

Spatial uncertainty and environmental geometry in navigation

Yul HR Kang^{1,2@}, Daniel M Wolpert^{1,3,4}, Máté Lengyel^{1,5}

¹Computational and Biological Learning Lab, Department of Engineering, University of Cambridge, Cambridge, UK

²Department of Biological and Experimental Psychology, Queen Mary University of London, London, UK

³Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY, USA

⁴Department of Neuroscience, Columbia University, New York, NY, USA

⁵Center for Cognitive Computation, Department of Cognitive Science, Central European University, Budapest, Hungary

@corresponding author: yul.hr.kang@gmail.com

Variations in the geometry of the environment, such as the shape and size of an enclosure, have profound effects on navigational behavior and its neural underpinning. Here, we show that these effects arise as a consequence of a single, unifying principle: to navigate efficiently, the brain must maintain and update the uncertainty about one's location. We developed an image-computable Bayesian ideal observer model of navigation, continually combining noisy visual and self-motion inputs, and a neural encoding model optimized to represent the location uncertainty computed by the ideal observer. Through mathematical analysis and numerical simulations, we show that the ideal observer accounts for a diverse range of sometimes paradoxical distortions of human homing behavior in anisotropic and deformed environments, including 'boundary tethering', and its neural encoding accounts for distortions of rodent grid cell responses under identical environmental manipulations. Our results demonstrate that spatial uncertainty plays a key role in navigation.

Determining our location in an environment is an essential component of successful navigation¹. The neural underpinning of this process involves some of the most intensively studied cell types in the brain, the place cells of the hippocampus and the grid cells of the medial entorhinal cortex^{2,3}. In particular, grid cells exhibit spatially ordered, periodic patterns of firing that have been hypothesized to provide a regular 'metric for space'⁴. However, both navigational behavior and regular, grid-like neural responses seem to go awry when the geometry of an enclosure is manipulated. For example, compared to isotropic, square-shaped environments, human spatial memory exhibits particular biases in an anisotropic environment with trapezoidal boundary geometry⁵. Under identical conditions, the tuning curves of grid cells in rats also become less regular and more anisotropic in shape⁶. When the size of the environment undergoes an unexpected change, such that it expands or compresses along an axis, homing behavior and neural responses do not only scale with the environment^{7,8} but also, paradoxically, show changes along the orthogonal, unchanged axis^{8–10}. Moreover, spatial patterns of homing behavior and grid cell responses dissociate depending on which wall along the changed axis was last approached^{11,12}.

Previous theories explained only a subset of these results, e.g. the effects of either environ-

mental shape^{5,13–17} or changes in size^{8,9,11,12}, but not both, or were only able to explain behavioral⁹ or neural data^{14–17}. In addition, they explained the apparent anomalies in behavioral response patterns as epiphenomena of deformations in neural response patterns, without addressing the potential adaptive value of these 'anomalies' from a normative perspective^{5,9,12}. Here we show instead that all these data can be accounted for by a single, unifying principle: that allocentric location is fundamentally uncertain, and this uncertainty thus pervades both behavioral and neural responses. While the importance of representing uncertainty about spatial location is well accepted in engineering solutions to navigation^{18,19} (even mobile phone navigation apps show a simple form of it by the shaded blue disc around the currently estimated location), its role in biological navigation has only rarely been considered^{20–25}. In particular, the fundamental effects of environmental geometry on spatial uncertainty have not been studied.

Formalizing the effects of environmental geometry on navigation is challenging because it can affect visual input, a main source of information for navigation, in complex ways that may not be well captured by simple, discrete point-like landmarks as used in some virtual-reality experiments^{26–28} and assumed to provide direct distance and bearing information in related modeling studies^{20,26,29}. Thus, we constructed the

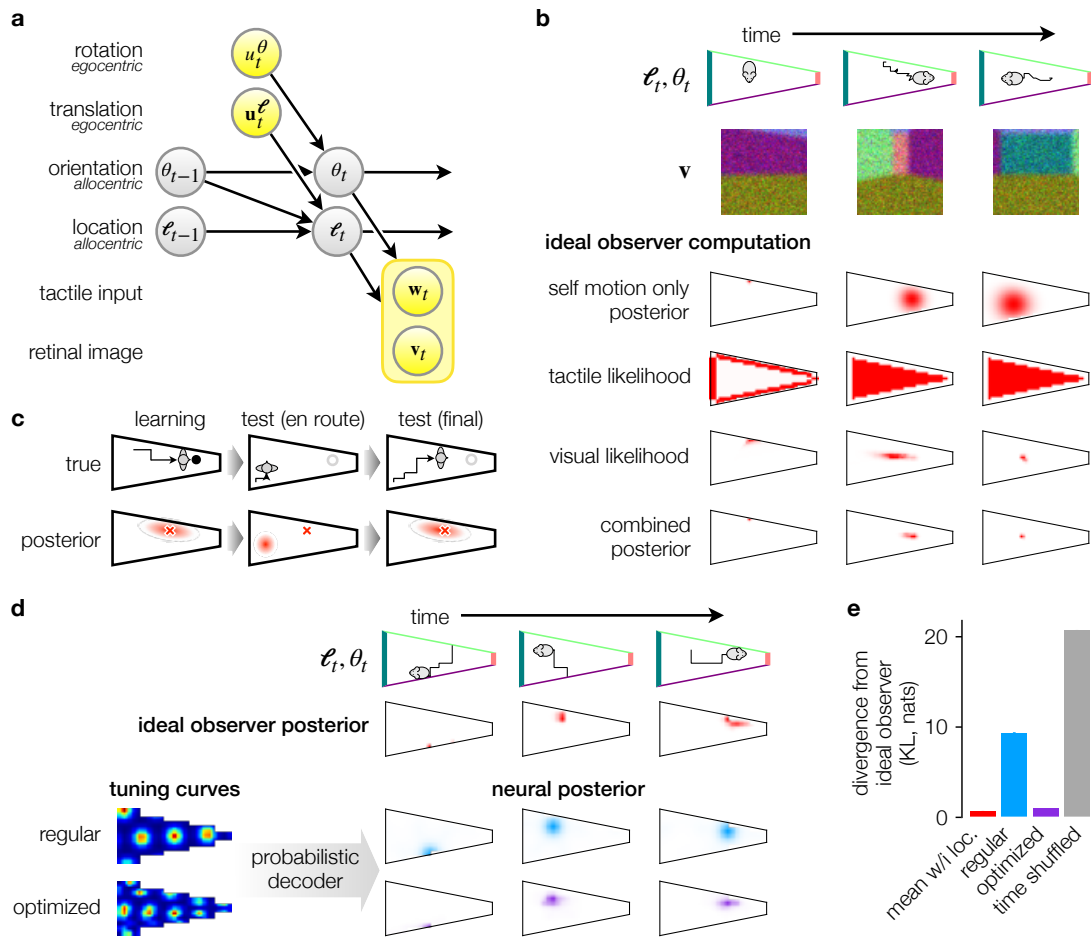


Fig. 1 | The BION model. **a.** Generative model of navigation. The subject's allocentric head direction (θ) and location (ℓ) are updated by egocentric self-motion (u^θ and u^ℓ for rotation and translation, respectively) and give rise to observations of visual input (v) and, potentially, tactile input (w). Yellow circles denote observed variables, gray circles denote hidden variables. **b.** *Row 1.* Top-down view of a rodent starting from a known location (1st column) and progressing along a trajectory (black line, 2nd and 3rd columns). *Row 2.* First-person view of the 3D environment (v) for the corresponding columns of the rodent shown in row 1. The view is corrupted by blur and Poisson noise (reduced for visualization). *Rows 3-6* show ideal observer inferences about location using different inputs (red). *Row 3.* Location posterior based only on self motion, without visual or tactile input, $P(\ell_t | u_{0:t}^\theta, u_{0:t}^\ell)$. *Row 4.* Location likelihood based only on tactile input, $P(w_t | \ell_t)$. *Row 5.* Location likelihood based only on visual input, $P(v_t | \ell_t)$. *Row 6.* Location posterior when all inputs are available, $P(\ell_t | v_{0:t}, w_{0:t}, u_{0:t}^\theta, u_{0:t}^\ell)$. **c.** Modeling behavior in the 'homing task'. Top row shows the participant location and bottom row shows the posterior distribution over locations (red shading). Learning the location of an object (black circle, 1st column) is modeled as encoding the mean (red \times) of the posterior in memory. During the test phase, the participant moves from a new location towards the remembered location of the now invisible object (2nd column). Navigational choices decrease the mean squared error between the remembered location of the object and the current location according to the current posterior, until no further decrease is possible (3rd column). **d.** Modeling neural responses. *Rows 1-2.* Location of a rodent exploring a trapezoid environment (*Row 1*) and corresponding ideal observer posteriors (*Row 2*). *Rows 3-4.* Grid cell responses (left) are fed into a probabilistic decoder to compute a neural posterior (right). With a regular grid (*Row 3*) the neural and ideal observer posteriors do not match. When the tuning curves are optimized by 'warping' (*Row 4*), the neural posterior matches the ideal observer posterior more accurately. **e.** Average (across-time) Kullback-Leibler divergence of ideal observer posteriors from neural posteriors decoded from regular tuning curves (blue) and optimized tuning curves (purple). For reference, the minimal achievable divergence is shown as the divergence of the ideal observer posterior from the average ideal observer posterior at each location (red). Chance level is shown as the divergence between time-shuffled ideal observer posteriors (gray). Error bars show s.e.m. across time points.

'Bayesian image-computable observer for navigation' (BION) model, an ideal observer that integrates information from successive retinal images and idiothetic (self-motion) cues to compute and continuously update a posterior probability distribution over pose (location and head direction). The posterior distribution formalizes uncertainty by expressing the probability with which the observer currently believes themselves to be in

each pose given past and current sensory information. This ideal observer reproduced empirically observed patterns of homing behavior in spatial memory experiments. By deriving analytical expressions for how uncertainty depends on environmental geometry in the BION model, we were also able to provide intuitions for how the fundamental principles of optics (the retinal projection) and optimal multi-sensory cue com-

bination explain the behavioral data. Encoding the posterior distributions of the ideal observer in neural population responses also allowed us to correctly predict how the tuning curves of entorhinal grid cells deform as a function of environmental geometry, and become tethered to boundaries when those boundaries are moved. These results demonstrate that spatial uncertainty plays a key role in navigation and in its neural underpinning.

Results

The BION model

In order to study spatial uncertainty, we formalized the task of estimating one's location in the environment as recursive probabilistic inference given noisy sensory inputs (Fig. 1, Methods). This formalization, based on widely used approaches to robotic navigation^{18,19}, treats allocentric pose (location and head direction) as hidden variables to which the brain has no direct access (Fig. 1a). Instead, sensory inputs deliver egocentric signals, such as a first-person view of the environment by visual inputs, and self-motion cues from a variety of senses (e.g. the vestibular system). Thus, the BION model infers allocentric pose from a stream of egocentric sensory inputs based on an internal model for how self-motion cues about translation and rotation (Fig. 1a, u^ℓ and u^θ) provide information about changes in pose (Fig. 1a, ℓ and θ for location and head direction, respectively), and how the retinal image depends on the current pose (Fig. 1a, v). As we expected environmental geometry to have complex effects on visual information that may not be well described by abstract Lidar-like readings of distance and bearing to point-like landmarks that were often used previously^{20,26,29,30}, we adopted an image-computable approach instead³¹. Thus, we considered a fully rendered retinal image through a pinhole projection (with Gaussian blur and Poisson noise). When modeling rodent navigation (but not human navigation in virtual reality), we also considered tactile input, providing egocentric information about whether the animal is next to a wall, and if so, on which side of the animal the wall is (Fig. 1a, w).

As the animal moves around in the environment (Fig. 1b, row 1), it receives a sequence of sensory inputs (Fig. 1b, row 2 shows visual input). Given the sequence of inputs received thus far, the ideal observer computes a posterior distribution, expressing the probability with which it currently believes itself to be at each location (and head direction, not shown; Fig. 1b, rows 3-6). When the ideal observer only has access to self-motion inputs, this results in the well-known inevitable

accumulation of errors over time by pure path integration (Fig. 1b, row 3). Tactile inputs provide information about being somewhere near the walls, but leave high levels of uncertainty when the animal is away from the walls (Fig. 1b, row 4). Compared to other senses, visual input can result in particularly complex and non-monotonic changes in uncertainty, especially along the line of sight, depending on the current pose of the animal relative to the geometry of the environment (Fig. 1b, row 5). Computing the Bayesian posterior with self-motion, tactile, and visual inputs allows the ideal observer to perform optimal multi-sensory cue combination. In particular, the posterior (Fig. 1b, row 6) is more precise than what each sensory modality alone can afford.

We used the BION model to make predictions about both behavioral patterns in human spatial memory experiments and grid cell recordings in rodents (Methods). To be able to maximally constrain the model, we selected a set of paradigms for which both behavioral and neural data was available. In addition, we considered experiments in which participants had already been familiarized with an environment so that they had minimal uncertainty about its layout (even though in the test phase of some experiments, the actual layout was changed unbeknownst to them). This allowed us to ignore uncertainty about the 'map' (see e.g. Ref. 20) and focus on the effects of uncertainty about location with respect to the map.

The behavioral paradigms we modeled required human participants to perform a 'homing task', i.e. to return to the location of a specific object whose location they had to previously encode in memory. These experiments use virtual reality environments so that visual input can be precisely controlled. We modeled both the learning and the test phases of these experiments (Fig. 1c). In the learning phase, we simulated trajectories that led participants to the object locations that they needed to remember and, at the end of each such trajectory, stored in memory the mean of the posterior distribution over locations (with noise to account for memory decay) as the remembered location (Fig. 1c, left). In the test phase, participants were removed from the original location of the object, and we modeled their sequence of navigational choices as greedily (albeit noisily) decreasing the mean squared error between their current location and the remembered location of the target object, until no further decrease was possible (Fig. 1c, center and right). Critically, due to the bias-variance trade-off inherent in minimizing the mean squared error³², the expected distance was generally not minimized simply where the mean of the current location poste-

rior matched the remembered target location, and thus participants' uncertainty (the covariance of their posterior) also affected the end point of their chosen routes.

For predicting grid cell responses, we simulated a population of grid cells (Fig. 1d, bottom left) in a randomly foraging animal (Fig. 1d, row 1), and used a probabilistic population decoder³³ to compute a 'neural' posterior over locations from their responses (Fig. 1d, right; Methods). In contrast to standard approaches to efficient coding that aim to minimize the uncertainty in the neural posterior³⁴, we required the uncertainty of the neural posterior to faithfully reflect that of the ideal observer (Fig. 1d, row 2), which prescribes the normatively appropriate level of uncertainty given noise and ambiguity in sensory inputs.

