and perceptual decisions. *Nat. Neurosci.* **14**, 143–146 (2011).
- Murray, J. D. et al. A hierarchy of intrinsic timescales across primate cortex. *Nat. Neurosci.* **17**, 1661–1663 (2014).
- O’Connell, R. G., Shadlen, M. N., Wong-Lin, K. & Kelly, S. P. Bridging neural and computational viewpoints on perceptual decision-making. *Trends Neurosci.* **41**, 838–852 (2018).
- Cohen, M. R. & Maunsell, J. H. R. Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron* **70**, 1192–1204 (2011).
- Kiani, R. & Shadlen, M. N. Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* **324**, 759–764 (2009).
- Meirhaeghe, N., Sohn, H. & Jazayeri, M. A precise and adaptive neural mechanism for predictive temporal processing in the frontal cortex. *Neuron* **109**, 2995–3011.e5 (2021).
- O’Shea, J., Muggleton, N. G., Cowey, A. & Walsh, V. Timing of target discrimination in human frontal eye fields. *J. Cogn. Neurosci.* **16**, 1060–1067 (2004).
- Ruff, D. A. & Cohen, M. R. Attention can either increase or decrease spike count correlations in visual cortex. *Nat. Neurosci.* **17**, 1591–1597 (2014).
- Hayhoe, M. & Ballard, D. Eye movements in natural behavior. *Trends Cogn. Sci.* **9**, 188–194 (2005).
- Hayhoe, M. M., McKinney, T., Chajka, K. & Pelz, J. B. Predictive eye movements in natural vision. *Exp. Brain Res.* **217**, 125–136 (2012).
- Mao, D. et al. Spatial modulation of hippocampal activity in freely moving macaques. *Neuron* **109**, 3521–3534.e6 (2021).
- Gibson, E. J., Owsley, C. J. & Johnston, J. Perception of invariants by five-month-old infants: differentiation of two types of motion. *Dev. Psychol.* **14**, 407–415 (1978).
- Yamins, D. L. K. & DiCarlo, J. J. Using goal-driven deep learning models to understand sensory cortex. *Nat. Neurosci.* **19**, 356–365 (2016).
- Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T. & Wang, X.-J. Task representations in neural networks trained to perform many cognitive tasks. *Nat. Neurosci.* **22**, 297–306 (2019).
- Hadjidimitrakakis, K., Bakola, S., Wong, Y. T. & Hagan, M. A. Mixed spatial and movement representations in the primate posterior parietal cortex. *Front. Neural Circuits* **13**, 15 (2019).
- Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* **22**, 1677–1686 (2019).
- Ebbesen, C. L. et al. More than just a “Motor”: recent surprises from the frontal cortex. *J. Neurosci.* **38**, 9402–9413 (2018).
- Zaksas, D. & Pasternak, T. Directional signals in the prefrontal cortex and in area MT during a working memory for visual motion task. *J. Neurosci.* **26**, 11726–11742 (2006).
- Kleinman, M., Chandrasekaran, C. & Kao, J. A mechanistic multi-area recurrent network model of decision-making. In *Advances in Neural Information Processing Systems* (eds Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P. S. & Vaughan, J. W.) Vol. 34, 23152–23165 (Curran Associates, Inc., 2021).
- Yang, G. R. & Molano-Mazón, M. Towards the next generation of recurrent network models for cognitive neuroscience. *Comput. Neurosci.* **70**, 182–192 (2021).
- Michaels, J. A., Schaffelhofer, S., Agudelo-Toro, A. & Scherberger, H. A goal-driven modular neural network predicts parietofrontal neural dynamics during grasping. *Proc. Natl. Acad. Sci. USA* **117**, 32124–32135 (2020).
- Pinto, L. et al. Task-dependent changes in the large-scale dynamics and necessity of cortical regions. *Neuron* **104**, 810–824.e9 (2019).
