

Geometry of abstract learned knowledge in the hippocampus

<https://doi.org/10.1038/s41586-021-03652-7>

Received: 5 February 2020

Accepted: 18 May 2021

Published online: 16 June 2021

 Check for updates

Edward H. Nieh^{1,6}, Manuel Schottdorf^{1,6}, Nicolas W. Freeman¹, Ryan J. Low¹, Sam Lewallen¹, Sue Ann Koay¹, Lucas Pinto^{1,5}, Jeffrey L. Gauthier¹, Carlos D. Brody^{1,2,3,7}✉ & David W. Tank^{1,2,4,7}✉

Hippocampal neurons encode physical variables^{1–7} such as space¹ or auditory frequency⁶ in cognitive maps⁸. In addition, functional magnetic resonance imaging studies in humans have shown that the hippocampus can also encode more abstract, learned variables^{9–11}. However, their integration into existing neural representations of physical variables^{12,13} is unknown. Here, using two-photon calcium imaging, we show that individual neurons in the dorsal hippocampus jointly encode accumulated evidence with spatial position in mice performing a decision-making task in virtual reality^{14–16}. Nonlinear dimensionality reduction¹³ showed that population activity was well-described by approximately four to six latent variables, which suggests that neural activity is constrained to a low-dimensional manifold. Within this low-dimensional space, both physical and abstract variables were jointly mapped in an orderly manner, creating a geometric representation that we show is similar across mice. The existence of conjoined cognitive maps suggests that the hippocampus performs a general computation—the creation of task-specific low-dimensional manifolds that contain a geometric representation of learned knowledge.

Since the discovery of place cells in the cornu ammonis 1 (CA1) of the dorsal hippocampus that increased their firing rates when rats moved through specific locations in a given environment¹, hippocampal neurons have also been shown to encode time^{17,18}, auditory frequency⁶, odours^{4,7} and taste⁵. Together, these studies support the view that the hippocampus constructs task-dependent cognitive maps^{8,19}, in which hippocampal neurons not only encode spatial position, but whichever environmental variable is relevant to the task at hand. Furthermore, functional MRI studies in humans have shown that the hippocampus can encode more cognitive variables, such as the sequential nature of a non-spatial task⁹ or social structures^{10,11}. Cognitive variables can be characterized by geometric properties such as adjacency and distance^{20–22}, suggesting that the neural encoding of these variables at the cellular level may also have geometric structure.

Neural activity can be described as a point in a high-dimensional coordinate system, in which each coordinate axis represents the activity of a single neuron. The underlying properties of the network and its inputs can confine neural trajectories to a subregion of this space—the neural manifold—which has been proposed to underlie motor movements^{23,24}, head direction cells²⁵ and hippocampal maps of physical variables¹³. The conceptual ideas in these studies suggest a general principle of hippocampal computation: the construction of organized maps of learned knowledge^{26,27} instantiated by neural manifolds. Here we examine how neurons in the dorsal CA1 integrate neural representations of cognitive and physical variables and whether low-dimensional manifolds underlie these representations.

Evidence accumulation in virtual reality

We used transgenic GCaMP6f-expressing mice ($n = 7$) performing an evidence-accumulation task in virtual reality^{14,28–30} (Supplementary Video 1) and two-photon calcium imaging to record the activity of neurons in dorsal CA1 at cellular resolution ($n = 3,144$ total neurons; 449 ± 64 neurons (mean \pm s.e.m.) simultaneously recorded per session) (Fig. 1a). The ‘accumulating towers’ task¹⁴ combines navigation with decision-making, such that position—a physical variable—has to be integrated with accumulated evidence^{14–16,31,32}—a cognitive variable that is not innate and can be inferred and calculated only after learning the task rules. Mice learned to traverse the stem of an immersive virtual-reality T-maze, while visual cues were presented randomly on the left and right walls. Turning to the side with more cues at the end of the maze resulted in the delivery of a liquid reward, while turning to the opposite side resulted in a time-out. Consistent with previously published results¹⁴, the behaviour showed characteristic psychometric curves (Fig. 1b) and mice used evidence (the number of right towers minus the number of left towers) from throughout the cue period (Fig. 1c).

Figure 1d, e illustrates two possibilities for how CA1 neurons may behave in the task. If the neurons behave similarly to previously described place cells that respond differently depending on context^{2,3,33,34}—for example, ‘splitter cells’ that encode turn direction—we would expect reliable place cell sequences that are specific to right- or left-turn trials (Fig. 1d). However, if individual CA1 neurons can encode a cognitive variable, such as the amount of accumulated evidence, in addition to the position in the maze, the cognitive map would comprise at least two independent axes—a position axis and an accumulated

¹Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA. ²Department of Molecular Biology, Princeton University, Princeton, NJ, USA. ³Howard Hughes Medical Institute, Princeton University, Princeton, NJ, USA. ⁴Bezos Center for Neural Dynamics, Princeton University, Princeton, NJ, USA. ⁵Present address: Department of Physiology, Feinberg School of Medicine, Northwestern University, Chicago, IL, USA. ⁶These authors contributed equally: Edward H. Nieh, Manuel Schottdorf. ⁷These authors jointly supervised this work: Carlos D. Brody, David W. Tank. ✉e-mail: brody@princeton.edu; dwtank@princeton.edu

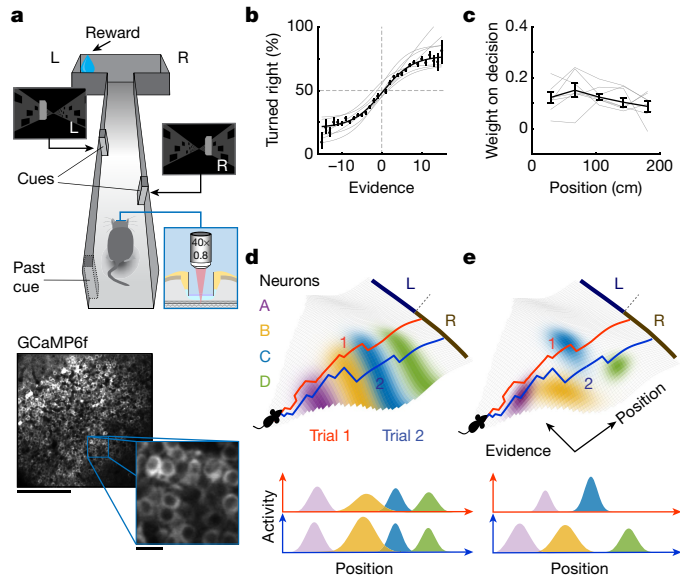


Fig. 1 | Imaging of CA1 neural activity in mice performing the accumulating towers task. **a**, Top, schematic of the task in which head-fixed mice navigate a virtual-reality evidence-accumulation T-maze task. Insets show example views from the perspective of the mice. Bottom, while mice ($n = 7$) perform the task, two-photon calcium imaging is used to record hippocampal CA1 neural activity. Scale bars, 200 μm (main image) and 20 μm (magnification). **b**, Psychometric curves of mice performing the towers task. Grey lines, $n = 7$ mice; black line, metamouse combining data across mice. Data are mean \pm binomial confidence interval. **c**, Logistic regression showing that mice use evidence (number of right towers minus number of left towers) from throughout the cue period. Grey lines, $n = 7$ mice; black line, metamouse combining data across mice. Data are mean \pm s.e.m. **d**, Firing fields of right-choice-selective place cells would not depend on evidence and would therefore divide a joint evidence-by-position ($E \times Y$) space into two halves (top). Two right-choice trials would generate the same neural sequence (bottom). **e**, Alternatively, if hippocampal neurons encoded evidence jointly with position, smaller firing fields that divide the evidence dimension would appear in $E \times Y$ space (top), and two right-choice trials could have different neural sequences that depend on the evidence values traversed (bottom).