In general, the ideal observer's posterior underwent dynamic changes over time (cf. Fig. 1b, row 6), such that it expressed less or more uncertainty about the animal's location (Fig. 1d, row 2, e.g. compare steps 1 and 2), sometimes being more isotropic, expressing approximately the same amount of uncertainty along each direction, while at other times becoming anisotropic, i.e. strongly elongated along a particular direction (compare time steps 2 and 3). Critically, the neural posteriors for grid cell responses with perfect hexagonal symmetry were translation-invariant and isotropic (except for trivial boundary effects), failing to exhibit the modulations in the uncertainty the ideal observer expressed (Fig. 1d, row 3). We thus created a 'warped population code'^{16,34} by warping the underlying grid on which grid cell tuning curves were defined. Specifically, we optimized this warping until the resulting neural posteriors maximally matched the ideal observer posteriors moment-by-moment (Fig. 1d, row 4 and Fig. 1e). The simulated grid cell responses obtained with the optimized warping constituted our neural predictions. Thus, rather than building a theory for the basic hexagonal symmetry of grid fields, as in most previous approaches^{35–41}, we focused on modeling how and why this perfect symmetry is broken under well-controlled experimental manipulations.

Deformations arise from inhomogeneous uncertainty

A critical insight of the BION model is that the same landmark provides different amounts of visual information about the subject's location, depending on its distance along, d , and eccentricity from, ϵ , the subject's line of sight (Fig. 2a, inset). For example, the landmark could be a corner of the enclosure, with d being the subject's distance from the corresponding wall when fac-

ing it. In particular, based on the basic laws of optics describing a pinhole projection with finite resolution, we were able to show analytically that the visual information (I , formally the Fisher information) about the distance (d) scales with the square of eccentricity (ϵ) and inversely with the fourth power of the distance itself (d):

$$I \propto \frac{\epsilon^2}{d^4} \quad (1)$$

Importantly, we were also able to show that the same expression holds for a more realistic scenario, with a fully rendered retinal image and a finite-sized retinotopic neural population whose signaling is corrupted by Poisson noise (Supp. Material). Indeed, numerical simulations of the visual likelihood of current location using the BION model confirmed these calculations (Fig. 2a).

Eq. 1 makes direct predictions about how visual information about location changes inside a square- versus a trapezoid-shaped environment, such as those used in several experiments^{5,6} (Fig. 2b). In general, because information falls steeply with d , overall information at any location is dominated by information obtained while looking in the direction of the closest wall. Thus, in both kinds of environments, information is higher near the walls. This is consistent with humans making smaller homing errors near walls⁵. More importantly, due to the fundamental symmetry of a square environment, information has a symmetrical distribution, and is also symmetric with respect to the axis along which spatial position needs to be determined (x or y in Fig. 2b). In contrast, in the trapezoid, information is asymmetric between its two sides, and (again consistent with empirical data⁵) overall smaller about position along the long (x-) axis (because distances are greater) than about position along its short (y-) axis (Fig. 2b).

In homing tasks (Fig. 2c, cf. Fig. 1c), total errors are dominated by the more uncertain direction, i.e. position about the long axis in the trapezoid (Fig. 2b, compare left two panels for the trapezoid, showing uncertainty about x and y separately, with the last panel, showing total uncertainty). Along this direction, on average, d is greater and thus information is smaller than along either direction in the square. Therefore, the BION model predicts homing performance to be worse in the trapezoid than in the square, which is what has been found experimentally⁵ (Fig. 2d, top, green vs. orange). Finally, the asymmetry between the two ends of the trapezoid results in ϵ being larger and thus information being greater near the broad end. As a consequence, homing performance is predicted to be relatively better near its broad end than near the narrow end,

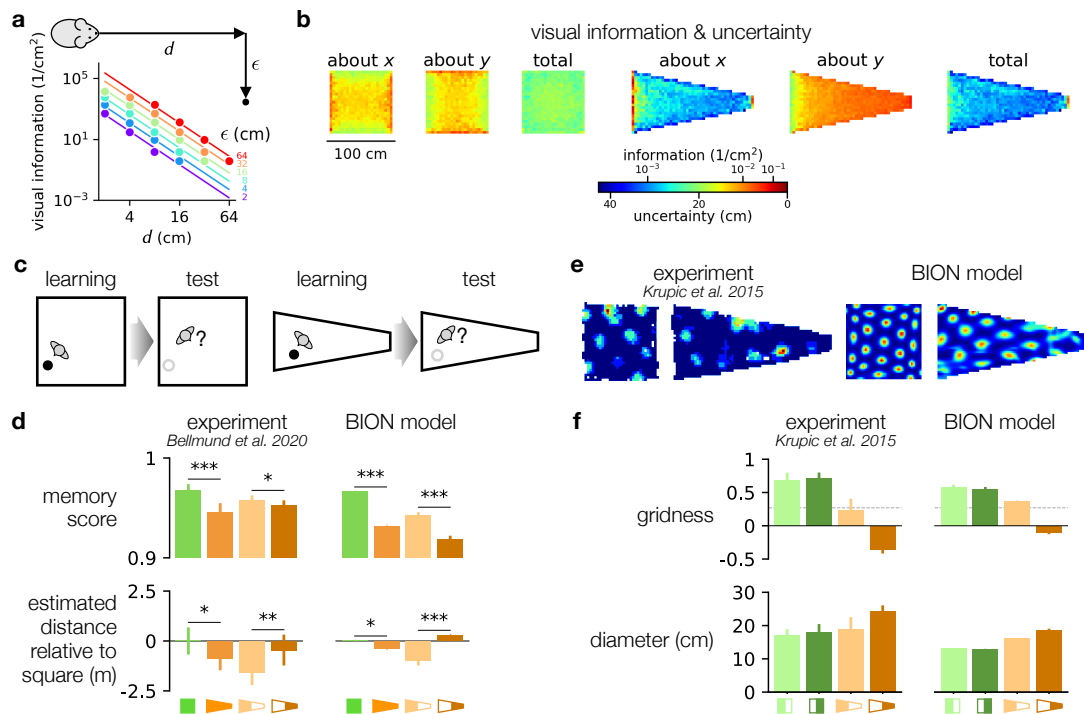


Fig. 2 | Deformations arise from inhomogeneous uncertainty. **a. Inset:** top-down view of an observer (rodent in this case) and a point landmark (solid black circle), at a given distance along the line of sight (d) and 'eccentricity' (ϵ , distance orthogonal to the line of sight) from the observer. **Main plot:** ideal observer's visual (Fisher) information about d as a function of the true distance (abscissa) and eccentricity of the landmark (colors), lines show analytic approximation, markers show numerical simulation. Note logarithmic scale for both axes as well as color scale. **b.** Visual information about x- and y-coordinates of location, and total information in the square and trapezoid environments. Warmer/cooler color indicates more/less information, and conversely, lower/higher uncertainty. **c.** Homing task. Participants memorized the location of an object during the learning phase (filled circle), to which they were asked to return (after the object was hidden, open circle) during the test phase. Two kinds of environments were used: square (left) and trapezoid (right). **d.** Homing performance from experimental data (left⁵) and BION model predictions (right). Each set of four columns shows square vs. trapezoid environments and broad vs. narrow part of the trapezoid environment (see filled areas in schematics at the bottom). **Top row:** memory score (rank of homing location's distance to target compared to all possible locations). **Bottom row:** estimated distance between target locations in each (half) of the environment compared to that in the square (true distances were matched). Error bars show s.e.m. across subjects. **e-f.** Example grid fields (**e**) and grid field characteristics (**f**) in experimental data (left⁶) and BION model predictions (right). **e.** Grid fields of the same cell in a square and a trapezoid environment. **f.** Gridness (top) and diameter of grid fields (bottom). Each set of four columns shows data for grid fields in the left and right halves of the square and trapezoid environments separately (see filled areas in schematics at the bottom). Error bars show s.e.m. across neurons.

again consistent with data⁵ (Fig. 2d, top, light vs. dark brown).

Note that path integration alone cannot explain these results, as it predicts an increase in uncertainty with time, not as a function of space. This may contribute to spatial information decreasing with d (and therefore the decrease of information away from the walls, and the overall increased uncertainty in the trapezoid than in the square environment that we described above) as long as there is an additional landmark-based mechanism that resets path integration near the walls. However, even such a hard reset mechanism cannot account for the scaling of information with ϵ that underlies the BION model's predictions about differences between the two halves of the trapezoid.

In another task used in the same experiment, participants had to estimate distances between ob-

jects. Overall, participants showed different degrees of estimation biases depending on the geometry of the environment (Fig. 2d, bottom left). Interestingly, however, the pattern of these estimation biases did not simply reflect the patterns seen in memory performance: while the biases in the trapezoid were overall more negative than in the square (in line with memory performance being poorer in the trapezoid than in the square), they were *less* negative (although memory performance was *more* diminished) in the narrow side of the trapezoid than in the broad side (Fig. 2d, left, compare bottom to top).

The ideal observer reproduced this pattern of results (Fig. 2d, bottom right). The relative underestimation in the trapezoid was due to regression to the mean: as visual information is always uncertain, it needs to be combined with the prior covering the extent of the environment. This biases

locations to be estimated as being closer to the middle of the environment (the mean of the prior) and thus overall distances between them to be smaller – and more so in the trapezoid because visual information is more uncertain there overall (Fig. 2b). The more subtle differences in underestimation between the two halves of the trapezoid environment were explained by the interplay of how environmental geometry changes both visual information (Eq. 1, i.e. the likelihood) and the prior (Supp. Material).

In the BION model, predicted grid field sizes depended on the overall level of spatial information, and their irregularities depended on the degree of warping that was necessary to achieve the level of spatial anisotropy in the uncertainties expressed by the ideal observer (Fig. 1d, and Fig. 2e, right). Compared to the square enclosure, uncertainty in the trapezoid was greater and more anisotropic (differing greatly along the two axes of the enclosure, Fig. 2b). Thus, the BION model predicted larger grid fields with lower levels of gridness for the trapezoid than for the square (Fig. 2f, right). For analogous reasons, the model also predicted larger grid fields with lower levels of gridness for the narrow than for the broad side of the trapezoid (Fig. 2f, right, orange). These predictions matched experimental measurements of grid fields in such square and trapezoid environments⁶ (Fig. 2e, left, f, left).

Scaling arises from model mismatch

The BION model also explains scaling of behavioral and neural response patterns in unfamiliar environments as a consequence of the mismatch between the internal model and the sensory input. In a typical human behavioral experiment⁸, participants were familiarized with a particular environment and their homing to a previously experienced target location was tested, after the environment had been expanded or compressed along one or both axes unbeknownst to them (Fig. 3a shows expansion along horizontal axis). Because the changing of the environment in the test phase was unknown to participants, we modeled the ideal observer's inferences as being computed in the familiar environment's frame of reference, resulting in a model mismatch between the internal model used by the ideal observer and the actual environment (Methods). Due to this model mismatch, the actual sensory inputs (that were generated by the test environment) led to higher uncertainty in the ideal observer (that interpreted them according to the familiar environment), particularly along the axis of the environment that was changed, and particularly for locations along the axis for which information was low

to begin with (because they were farther from the nearest wall, see above). Thus, in line with experimental data⁸ (Fig. 3b, top), the BION model (Fig. 3b, bottom) predicts response variability to be higher along the changed axis especially when the distance to the closest wall along that axis is greater (Fig. 3b, red vs. blue).

Unfamiliar environments do not only lead to increased response variability, but also to systematic biases in homing responses. These biases have been characterized in an experiment in which participants were required to return to a previously visited target without visual input after they were led away from it (in the presence of visual input) along an experimentally controlled path (Fig. 3c, gray trajectories in left panels followed by orange or black arrows in right panels). Critically, participants performed this task for three consecutive trials with a fixed ('familiar') environmental geometry, and then on the fourth trial they performed it either in the same environment, or in an unfamiliar 'test' environment that was scaled (expanded or compressed) along one of its axes (Fig. 3c shows the case when the test environment is expanded along the x-axis). Participants were specifically led to believe that the test environment was the same as the familiar one although it was in fact scaled – corresponding to a model mismatch. Response biases (the length of the return journey in the scaled environment relative to that in the familiar environment with a matched distance to the target) were measured separately for return journeys along the changed (Fig. 3c, top, orange) or unchanged axis of the environment (Fig. 3c, bottom, black arrow).

Responses were negatively biased (i.e. undershot) along an expanded axis (Fig. 3d, top, orange bars) and positively biased (i.e. overshoot) along a compressed axis (Fig. 3d, top, green bars). The BION model captured this because, during the outward journey, the visual image of the far wall along the expanded axis suggested that this wall was still at a great distance, which in turn made the outward journey seem shorter than it actually was, leading to an undershoot in simulated responses (Fig. 3d, bottom, orange and Fig. 3e, top) – and vice versa when the far wall was seen along the compressed axis (Fig. 3d, bottom, green). Interestingly, even responses along the unchanged axes were biased, and in the opposite direction: positively (overshooting), when the environment was expanded, and negatively (undershooting), when it was compressed (Fig. 3d, top, black bars). The BION model also captured this behavior simply because looking along the unchanged direction showed the farther expanded wall (that was seen more during the