- Musall, S., Urai, A. E., Sussillo, D. & Churchland, A. K. Harnessing behavioral diversity to understand neural computations for cognition. *Comput. Neurosci.* **58**, 229–238 (2019).
- Alefantis, P. et al. Sensory evidence accumulation using optic flow in a naturalistic navigation task. *J. Neurosci.* **42**, 5451–5462 (2022).
- Lakshminarasimhan, K. J. et al. A dynamic Bayesian observer model reveals origins of bias in visual path integration. *Neuron* <https://doi.org/10.1016/j.neuron.2018.05.040> (2018).
- Lakshminarasimhan, K. J. et al. Tracking the mind’s eye: primate gaze behavior during virtual visuomotor navigation reflects belief dynamics. *Neuron* <https://doi.org/10.1016/j.neuron.2020.02.023> (2020).
- Stavropoulos, A., Lakshminarasimhan, K. J., Laurens, J., Pitkow, X. & Angelaki, D. Influence of sensory modality and control dynamics on human path integration. *eLife* **11**, e63405 (2022).
- Noel, J.-P. et al. Coding of latent variables in sensory, parietal, and frontal cortices during closed-loop virtual navigation. *eLife* **11**, e80280 (2022).
- Lakshminarasimhan, K. J., Avila, E., Pitkow, X. & Angelaki, D. E. Dynamical latent state computation in the male macaque posterior parietal cortex. *Nat. Commun.* **14**, 1832 (2023).
- Ballard, D. H., Hayhoe, M. M. & Pelz, J. B. Memory representations in natural tasks. *J. Cogn. Neurosci.* <https://doi.org/10.1162/jocn.1995.7.1.66> (1995).
- Johansson, R., Holsanova, J., Dewhurst, R. & Holmqvist, K. Eye movements during scene recollection have a functional role, but they are not reinstatements of those produced during encoding. *J. Exp. Psychol. Hum. Percept. Perform.* <https://doi.org/10.1037/a0026585> (2012).
- Johansson, R. & Johansson, M. Look here, eye movements play a functional role in memory retrieval. *Psychol. Sci.* <https://doi.org/10.1177/0956797613498260> (2014).

37. Spivey, M. J. & Geng, J. J. Oculomotor mechanisms activated by imagery and memory: eye movements to absent objects. *Psychol. Res.* <https://doi.org/10.1007/s004260100059> (2001).
38. Chow, H. M. & Sperling, M. Eye movements during optic flow perception. *Vis. Res.* **204**, 108164 (2023).
39. Gold, J. I. & Shadlen, M. N. Representation of a perceptual decision in developing oculomotor commands. *Nature* **404**, 390–394 (2000).
40. Loetscher, T., Bockisch, C. J., Nicholls, M. E. R. & Brugger, P. Eye position predicts what number you have in mind. *Curr. Biol.* **20**, R264–R265 (2010).
41. Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M. & Sedivy, J. C. Integration of visual and linguistic information in spoken language comprehension. *Science* **268**, 1632–1634 (1995).
42. Crane, B. T. & Demer, J. L. Latency of voluntary cancellation of the human vestibulo-ocular reflex during transient yaw rotation. *Exp. Brain Res.* **127**, 67–74 (1999).
43. Lisberger, S. G. Visual tracking in monkeys: evidence for short-latency suppression of the vestibuloocular reflex. *J. Neurophysiol.* **63**, 676–688 (1990).
44. Mahon, B. Z. What is embodied about cognition? *Lang. Cogn. Neurosci.* **30**, 420–429 (2015).
45. Goehring, T., Keshavarzi, M., Carlyon, R. P. & Moore, B. C. J. Using recurrent neural networks to improve the perception of speech in non-stationary noise by people with cochlear implants. *J. Acoust. Soc. Am.* **146**, 705–718 (2019).
46. Martinez, J., Black, M. J. & Romero, J. On Human Motion Prediction Using Recurrent Neural Networks. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, 2017).