evidence axis³⁵. If so, we would expect each right-choice trial to evoke different neural sequences, depending on the time courses of evidence that the mice encountered throughout the maze (Fig. 1e). Notably, in this second scenario, firing fields evaluated in a single dimension—such as position—would exist, but would appear unreliable across trials with different amounts of accumulated evidence (Fig. 1e, bottom). Unreliability could appear as either missing activity in the place field of the cell or variability in the position at which the cell is active.

Joint encoding of position and evidence

To distinguish between these two possibilities, we examined how neural activity depended on known behavioural variables such as position, choice and evidence. We first calculated $\Delta F/F$ for each identified hippocampal CA1 neuron following established methods^{15,36,37}. We then measured the mutual information between the neural activity of each cell and the position of the mouse along the stem of the T-maze (0 to 300 cm) and compared it to a shuffled dataset in which the activity of each cell was circularly shifted within each trial. CA1 neurons exhibited choice-specific place cell sequences when activity was sorted by the position of peak activity (Fig. 2a). However, the response of individual cells in these sequences was more variable and unreliable on a trial-by-trial basis in comparison to a simpler alternation task (Extended Data Fig. 1a, b). This is against the prediction of choice-specific cell maps

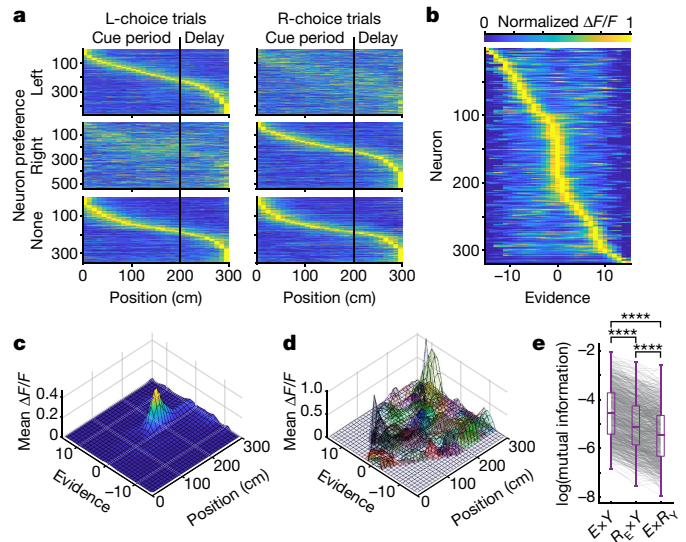


Fig. 2 | CA1 neurons jointly encode the position of the mice and accumulated evidence in an evidence-accumulation task. **a**, Choice-specific place cell sequences, divided into left-choice-preferring (top), right-choice-preferring (middle) and non-preferring (bottom) cells. Cells are shown in the same order within each row group. $\Delta F/F$ was normalized within each neuron. **b**, CA1 neurons have firing fields in accumulated evidence space (number of right towers minus number of left towers). **c**, Example of the average neural activity of a single neuron in joint evidence-by-position ($E \times Y$) space. **d**, Twenty-five neurons with significant information in $E \times Y$ space. Each colour represents one cell, and surfaces represent neural activity that exceeds 2σ above the shuffled means (Extended Data Fig. 2a). **e**, Mutual information of cells found to have significant information in $E \times Y$ space is significantly greater than mutual information in two-dimensional spaces in which either evidence (R_E) or position (R_Y) has been randomized. Two-tailed paired Student's t -tests, Bonferroni correction, $n = 917$ neurons; $E \times Y$ versus $R_E \times Y$, **** $P < 0.0001$; $E \times Y$ versus $E \times R_Y$, **** $P < 0.0001$; $R_E \times Y$ versus $E \times R_Y$, **** $P < 0.0001$. For box plots, boundaries, 25–75th percentiles; midline, median; whiskers, minimum–maximum.

(Fig. 1d), but is consistent with maps in which evidence and position are jointly encoded (Fig. 1e). We next measured the mutual information between accumulated evidence and the neural activity of each cell and found that CA1 neurons formed firing fields in evidence space that spanned small segments of evidence values (Fig. 2b and Extended Data Fig. 1c), consistent with Fig. 1e.

To directly test the hypothesis that CA1 neurons encode evidence and position jointly (Fig. 1e), we measured the amount of mutual information between neural activity and occupancy in a two-dimensional evidence-by-position ($E \times Y$) space and compared this to the amount of mutual information if cells encoded position or evidence independently. The neural activity of an example neuron with significant mutual information between activity and occupancy in $E \times Y$ space is shown in Fig. 2c, and 25 of these neurons from a single imaging session are shown in Fig. 2d and Extended Data Fig. 2a. For these neurons that jointly encode position and evidence, mutual information in $E \times Y$ space was significantly greater than in two-dimensional spaces in which either evidence or position values were shuffled (Fig. 2e and Extended Data Fig. 2b, c).

Geometric representation by a neural manifold

Although the mutual information metric has historically been used to measure spatial information in single hippocampal neurons³⁸, it relies on the manual selection of predetermined behavioural variables. We therefore turned to the unsupervised extraction of neural manifolds

using a principled method: manifold inference from neural dynamics (MIND)¹³. Whereas most nonlinear dimensionality reduction techniques focus on the geometric properties of the cloud of population-state data, MIND constructs a set of latent variables with a specific emphasis on incorporating temporal dynamics. It is therefore particularly suited to find low-dimensional representations in data with sequential activity.

We first used the distance metric in MIND to estimate the dimensionality of the neural manifold in the hippocampus during the accumulating towers task. We calculated distances from estimated transition probabilities between observed population activity states and counted the cumulative number of population states that fell within spheres of growing radii r , where r is an estimate for the inner distance^{23,39}. If the manifold has d dimensions, we expect the number of states to grow as r^d . We found that the number of states grows approximately as $d = 5.4$ (4.8–6.0; 95% bootstrapped confidence interval), strongly indicating a low, approximately four-to-six-dimensional latent geometry (Fig. 3a). Notably, the dimensionality estimate of a simpler task, in which visual cues appeared only on one side of the maze, was significantly lower (Extended Data Fig. 3).