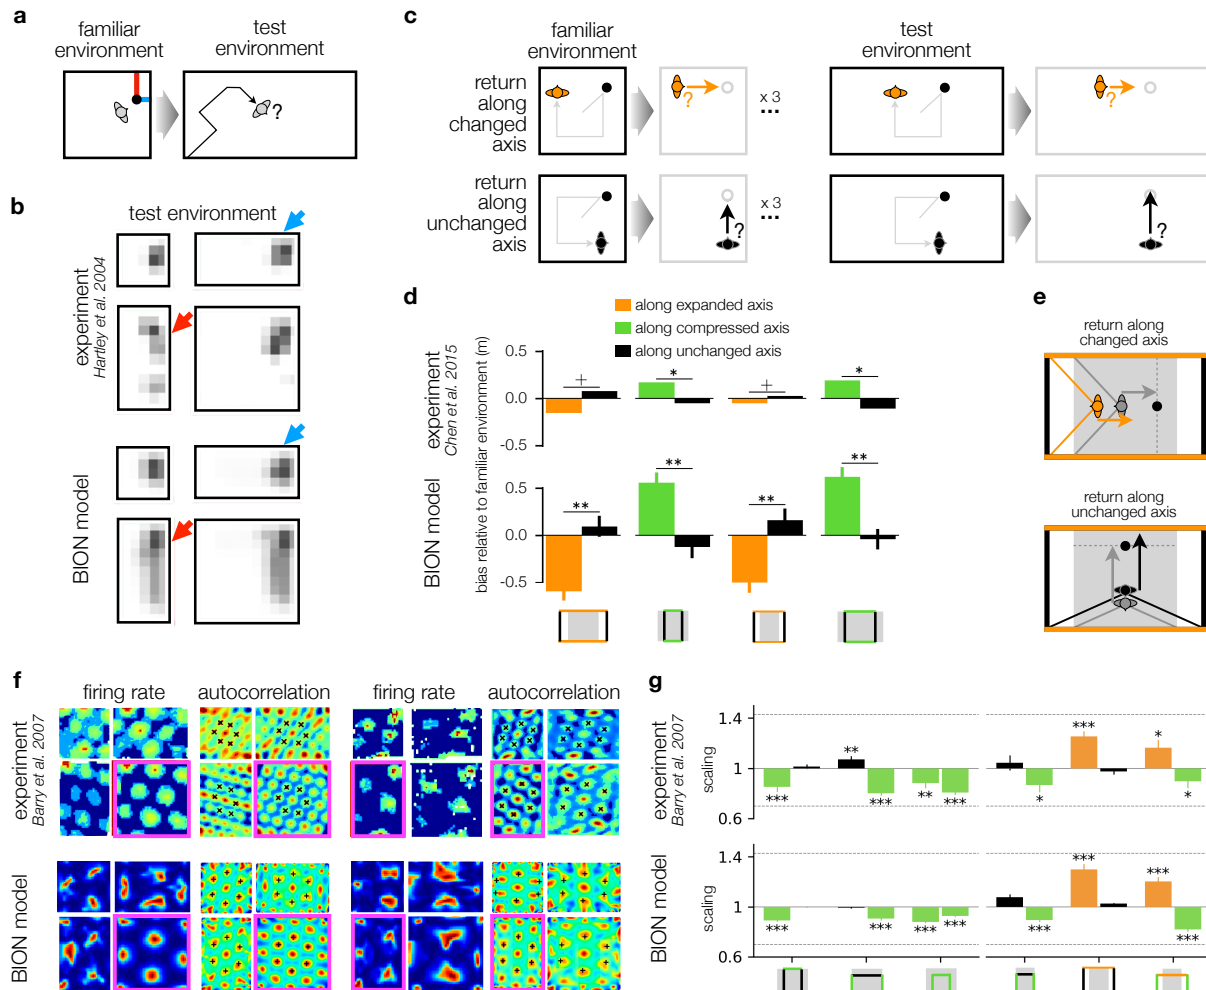


Fig. 3 | Scaling arises from model mismatch. **a.** Homing in scaled environments⁸. Human participants first learned a target location (black circle) in a familiar environment (e.g. a square room) and were asked to return to that location (black path) with the target now invisible in a test environment that was either the same as the familiar environment or, unbeknownst to them, scaled (e.g. a rectangular room as shown). Blue and red lines show distances to the nearest walls along the horizontal and vertical axes, respectively. **b.** Homing location distributions for the square familiar environment and target object location shown in **a** in experimental data (top⁸) and the BION model (bottom). Each sub-panel shows a different test environment. Blue/red arrows show conditions where the target object's distance to the nearest wall along the axis of scaling was comparatively smaller/larger (see blue/red lines in **a**). **c.** Axis-dependent homing in scaled environments⁹. Participants first walked from the center of an environment to a visible target and were asked to remember its location (top row, black circle). The target then disappeared and participants were guided across the environment along a random, multistep path (gray 'outbound' path). Participants were then asked to return to the location of the target (orange and black arrows) in the absence of any visual input (gray bounding box). In each block, three trials were performed in a 'familiar' environment (left) and then a single trial was performed in a scaled test environment (right). The beginning of the return path (i.e. the endpoint of the outbound path) was always such that the direction to the target was either along the axis that was changed in that trial (top, orange arrows), or along the unchanged axis (bottom, black arrows). **d.** Homing biases in test environments along the changed (orange and green bars) and unchanged (black bars) axes observed in the data (top⁹) and the BION model (bottom). Schematics in the bottom show familiar environment (gray) and test environment boundaries (black and colored lines) for reference. Error bars show s.e.m. across simulated subjects (not drawn for the experimental data because only the significance of the comparison and mean differences were available). **e.** Schematic of how biases arise when the test environment is scaled (orange) from the familiar environment (gray). When the return direction is along the expanded axis (top), the angle subtended by the corners of the far wall (orange), as seen during the outbound journey in the test environment, matches the angle at a location in the familiar environment (gray) that is closer to the target location (black circle). This results in an undershoot during the return journey (negative bias, orange vs. gray arrows). Conversely, for homing along the unchanged axis (bottom), the angle subtended by the corners of the far wall in the test environment (black) suggests being farther from the target in the familiar environment (gray), hence leading to an overshoot (positive bias, black vs. gray arrows). **f.** Effect of environmental scaling on rodent grid cell responses¹⁰. Rats were first familiarized with one room (pink outline, square and rectangular room in left and right panels, respectively) and then were introduced to scaled rooms (three adjacent rectangles). Grid cell firing rates are shown on the left and autocorrelations on the right for each room shape. **g.** Grid field scaling ratio along the changed axis (green and orange for compressed and expanded, respectively) and along the unchanged axis (black) in the data (top¹⁰) and the BION model (bottom). Schematics in the bottom show familiar environment (gray) and test environment boundaries (black and colored lines) for reference. Error bars show s.e.m. across neurons.

outward journey) subtend a large angle, making it appear as if it was closer, and thus the outward

journey longer than it actually was, leading to an overshoot in simulated responses (Fig. 3e, bottom) – and again vice versa in the compressed environment (Fig. 3d, bottom, black bars).

Although the geometrical intuitions above only apply to the effect of visual inputs, and are most straightforward assuming simple point-like landmarks (the corners of the enclosure), the simulation results we show were obtained by the full ideal observer that uses a rendered retinal image as visual input, and combines it with self-motion inputs (Fig. 1a-b). This shows that the basic laws of optics determining the retinal projection can have significant behavioral effects in a Bayes-optimal observer during navigation – and thus explain the experimentally observed pattern of results.

Recordings of rodent grid cell responses in analogous scaled rectangular enclosures found that the regularity (specifically, isotropy) of grid fields depended on whether a given enclosure shape was familiar (pink outlined environments in Fig. 3f) or encountered after scaling, rather than on the shape itself. For example, grid fields in the identical square environment were isotropic when the square was familiar (Fig. 3f, top left, firing rate and autocorrelation map inside pink outline), but became more anisotropic when the familiar environment was a rectangle and the square environment was encountered as a scaled version of this rectangle (Fig. 3f, top right). Conversely, grid fields were isotropic in the rectangular environment in the second case but anisotropic in the first case. These results were naturally reproduced by the BION model, simply because the distortions of uncertainty due to model mismatch, which explained human behavior, required increased warping of grid fields (Fig. 3f, bottom).

In the same set of experiments, a finer grained analysis revealed systematic distortions of grid fields along the changed vs. unchanged axes of these environments (Fig. 3g, top). In particular, in line with the biases seen in human behavior (Fig. 3d, top), grid fields compressed and expanded respectively along the compressed and expanded axes of the environment (Fig. 3g, top, green and orange bars), and showed opposite scaling along the unchanged direction (Fig. 3g, top, black bars). Once again, the same distortions in the BION model's spatial uncertainty that explained human behavior, also required grid field warpings that reproduced these neural response patterns (Fig. 3g, bottom).

Tethering arises from cue conflict

Unexpected changes in environmental geometry have also revealed an intriguing 'tethering' effect in both behavioral and neural responses. In behavioral experiments, homing end point locations corresponding to different targets in the changed geometry have been compared to the original locations of those targets in the familiar geometry (Fig. 4a). The relationship between these two sets of target locations was characterized by a 'scale' factor, measuring the ratio of distances between homing locations and distances between original target locations, and an overall 'shift' in homing locations away from the wall from which participants started their return journeys ('starting wall') in a given trial with the changed geometry. As expected, the scaling of homing locations reflected the scaling of the environment: when the environment was compressed or expanded, homing locations were respectively compressed or expanded (Fig. 4b, top left). However, there was also a systematic shift in homing locations (Fig. 4b, bottom left): when the environment was compressed, they shifted forwards, while in expanded environments, they shifted backwards, keeping their distance to the starting wall similar to that in the familiar environment in both cases (i.e. as if they were 'tethered' to the starting wall)¹².

The BION model readily explained these results as a fundamental feature of navigation in the face of spatial uncertainty. According to the model, interpreting visual and self-motion inputs in the changed geometry under the internal model of the familiar geometry leads to cue conflict between these two inputs. That is, when path integration using self-motion inputs on a return journey suggests that some distance has been traveled since leaving the starting wall, visual inputs are in conflict with this, because the opposite wall appears nearer or farther depending on whether the environment has been compressed or expanded. In other words, path integration anchors responses to the starting wall, while visual inputs tend to anchor them to the opposite wall (Fig. 4c, red and blue distributions peaking at the ends of corresponding arrows). Critically, the ideal observer combines these two cues according to their respective uncertainties, such that the final estimated spatial location depends more on the input with less uncertainty. In general, uncertainty from path integration grows with distance traveled from the starting wall (Fig. 1b), while uncertainty in visual inputs decreases because of its inverse scaling with the distance from the opposite wall (Eq. 1). Indeed, when we simulated the ideal observer with each input sepa-

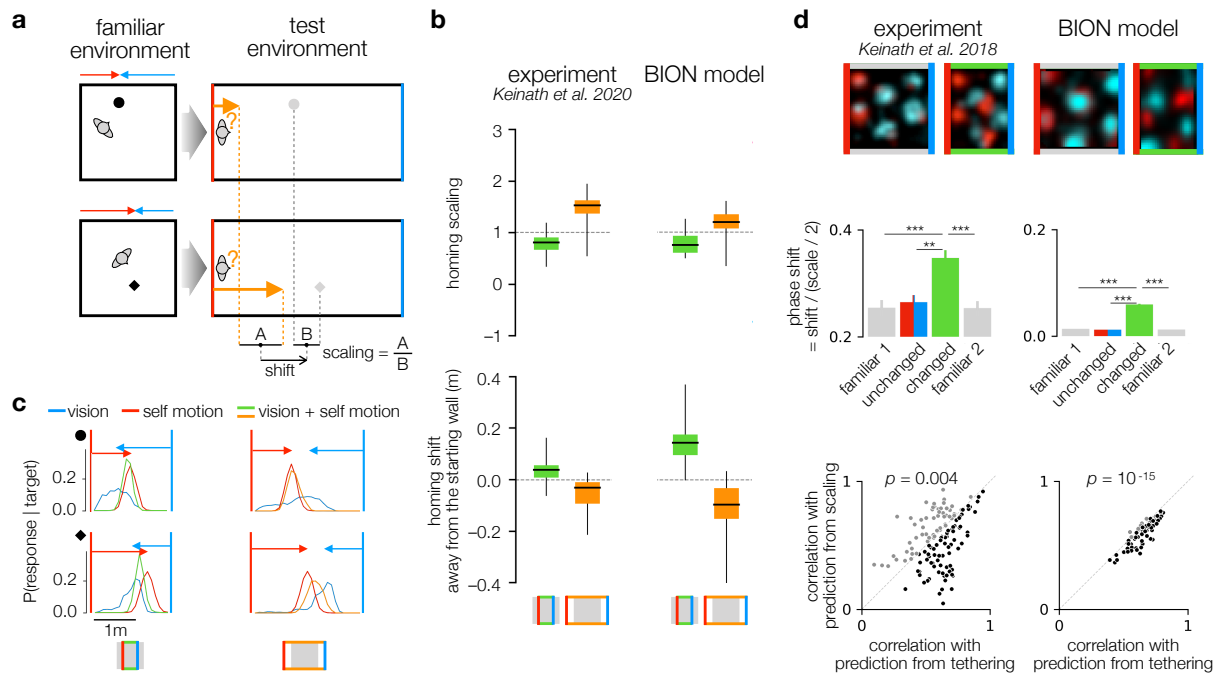


Fig. 4 | Tethering arises from cue conflict. **a.** Distance-dependent homing in scaled environments¹². *Left.* Participants learned locations of four nameable objects (only two shown here for simplicity, black circle and diamond) in a square familiar environment. The objects were at different distances from the walls (red and blue arrows indicate the distance from two opposite walls along the axis that is to be changed). *Right.* In a test trial, participants were asked to return to the location of an object cued by its name in an expanded (as shown) or compressed test environment, starting from one of the four walls (for this trial, the starting wall is red and the opposite wall is blue for illustration purposes only). Gray circles show the 'equivalent' location of each target object (hidden from participants), i.e. the location that is at the same distance from the center of the enclosure in the test as in the familiar environment (note that this preserves distances between objects). Schematics in the bottom give intuitive explanation for main experimental measures (actual calculations used linear regression across 4 objects \times 4 starting walls with the predicted shift appropriately changed depending on the starting wall¹²). The scaling of homing responses measures the ratio of the distance (A) between the two homing locations (orange) and the distance (B) between the target locations (or equivalent target locations, gray). The shift of responses measures the difference between the average of the homing locations (orange arrows) and the average of equivalent target locations (gray), with a sign that is defined to be positive when the shift is away from the starting wall. **b.** Distribution of scaling and shift in the data (left; Fig. 6d in Ref. 12, with an opposite definition of the sign of shift as we use here) and the BION model (right) is consistent with tethering. Box-and-whisker plots show range (whiskers), interquartile range (box), and median (horizontal line) across subjects. **c.** Distribution of homing locations simulated by the BION model, when the test environment is compressed (left column) and expanded (right column), for a target nearer the starting wall (top row) and farther from the starting wall (bottom row). Red and blue vertical lines indicate the starting and opposite walls, corresponding arrows show the distance of the object to these walls in the familiar environment shown in **a**. Curves of different colors indicate the distribution of homing responses using both visual and self motion (green and orange), vision alone (blue, mostly preserving the original distance to the blue wall), and self motion alone (red, mostly preserving the original distance to the red wall). Schematics in the bottom show familiar environment (gray) and test environment boundaries (colored lines) for reference. **d.** *Top row.* Grid fields in the data (left¹¹) and the BION model (right) show tethering analogous to that seen in behavior. The familiar, training environment is a square and the test environment is compressed along the horizontal axis. Red and cyan shading show firing rates conditioned on the last approached wall being the west and east wall, respectively (walls are colored in the figure for visualization only). Note that in the square environment, the cyan shading overlaps the red shading because the phases of the grid fields are the same regardless of the last wall approached. *Middle row.* Phase shift of the grid fields relative to the scale of the grid (distance between adjacent peaks) in the familiar environment (gray, before [1] and after [2] experiencing the novel environment) and along the unchanged (red/blue) and changed (green) axes of the compressed environment. Note that, due to the periodicity of grid fields, their shifts can only be measured up to half the distance between them, i.e. half their scale. Error bars show s.e.m. across neurons. *Bottom row.* Pearson's correlation between the observed tuning curves conditioned on the last wall being on the changed axis and the tuning curves predicted from perfect scaling (ordinate) vs. perfect shifting (i.e. boundary tethering, abscissa). Each point is a neuron, colored black if tethering correlation is bigger than scaling correlation, and gray otherwise. P-values show the significance of a paired t-test in favor of tethering vs. scaling across neurons (comparing Fisher-transformed correlations).

ately, response distributions based on only self-motion input became slightly broader for targets farther from the starting wall (Fig. 4c, red distributions in bottom compared to top). In contrast, response distributions based on only visual inputs were substantially broader for targets nearer the starting wall compared to more distant tar-

gets (Fig. 4c, blue distributions in top compared to bottom), as nearer targets are farther from the opposite wall. Therefore, responses to targets near the starting wall were dominated by path integration, and were thus mostly tethered to (i.e. kept their distance from) the starting wall (Fig. 4c, top row, green and orange response distributions

largely overlap with red distributions). In comparison, homing to targets near the opposite wall relied more on visual inputs, whose reliability was higher there, and thus these responses were tethered relatively more to the opposite wall (Fig. 4c, bottom row, green and orange response distributions interpolate between red and blue distributions).