47. Mi, Y., Katkov, M. & Tsodyks, M. Synaptic correlates of working memory capacity. *Neuron* **93**, 323–330 (2017).
48. Schrimpf, M. et al. Integrative benchmarking to advance neurally mechanistic models of human intelligence. *Neuron* **108**, 413–423 (2020).
49. Wolbers, T., Wiener, J. M., Mallot, H. A. & Büchel, C. Differential recruitment of the hippocampus, medial prefrontal cortex, and the human motion complex during path integration in humans. *J. Neurosci.* **27**, 9408–9416 (2007).
50. Chrastil, E. R., Sherrill, K. R., Hasselmo, M. E. & Stern, C. E. Which way and how far? Tracking of translation and rotation information for human path integration. *Hum. Brain Mapp.* <https://doi.org/10.1002/hbm.23265> (2016).
51. Wolbers, T. & Wiener, J. M. Challenges for identifying the neural mechanisms that support spatial navigation: the impact of spatial scale. *Front. Hum. Neurosci.* **8**, 571 (2014).
52. Barbas, H. & Mesulam, M.-M. Organization of afferent input to subdivisions of area 8 in the rhesus monkey. *J. Comp. Neurol.* **200**, 407–431 (1981).
53. Huerta, M. F. & Kaas, J. H. Supplementary eye field as defined by intracortical microstimulation: connections in macaques. *J. Comp. Neurol.* **293**, 299–330 (1990).
54. Leichnetz, G. R. Connections of the medial posterior parietal cortex (area 7m) in the monkey. *Anat. Rec.* **263**, 215–236 (2001).
55. Stanton, G. B., Bruce, C. J. & Goldberg, M. E. Topography of projections to posterior cortical areas from the macaque frontal eye fields. *J. Comp. Neurol.* **353**, 291–305 (1995).
56. Seung, H., Lee, D., Reis, B. & Tank, D. Stability of the memory of eye position in a recurrent network of conductance-based model neurons. *Neuron* **26**, 259–271 (2000).
57. Seung, H. S. How the brain keeps the eyes still. *Proc. Natl. Acad. Sci. USA* **93**, 13339–13344 (1996).
58. Chavis, D. A. & Pandya, D. N. Further observations on corticofrontal connections in the rhesus monkey. *Brain Res.* [https://doi.org/10.1016/0006-8993\(76\)90089-5](https://doi.org/10.1016/0006-8993(76)90089-5) (1976).
59. Chrastil, E. R., Nicora, G. L. & Huang, A. Vision and proprioception make equal contributions to path integration in a novel homing task. *Cognition* <https://doi.org/10.1016/j.cognition.2019.06.010> (2019).
60. Churan, J., Paul, J., Klingenhoefer, S. & Bremmer, F. Integration of visual and tactile information in reproduction of traveled distance. *J. Neurophysiol.* **118**, 1650–1663 (2017).
61. Churan, J., von Hopffgarten, A. & Bremmer, F. Eye movements during path integration. *Physiol. Rep.* **6**, e13921 (2018).
62. Rosenblum, L., Grewe, E., Churan, J. & Bremmer, F. Influence of tactile flow on visual heading perception. *Multisens. Res.* **35**, 291–308 (2022).
63. Rosenblum, L., Kreß, A., Arikan, B. E., Straube, B. & Bremmer, F. Neural correlates of visual and tactile path integration and their task related modulation. *Sci. Rep.* **13**, 9913 (2023).
64. Freedman, E. G. & Sparks, D. L. Eye-head coordination during head-unrestrained gaze shifts in rhesus monkeys. *J. Neurophysiol.* **77**, 2328–2348 (1997).
65. Grasso, R., Glasauer, S., Takei, Y. & Berthoz, A. The predictive brain: anticipatory control of head direction for the steering of locomotion. *Neuroreport* **7**, 1170–1174 (1996).
66. Grasso, R., Prévost, P., Ivanenko, Y. P. & Berthoz, A. Eye-head coordination for the steering of locomotion in humans: an anticipatory synergy. *Neurosci. Lett.* **253**, 115–118 (1998).