To validate this estimate, we next embedded the manifold into d -dimensional Euclidean spaces and assessed how well these embedded manifolds described neural data using cross-validation on held-out trials (Extended Data Fig. 4). Figure 3b shows a small portion of the activity from 40 neurons and the reconstruction of that same data from the five latent variables obtained after embedding the manifold into a five-dimensional Euclidean space. We measured the average cross-validated correlation coefficient between the neural data and the reconstruction of the same data from manifolds embedded into two- to seven-dimensional Euclidean spaces. Consistent with the dimensionality estimate in Fig. 3a, we find that the reconstruction performance saturates at around five to six dimensions (Fig. 3c). Using a linear dimensionality reduction technique—principal component analysis—comparable decoding indices for embedding into 4, 5 and 6 dimensions are reached using 29, 40 and 47 principal components, respectively. This reveals that hippocampal activity is constrained to an intricately shaped low-dimensional manifold that can only be identified with nonlinear dimensionality reduction techniques.

If the neural manifold accurately represents the cognitive map of the task that individual neurons encode, two key predictions should hold true. First, individual neurons should have firing fields that tile the latent space and, second, important variables in the task—such as position and evidence—should be organized in an orderly manner. The activity of a representative neuron plotted as a heat map on a three-dimensional embedding of the manifold is shown in Fig. 3d, demonstrating a localized firing field on the manifold. Plotting the activity of multiple neurons on the same manifold reveals that the manifold is tiled with multiple firing fields (Fig. 3e and Supplementary Video 2). Furthermore, the manifold structure implies the coordinated activity of the entire neural population, such that activity of a single neuron can be well-predicted by activity from the rest of the population (Extended Data Fig. 4).

The second key prediction of our hypothesis is the orderly organization of important task variables on the manifold. Figure 3f reveals that both position and evidence appear organized as gradients in the latent space, in that the trajectory of the neural state typically progresses along a position direction in the course of a trial, while splitting along an independent, but integrated, evidence direction (Supplementary Video 3)—a structure that is fundamentally different from the visual inputs that the mouse experiences in the towers task (Extended Data Fig. 5). We then used Gaussian process regression to decode position and evidence from the manifold and found that both variables can be decoded with similar accuracy as from neural data (Fig. 3g). In addition, other behavioural variables such as velocity and view angle could also be decoded from the manifold, as well as binary task variables such as the choice on the previous trial and

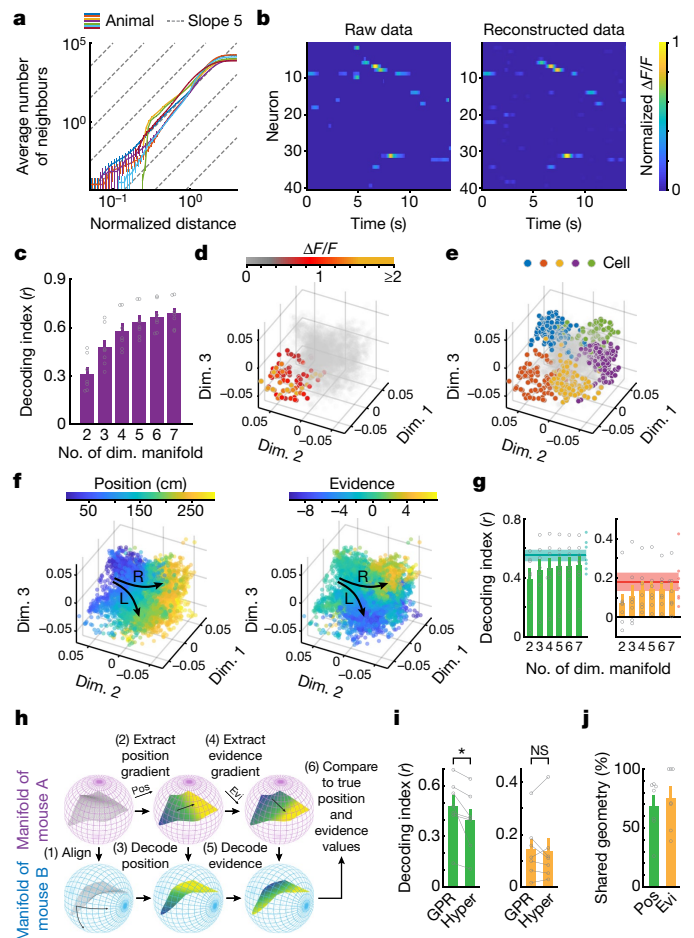


Fig. 3 | Geometric representation of task variables on low-dimensional neural manifolds. **a**, The mean cumulative number of neighbouring neural states as a function of the geodesic distance. $n = 7$ mice; data are mean \pm 95% bootstrapped confidence intervals. Grey dashed lines illustrate the slope expected for a five-dimensional manifold. **b**, Examples of raw (left) and manifold-reconstructed (right) neural activity from 40 neurons. $\Delta F/F$ is normalized to the maximum $\Delta F/F$ in the window shown. **c**, Reconstruction of held-out neural data from d -dimensional (dim.) embeddings of the neural manifold. The decoding index is the correlation coefficient between the predicted and true $\Delta F/F$ data in held-out trials. **d**, Each point is a location in the three-dimensional embedding of the manifold at one time point. Coloured points represent $\Delta F/F$ values that are 3σ above the mean activity for one example cell. **e**, Firing field of five cells plotted on the manifold. **f**, Position (left) and evidence (right) plotted on the three-dimensional embedding of the manifold. Black arrows represent two hypothetical trajectories that would traverse through position and evidence space. **g**, Decoding position (left) and evidence (right) from d -dimensional embeddings of the manifold using Gaussian process regression (GPR). The decoding index is the correlation coefficient between the predicted and true position or evidence values. The shaded area and line represent the mean decoding index \pm s.e.m. using GPR on the top 10% of neurons with the highest mutual information for position or evidence. **h**, Schematic of the hyperalignment (hyper) procedure (Methods). **i**, Decoding index of position (left) and evidence (right) using the best hyperaligned manifold from another mouse versus decoding with GPR using the manifold of the mouse. Five-dimensional embedding; $n = 7$ mice; two-tailed Wilcoxon signed-rank test; position, $*P = 0.016$; evidence, $P = 0.81$. NS, not significant. **j**, The majority of manifold geometry is shared between the best pairs of mice. **c, g, i, j**, Data are mean \pm s.e.m.; $n = 7$ mice.

whether the previous trial was correct (Extended Data Fig. 6 and Supplementary Discussion).