The opposite tethering of near and far targets to the starting and opposite walls in the BION model, respectively, explains the experimentally observed patterns of scaling of response distributions in compressed and expanded environments (Fig. 4b, top right). The model also predicted the experimentally observed overall tethering of responses to the starting wall (Fig. 4b, bottom right) because, compared to other experiments, self-motion inputs were more informative than visual inputs in this experiment, as subjects walked a short, straight path from a known location (center of an identifiable wall). This contrasts with some of the other experiments we modeled above, in which visual inputs were relatively more informative because subjects started from a random location uniformly sampled within the environment (Fig. 2c) or walked a long meandering path before the homing phase (Fig. 3c).

A re-analysis of previous experimental data^{10,42} revealed that, similar to human homing responses, rodent grid fields are also tethered to environmental boundaries under analogous conditions (i.e. unexpected environmental scaling). In particular, the distance of a given grid field to the wall that the animal most recently approached (i.e. was within 12 cm of it) is preserved between familiar and compressed test environments¹¹ (Fig. 4d, top left). In order for the neural posterior to follow the overall tethering of the ideal observer's posterior to the last wall approached, which was apparent in the behavioral data (Fig. 4b, right), the BION model also predicted this tethering of neural tuning curves (Fig. 4d, right; Supp. Material). Further quantifying tethering as a relative phase shift of grid fields confirmed that it was significantly larger along the changed axis than along the unchanged axis of the compressed environment, and also larger than in the familiar environment before and after testing in the compressed environment both in experiments and in the model (Fig. 4d, middle left and right). Finally, empirical changes in tuning curves in the compressed environment were also compared to hypothetical deformations based on pure scaling vs. pure tethering. Although, overall, these two kinds of deformations were found to provide comparable fits to the empirical changes, tuning curves were significantly better correlated

with tethering-based predictions in experiments as well as in the model (Fig. 4d, bottom left and right).

Discussion

Our work provides a unifying, normative theory of how both behavioral and neural response patterns depend on environmental geometry. The key insight is that, due to the basic laws of optics, visual inputs have a fundamental effect on spatial uncertainty, which in turn is reflected in behavioral responses and encoded in neural responses as deformation, scaling, and tethering by environmental boundaries. We quantify these effects through numerical simulations of the BION model, and provide mathematical insights into their main qualitative features by analytical derivations.

Previous models that attempted to link behavior or grid cell properties to environmental geometry only explained a subset of the phenomena that we predict, and typically accounted for behavioral effects as mere epiphenomena of particular neural representations that were themselves often posited to be based on idiosyncratic features of the environment^{8,9,11,12}. In particular, based on observations of place cell responses, the 'boundary vector hypothesis' postulated a special status for environmental boundaries, and their representation by specialized cell types⁷ – a prediction confirmed by later experiments^{43–45}. These cells were in turn suggested to underlie some of the behavioral⁸ and neural^{11,46} effects seen during environmental deformations. In contrast, the BION model does not rely on explicit notions of walls or boundaries, or externally provided identifying tags for them. Instead, it provides a normative basis for the special status of environmental boundaries in navigation by only using 'raw' visual inputs. Thus, it explains the behavioral and neural effects of boundaries from first principles, as consequences of the higher visual information close to such boundaries (Fig. 2a, Eq. 1). In line with this, the dependence of hippocampal place cell responses on environmental boundaries has recently also been hypothesized to reflect such information-theoretic principles⁴⁷. Taken together, these results raise the possibility that the very existence of boundary vector cells⁴³ (and related boundary⁴⁴ and border cells⁴⁵) may also derive from the same normative principle.

The BION model starts from the fundamental principle (long-recognized in engineering applications) that representing spatial uncertainty is key for efficient and reliable navigation because allocentric location can only be inferred from sen-

sory inputs and motor outputs that are inherently noisy and ambiguous. This is a fundamental difference from dominant theories of grid cell responses that either make the simplifying assumption that location is already given to the model as an input, without uncertainty^{13,17,38,46,48,49}, or perform localization from sensory inputs without considering the unavoidable uncertainty in the process^{39,50}. Conversely, models that formalized navigation in a Bayesian inference framework have not incorporated an image-computable ideal observer. Thus, these models did not study how environmental geometry affects spatial uncertainty and, in turn, homing behavior or grid cell responses^{16,20,21,25,51–56}.

An alternative theoretical account of some forms of grid field deformations is based on the successor representation (SR)⁵⁷. Specifically, according to SR-theory, place cell responses encode a predictive map, the (time-discounted) probability with which a given location in the environment will be reached from the animal's current location, while grid fields provide a set of basis functions for efficiently representing this map¹³. Because these transition probabilities depend on the geometry of the environment, they inherit the asymmetry of the environment. Therefore, SR predicts an overall asymmetry of grid fields between the two halves of a trapezoid environment¹³, and lower gridness in the narrow side⁴⁶. Based on visual uncertainty alone (Fig. 2b), the BION model also predicts this overall asymmetry, along with more fine-grained differences in the gridness and diameter of grid fields (Fig. 2f). While SR-based theories have also made predictions about grid field diameters (or equivalently, radial frequencies, Fig. 1j in Ref. 5), these seem to be in conflict with experimental data that shows larger diameters in the narrow than in the broad half of the trapezoid⁶ that the BION model successfully predicted (Fig. 2f). In addition, unlike the BION model, SR does not explain experimental findings when the environment is changed (Figs. 3 and 4).

The recently proposed 'default representation' (DR) is similar to SR, in that it also encodes a predictive map without considering spatial uncertainty⁴⁹. DR mainly differs from SR in that it is based on a 'default' behavioral policy that depends on basic physical constraints of the environment rather than the current task-dependent behavioral policy. Therefore, although the predictions of DR have not been spelled out for the experiments we considered here, they can be expected to be similar to those of SR, as long as the behavior of the animal is constrained mainly by environmental geometry (as in random pellet chasing, but see below).

Critically, SR-like and our uncertainty-based mechanisms need not be mutually exclusive. For example, there are changes in place cell responses (albeit not in grid cell responses^{58–60}) that happen in the absence of any change in visual inputs, when only the behavioral policy of the animal is changed⁶¹. Such changes may be naturally accounted for by SR, as changing the behavioral policy will also change the transition probabilities between locations, but might not be readily predicted by our uncertainty-based account. Conversely, the grid field deformations described for unexpected environmental changes that the BION model naturally captured (Figs. 3 and 4) seem difficult to reconcile with a purely SR-based account. Moreover, studies of grid field deformations using a 'hairpin maze' have found profound differences in grid cell responses when subjects navigated the same environment with vs. without walls, even though their trajectories (and hence transition probabilities) were constrained to be similar^{58,60}. These results challenge a purely SR-based (but not DR-based) account of grid cell firing⁴⁹, while they are consistent with an uncertainty-based mechanism given the critical role of visual inputs in determining spatial uncertainty, as revealed by our analyses. Taken together, we propose that SR (or DR) may be most relevant for place cell responses, while grid cell deformations may be driven primarily by uncertainty-based mechanisms.

Our work focuses on how a probabilistic estimate of location (and head direction) is computed from noisy and ambiguous sensory inputs. A complete understanding of navigational behavior also requires the formalization of how these estimates, and the uncertainty about them, are used for route planning⁶² – something that we only modeled phenomenologically (Fig. 1c). Indeed, recent work using a more formal model of the planning process²⁰ has been able to account for apparent violations of the Bayes-optimal combination of landmark and self motion cues^{27,28} and reconcile them with earlier results suggesting optimal cue combination²⁶. These results lend further support to our core suggestion: that Bayesian principles play a fundamental role in navigation.

Acknowledgements

We thank T. Durrant for technical help with digitizing heatmaps. This work was supported by Wolfson College, University of Cambridge (Junior Research Fellowship to Y.K.), the National Institutes of Health (R01NS117699 to D.M.W), the Air Force Office of Scientific Research (FA9550-22-1-0337 to D.M.W), the Wellcome Trust (Investiga-

tor Award in Science 212262/Z/18/Z to M.L.), and the Human Frontiers Science Programme (Research Grant RGP0044/2018 to M.L.). For the purpose of open access, the authors have applied a CC-BY public copyright license to any author accepted manuscript version arising from this submission.

References

- Gallistel, C. R. *The organization of learning*. (The MIT Press, 1990).
- O'Keefe, J. & Dostrovsky, J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research* **34**, 171–175 (1971).
- Fyhn, M., Molden, S., Witter, M. P., Moser, E. I. & Moser, M.-B. Spatial Representation in the Entorhinal Cortex. *Science* **305**, 1258–1264 (2004).
- Moser, E. I. & Moser, M.-B. A metric for space. *Hippocampus* **18**, 1142–1156 (2008).
- Bellmund, J. L. *et al.* Deforming the metric of cognitive maps distorts memory. *Nature Human Behaviour* 1–12 (2020).
- Krupic, J., Bauza, M., Burton, S., Barry, C. & O'Keefe, J. Grid cell symmetry is shaped by environmental geometry. *Nature* **518**, 232 (2015).
- O'Keefe, J. & Burgess, N. Geometric determinants of the place fields of hippocampal neurons. *Nature* **381**, 425–428 (1996).
- Hartley, T., Trinkl, I. & Burgess, N. Geometric determinants of human spatial memory. *Cognition* **94**, 39–75 (2004).
- Chen, X., He, Q., Kelly, J., Fiete, I. & McNamara, T. Bias in Human Path Integration Is Predicted by Properties of Grid Cells. *Current Biology* **25**, 1771–1776 (2015).
- Barry, C., Hayman, R., Burgess, N. & Jeffery, K. J. Experience-dependent rescaling of entorhinal grids. *Nature Neuroscience* **10**, nn1905 (2007).
- Keinath, A. T., Epstein, R. A. & Balasubramanian, V. Environmental deformations dynamically shift the grid cell spatial metric. *eLife* **7**, e38169 (2018).
- Keinath, A. T., Rehnitz, O., Balasubramanian, V. & Epstein, R. A. Environmental deformations dynamically shift human spatial memory. *Hippocampus* (2020).
- Stachenfeld, K. L., Botvinick, M. M. & Gershman, S. J. The hippocampus as a predictive map. *Nature Neuroscience* **20**, 1643 (2017).
- Krupic, J., Bauza, M., Burton, S., Lever, C. & O'Keefe, J. How environment geometry affects grid cell symmetry and what we can learn from it. *Philosophical Transactions of the Royal Society B: Biological Sciences* **369**, 20130188 (2014).
- Mok, R. M. & Love, B. C. A non-spatial account of place and grid cells based on clustering models of concept learning. *Nature Communications* **10**, 5685 (2019).
- Evans, T. & Burgess, N. Replay as structural inference in the hippocampal-entorhinal system. *bioRxiv* (2020).
- Wang, W. & Wang, W. Place cells and geometry lead to a flexible grid pattern. *Journal of Computational Neuroscience* 1–12 (2021).
- Durrant-Whyte, H. & Bailey, T. Simultaneous localization and mapping: part I. *IEEE Robotics & Automation Magazine* **13**, 99–110 (2006).
- Bailey, T. & Durrant-Whyte, H. Simultaneous localization and mapping (SLAM): part II. *IEEE Robotics & Automation Magazine* **13**, 108–117 (2006).
- Kessler, F., Frankenstein, J. & Rothkopf, C. A. A Dynamic Bayesian Actor Model explains Endpoint Variability in Homing Tasks. *bioRxiv* (2022).
- Cheung, A. Probabilistic Learning by Rodent Grid Cells. *PLOS Computational Biology* **12**, e1005165 (2016). Uses self-motion and boundary information. Hence doesn't consider visual uncertainty, nor does visual active sensing.
- Lakshminarasimhan, K. J. *et al.* Tracking the Mind's Eye: Primate Gaze Behavior during Virtual Visuomotor Navigation Reflects Belief Dynamics. *Neuron* **106**, 662–674.e5 (2020).
- Madl, T., Franklin, S., Chen, K., Montaldi, D. & Trapp, R. Bayesian Integration of Information in Hippocampal Place Cells. *PLoS ONE* **9**, e89762 (2014).
- Schröder, T. N. *et al.* Environmental anchoring of grid-like representations minimizes spatial uncertainty during navigation. *bioRxiv* (2020).
- Kanitscheider, I. & Fiete, I. Emergence of dynamically reconfigurable hippocampal responses by learning to perform probabilistic spatial reasoning. *bioRxiv* 231159 (2017).
- Nardini, M., Jones, P., Bedford, R. & Braddick, O. Development of Cue Integration in Human Navigation. *Current Biology* 689–693 (2008).
- Zhao, M. & Warren, W. H. How You Get There From Here. *Psychological Science* **26**, 915–924 (2015).
- Chen, X., McNamara, T. P., Kelly, J. W. & Wolbers, T. Cue combination in human spatial navigation. *Cognitive Psychology* **95**, 105–144 (2017).