67. Bernardin, D. et al. Gaze anticipation during human locomotion. *Exp. Brain Res.* **223**, 65–78 (2012).
68. Robins, R. K. & Hollands, M. A. The effects of constraining vision and eye movements on whole-body coordination during standing turns. *Exp. Brain Res.* **235**, 3593–3603 (2017).
69. Hollands, M., Khobkhun, F., Ajijmaporn, A., Robins, R. & Richards, J. The effects of constraining head rotation on eye and whole-body coordination during standing turns at different speeds. *J. Appl. Biomech.* **38**, 301–311 (2022).
70. Ahmad, S., Huang, H. & Yu, A. J. Cost-sensitive Bayesian control policy in human active sensing. *Front. Hum. Neurosci.* **8**, 955 (2014).
71. Hoppe, D. & Rothkopf, C. A. Learning rational temporal eye movement strategies. *Proc. Natl. Acad. Sci. USA* **113**, 8332–8337 (2016).
72. Hoppe, D. & Rothkopf, C. A. Multi-step planning of eye movements in visual search. *Sci. Rep.* **9**, 144 (2019).
73. Yang, S. C.-H., Lengyel, M. & Wolpert, D. M. Active sensing in the categorization of visual patterns. *eLife* **5**, e12215 (2016).
74. Zhu, S., Lakshminarasimhan, K. J., Arfaei, N. & Angelaki, D. E. Eye movements reveal spatiotemporal dynamics of visually-informed planning in navigation. *eLife* **11**, e73097 (2022).
75. Adams, R. A., Aponte, E., Marshall, L. & Friston, K. J. Active inference and oculomotor pursuit: the dynamic causal modelling of eye movements. *J. Neurosci. Methods* **242**, 1–14 (2015).
76. de Xivry, J.-J. O., Coppe, S., Blohm, G. & Lefevre, P. Kalman filtering naturally accounts for visually guided and predictive smooth pursuit dynamics. *J. Neurosci.* **33**, 17301–17313 (2013).
77. Deravet, N., Blohm, G., de Xivry, J.-J. O. & Lefèvre, P. Weighted integration of short-term memory and sensory signals in the oculomotor system. *J. Vis.* **18**, 16–16 (2018).
78. Molano-Mazón, M. et al. Recurrent networks endowed with structural priors explain suboptimal animal behavior. *Curr. Biol.* **33**, 622–638.e7 (2023).
79. Andalman, A. S. et al. Neuronal dynamics regulating brain and behavioral state transitions. *Cell* **177**, 970–985.e20 (2019).
80. Rajalingham, R., Piccato, A. & Jazayeri, M. Recurrent neural networks with explicit representation of dynamic latent variables can mimic behavioral patterns in a physical inference task. *Nat. Commun.* **13**, 5865 (2022).
81. Kumar, A., Wu, Z., Pitkow, X. & Schrater, P. Belief dynamics extraction. *Cogn. Sci.* **2019**, 2058–2064 (2019).
82. Lee, D. D., Ortega, P. A. & Stocker, A. A. Dynamic belief state representations. *Curr. Opin. Neurobiol.* <https://doi.org/10.1016/j.conb.2014.01.018> (2014).

83. Reddy, S., Dragan, A. D. & Levine, S. Where do you think you're going?: inferring beliefs about dynamics from behavior. In *Proc. Advances in Neural Information Processing Systems*. https://proceedings.neurips.cc/paper_files/paper/2018/hash/6f2268bd1d3d3ebaabb04d6b5d099425-Abstract.html (2018).
84. Sohn, H., Narain, D., Meirhaeghe, N. & Jazayeri, M. Bayesian Computation through Cortical Latent Dynamics. *Neuron* <https://doi.org/10.1016/j.neuron.2019.06.012> (2019).
85. Zhang, R., Pitkow, X. & Angelaki, D. E. Inductive biases of neural network modularity in spatial navigation. *Sci. Adv.* **10**, eadk1256 (2024).