If these geometric objects are task-specific, rather than mouse-specific, there should be a high degree of similarity across mice

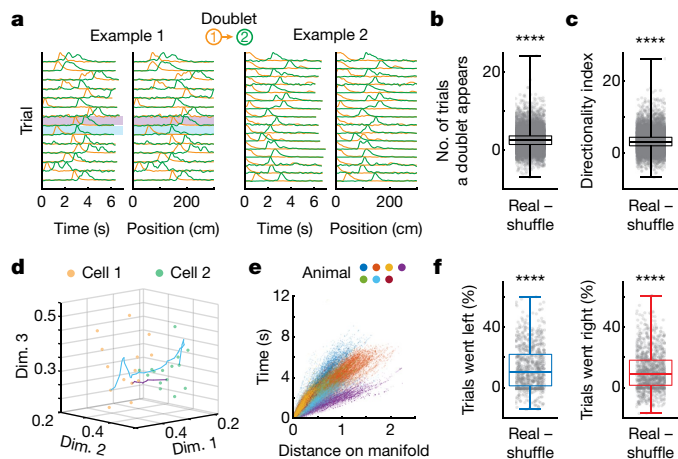


Fig. 4 | Sequential activity of CA1 neurons in single trials is predictive of behaviour and explained by the manifold. **a**, Two examples of doublets, in which two neurons consistently fire one after the other. In example 1, activity does not appear to be tied to time (left) or position (right). Highlighted trials in cyan and purple are the same trials plotted in **d**. In example 2, activity in both neurons appears to be related to time and/or position in the trial. **b**, Doublets appear more frequently in real data than in a shuffled dataset. $n = 16,088$ doublets; two-tailed paired Student's *t*-test, real versus shuffled data, **** $P < 0.0001$. **c**, Doublets are asymmetric. The directionality index is defined as the number of times cell 1 fires before cell 2 in a trial minus the number of times cell 2 fires before cell 1 in a trial. $n = 16,088$ doublets; two-tailed paired Student's *t*-test; real versus shuffled data, **** $P < 0.0001$. **d**, Example showing how events from cell 1 (orange) and cell 2 (green) of a doublet are separated in manifold space. Cyan and purple lines each represent a trial trajectory between when cell 1 and cell 2 fire. **e**, The amount of time between when cell 1 and cell 2 fire plotted against the distance in manifold space. **f**, Left- and right-choice-predictive doublets (blue and red box plots, respectively) are significantly more predictive of upcoming choice than doublets generated from shuffled data in which trial IDs were shuffled. Left-predictive, $n = 922$ doublets, two-tailed paired Student's *t*-test, real versus shuffled data, **** $P < 0.0001$; right-predictive, $n = 1,227$ doublets, two-tailed paired Student's *t*-test, real versus shuffled data, **** $P < 0.0001$. For box plots, boundaries, 25–75th percentiles; midline, median; whiskers, minimum–maximum.

performing the same task¹². To test this hypothesis, we trained a model to predict position and evidence from manifold coordinates in one mouse and used the model to decode these variables in another mouse, after aligning their manifolds in the five-dimensional embedding space (Fig. 3h, i). We found that the majority of the geometric structure was shared across mice (Fig. 3j).

Sequential neural activity encodes behaviour

If the manifold is a good representation of hippocampal neural activity, then each trial in the accumulating towers task has a corresponding trajectory within the manifold, leading to the emergence of trial-specific sequences of active cells. To detect sequences, we identified pairs of cells that consistently fired one after the other without any restrictions on the time and/or place in the maze that each cell fired (Extended Data Fig. 7a) and termed a pair of cells as a 'doublet' if one cell fired after the other significantly more often than in a shuffled dataset, in which the activity of each neuron was circularly shifted within each trial (Fig. 4a and Extended Data Fig. 7b). To test whether these doublets appear more often than would be expected from independently behaving choice- and position-selective cells, we shuffled the trial IDs of each cell independently within left- and right-choice trials to remove pairwise correlations while preserving the place and side structure seen in Fig. 2a (Extended Data Fig. 8). The number of trials in which

doublets appeared was significantly greater than in the shuffled dataset (Fig. 4b and Extended Data Fig. 9a). Furthermore, given the mostly unidirectional trajectories of the task in conceptual $E \times Y$ space (Extended Data Fig. 8a, b), we found that doublets were asymmetric (Fig. 4c and Extended Data Fig. 9b).

Next, we used the latent dimensions from the five-dimensional embedding of the manifold to reconstruct the neural activity of all cells and extracted doublets from this reconstructed data. Even though doublets are very rare (on average, a given doublet is only active in $3.6 \pm 0.01\%$ (mean \pm s.e.m.) of trials; $n = 16,088$ doublets), the manifold predicted the presence of doublets with a 0.87 ± 0.02 true-positive rate and 0.14 ± 0.01 false-positive rate ($n = 7$ mice; mean \pm s.e.m.) (Extended Data Fig. 9d). Furthermore, we found that the manifold could also predict the precise timing of doublet events—the correlation between the timing of a doublet and the distance traversed on the manifold was significantly greater than the correlation in a shuffled dataset in which manifold path lengths were taken from different trials with the same time interval (Fig. 4d, e) (two-tailed Wilcoxon signed-rank test, $n = 7$ mice, $P = 0.031$).

As the manifold encodes sequential activity well and given that behavioural variables are geometrically represented on the manifold (Fig. 3f, g), we would expect sequences to encode information about the behaviour of the mouse, specifically the upcoming choice. First, we identified doublets that were significantly choice-predictive by comparing the probability that the mouse turns left or right in trials in which a doublet appears to the same probability in a shuffled dataset in which choices in each trial were shuffled. Next, we found that these choice-predictive doublets were significantly more predictive than the same doublets drawn from the shuffled dataset in which trial IDs were shuffled (Fig. 4f and Extended Data Fig. 9c). Taken together, these sequences are informative beyond independently behaving cells, suggesting population activity that is consistent with movement along the low-dimensional manifold.

Discussion

By combining large-scale calcium imaging with a behavioural task in which animals accumulate abstract evidence during navigation, we show how the coordinated activity of neurons in the dorsal CA1 region of the hippocampus gives rise to a task-specific geometric representation of a cognitive process. The neural population manifests this geometric representation by having firing fields within a low-dimensional nonlinear manifold, along which key task variables—both continuous and discrete—have an orderly arrangement. Previous rodent studies have shown the existence of low-dimensional manifolds in the hippocampus representing spatial position^{12,13}, and functional MRI studies in humans have shown that more abstract variables, such as social structures^{10,11}, can be decoded from the hippocampus. One possibility was that different sets of hippocampal neurons could have encoded these variables separately, similar to the specialized coding of sensory, motor and cognitive variables by dopamine neurons in the ventral tegmental area in the same task¹⁶. However, we found that the majority of task-responsive neurons encoded position and evidence jointly (Fig. 2), leading to population dynamics that also reflect this joint neural code (Figs. 3, 4).

The formation of a conjoined geometric representation of physical and abstract task variables, within neural manifolds in the hippocampus, could serve as a common organizing principle across two roles of the hippocampus—storing declarative memory and generating cognitive maps—that have historically been studied separately^{21,26,40}. Low-dimensional manifolds could serve as the substrate on which relational networks for both declarative and spatial memories are stored²⁷. In addition, our work suggests that the fast replay sequences seen in human non-spatial tasks⁹ could be organized by the geometric structure of the neural manifold, analogous to the process by which neural

sequences during ongoing behaviour are evoked from trajectories through the manifold (Fig. 4). Finally, recent computational work has focused on how representations of knowledge in a reinforcement learning⁴⁰ or predictive coding²⁷ context can be used to guide behaviour. There are intriguing parallels between the latent structure identified in these models and the latent variable structure we have uncovered in our studies. However, future work is required to provide a quantitative understanding of how our experimental results relate to these learning models.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-021-03652-7>.