29. Doeller, C. F., King, J. A. & Burgess, N. Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. *Proceedings of the National Academy of Sciences* **105**, 5915–5920 (2008).
30. Pickup, L. C., Fitzgibbon, A. W. & Glennerster, A. Modelling human visual navigation using multi-view scene reconstruction. *Biological Cybernetics* **107**, 449–464 (2013).
31. Burge, J. Image-Computable Ideal Observers for Tasks with Natural Stimuli. *Annual Review of Vision Science* **6**, 1–27 (2020).
32. Dayan, P. & Abbott, L. F. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems* (MIT Press, 2005).
33. Ma, W. J., Beck, J. M., Latham, P. E. & Pouget, A. Bayesian inference with probabilistic population codes. *Nat Neurosci* **9**, 1432–1438 (2006).
34. Ganguli, D. & Simoncelli, E. P. Efficient sensory encoding and Bayesian inference with heterogeneous neural populations. *Neural computation* **26**, 2103–34 (2014).
35. Burgess, N., Barry, C. & O’keefe, J. An oscillatory interference model of grid cell firing. *Hippocampus* **17**, 801–812 (2007).
36. Remme, M. W., Lengyel, M. & Gutkin, B. S. Democracy-independence trade-off in oscillating dendrites and its implications for grid cells. *Neuron* **66**, 429–437 (2010).
37. Burak, Y. & Fiete, I. R. Accurate path integration in continuous attractor network models of grid cells. *PLoS computational biology* **5**, e1000291 (2009).
38. Sreenivasan, S. & Fiete, I. Grid cells generate an analog error-correcting code for singularly precise neural computation. *Nature neuroscience* **14**, 1330–1337 (2011).
39. Banino, A. *et al.* Vector-based navigation using grid-like representations in artificial agents. *Nature* **557**, 429–433 (2018).
40. Sorscher, B., Mel, G. C., Ocko, S. A., Giocomo, L. M. & Ganguli, S. A unified theory for the computational and mechanistic origins of grid cells. *Neuron* (2022).
41. Schaeffer, R., Khona, M. & Fiete, I. No free lunch from deep learning in neuroscience: A case study through models of the entorhinal-hippocampal circuit. *bioRxiv* (2022).
42. Stensola, H. *et al.* The entorhinal grid map is discretized. *Nature* **492**, 72–78 (2012).
43. Lever, C., Burton, S., Jeewajee, A., O’Keefe, J. & Burgess, N. Boundary Vector Cells in the Subiculum of the Hippocampal Formation. *Journal of Neuroscience* **29**, 9771–9777 (2009).
44. Savelli, F., Yoganarasimha, D. & Knierim, J. J. Influence of boundary removal on the spatial representations of the medial entorhinal cortex. *Hippocampus* **18**, 1270–82 (2008).
45. Solstad, T., Boccara, C. N., Kropff, E., Moser, M.-B. & Moser, E. I. Representation of Geometric Borders in the Entorhinal Cortex. *Science* **322**, 1865–1868 (2008).
46. Cothi, W. & Barry, C. Neurobiological successor features for spatial navigation. *Hippocampus* **30**, 1347–1355 (2020).
47. Tanni, S., Cothi, W. d. & Barry, C. State transitions in the statistically stable place cell population correspond to rate of perceptual change. *Current Biology* (2022).
48. Monsalve-Mercado, M. M. & Leibold, C. Effect of boundaries on grid cell patterns. *Physical Review Research* **2**, 043137 (2020).
49. Piray, P. & Daw, N. D. Linear reinforcement learning in planning, grid fields, and cognitive control. *Nature communications* **12**, 1–20 (2021).
50. Burak, Y. & Fiete, I. R. Accurate Path Integration in Continuous Attractor Network Models of Grid Cells. *PLoS Computational Biology* (2009).
51. Whittington, J. C. *et al.* The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. *Cell* (2020).
52. Madl, T., Franklin, S., Chen, K., Montaldi, D. & Trapp, R. Bayesian Integration of Information in Hippocampal Place Cells. *PLoS ONE* **9**, e89762 (2014).
53. Vertes, E. & Sahani, M. A neurally plausible model learns successor representations in partially observable environments. In *Advances in Neural Information Processing Systems*, 13692–13702 (2019).
54. Lakshminarasimhan, K. J. *et al.* Tracking the Mind’s Eye: Primate Gaze Behavior during Virtual Visuomotor Navigation Reflects Belief Dynamics. *Neuron* **106**, 662–674.e5 (2020).
55. Sanders, H., Wilson, M. A. & Gershman, S. J. Hippocampal remapping as hidden state inference. *eLife* **9**, e51140 (2020).
56. Plitt, M. H. & Giocomo, L. M. Experience-dependent contextual codes in the hippocampus. *Nature Neuroscience* 1–10 (2021).
57. Dayan, P. Improving generalization for temporal difference learning: The successor representation. *Neural computation* **5**, 613–624 (1993).
58. Derdikman, D. *et al.* Fragmentation of grid cell maps in a multicompartment environment. *Nature neuroscience* **12**, 1325–1332 (2009).

59. Howard, L. *et al.* The Hippocampus and Entorhinal Cortex Encode the Path and Euclidean Distances to Goals during Navigation. *Current Biology* **24**, 1331–1340 (2014).
60. He, Q. & Brown, T. I. Environmental Barriers Disrupt Grid-like Representations in Humans during Navigation. *Current Biology* (2019).
61. Markus, E. *et al.* Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *The Journal of Neuroscience* **15**, 7079–7094 (1995).
62. Mattar, M. G. & Lengyel, M. Planning in the brain. *Neuron* **110**, 914–934 (2022).

Environmental geometry	Behavioral (Human)	Neural (Rodent)
Square vs. trapezoid	Bellmund et al. 2020 ⁵	Krupic et al. 2015 ⁶
Expanded & compressed rect-angles	Chen et al. 2015 ⁹ , Hartley et al. 2004 ⁸ , Keinath et al. 2020 ¹²	Barry et al. 2007 ¹⁰ , Keinath et al. 2018 ¹¹

Table 1 | Choice of experimental tasks.

Methods

Choice of experimental tasks

We chose experimental tasks with which we can test both behavioral and neural effects of uncertainty as a function of the environmental geometry. Typically, human experiments yielded suitable behavioral measures (e.g., endpoints in homing tasks⁵) and rodent experiments yielded suitable neural measures (e.g., electrophysiological recordings from grid cells while the animal chased a randomly dropped pellets). Therefore, we chose pairs of experiments that employed analogous environmental geometry and compared their results with our model (Table 1).

BION ideal observer model

Generative model

We consider an observer (a modeled human or nonhuman subject) that navigates an environment while trying to find out its own location within the environment. In experiments we modeled, the subjects were familiarized with the environment, or believed that they were in a familiar environment, so we assume that the observer's uncertainty about the configuration of the environment (also known as a 'map') to be negligible. The observer still has uncertainty about where it is, due to noise in its sensory inputs (visual, self-motion, and tactile) and its movements.

Formally, at each time step t , the observer intends to turn by u^θ , which results in a noisy rotation \tilde{u}^θ in its allocentric head direction θ :

$$\tilde{u}_t^\theta \sim \text{vM}(u_t^\theta \Delta t, \kappa \Delta t) \quad (2)$$

$$\theta_t = \theta_{t-1} + \tilde{u}_t^\theta \quad (3)$$

where $\text{vM}(\cdot)$ denotes von Mises distribution. Likewise, the observer intends to move forward by u^x , which results in a noisy forward motion \tilde{u}^x in its allocentric location ℓ , except when it meets a wall $\partial\mathcal{E}$, in which case it stops at the nearest intersection of its trajectory and the wall:

$$\tilde{u}_t^x \sim \text{gamma}(u_t^x \Delta t, (\sigma^x)^2 u_t^x \Delta t) \quad (4)$$

$$\tilde{\mathbf{u}}^\ell := (\tilde{u}^x, 0)^\top \quad (5)$$

$$\ell_t' = \ell_{t-1} + \mathbf{R}(\theta_t) \tilde{\mathbf{u}}_t^\ell \quad (6)$$

$$\mathcal{B}(\ell, \ell') := \overline{\ell\ell'} \cap \partial\mathcal{E} \quad (7)$$

$$\ell_t = \begin{cases} \ell_t' & \text{if } \ell_t' \in \mathcal{E} \\ \text{argmin}_{\ell \in \mathcal{B}(\ell_{t-1}, \ell_t')} d(\ell_{t-1}, \ell) & \text{if } \ell_t' \notin \mathcal{E} \end{cases} \quad (8)$$

The observer receives egocentric visual input on a flat rectangular array of retinotopic neurons through a pinhole projection of the 3D scene. The individual neuron's signal V is subject to Poisson noise with the tuning curve (i.e., expected firing rate) given by a blurred first person view \bar{v} , which is a function of the observer's location and head direction:

$$V_t^i \sim \text{Poisson}(\bar{v}_i(\ell_t, \theta_t) \Delta t) \quad (9)$$

Finally, the observer receives egocentric tactile input W if it is within distance x_w from a wall, with reliability p_w :

$$W_\gamma \sim \begin{cases} \text{Bernoulli}(\frac{1}{2} + \frac{1}{2}p_w) & \text{if } |\mathcal{B}(\ell, \ell + x_w \mathbf{u}(\theta + \gamma))| > 0 \\ \text{Bernoulli}(\frac{1}{2} - \frac{1}{2}p_w) & \text{otherwise} \end{cases} \quad (10)$$

There were a few differences between the generative models we used for different experiments: (1) For human experiments, which were performed in VR, we did not use tactile input (Eq. 10). (2) For the

scaling experiments done in large environments^{8,9}, we used isotropic Gaussian distribution instead of gamma distribution for the translation (Eq. 4).

A recent theory used a normative model to explain undershoot in navigation to a remembered target location by having a prior that places higher probability on slower speeds (i.e. a slowness prior rather than imperfect integration)²². In contrast, in our model we assumed that the perception of the control signal is well calibrated and is observed, while the resulting movement is noisy, following the convention of some of the robotic navigation literature¹⁹. That said, as our model specifies $P(\Delta(\ell, \theta) | (\mathbf{u}^\ell, u^\theta))$, it would be possible to emulate a slowness prior. However, we did not include slowness prior in our analyses because (1) in the experiments we modeled, vestibular and/or afferent copy of motor signals (in addition to optic flow) were available, which would have made the perception of self motion more reliable, compared to the experiment modeled by Ref. 22 where only optic flow was available; (2) the main results we replicated were differences between conditions (e.g., difference of biases between expansion vs. compression of environments) rather than the bias relative to the true distance, so that any undershoot would be subtracted in the comparison.

Bayesian filtering

The ideal observer is given by Bayesian inversion of the generative model above. Specifically, we used Bayesian filtering, where the observer infers the posterior distribution over its allocentric pose (location and head direction) given the full history of noisy sensory inputs and control signal.

$$P(\mathbf{z}_T = \mathbf{z} | \mathbf{u}_{0:T}, \mathbf{o}_{0:T}) \propto \int \underbrace{P(\mathbf{o}_T | \mathbf{z}_T = \mathbf{z})}_{\text{observation}} \underbrace{P(\mathbf{z}_T = \mathbf{z} | \mathbf{u}_T, \mathbf{z}_{T-1} = \mathbf{z}')}_{\text{prediction}} \underbrace{P(\mathbf{z}_{T-1} = \mathbf{z}' | \mathbf{u}_{0:T-1}, \mathbf{o}_{0:T-1})}_{\text{dynamic prior}} d\mathbf{z}' \quad (11)$$

where the pose is the latent state $\mathbf{z} := (\ell, \theta)$ that consists of the allocentric location ℓ and head direction θ , and observation $\mathbf{o} := (\mathbf{v}, \mathbf{w})$ consists of egocentric visual input \mathbf{v} and tactile input \mathbf{w} , which we assume to be independent given \mathbf{z} . Hence, $P(\mathbf{o} | \mathbf{z}) = P(\mathbf{v} | \mathbf{z}) P(\mathbf{w} | \mathbf{z})$. For simplicity, we assume that the control signal \mathbf{u} (= intended movement) is known by the observer without noise, but that the actual movement is noisy¹⁹.

Visual likelihood

A key novel component of the BION model is its visual likelihood. It is computed as the likelihood of the egocentric visual observation given an allocentric pose. This is achieved by inverting the generative model of the visual input (Eq. 9). Specifically:

$$P(V = \mathbf{v} | \mathbf{z}) = \prod_i \frac{(\bar{v}^i(\mathbf{z}) \Delta t)^{r^i \Delta t}}{(r^i \Delta t)!} e^{-\bar{v}^i(\mathbf{z}) \Delta t} \quad (12)$$

where i indexes neurons with a retinotopic arrangement ('retinotopic neural population') and $\mathbf{z} := (\ell_t, \theta_t)$. This is an extension of a classical Bayesian decoder (Eq. 3.30 of Ref. 32) to the 3D case, by using tuning curves obtained from the projection of a 3D scene onto the retinal plane, $\bar{v}^i(\mathbf{z})$. This extension has an important implication for the amount of visual information. While the assumption of independent Poisson noise of pixels is an obvious simplification, it still yields powerful predictions when combined with the law of optics of the projection from the 3D world onto the 2D retina (Supp. Material and Results). The extension also has implications for numerical computation. To see this, note that expressions like the above is typically evaluated using logarithm for numerical stability as:

$$\log P(V = \mathbf{v} | \mathbf{z}) = \sum_i [r^i \log \bar{v}^i(\mathbf{z}) - \bar{v}^i(\mathbf{z}) \cdots] \quad (13)$$

Here, unlike Eq. 3.31 of Ref. 32, we cannot ignore the second term, because we cannot assume it will be independent of \mathbf{z} . That is because in case of Ref. 32, the tuning curves are arranged evenly in the space of the stimulus, whereas in our case, the space of the stimulus, \mathbf{z} , cannot be linearly mapped onto the retinal space (on which the tuning curves are approximately evenly arranged). As a result, for example, some places \mathbf{z} may simply be darker than other places and evoke fewer spikes across the retinal space, making the second term different between different \mathbf{z} .

BION neural representation model

We use the readout $P^g(\ell)$ of a Bayesian decoder of the animal's location ℓ given Poissonian firing rates \mathbf{r} of a grid cell population as 'neural posterior'. Our hypothesis is that the tuning curves of the grid cells are warped such that this neural posterior represents the ideal observer's posterior P^b as closely as possible. Formally, the neural posterior has a parametric form obtained by inverting the generative model of the firing rates:

$$r_i \sim \text{Poisson}(f_i^{\mathcal{E}^*}(\mathbf{z})T) \quad (14)$$

$$P^g(\ell; \mathbf{r}, \Theta^w) \propto P(\ell \mid \mathbf{r}; f^{\tilde{\mathcal{E}}}(\mathbf{z}))^{\alpha'} \quad (15)$$

$$\propto \left[P(\mathbf{r} \mid \ell; f^{\tilde{\mathcal{E}}}(\mathbf{z})) \underbrace{P(\mathbf{r}; f^{\tilde{\mathcal{E}}}(\mathbf{z}))}_{\perp \ell} \right]^{\alpha'} \quad (16)$$

$$\log P^g(\ell; \mathbf{r}, \Theta^w) = \alpha' \sum_i \log P(r_i \mid \ell; f_i^{\tilde{\mathcal{E}}}(\mathbf{z})) + (\text{const}) \quad (17)$$

where i indexes neurons, $f^{\mathcal{E}^*}$ is the 'encoding' tuning curve, which gives rise to the firing rate $\mathbf{r} := \{r_i\}$, and $f^{\tilde{\mathcal{E}}}$ is the 'decoding' tuning curve. Here $\tilde{\mathcal{E}}$ is the familiar environment, and \mathcal{E}^* is the test environment, which may or may not be the same as $\tilde{\mathcal{E}}$. α' is a parameter that determines the sharpness of the posterior, which ensures that the neural posterior does not become infinitely sharp even with infinite number of neurons. It is necessary because we assume that there is nonzero uncertainty in the ideal observer's posterior due to noisy observation (see Results).