86. Abdou, M. et al. Can language models encode perceptual structure without grounding? A case study in color. Preprint at <https://doi.org/10.48550/arXiv.2109.06129> (2021).
87. Bender, E. M. & Koller, A. Climbing towards NLU: on meaning, form, and understanding in the age of data. In *Proc. 58th Annual Meeting of the Association for Computational Linguistics* 5185–5198 (Association for Computational Linguistics, 2020).
88. Chalmers, D. J. Could a large language model be conscious? Preprint at <https://doi.org/10.48550/arXiv.2303.07103> (2023).
89. Harnad, S. The symbol grounding problem. *Phys. Nonlinear Phenom.* **42**, 335–346 (1990).
90. Lake, B. M. & Murphy, G. L. Word meaning in minds and machines. *Psychol. Rev.* **130**, 401–431 (2023).
91. LeCun, Y. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. Open Review. Vol. 62 (2022).
92. Piantadosi, S. T. & Hill, F. Meaning without reference in large language models. Preprint at <https://doi.org/10.48550/arXiv.2208.02957> (2022).
93. Einstein, A. Über das Relativitätsprinzip und die aus demselben gezogene Folgerungen, Jahrbuch der Radioaktivität und Elektronik. In *Proc. Collected Papers of Albert Einstein* (Princeton University Press, 1907).
94. Angelaki, D. E. & Dickman, J. D. Spatiotemporal processing of linear accelerations: primary afferent and central vestibular neuron responses. *J. Neurophysiol.* <https://doi.org/10.1152/jn.2000.84.4.2113> (2000).
95. Fernandez, C., Goldberg, J. M. & Abend, W. K. Response to static tilts of peripheral neurons innervating otolith organs of the squirrel monkey. *J. Neurophysiol.* <https://doi.org/10.1152/jn.1972.35.6.978> (1972).
96. Fernandez, C. & Goldberg, J. M. Physiology of peripheral neurons innervating otolith organs of the squirrel monkey. I. Response to static tilts and to long duration centrifugal force. *J. Neurophysiol.* <https://doi.org/10.1152/jn.1976.39.5.970> (1976).
97. Erkelens, C. J., Steinman, R. M. & Collewijn, H. Ocular vergence under natural conditions. II. Gaze shifts between real targets differing in distance and direction. *Proc. R. Soc. Lond. B Biol. Sci.* **236**, 441–465 (1989).
98. Lakshminarasimhan, K. J. Belief embodiment through eye movements facilitates memory-guided navigation, kaushik-l/firefly-rnn. <https://doi.org/10.5281/zenodo.7434801> (2025).
99. Stavropoulos, A. Belief embodiment through eye movements facilitates memory-guided navigation, AkisStavropoulos/eye_movement_analysis: code for eye movement analysis (AS, 2025) (v1.0). Zenodo, <https://doi.org/10.5281/zenodo.15646521> (2025).

Acknowledgements

The authors thank Jing Lin and Jian Chen for their technical support, and Vicky Stavropoulou, Evangelia Pappou and Babis Stavropoulos for their insights and overall support. This work was supported by NIH grant 1R01 DC004260, 1R01 NS127122, NSF NeuroNex 1707400, NIH CRCNS 1R01 NS120407-01, 1U19 NS118246, and Simons Collaboration on the Global Brain grant 324143 to D.E.A. K.L. was supported by a NARSAD Young Investigator Grant from the Brain & Behavior Research Foundation. The authors also thank the Kavli Foundation and the Gatsby Charitable Foundation GAT3780 for their support.

Author contributions

A.S. performed experiments, data analysis, and model development. K.J.L. performed experiments, data analysis, and model development. D.E.A. supervised the work.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-025-66080-5>.

Correspondence and requests for materials should be addressed to Akis Stavropoulos.

Peer review information *Nature Communications* thanks Frank Bremmer, Alessio Fracasso, and the other anonymous reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025