- O'Keefe, J. & Dostrovsky, J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* **34**, 171–175 (1971).
- Frank, L. M., Brown, E. N. & Wilson, M. Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron* **27**, 169–178 (2000).
- Wood, E. R., Dudchenko, P. A., Robitsek, R. J. & Eichenbaum, H. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron* **27**, 623–633 (2000).
- Eichenbaum, H., Kuperstein, M., Fagan, A. & Nagode, J. Cue-sampling and goal-approach correlates of hippocampal unit activity in rats performing an odor-discrimination task. *J. Neurosci.* **7**, 716–732 (1987).
- Herzog, L. E. et al. Interaction of taste and place coding in the hippocampus. *J. Neurosci.* **39**, 3057–3069 (2019).
- Aronov, D., Nevers, R. & Tank, D. W. Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature* **543**, 719–722 (2017).
- Taxidis, J. et al. Differential emergence and stability of sensory and temporal representations in context-specific hippocampal sequences. *Neuron* **108**, 984–998.e9 (2020).
- O'Keefe, J. & Nadel, L. *The Hippocampus as a Cognitive Map* (Clarendon, 1978).
- Schuck, N. W. & Niv, Y. Sequential replay of nonspatial task states in the human hippocampus. *Science* **364**, eaaw5181 (2019).
- Tavares, R. M. et al. A map for social navigation in the human brain. *Neuron* **87**, 231–243 (2015).
- Park, S. A., Miller, D. S., Nili, H., Ranganath, C. & Boorman, E. D. Map making: constructing, combining, and inferring on abstract cognitive maps. *Neuron* **107**, 1226–1238.e8 (2020).
- Rubin, A. et al. Revealing neural correlates of behavior without behavioral measurements. *Nat. Commun.* **10**, 4745 (2019).
- Low, R. J., Lewallen, S., Aronov, D., Nevers, R. & Tank, D. W. Probing variability in a cognitive map using manifold inference from neural dynamics. Preprint at <https://doi.org/10.1101/418939> (2018).
- Pinto, L. et al. An accumulation-of-evidence task using visual pulses for mice navigating in virtual reality. *Front. Behav. Neurosci.* **12**, 36 (2018).
- Koay, S. A., Thiberge, S., Brody, C. D. & Tank, D. W. Amplitude modulations of cortical sensory responses in pulsatile evidence accumulation. *eLife* **9**, e60628 (2020).
- Engelhard, B. et al. Specialized coding of sensory, motor and cognitive variables in VTA dopamine neurons. *Nature* **570**, 509–513 (2019).
- MacDonald, C. J., Lepage, K. Q., Eden, U. T. & Eichenbaum, H. Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron* **71**, 737–749 (2011).
- Pastalkova, E., Itskov, V., Amarasingham, A. & Buzsáki, G. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321**, 1322–1327 (2008).
- Tolman, E. C. Cognitive maps in rats and men. *Psychol. Rev.* **55**, 189–208 (1948).
- Bellmund, J. L. S., Gärdenfors, P., Moser, E. I. & Doeller, C. F. Navigating cognition: spatial codes for human thinking. *Science* **362**, eaat6766 (2018).
- Eichenbaum, H. in *Behavioral Neuroscience of Learning and Memory* (eds Clark, R. E. & Martin, S. J.) 101–117 (Springer, 2018).
- Constantinescu, A. O., O'Reilly, J. X. & Behrens, T. E. J. Organizing conceptual knowledge in humans with a gridlike code. *Science* **352**, 1464–1468 (2016).
- Gallego, J. A., Perich, M. G., Miller, L. E. & Solla, S. A. Neural manifolds for the control of movement. *Neuron* **94**, 978–984 (2017).
- Russo, A. A. et al. Motor cortex embeds muscle-like commands in an untangled population response. *Neuron* **97**, 953–966.e8 (2018).
- Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A. & Fiete, I. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nat. Neurosci.* **22**, 1512–1520 (2019).
- Eichenbaum, H. & Cohen, N. J. Can we reconcile the declarative memory and spatial navigation views on hippocampal function? *Neuron* **83**, 764–770 (2014).
- Recanatesi, S. et al. Predictive learning as a network mechanism for extracting low-dimensional latent space representations. *Nat. Commun.* **12**, 1417 (2021).
- Aronov, D. & Tank, D. W. Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system. *Neuron* **84**, 442–456 (2014).
- Dombeck, D. A., Harvey, C. D., Tian, L., Looger, L. L. & Tank, D. W. Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. *Nat. Neurosci.* **13**, 1433–1440 (2010).
- Harvey, C. D., Coen, P. & Tank, D. W. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**, 62–68 (2012).
- Brunton, B. W., Botvinick, M. M. & Brody, C. D. Rats and humans can optimally accumulate evidence for decision-making. *Science* **340**, 95–98 (2013).
- Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
- Gill, P. R., Mizumori, S. J. Y. & Smith, D. M. Hippocampal episode fields develop with learning. *Hippocampus* **21**, 1240–1249 (2011).
- McKenzie, S. et al. Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. *Neuron* **83**, 202–215 (2014).
- Howard, M. W., Luzardo, A. & Tiganj, Z. Evidence accumulation in a Laplace domain decision space. *Comput. Brain Behav.* **1**, 237–251 (2018).
- Pnevmatikakis, E. A. et al. Simultaneous denoising, deconvolution, and demixing of calcium imaging data. *Neuron* **89**, 285–299 (2016).
- Pnevmatikakis, E. A. & Giovannucci, A. NoRMCorre: an online algorithm for piecewise rigid motion correction of calcium imaging data. *J. Neurosci. Methods* **291**, 83–94 (2017).
- Skaggs, W. E., McNaughton, B. L. & Gothard, K. M. An information-theoretic approach to deciphering the hippocampal code. In *Advances in Neural Information Processing Systems 5* (eds Hanson, S. J. et al.) 1030–1037 (Morgan-Kaufmann, 1993).
- Grassberger, P. & Procaccia, I. Measuring the strangeness of strange attractors. *Physica D* **9**, 189–208 (1983).
- Stachenfeld, K. L., Botvinick, M. M. & Gershman, S. J. The hippocampus as a predictive map. *Nat. Neurosci.* **20**, 1643–1653 (2017).

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

© The Author(s), under exclusive licence to Springer Nature Limited 2021

Methods

Animals and stereotaxic surgery

All procedures performed in this study were approved by the Institutional Animal Care and Use Committee at Princeton University and were performed in accordance with the Guide for the Care and Use of Laboratory Animals (National Research Council, 2011). Male and female mice aged 2–18 months expressing GCaMP6f were used for chronic expression of the calcium indicator.