Our strategy is to first optimize the tuning curves to match the neural posterior with the ideal observer's posterior, and then to compare them with the experimental data. For this comparison, for the neural model, we chose, $\mathbf{z} := (\ell, h)$ where ℓ is the location of the animal and h is the last wall approached, to match the way the tuning curves are reported in Ref. 11 (Fig. 4d). Specifically, we fit $f^{\mathcal{E}^*}$ (observable), parameterized with Θ^w to be a warping with a mesh. When the test environment is familiar, $f^{\tilde{\mathcal{E}}}$ is constrained to be the same as $f^{\mathcal{E}^*}$, which is fit. Otherwise, $f^{\tilde{\mathcal{E}}}$ is fixed to the value fit when the particular $\tilde{\mathcal{E}}$ was the test environment, and $f^{\mathcal{E}^*}$ is fit.

In practice, we assume that head direction uncertainty is small (because salient orientation cues were available in all experiments) and that tactile reliability is high, so we treat h as observed. Therefore, $f^{\mathcal{E}^*}$ and $f^{\tilde{\mathcal{E}}}$ depend on h , but we only need to compute the posterior over the location ℓ .

Also, for simplicity of optimization, we assume the limit of an infinite-sized neural population. (Note that this contrasts with our assumption of a finite retinotopic neural population (Eq. 12), where we used the Poisson noise as a proxy for noise and ambiguity inherent in the egocentric visual input.) This assumption allows us to express the right hand side of Eq. 17 with $f^{\mathcal{E}^*}$ instead of \mathbf{r} :

$$\alpha' \sum_i \log P(r_i \mid \ell; f_i^{\tilde{\mathcal{E}}}(\mathbf{z})) = \alpha' \sum_i \sum_{j=1}^{N_i} \log P(r_j \mid \ell; f_i^{\tilde{\mathcal{E}}}(\mathbf{z})) \quad (18)$$

$$= \alpha' \sum_i \sum_{j=1}^{N_i} \left[r_j T \log(f_i^{\tilde{\mathcal{E}}}(\mathbf{z})T) - f_i^{\tilde{\mathcal{E}}}(\mathbf{z})T - \underbrace{\log(r_j T!)}_{\perp \mathbf{z}} \right] \quad (19)$$

$$\log P^g(\ell; \mathbf{z}, \Theta) = \sum_i \alpha_i \left[f_i^{\mathcal{E}^*}(\mathbf{z})T \log(f_i^{\tilde{\mathcal{E}}}(\mathbf{z})T) - f_i^{\tilde{\mathcal{E}}}(\mathbf{z})T \right] + \underbrace{\dots}_{\perp \mathbf{z}} \quad (20)$$

where N_i is the number of neurons with the same tuning curve. We suppress the Poisson variability and use $f_i^{\mathcal{E}^*}$ in place of r_i , by setting the population size $N_i \rightarrow \infty$ and setting the contribution of an individual neuron $\alpha' := \alpha_i/N_i \rightarrow +0$ while keeping α_i finite. For simplicity, we treat the N_i neurons that share the same tuning curve $f_i^{\mathcal{E}^*}$ as a single deterministic neuron i with a continuous firing rate $f_i^{\mathcal{E}^*}$ from here. We then fit $\alpha := \{\alpha_i\}$, to account for the difference between the contribution of each neuron to the decoded neural posterior. Since our hypothesis is that warping is the dominant factor determining the posterior, we constrain α_i to be the same for all neurons within a grid module. We denote all parameters controlling the neural posterior together as $\Theta := (\Theta^w, \alpha)$. We fix $T = 1$ for simplicity.

To fit the warping, we use gradient descent to find:

$$\hat{\Theta} = \underset{\Theta}{\operatorname{argmin}} \mathcal{L}^{\text{fit}} + \mathcal{L}^{\text{reg}} \quad (21)$$

where

$$\mathcal{L}^{\text{fit}} := \langle D_{\text{KL}}[P_t^b \parallel P_{r_i}^g] \rangle_t \quad (22)$$

is the divergence of the neural posterior from the ideal observer's posterior, averaged across time. \mathcal{L}^{reg} is a regularization loss that penalizes deviation of the angle between neighboring nodes of the mesh from 90° . It is defined as:

$$\mathcal{L}^{\text{reg}} := w^{\text{reg}} \sum \left(\frac{1 - \sin \vartheta}{2} \right)^\beta \quad (23)$$

Here, ϑ is an angle subtended by neighboring points in the mesh controlling the warp, w^{reg} and β are hyperparameters that control the weight and stiffness of the regularization, respectively.

It turns out that we can further simplify the calculation of the \mathcal{L}^{fit} , because it is equivalent to minimizing the divergence of the posterior (in our case, defined over believed location $\tilde{\ell}$), given a state that parameterizes the neural firing rate (in our case, \mathbf{z}), *averaged across time*:

$$\langle D_{\text{KL}}[P^b(\tilde{\ell}; \mathbf{z}_t, t) \parallel P^g(\tilde{\ell}; \mathbf{z}_t, \Theta)] \rangle_t = \frac{1}{T} \sum_t \sum_{\tilde{\ell}} P^b(\tilde{\ell}; \mathbf{z}_t, t) \left(\log P^b(\tilde{\ell}; \mathbf{z}_t, t) - \log P^g(\tilde{\ell}; \mathbf{z}_t, \Theta) \right) \quad (24)$$

$$= C - \frac{1}{T} \sum_{\mathbf{z}} \sum_{\tilde{\ell}} \log P^g(\tilde{\ell}; \mathbf{z}, \Theta) \cdot \sum_{t: \mathbf{z}_t = \mathbf{z}} P^b(\tilde{\ell}; \mathbf{z}_t, t) \quad (25)$$

$$= C - \frac{1}{T} \sum_{\mathbf{z}} \sum_{\tilde{\ell}} \log P^g(\tilde{\ell}; \mathbf{z}, \Theta) M(\mathbf{z}) \overline{P^b}(\tilde{\ell}; \mathbf{z}) \quad (26)$$

Here, C is a constant term that is independent of Θ , $M(\mathbf{z}) := |\{t : \mathbf{z}_t = \mathbf{z}\}|$ is the number of visits to \mathbf{z} , and $\overline{P^b}(\mathbf{z}) := \frac{1}{M(\mathbf{z})} \sum_{t: \mathbf{z}_t = \mathbf{z}} P^b(\tilde{\ell}; \mathbf{z}_t, t)$ is the ideal observer's posterior at \mathbf{z} averaged across time, which is also independent of Θ . These two terms can be pre-computed, which greatly accelerates the optimization of Θ when $|\{\mathbf{z}\}| \ll T$, as is the case for our simulations.

To gauge the quality of the fit, we show optimized D_{KL} along with controls in Fig. 1e, which are defined as follows:

$$D_{\text{KL}}^{\text{optim}} := \langle D_{\text{KL}}[P^b(\tilde{\ell}; \mathbf{z}_t, t) \parallel P^g(\tilde{\ell}; \mathbf{z}_t, \hat{\Theta})] \rangle_t \quad (27)$$

$$D_{\text{KL}}^{\text{within}} := \frac{1}{\sum_{\mathbf{z}'} M(\mathbf{z}')} \sum_{\mathbf{z}'} M(\mathbf{z}') D_{\text{KL}}[\langle P^b(\tilde{\ell}; \mathbf{z}_t = \mathbf{z}') \rangle_t \parallel \langle P^g(\tilde{\ell}; \mathbf{z}_t = \mathbf{z}', \hat{\Theta}) \rangle_t] \quad (28)$$

$$D_{\text{KL}}^{\text{reg}} := \langle D_{\text{KL}}[P^b(\tilde{\ell}; \mathbf{z}_t, t) \parallel P^g(\tilde{\ell}; \mathbf{z}_t, \Theta^{\text{reg}})] \rangle_t \quad (29)$$

$$D_{\text{KL}}^{\text{shuffle}} := \langle D_{\text{KL}}[P^b(\tilde{\ell}; \mathbf{z}_{\tilde{t}(t)}, \tilde{t}(t)) \parallel P^g(\tilde{\ell}; \mathbf{z}_t, \hat{\Theta})] \rangle_t \quad (30)$$

where Θ^{reg} is the parameter of the warping that gives regular grid fields, and $\tilde{t}(t)$ is the t -th time step after shuffling.

For results in Fig. 4d, we only fit $f^{\mathcal{E}^*}$ while fixing $f^{\mathcal{E}}$ to be regular grid fields. This was because the familiar environment was always a square, in which the effect of warping was small.

Homing task

Overview

Homing experiments^{5,8,9} consisted of a learning phase followed by a test phase. In the learning phase, subjects navigated to a visible target location. In the test phase, subjects had to navigate back to the hidden target location from memory, starting from a designated location that differed from the target location. In Ref. 9, subjects had to follow a designated path (that was randomly chosen by the computer program every trial) before they navigate back to the hidden target location.

In both phases, at each time step, a control signal \mathbf{u} is chosen among four options: staying still, moving forward, turning right, and turning left. Formally:

$$\mathbf{u} := (u^x, u^\theta) \quad (31)$$

$$\in \mathcal{C} \quad (32)$$

$$:= \{\mathbf{u}^0, \dots, \mathbf{u}^C\} \quad (33)$$

$$:= \{(0, 0), (\Delta\ell, 0), (\Delta\ell, -\Delta\theta), (\Delta\ell, +\Delta\theta)\} \quad (34)$$

Learning phase

We modeled the observer to be from the starting location toward the true target location in as straight as possible trajectory at each moment (or following a designated path⁹). The control signals are sampled passively, i.e., without regard to the observer's beliefs. If there are more than two possible such control signals at a given moment, one is chosen at random. When the observer arrives at the true target location, its belief about its location is stored as the mean believed target location $\bar{\ell}$, given all observations \mathbf{o} and control signals \mathbf{u} (which include translation and rotation):

$$\tilde{p}^{\text{learn}}(\tilde{\ell}^{\text{learn}}) \propto P(\tilde{\ell}^{\text{learn}} | \mathbf{o}_{0:T^{\text{learn}}}, \mathbf{u}_{0:T^{\text{learn}}}, \tilde{\mathcal{E}})^{\beta_{\text{learn}}} \quad (35)$$

$$\bar{\ell} = E(\tilde{\ell}^{\text{learn}}) \quad (36)$$

where $0 < \beta_{\text{learn}} \leq 1$ is the inverse temperature parameter that models forgetting (the lower, the more forgetting).

Test phase: position estimation

The observer chooses its own control signal. For this, it first computes the counterfactual expected value from executing each candidate control:

$$v_j := \sum_{\tilde{\ell}^{\text{learn}}} \sum_{\tilde{\ell}^{\text{test}}} v(\tilde{\ell}^{\text{learn}}, \tilde{\ell}^{\text{test}}) \tilde{p}^{\text{learn}}(\tilde{\ell}^{\text{learn}}) \tilde{p}^j(\tilde{\ell}^{\text{test}}) \quad (37)$$

where the value is defined to be the negative squared Euclidean distance:

$$v(\tilde{\ell}^{\text{learn}}, \tilde{\ell}^{\text{test}}) := -\|\tilde{\ell}^{\text{learn}} - \tilde{\ell}^{\text{test}}\|_2^2 \quad (38)$$

The counterfactual location posterior when staying is defined to be the same as the current posterior:

$$\tilde{p}^0(\tilde{\ell}^{\text{test}}) := P(\tilde{\ell}^{\text{test}} | \mathbf{o}_{0:t}^{\text{test}}, \mathbf{u}_{0:t}^{\text{test}}, \tilde{\mathcal{E}}) \quad (39)$$

In general, the counterfactual location distribution after executing the j -th candidate control \mathbf{u}^j is defined in the believed environment space:

$$\tilde{p}^j(\tilde{\ell}^{\text{test}}) := P(\tilde{\ell}^{\text{test}} | \mathbf{o}_{0:t}^{\text{test}}, \mathbf{u}_{0:t}^{\text{test}}, \mathbf{u}^j, \tilde{\mathcal{E}}) \quad (40)$$

with a temperature parameter σ_v^2 , which effectively scales the distance in Eq. 38:

$$P(\mathbf{u}_t^{\text{test}} = \mathbf{u}^j) \propto e^{v_j / \sigma_v^2} \quad (41)$$

The final response is made when the observer chooses to stay, i.e., $\mathbf{u}_t^{\text{test}} = \mathbf{u}^0$, and the observer's true location at that time is recorded as the responded location.

Test phase: distance estimation

In addition to the homing task, subjects in Ref. 5 also performed a distance estimation task. In this task, subjects were asked to indicate the distance between two target objects they have seen during the learning phase. During the test phase, subjects were asked to walk the distance they want to indicate in a virtual circular enclosure (to which they were previously familiarized). Specifically, they were asked to walk the same distance as the distance between a pair of objects, starting from a root of a static visual arrow placed randomly in the familiar circular environment, following the direction of the arrow, which always pointed to the center.

In our simulation, the observer is driven from a random starting location toward the root of the arrow in as straight as possible trajectory at each moment, until it arrives at the root (similar to the learning phase of the homing task Eq. 35). The observer starts this trajectory with a uniform prior belief over locations and heading directions, and ends it with a distribution of believed position of the root $P(\tilde{\ell}_{\text{root}})$. Note that, since the observer continues onto the estimation as soon as it reaches the root of the arrow, we assume it does not forget the position of the root (i.e., retains the full posterior without exponentiation by β_{learn} , unlike Eq. 35 and Eq. 36).

At each time step from then on, the observer decides whether to walk another step in the direction of the arrow or stay, using the counterfactual value of the predicted location after each action (walk or stay), as in the test phase of the homing task Eq. 37-Eq. 41. Here, the value v_u^{est} of a control signal $u_t = u \in \{\text{walk}, \text{stay}\}$ is defined as the negative expected squared difference between the predicted distance from the root of the arrow and the believed distance between the given pair of objects (j, k):

$$v_u^{\text{est}} = \sum_{\ell} \sum_{\ell'} \tilde{P}(\tilde{\ell}_{\text{root}} = \ell) \tilde{P}^u(\ell' | \tilde{\ell}_{\text{root}} = \ell, u_{0:t-1}, \mathbf{o}_{0:t-1}) v^{\text{est}}(\|\ell - \ell'\|_2) \quad (42)$$

$$v^{\text{est}}(\|\ell - \ell'\|_2) = -(\|\ell - \ell'\|_2 - \|\bar{\ell}_j - \bar{\ell}_k\|_2)^2 \quad (43)$$

Then the choice is made based on the softmax over u , with the temperature parameter σ_v , which effectively scales the difference between distances in Eq. 43, as it does in the homing task (Eq. 41). The estimation trial ends when the observer chooses to stay, and the observer's true distance from the root of the arrow is recorded as the responded distance.

Supplementary Material

Visual uncertainty

A point landmark and a single neuron

While all results are obtained from numerical simulation using Eq. 11, we also derived simple analytic relation between Fisher information and the environmental geometry to gain analytic insights, and checked its correspondence with the numerical simulation (Fig. 2a). Here we start by deriving the Fisher information I_ϵ^ℓ of the apparent position $\tilde{\epsilon}$ about the true distance ℓ along an axis of interest, assuming that the axis of interest is the same as the optical axis of a pinhole projection (the line of sight of our ideal observer). We further assume that the distance of the point landmark from the axis, ϵ , as well as blur σ (on the image plane at a focal length of 1 without loss of generality) is known, which gives the distribution of the apparent position $\tilde{\epsilon} \sim \mathcal{N}(\epsilon', \sigma)$ with $\epsilon' := \epsilon/\ell$. Then:

$$I_\epsilon^\ell(\ell) = \left\langle -\frac{\partial^2}{\partial \ell^2} \ln p(\tilde{\epsilon} | \ell) \right\rangle_{\tilde{\epsilon}} \quad (\text{S1})$$

$$= - \int_{-\infty}^{\infty} d\tilde{\epsilon} p(\tilde{\epsilon} | \ell) \frac{\partial^2}{\partial \ell^2} \ln p(\tilde{\epsilon} | \ell) \quad (\text{S2})$$

$$= \frac{1}{2} C \int_{-\infty}^{\infty} d\tilde{\epsilon} e^{-\frac{1}{2} \left(\frac{\tilde{\epsilon} - \epsilon'}{\sigma} \right)^2} \frac{\partial^2}{\partial \ell^2} \left(\frac{\tilde{\epsilon} - \epsilon'}{\sigma} \right)^2 \quad (\text{S3})$$

where $C = \frac{1}{\sigma\sqrt{2\pi}}$ is the normalization factor. The partial derivative evaluates to:

$$\frac{\partial^2}{\partial \ell^2} \left(\frac{\tilde{\epsilon} - \epsilon'}{\sigma} \right)^2 = 2 \frac{\partial}{\partial \ell} \left(\frac{\tilde{\epsilon} - \epsilon\ell^{-1}}{\sigma} \right) \left(\frac{\epsilon}{\sigma} \ell^{-2} \right) \quad (\text{S4})$$

$$= \frac{2}{\sigma^2} \frac{\partial}{\partial \ell} (\tilde{\epsilon}\ell^{-2} - \epsilon^2\ell^{-3}) \quad (\text{S5})$$

$$= \frac{2}{\sigma^2} (-2\tilde{\epsilon}\ell^{-3} + 3\epsilon^2\ell^{-4}) \quad (\text{S6})$$

$$= \underbrace{\frac{2}{\sigma^2} (-2\epsilon\ell^{-3})}_{-4\epsilon\ell^{-3}/\sigma^2} \left(\tilde{\epsilon} - \epsilon' + \epsilon' + \underbrace{\frac{3\epsilon^2\ell^{-4}}{-2\epsilon\ell^{-3}}}_{-\frac{1}{2}\epsilon'} \right) \quad (\text{S7})$$

Substituting back to Eq. S3:

$$I_\epsilon^\ell(\ell) = \frac{1}{2} C \left(-\frac{4\epsilon\ell^{-3}}{\sigma^2} \right) \int_{-\infty}^{\infty} d\tilde{\epsilon} e^{-\frac{1}{2} \left(\frac{\tilde{\epsilon} - \epsilon'}{\sigma} \right)^2} \left((\tilde{\epsilon} - \epsilon') - \frac{1}{2}\epsilon' \right) \quad (\text{S8})$$

To evaluate the integral, note that:

$$\frac{\partial}{\partial \tilde{\epsilon}} e^{-\frac{1}{2} \left(\frac{\tilde{\epsilon} - \epsilon'}{\sigma} \right)^2} = -\frac{1}{\sigma} \left(\frac{\tilde{\epsilon} - \epsilon'}{\sigma} \right) e^{-\frac{1}{2} \left(\frac{\tilde{\epsilon} - \epsilon'}{\sigma} \right)^2} \quad (\text{S9})$$

$$= -\frac{1}{\sigma^2} (\tilde{\epsilon} - \epsilon') e^{-\frac{1}{2} \left(\frac{\tilde{\epsilon} - \epsilon'}{\sigma} \right)^2} \quad (\text{S10})$$

Substituting back to Eq. S8:

$$I_\epsilon^\ell(\ell) = \frac{1}{2} C \left(-\frac{4\epsilon\ell^{-3}}{\sigma^2} \right) \left(-\sigma^2 \underbrace{\left[e^{-\frac{1}{2} \left(\frac{\tilde{\epsilon} - \epsilon'}{\sigma} \right)^2} \right]_{\tilde{\epsilon}=-\infty}^{\infty}}_{=0} - \frac{1}{2}\epsilon' \underbrace{\int_{-\infty}^{\infty} d\tilde{\epsilon} e^{-\frac{1}{2} \left(\frac{\tilde{\epsilon} - \epsilon'}{\sigma} \right)^2}}_{=1/C} \right) \quad (\text{S11})$$

$$= \frac{\epsilon^2}{\sigma^2 \ell^4} \quad (\text{S12})$$

The result's unit is $(\text{length})^{-2}$, as it should be from Eq. S1, where $\frac{\partial^2}{\partial \ell^2}$ (whose unit is $(\text{length})^{-2}$) is applied to $((\tilde{\epsilon} - \epsilon')/\sigma)^2$, which is dimensionless since $\tilde{\epsilon}$, ϵ' , and σ are dimensionless ratios between lengths (e.g., $\epsilon' := \epsilon/\ell$).

An array of Poissonian neurons

Now we derive the Fisher information $I_v^\ell(\ell)$ about ℓ given the firing rates \mathbf{r} of an array of N neurons with independent Poisson spikes, which is defined as:

$$I_v^\ell(\ell) := - \int_{\mathbf{r}} d\mathbf{r} p(\mathbf{r} | \ell) \frac{\partial^2}{\partial \ell^2} \ln p(\mathbf{r} | \ell) \quad (\text{S13})$$

While the results of this step has been known³², we describe them here for completeness. Here, our aim is to express $I_v^\ell(\ell)$ as a function of each neuron i 's tuning curve $f_i(\ell)$, which is defined as the mean firing rate given the true stimulus ℓ , i.e.:

$$f_i(\ell) := \int_{r_i} dr_i p(r_i | \ell) r_i \quad (\text{S14})$$

First consider $p(\mathbf{r} | \ell)$ in Eq. S13. Since we assume that each neuron's spikes follow independent Poisson distribution, the likelihood of observing a certain vector of firing rates \mathbf{r} is simply:

$$p(\mathbf{r} | \ell) = \prod_{i=1}^N \frac{(f_i(\ell) T)^{r_i T}}{(r_i T)!} \quad (\text{S15})$$

where T is the duration during which the spikes are observed. Now we can evaluate the partial derivative in Eq. S13 using Eq. S15:

$$\frac{\partial^2}{\partial \ell^2} \ln p(\mathbf{r} | \ell) = \frac{\partial^2}{\partial \ell^2} \ln \left(\prod_{i=1}^N \frac{(f_i(\ell) T)^{r_i T}}{(r_i T)!} \right) \quad (\text{S16})$$

$$= \frac{\partial^2}{\partial \ell^2} \left(\sum_{i=1}^N (r_i T \ln(f_i(\ell) T) - f_i(\ell) T - \ln(r_i T)!) \right) \quad (\text{S17})$$

$$= \frac{\partial}{\partial \ell} \left(T \sum_{i=1}^N \left(\frac{T f_i'(\ell)}{f_i(\ell) T} r_i - f_i'(\ell) \right) \right) \quad (\text{S18})$$

$$= T \sum_{i=1}^N \left(\left(\frac{f_i''(\ell)}{f_i(\ell)} - \frac{f_i'(\ell)^2}{f_i(\ell)^2} \right) r_i - f_i''(\ell) \right) \quad (\text{S19})$$

Finally, substitute Eq. S19 back into Eq. S13 and simplify using Eq. S14:

$$I_v^\ell(\ell) = - \int_{\mathbf{r}} d\mathbf{r} p(\mathbf{r} | \ell) T \sum_{i=1}^N \left(\left(\frac{f_i''(\ell)}{f_i(\ell)} - \frac{f_i'(\ell)^2}{f_i(\ell)^2} \right) \underbrace{r_i}_{\text{use Eq. S14}} - f_i''(\ell) \right) \quad (\text{S20})$$

$$= -T \sum_{i=1}^N \left(\left(-\frac{f_i'(\ell)^2}{f_i(\ell)^2} f_i(\ell) + \frac{f_i''(\ell)}{f_i(\ell)} f_i(\ell) - f_i''(\ell) \right) \right) \quad (\text{S21})$$

$$= T \sum_{i=1}^N \frac{f_i'(\ell)^2}{f_i(\ell)} \quad (\text{S22})$$

Therefore, as aimed, we expressed $I_v^\ell(\ell)$ as a function of $f_i(\ell)$.

A retinal image

We are finally ready to derive the Fisher information $I_v^\ell(\ell)$ about the 3D distance ℓ from neurons with a uniform 2D retinotopic array of tuning curves. Assume a Gaussian tuning curve

$$f_i(\ell) := r_{\max} e^{-\frac{1}{2} \left(\frac{\epsilon_i - \ell}{\sigma} \right)^2} \quad (\text{S23})$$

then

$$f_i'(\ell) = -\frac{1}{2} \cdot 2 \left(\frac{\epsilon_i - \ell}{\sigma} \right) \left(\frac{\epsilon_i}{\sigma} - \frac{\ell}{\sigma} \right) f_i(\ell) \quad (\text{S24})$$

$$= -\frac{\epsilon_i - \ell}{\sigma^2} (\epsilon_i - \ell) f_i(\ell) \quad (\text{S25})$$

Substitute Eq. S23 and Eq. S25 into Eq. S22:

$$I_v^\ell(\ell) = \frac{\epsilon^2 T}{\ell^4} \sum_{i=1}^N \left(\frac{(\epsilon_i - \epsilon \ell^{-1})^2}{\sigma^4} f_i(\ell) \right) \quad (\text{S26})$$

Compare this with Ref. 32, Eqs. 3.46-48 to approximate the sum with an integral to obtain:

$$I_v^\ell(\ell) = \frac{\epsilon^2}{\ell^4} \cdot \frac{(2\pi)^{D/2} \rho \sigma^{D-2} r_{\max} T}{D} \quad (\text{S27})$$

$$= \pi \rho r_{\max} T \cdot \frac{\epsilon^2}{\ell^4} \quad (\text{S28})$$

where ρ is the number of neurons in a unit range of $\epsilon' := \epsilon/\ell$, and D is the number of dimensions of the space where the neurons are arranged, which is 2 for us. This expression is used for the lines in Fig. 2a, and its simplified form is presented as Eq. 1. Note that here we used ℓ instead of d to avoid confusion during integration.

Intriguingly, as explained in Ref. 32, the Fisher information is independent of the width of the tuning curve σ , because the number of neurons responding to the stimulus scales with $\rho \sigma^D$. The difference from Ref. 32 is the factor of ϵ^2/ℓ^4 , which arises because our quantity of interest ℓ is inversely related to the center of the tuning curve ϵ/ℓ (rather than linearly related, which would have given the center of ℓ). Note that the resulting expression for $I_v^\ell(\ell)$ has the same unit as $I_\epsilon^\ell(\ell)$, and is proportional to it without the σ term, scaled with the number of spikes, $\rho r_{\max} T$.

Bias in distance estimation – analytic intuition

Our numerical simulation replicated different biases in distance estimation between environments and between sides of the trapezoid (Fig. 2d, bottom). Here, we provide analytic intuition for the results with a simplified version of the ideal observer model. We derive a relation that explains how strongly the bias in distance estimation along a given axis (x or y) depends on the (half of the) sizes of the environment along that axis ('depth', L) and the axis orthogonal to it ('eccentricity', ϵ). Then we use this relation to show how the relative sizes of the environments and its sides (Fig. S1) translate into the difference in biases.

First, we express $\hat{\ell}$, the estimated location, obtained with Bayesian inference after Laplace approximation of the prior and the visual likelihood, as a function of L and ϵ . Here, without loss of generality, we adopt a convention where $\hat{\ell} = 0$ means that the estimated location matches the mean prior location:

$$\hat{\ell} = f(\ell)\ell \quad 0 < \ell < L \quad (\text{S29})$$

$$\text{where } f(\ell) := \frac{1}{1 + g(\ell)} \quad 0 < f(\ell) < 1 \quad (\text{S30})$$

$$\text{and } g(\ell) := \frac{(L - \ell)^4}{\epsilon^2} \frac{\lambda}{L^2} \quad 0 < g(\ell) \quad (\text{S31})$$

Here, $L - \ell$ is the distance to the wall in front, $\frac{\epsilon^2}{(L - \ell)^4}$ is the precision of the likelihood (following Eq. S28), $\frac{\lambda}{L^2}$ is the precision of the prior (which is lower in a larger environment), where λ is a free parameter that indicates the relative informativeness of the prior compared to the likelihood.

Then, we express how much g depends on L vs. ϵ in preparation for a later comparison on how much the estimated location $\hat{\ell}$ depends on them:

$$\frac{\partial g}{\partial L} = 2 \frac{g(\ell)}{L - \ell} \left(2 - \frac{L - \ell}{L} \right) \quad (\text{S32})$$

$$\frac{\partial g}{\partial \epsilon} = -\frac{2}{\epsilon} g(\ell) \quad (\text{S33})$$

which means that, since $L - \ell \in [0, L]$ and $2 - \frac{L - \ell}{L} \in [0, 2]$, when $L = \epsilon$:

$$\left| \frac{\partial g}{\partial L} \right| > \left| \frac{\partial g}{\partial \epsilon} \right| \quad (\text{S34})$$

In other words, g changes more with L than with ϵ . Now we express the estimated distance \hat{d} between two target locations ℓ_1 and ℓ_2 as a function of f , and in turn, of g :

$$d = \ell_1 - \ell_2 \quad \ell_1 > \ell_2 \Rightarrow f(\ell_1) > f(\ell_2) \quad (\text{S35})$$

$$\hat{d} = \hat{\ell}_1 - \hat{\ell}_2 \quad (\text{S36})$$

$$= f(\ell_1)\ell_1 - f(\ell_2)\ell_2 \quad (\text{S37})$$

Then we compare how the change in L vs. ϵ affects \hat{d} using the relations derived above:

$$\frac{\partial \hat{d}}{\partial g} = -f^2(\ell_1)\ell_1 + f^2(\ell_2)\ell_2 \quad (\text{S38})$$

$$\frac{\partial \hat{d}}{\partial L} = \frac{\partial g}{\partial L} \frac{\partial \hat{d}}{\partial g} < 0 \quad (\text{S39})$$

$$\frac{\partial \hat{d}}{\partial \epsilon} = \frac{\partial g}{\partial \epsilon} \frac{\partial \hat{d}}{\partial g} > 0 \quad (\text{S40})$$

In words, the estimated distance between two target locations ℓ_1 and ℓ_2 is shorter when the environment is deeper (when the wall is farther and visual likelihood is less precise) due to stronger regression to mean, and longer when the environment is wider due to weaker regression to mean. From Eq. S34, the absolute values of the above two quantities can be compared as:

$$\left| \frac{\partial \hat{d}}{\partial L} \right| > \left| \frac{\partial \hat{d}}{\partial \epsilon} \right| \quad (\text{S41})$$

Therefore, the estimated distance relies more on the depth of the environment than on its width.