$n = 5$ triple transgenic mice expressing GCaMP6f under the CaMKII α promoter were obtained by crossing *Ai93-D;CaMKII α -tTA* mice (*Igsl^{tm93.1(tetO-GCaMP6f)HzeTg}(Camk2a-tTA)1Mmay/J*), Jackson Laboratories, 024108) and *Emx1-IRES-cre* mice (B6.129S2-*Emx1^{tm1(cre)Krl}/J*), Jackson Laboratories, 005628). These mice are also referred to as Ai93 \times EMX1.

$n = 10$ Thy1-GCaMP6f mice (C57BL/6J-Tg(Thy1-GCaMP6f)GP5.3Dkim/J), Jackson Laboratories, 028280) were used. These mice are also referred to as GP5.3.

Behaviourally, no differences have been observed in Ai93 \times EMX1 and GP5.3 mice¹⁴. In terms of calcium imaging, Ai93 \times EMX1 mice have higher expression levels of GCaMP6f than GP5.3 mice and therefore higher signal-to-noise ratios, which results in the use of different activity thresholds to identify calcium events (described in 'Identification and analysis of sequences'). Some mice were used in multiple behavioural experiments that we analysed. For example, the one-side cues task is a training stage in the shaping procedure for the accumulating towers task. For analyses and statistics other than those in Fig. 1b, c (described in 'Psychometric curves' and 'Logistic regression analysis'), one imaging session for each mouse was selected based on the performance of the mouse in the task during the session, the number of active cells identified by the automated cell-finding algorithm, the amount of noise in the $\Delta F/F$ signal and the quality of motion correction.

Mice underwent surgical procedures as previously described^{29,41} to acquire optical access to the hippocampus. Surgery was performed on mice under aseptic conditions and body temperature was maintained with a heating pad (Harvard Apparatus). Mice were anaesthetized with isoflurane (2.5% for induction, 1–1.5% for maintenance) and given a pre-operative dose of meloxicam subcutaneously for analgesia (1 mg kg⁻¹) and a postoperative dose 24 h later. After asepsis, the skull was exposed, and the periosteum was removed.

A custom lightweight titanium headplate was attached to the skull with adhesive cement (C&B Metabond, Parkell). A craniotomy in the left hemisphere centred over the CA1 (mediolateral, -1.8 mm from the midline; anteroposterior, 2.0 mm posterior from bregma) was made using a pneumatic drill. A small volume of overlying cortical tissue was aspirated to expose the external capsule; superficial fibres were then removed until the alveus became visible. A thin layer of Kwik-Sil (WPI) was injected into the resected area, and a metal cannula (316 S/S Hypo Tub 12T GA. 0.1080/0.1100 inches outer diameter \times 0.0890/.0930 inches inner diameter \times 0.060 inches long; cut and deburred) with a coverglass (2.5 mm diameter, Erie Scientific) attached to the bottom (NOA81 adhesive, Norland) was implanted on top of the Kwik-Sil, so that the Kwik-Sil served as a stabilizing medium between the glass and brain tissue. Another layer of adhesive cement was added to attach the cannula to the skull and the headplate. Mice were allowed to recover for at least 5 days before starting water restriction for behavioural training. Mice were extensively handled during the restriction process to familiarize them to experimenters. Mice were allotted daily volumes of 1–2 ml of liquid per day, delivered either during behavioural sessions or supplemented after sessions. Mice were examined daily to ensure that there were no signs of dehydration and that a body mass of at least 80% of the initial value was maintained.

Behavioural training

The mice were trained to perform the accumulating towers task in a virtual-reality environment, as previously described^{14–16,42}. In brief,

mice were headfixed so that they could run comfortably on an 8-inch (20 cm) Styrofoam ball suspended by compressed air. Ball movements were measured with optical flow sensors (ADNS3080) via an Arduino Due, and the virtual-reality environment was projected onto a coated Styrofoam screen (approximately 270° horizontal and 80° vertical visual field) using a DLP projector (Mitsubishi). The virtual environment was generated using ViRMEn software²⁸. Rewards were delivered by a solenoid valve (NResearch), controlled by a NI-DAQ card (PCI-6229, National Instruments). This virtual-reality system has previously been used^{14–16,42} and was designed, by choice of material and size of the spherical treadmill, to minimize the amount of effort to turn the floating ball, such that the moment of inertia of a mouse pushing back the ball (2.78×10^{-4} kg m²) is comparable to the moment of inertia of a mouse pushing itself (2.68×10^{-4} kg m²).

Mice were trained to run down a 330-cm virtual T-maze (30-cm start region, 200-cm cue region and 100-cm delay region). As mice ran through the cue region, tall, high-contrast visual cues (towers, 6 cm tall and 2 cm wide) were shown along either wall. After the delay period, mice were presented with a liquid reward for turning into the arm on the side where more towers had been shown (4–8 μ l of 10% v/v sweet condensed milk or 10% w/v sucrose). Rewarded trials were followed by a 3-s inter-trial interval, and error trials were followed by an audio error cue and a 12-s inter-trial interval. When rewards or error cues were delivered, the visual display froze for the first second after which the display was then blacked out. The trial length for the seven experimental mice was 6.3 ± 0.8 s (mean \pm s.e.m.) (cue, 2.4 ± 0.4 s; delay, 1.9 ± 0.2 s).

Tower positions were drawn randomly from spatial Poisson processes with means of 7.7 and 2.3 towers per metre on the rewarded and unrewarded sides, respectively. Towers were transient, appearing when mice were 10 cm away from their locations and disappeared after 200 ms. Each session started with at least 10 trials of a visually guided version of the task as warm-up before proceeding to the main task. Behavioural sessions lasted $48:16 \pm 03:44$ (mm:ss; mean \pm s.e.m.; $n = 7$ mice). For analyses, trials in which mice turned around 180° or backtracked to before halfway in the delay region were not included. Detailed methods for the shaping procedures involved in training mice to perform the task, as well as performance and behavioural analyses, have previously been published¹⁴.

A different set of mice learned a simplified version of this task ('alternation task'), in which no towers were presented in the T-maze. In one version of the alternation task ($n = 2$) (Extended Data Fig. 1), the walls were textured differently along the long stem and large distal cues were added, as previously described⁴¹. The maze itself was also slightly longer (340 cm instead of 300 cm). In a second version ($n = 7$) (Extended Data Fig. 6), the maze was identical to the accumulating towers task, except no towers were ever shown. In both cases, mice simply needed to alternate between left and right turns to be rewarded. Visual guides were also present in the arm in which the reward would be located.