Finally, we use Eq. S41 to compare the estimated distances in square vs. trapezoid, and in the broad vs. narrow sides of the trapezoid, which we denote as \hat{S} (square), \hat{T} (trapezoid), \hat{B} (broad side of the trapezoid), and \hat{N} (narrow side of the trapezoid). We further denote the size of the environments and the halves of the trapezoid as X and Y as shown in Fig. S1.

Using these relations, we begin by comparing the estimated distance along the x axis between the square and trapezoid environments. Here, X_S and X_T are the depths. (To be exact, the prior mean x location of the trapezoid is off center, but we ignore it for simplicity and use X_T for L , since the following holds as long as $X_S \ll X_T$.) Since $X_S \ll X_T$, we can infer that $\hat{S} \gg \hat{T}$ (larger depth results in shorter estimated distance; Eq. S39). To disambiguate this comparison with others, we denote it as $\hat{S}_L^X \gg \hat{T}_L^X$, which means we used X_S and X_T as the depth (L) of the environment.

Likewise, we can compare the widths (ϵ) of the environments as $Y_S < Y_T$, from which we infer that $\hat{S}_\epsilon^Y > \hat{T}_\epsilon^Y$ (larger width results in longer estimated distance; Eq. S40) using the same convention as above.

We can also compare the estimated distance along the y axis. Then Y_S and Y_T would be the depths, and X_S and X_T the widths. We can summarize the comparisons along the x and y axes as:

$$\text{square vs. trapezoid} \begin{cases} \text{along x axis} \begin{cases} X_S \ll X_T \Rightarrow \hat{S}_L^X \gg \hat{T}_L^X \\ Y_S > Y_T \Rightarrow \hat{S}_\epsilon^Y > \hat{T}_\epsilon^Y \end{cases} \\ \text{along y axis} \begin{cases} Y_S > Y_T \Rightarrow \hat{S}_L^Y < \hat{T}_L^Y \\ X_S \ll X_T \Rightarrow \hat{S}_\epsilon^X \ll \hat{T}_\epsilon^X \end{cases} \end{cases} \quad (\text{S42})$$

Thus, the estimated distance in the square should be longer along the x axis, while it should be shorter along the y axis. To resolve this conflict, we compare the relative importance of the two axes. According to Eq. S41, the influence of the depth should dominate the influence of the width when $L = \epsilon$. This allows us to compare, e.g., $|\hat{S}_L^X - \hat{T}_L^X|$ and $|\hat{S}_\epsilon^X - \hat{T}_\epsilon^X|$, where the same set of sizes (X_S and X_T) were used. This, combined with the relative size difference of the two environments along the two axes means that:

$$\begin{aligned} |\hat{S}_L^X - \hat{T}_L^X| > |\hat{S}_\epsilon^X - \hat{T}_\epsilon^X| &\Rightarrow \hat{S} \gg \hat{T} \\ |\hat{S}_L^Y - \hat{T}_L^Y| > |\hat{S}_\epsilon^Y - \hat{T}_\epsilon^Y| &\Rightarrow \hat{S} < \hat{T} \end{aligned} \Rightarrow \hat{S} > \hat{T} \quad (\text{S43})$$




Fig. S1 | Symbols used to explain the analytic intuition for the bias in distance estimation. Schematics of the environments where their relative sizes roughly follow that of Ref. 5, where the distance estimation experiments are reported. **a.** Whole square. **b.** Whole trapezoid. **c.** Broad and narrow side of the trapezoid, split in half. Because the subject does not know whether they are in the broad or narrow half, their prior along the x axis covers the whole environment. Hence, we define $X_B := X_N := X_T$ so that any of them can substitute L from Eq. S29 to Eq. S41.

which agrees with the experimental results (Fig. 2d bottom, green vs. orange).

Likewise, we compare the broad vs. narrow sides of the trapezoid. Note that the two sides share the same size along the x axis (which is X_T) for the purpose of determining the precision of the prior and likelihood:

$$\text{broad vs. narrow side of the trapezoid} \left\{ \begin{array}{l} \text{along x axis} \left\{ \begin{array}{l} X_B = X_N = X_T \Rightarrow \hat{B}_L^X = \hat{N}_L^X \\ Y_B > Y_N \Rightarrow \hat{B}_\epsilon^Y > \hat{N}_\epsilon^Y \end{array} \right. \\ \text{along y axis} \left\{ \begin{array}{l} Y_B > Y_N \Rightarrow \hat{B}_L^Y < \hat{N}_L^Y \\ X_B = X_N = X_T \Rightarrow \hat{B}_\epsilon^X = \hat{N}_\epsilon^X \end{array} \right. \end{array} \right. \quad (\text{S44})$$

Again, we compare them with Eq. S41:

$$\left. \begin{array}{l} |\hat{B}_L^X - \hat{N}_L^X| = |\hat{B}_\epsilon^X - \hat{N}_\epsilon^X| = 0 \Rightarrow \hat{B} = \hat{N} \\ |\hat{B}_L^Y - \hat{N}_L^Y| > |\hat{B}_\epsilon^Y - \hat{N}_\epsilon^Y| \Rightarrow \hat{B} < \hat{N} \end{array} \right\} \Rightarrow \hat{B} < \hat{N} \quad (\text{S45})$$

which, again, agrees with the experimental results (Fig. 2d bottom, light vs. dark orange). Therefore, we can see that the difference in the estimated distance (Eqs. S43 and S45) can be explained by the effect of the size of the environments on the precision of the prior and likelihood in Bayesian inference.

Neural tethering – analytic intuition

Our goal is to show that, if the ideal observer's posterior is tethered to the last wall approached, then a neural posterior, if it is encoded by tuning curves tethered to the last wall approached, also matches the ideal observer's posterior. Since tuning curves define mean firing rates, we first consider the mean posterior \bar{P} , defined as the ideal observer posterior averaged over time:

$$\bar{P}(\tilde{\ell}|\ell, \theta, \mathcal{E}) := \left\langle P(\tilde{\ell}|\tilde{u}_{0:t}, \mathbf{v}_{0:t}, \mathcal{E}) \right\rangle_{t: (\ell_t^*, \theta_t^*) = (\ell, \theta)} \quad (\text{S46})$$

For simplicity, here we only consider the 1D location $\ell \in \mathcal{L} = [0, L]$ along the changed axis, and two opposite head directions along that axis, $\theta \in \Omega = \{0, \pi\}$, where $\theta = 0$ means the observer faces toward L . We also assume that the momentary posterior of the ideal observer matches its average posterior in either environment $\mathcal{E} \in \{\tilde{\mathcal{E}}, \mathcal{E}^*\}$, which allows us to parameterize the ideal observer posterior P^b with location, head direction, and environment, as follows:

$$\begin{aligned} P^b(\tilde{\ell}; \ell_t^*, \theta_t^*, \mathcal{E}) &:= \bar{P}(\tilde{\ell}|\ell_t^*, \theta_t^*, \mathcal{E}) \\ &= P(\tilde{\ell}|\tilde{u}_{0:t}, \mathbf{v}_{0:t}, \mathcal{E}) \quad \forall t \end{aligned} \quad (\text{S47})$$

Simplifying the approximate patterns seen in behavioral data, assume that the ideal observer's belief P^b is tethered to the last wall in the new environment \mathcal{E}^* , and that the last wall is always behind the observer (i.e., when $\theta = 0$, the last wall is at 0, and when $\theta = \pi$, the last wall is at L^{new}). For $\theta^* = 0$, this is simply:

$$\underbrace{P^b(\tilde{\ell}; \ell^*, \theta^* = 0, \mathcal{E}^*)}_{P^b_{\text{new}}(\theta^*=0)} \propto \underbrace{P^b(\tilde{\ell}; \ell^*, \theta^* = 0, \tilde{\mathcal{E}})}_{P^b_{\text{familiar}}(\theta^*=0)} \quad \forall \ell^* \in [0, L^{\text{new}}] \quad (\text{S48})$$

We also assume an exact match between the ideal observer and neural posteriors in $\tilde{\mathcal{E}}$:

$$\underbrace{P^b(\tilde{\ell}; \ell^*, \theta^*, \tilde{\mathcal{E}})}_{P^b_{\text{familiar}}} = \underbrace{P^g\left(\tilde{\ell} \middle| \underbrace{f(\ell^*, \theta^*, \tilde{\mathcal{E}})}_{f_{\text{familiar}}}\right)}_{P^g_{\text{familiar}}} \quad \forall (\ell^*, \theta^*) \quad (\text{S49})$$

Note that we are considering the limit of a large neural population size, which gives $r_t = f(\ell_t^*, \theta_t^*, \mathcal{E}) \forall t$, and the limit of good coverage of the environment by the observer's trajectory, which gives $\{(\ell_t^*, \theta_t^*)\} = \mathcal{L} \times \Theta$. Here, for simplicity, we assume $L^{\text{new}} < L^{\text{familiar}}$, i.e., that the observer is in a new environment \mathcal{E}^* compressed compared to the familiar environment $\tilde{\mathcal{E}}$, such that tethered beliefs can cover the whole length of \mathcal{E}^* .

Then, having tethered tuning curve, i.e.,

$$\underbrace{f(\ell^*, \theta^* = 0, \mathcal{E}^*)}_{f_{\text{new}}(\theta^*=0)} = \underbrace{f(\ell^*, \theta^* = 0, \tilde{\mathcal{E}})}_{f_{\text{familiar}}(\theta^*=0)} \quad \forall \ell^* \in [0, L^{\text{new}}] \quad (\text{S50})$$

which implies

$$\underbrace{P^g\left(\tilde{\ell} \middle| \underbrace{f(\ell^*, \theta^* = 0, \tilde{\mathcal{E}})}_{f_{\text{familiar}}(\theta^*=0)}\right)}_{P^g_{\text{familiar}}(\theta^*=0)} \propto \underbrace{P^g\left(\tilde{\ell} \middle| \underbrace{f(\ell^*, \theta^* = 0, \mathcal{E}^*)}_{f_{\text{new}}(\theta^*=0)}\right)}_{P^g_{\text{new}}(\theta^*=0)} \quad \forall \ell^* \in [0, L^{\text{new}}] \quad (\text{S51})$$

is, by transitivity across Eqs. S48, S49 and S51, sufficient to give

$$\underbrace{P^g\left(\tilde{\ell} \middle| \underbrace{f(\ell^*, \theta^* = 0, \mathcal{E}^*)}_{f_{\text{new}}(\theta^*=0)}\right)}_{P^g_{\text{new}}(\theta^*=0)} = \underbrace{P^b(\tilde{\ell}; \ell^*, \theta^* = 0, \mathcal{E}^*)}_{P^b_{\text{new}}(\theta^*=0)} \quad \forall \ell^* \in [0, L^{\text{new}}] \quad (\text{S52})$$

Also, in the limits of good coverage that we're assuming, it is also natural to assume mirror symmetry in $\tilde{\mathcal{E}}$, i.e.,

$$\underbrace{P^b(\tilde{\ell}; \ell^*, \theta^* = 0, \tilde{\mathcal{E}})}_{P^b_{\text{familiar}}(\theta^*=0)} = \underbrace{P^b(\tilde{\ell}; L^{\text{familiar}} - \ell^*, \theta^* = \pi, \tilde{\mathcal{E}})}_{P^b_{\text{familiar}}(\theta^*=\pi)} \quad (\text{S53})$$

and the same for the neural posterior and tuning curves. Then, by mirror symmetry, results analogous to Eq. S52 also hold when $\theta^* = \pi$. That is, tethering of the ideal observer's belief...

$$\underbrace{P^b(\tilde{\ell}; L^{\text{new}} - \ell^*, \theta^* = \pi, \mathcal{E}^*)}_{P^b_{\text{new}}(\theta^*=\pi)} \propto \underbrace{P^b(\tilde{\ell}; L^{\text{familiar}} - \ell^*, \theta^* = \pi, \tilde{\mathcal{E}})}_{P^b_{\text{familiar}}(\theta^*=\pi)} \quad \forall \ell^* \in [0, L^{\text{new}}] \quad (\text{S54})$$

... and tethering of tuning curves...

$$\underbrace{f(L^{\text{new}} - \ell^*, \theta^* = 0, \mathcal{E}^*)}_{f_{\text{new}}(\theta^*=0)} = \underbrace{f(L^{\text{familiar}} - \ell^*, \theta^* = 0, \tilde{\mathcal{E}})}_{f_{\text{familiar}}(\theta^*=0)} \quad \forall \ell^* \in [0, L^{\text{new}}] \quad (\text{S55})$$

... which imply...

$$\underbrace{P^g\left(\tilde{\ell} \middle| \underbrace{f(L^{\text{familiar}} - \ell^*, \theta^* = \pi, \tilde{\mathcal{E}})}_{f_{\text{familiar}}(\theta^*=\pi)}\right)}_{P^g_{\text{familiar}}(\theta^*=\pi)} \propto \underbrace{P^g\left(\tilde{\ell} \middle| \underbrace{f(L^{\text{new}} - \ell^*, \theta^* = \pi, \mathcal{E}^*)}_{f_{\text{new}}(\theta^*=\pi)}\right)}_{P^g_{\text{new}}(\theta^*=\pi)} \quad \forall \ell^* \in [0, L^{\text{new}}] \quad (\text{S56})$$

... are, by transitivity across Eqs. S49, S54 and S56, sufficient for the neural posterior to match the ideal observer's posterior given $\theta^* = \pi$:

$$\underbrace{P^g \left(\tilde{\ell} \mid \underbrace{f(L^{\text{new}} - \ell^*, \theta^* = 0, \mathcal{E}^*)}_{f_{\text{new}}(\theta^* = \pi)} \right)}_{P_{\text{new}}^g(\theta^* = \pi)} = \underbrace{P^b \left(\tilde{\ell}; L^{\text{new}} - \ell^*, \theta^* = \pi, \mathcal{E}^* \right)}_{P_{\text{new}}^b(\theta^* = \pi)} \quad \forall \ell^* \in [0, L^{\text{new}}] \quad (\text{S57})$$

Therefore, tuning curves tethered to the last wall in \mathcal{E}^* is sufficient for the neural likelihood to be tethered the same way, and when the likelihood is sharp enough, for the neural posterior to match the ideal observer posterior tethered to the last wall approached in either direction, as in Eqs. S52 and S57.