Two-photon cellular-resolution calcium imaging

The two-photon calcium imaging setup was identical to a previously published design¹⁵. Two-photon illumination was achieved with a Ti:Sapphire laser (Chameleon Vision II, Coherent) operating at 920 nm. Fluorescence was acquired using a 40 \times /0.8 NA objective (LUMPLFLN40X/W, Olympus) and GaAsP photomultiplier tubes (H10770PA-40, Hamamatsu) after passing through a dichroic (FF670-SDi01, Semrock), an infrared filter (FF01-720sp, Semrock), reflected by a second dichroic (FF562-Di03, Semrock) and passing through a final band-pass filter (FF01-520/60, Semrock). The output signal of the photomultiplier tube was amplified (Variable High Speed Current Amplifier; 59-179, Edmund Optics) and digitized (PXIe-7961R FlexRIO, National Instruments). The microscope was controlled by ScanImage (Vidrio Technologies) software using additional analogue output units (PXIe-6341, National Instruments) for the laser power control and the control of the scanners. Double-distilled water was

Article

used as the immersion medium for the objective. The average beam power measured at the front of the objective was 60–160 mW. The region between the objective and imaging site was shielded from external sources of light using a black rubber tube. Horizontal scans of the laser were achieved using a resonant galvanometer (Thorlabs). Typical fields of view measured approximately $500 \times 500 \mu\text{m}$ and data were acquired at 30 Hz.

Data processing and cell identification

All imaging data were corrected for non-rigid brain motion using custom MATLAB code based on a technique similar to NoRMCorre, in which the image is divided into multiple overlapping patches and a rigid translation is estimated for each patch and frame by aligning against a template³⁷. The set of transitions are then upsampled to create a smooth motion field that is applied to a set of smaller overlapping patches, and the registered frame is then used to update the template by calculating a running mean of past registered frames.

After correcting for motion, fluorescence traces (downsampled to 15 Hz) corresponding to individual cells were extracted using a constrained non-negative matrix factorization algorithm (CNMF)³⁶. Initialization of the spatial components for CNMF was done as previously published, as was classification of identified components into cell-like and non-cell-like categories¹⁵. Automated classification was followed by manual re-classification of a subset of components and artefact rejection. $\Delta F/F$ for each cell was calculated using the modal value of fluorescence in 3-min long windows as baseline fluorescence. An important note is that CNMF can only identify cells with calcium activity during the imaging session, hence cells that were silent for the entire imaging session were not included in reported cell numbers.

Psychometric curves

Psychometric curves (Fig. 1b) were plotted using previously described methods¹⁴. In brief, psychometric curves were fit using a four-parameter sigmoid, $p_R(\Delta) = p_0 + B[1 + e^{-(\Delta - \Delta_0)/\lambda}]^{-1}$, in which Δ is the difference between the number of right and left towers. The binomial confidence interval was calculated using Jeffrey's method^{14,15}.

Logistic regression analysis

Logistic regression (Fig. 1c) was performed using previously described methods¹⁴. In brief, we modelled choices of the mice in each trial with logistic regression in which the factors are the evidence (number of right towers minus number of left towers) in five equally sized regions in the cue period.

For both the psychometric curves (Fig. 1b) and the logistic regression analysis (Fig. 1c), all sessions in which mice ($n = 7$) performed above 60% correct were included ($n = 109$ total sessions).

Mutual information analysis

For each cell, we evaluated a previously defined mutual information metric³⁸, $I = \int_x \lambda(x) \log_2 \frac{\lambda(x)}{\lambda} p(x) dx$, in which I is the mutual information rate of the cell in bits per second, x is the spatial location of the mouse, $\lambda(x)$ is the mean $\Delta F/F$ of the cell at location x , $p(x)$ is the probability density of the mouse occupying location x and $\lambda = \int_x \lambda(x) p(x) dx$ is the overall mean $\Delta F/F$ of the cell.

To obtain $\lambda(x)$, we first denoised $\Delta F/F$ by smoothing with a Gaussian filter with a length of 5 bins and thresholded the result so that values less than 2 robust σ across the time series were set to 0. $\lambda(x)$ was then calculated bin-wise by collecting all smoothed and thresholded $\Delta F/F$ values in their respective bins across the entire session and taking the mean. $\lambda(x)$ was then smoothed by convolution with a Gaussian filter with a length of 5 bins and a σ of 1 bin. $p(x)$ was calculated similarly by counting the number of frames that the mouse spent in each bin across trials and normalized to have a sum of 1.

For position data, 10-cm bins from 0 cm to 300 cm were used. For evidence data, 31 bins (-15 to 16 , number of right towers minus number of left towers) were used. For multidimensional spaces in which we randomized one of the dimensions ($R_E \times Y$ and $E \times R_Y$ in Fig. 2e), the randomized variables (R_E or R_Y) were created by using uniform random sampling with replacement from the joint distribution of discrete evidence (E) and position (Y) values. More specifically, for the $R_E \times Y$ space, in which Y is the non-randomized dimension, we first found the distribution of E values present in the data for each Y value. This created 30 separate E distributions with respect to Y . The R_E value for each frame was generated by randomly sampling from the sole E distribution that corresponded to the non-randomized Y value for that frame. This procedure was performed to control for the non-uniformity of the joint $E \times Y$ distribution in which specific combinations of E and Y values can have greatly different probabilities. A similar procedure was followed for generating the $E \times R_Y$ variable.

To determine significance, the mutual information value of each cell was compared with the mean mutual information value of a shuffled dataset (100 shuffles), in which the $\Delta F/F$ of each cell was circularly shifted by a random interval within each trial, which disrupts the relationship between position and neural activity, but maintains neural activity patterns. Only cells that had mutual information values greater than 2σ above the average mutual information of the shuffle distribution were considered statistically significant. Cells with statistically significant mutual information between neural activity and position in left-choice trials, but not right-choice trials were categorized as left-choice preferring, whereas cells with statistically significant mutual information between neural activity and position in right-choice trials, but not left-choice trials were categorized as right-choice preferring. Those that were significant for both left and right choice trials were categorized as non-preferring. Similar tests were done for mutual information between neural activity and evidence, with the addition that cells for which training and test sets were not correlated (described below) were rejected.

For one-dimensional sequence plots (Fig. 2a, b and Extended Data Figs. 1c, 3c), $\lambda(x)$ values were sorted and normalized on the basis of their peak mean $\Delta F/F$ values. For the cross-validation procedure for evidence fields (Extended Data Fig. 1c), trials were ranked on the basis of the maximum $\Delta F/F$ value of the cell in a given trial. Odd-ranked and even-ranked trials were assigned to the training and test sets, respectively. $\lambda(x)$ was recalculated on the training and test sets and smoothed as described above. Only cells with significantly correlated $\lambda(x)$ between the training and test sets ($P < 0.05$) were used in the sequence plots. The training set was sorted on the basis of the peak mean $\Delta F/F$ values and plotted. This same sorting index was then used to plot the test set.

For Extended Data Fig. 2d, cells were considered to encode both evidence and position if they had significant mutual information in $E \times Y$ space, as described above. Of the remaining cells, cells were considered to encode only position if they were significant in $R_E \times Y$ space (16%) and only evidence if they were significant in $E \times R_Y$ space (6%). For Extended Data Fig. 2e, distributions of mutual information in $R_E \times Y$ and $E \times R_Y$ space were calculated from 50 different shuffles, in which either E or Y were shuffled. Of the $E \times Y$ cells described above, 89.9% had mutual information values in $E \times Y$ space greater than 2σ above the mean of both shuffled distributions, 9.8% had mutual information values that were greater than only the $E \times R_Y$ distribution and 0.3% had mutual information values that were greater than only the $R_E \times Y$ distribution.

Counting the number of place fields

To estimate the number of place fields in $E \times Y$ space, we followed a heuristic to count peaks derived from previous studies^{43,44}. Using the neural activity maps for each neuron in $E \times Y$ space (Fig. 2c, d and Extended Data Fig. 2a) obtained as described above (see 'Mutual information analysis'), we considered all bins that were 2σ above the shuffled mean as candidate place fields in the $E \times Y$ space. We then joined all bins with

adjacent significant bins, and if a connected component exceeded $3 \times 3 = 9$ bins, we counted the connected component as a place field. The distribution of the place field counts is shown in Extended Data Fig. 2g. Note that a very small number of cells (3%, $n = 31$ out of 917 cells) had significant firing fields above the shuffled control that were smaller than 9 bins. These appear in the histogram as '0'. Cells had approximately 1.7 ± 0.3 (mean \pm s.e.m.) firing fields, with 53% ($n = 490$ out of 917) of cells having more than one firing field.

Manifold inference from neural dynamics

To infer latent dimensions from neural dynamics, we adopted a previously developed procedure¹³ for calcium-imaging data. We first smoothed the raw $\Delta F/F$ traces with an 11-bin Gaussian filter and thresholded at 4σ , for which we estimated the robust σ across the entire time series, but individually for every neuron. We restricted our analysis to cells that had at least one transient in the recording session, and imaging frames that had at least one active cell, as well as the portion of the maze represented by 0–300 cm (cue and delay periods). We then followed the previously published procedure¹³ to calculate the distances between pairs of population activity vectors, extracting a set of latent variables from these distances with multidimensional scaling, and learning a map between latent space and network activity with local linear embedding (LLE).

In brief, we first learned a generative model of transition probabilities from population activity $s(t) = [s_1(t), \dots, s_N(t)]$ of N neurons at time $0 < t < T$, to the activity $s(t + \Delta t)$ using the previously developed random forest method¹³ with a few modifications. First, when splitting the neural state space into regions using a set of hyperplanes organized in a decision tree, we assessed 20 random hyperplane orientations at every node of the tree and selected the orientation that best split the data. This improved performance with the large numbers of neurons typically encountered in calcium imaging. Second, we set the minimum number of leaves in each random tree to 500. Third, to define transitions, we considered all states $\Delta t = 67$ ms apart (one frame at a 15-Hz frame rate). Fourth, we fit manifolds to all data points, not only a subset of landmarks. All other hyperparameters were chosen as previously described¹³. The random forest model provides us with a set of transition probabilities $p(s(t + \Delta t)|s(t))$ that can be translated into a local distance $\delta(s(t + \Delta t), s(t))$ under a diffusion approximation, in which the transition probability p decreases with distance δ as $p \propto \exp(-\delta^2)$. Similar to isomap⁴⁵, we then calculated the global distance between two states as the length of the shortest path from one to the other via any intermediate, connected states. The pairwise geodesic distances of l points $\rho(i, j)$, in which $0 < i, j \leq l$, then yields a matrix of size $l \times l$ that was embedded using multidimensional scaling with Sammon's nonlinear mapping. This yielded latent variables to describe population data. The mapping from latent space to neural activity and back was then achieved with LLE¹³.

Manifold inference on video files

To construct a low-dimensional representation of the task itself, we applied the algorithm described above (see 'Manifold inference from neural dynamics') to the visual input that the mice received in a typical experimental session, more specifically to the blue channel across all RGB pixels in each frame of the video files displaying the field of view of the mice. This corresponds to a vectorized time series of $1,792 \times 1,088 = 1,949,696$ pixels as a function of time. To make this analysis computationally viable, we first downsampled the videos $17\times$ from the original $1,792 \times 1,088$, restricted our analysis to trials shorter than 30 s and frames with positions between 0 and 350 cm, and simplified the hyperparameters, in comparison to the analysis of neural data by using only two random hyperplane orientations and 1,000 landmarks. All other parameters were identical to the analysis of neural data. The results are shown in Extended Data Fig. 5, where Extended Data Fig. 5a shows the mean luminance of the blue channel, after averaging across all pixels.

Dimensionality estimation

To estimate the dimensionality of the latent manifold, we analysed the geometric properties of the geodesic distance matrix $\rho(i, j)$. We specifically studied the statistics of nearest neighbour distances. Suppose that the neural states were confined to a two-dimensional sheet in high-dimensional neural state space. Within the sheet, the cumulative number of points N within distance r will increase quadratically with distance r , as more points on the sheet will fall within the neighbourhood, thus recovering the two-dimensional sheet structure. Using this variation of the correlation dimension that can also be used for complex attractor geometries^{12,39}, we found a wide range of values for which the number of points scaled like a power law.

We fit this power law by minimizing the quadratic error to the model function $N(r) = cr^d$, in which N is the total number of neighbours, r is the distance, and c and d are fit parameters. We fit this function over three orders of magnitude, for $10^3 < N < 10^6$. The average across the seven mice yielded $d = 5.4$ (4.8–6.0; 95% bootstrapped confidence intervals). These numbers are consistent with an approximately 4–6-dimensional manifold, embedded in an approximately 450-dimensional neural state space (Fig. 3a). For the illustrations in Fig. 3a and Extended Data Fig. 3a, we normalized the distance by the average length of a trial along the manifold for each mouse.

Reconstructing neural data from embedded manifolds

To assess the quality of the dimensionality reduction performed with MIND, we measured how well the neural data can be reconstructed from the d latent variables after embedding the manifold into d dimensions (Extended Data Fig. 4). This provides us with an estimate of the minimum number of dimensions required for the reconstruction quality to saturate. This number should be comparable to the intrinsic dimensionality of the manifold, and thus provided us with a separate measurement of the dimensionality of the manifolds.

Measuring how well the coordinated activity of neurons is predicted by the manifold

To this end, we held out a random trial, fit a manifold to the remaining data, and embedded this manifold into 2–7 dimensions using the methods described above (see 'Manifold inference from neural dynamics'). After fitting the manifold to the training data, we first projected the held-out trial onto the manifold to obtain d coordinates for every time point and then reconstructed neural activity from these d numbers in the test dataset using LLE¹³. We then thresholded the LLE estimate to capture the thresholding nonlinearity of calcium imaging. The thresholding cut-off was estimated from the training data for the best reconstruction. To assess the similarity between the raw data and the reconstruction, we then measured the correlation coefficient between the reconstructed neural data and the real data. These data are a vectorized time-series of the form neurons \times time. To perform an element-wise comparison, we concatenate all columns into a single vector and calculate the correlation coefficient. This number was averaged across the 10 held-out trials to form the decoding index, and the process was repeated for all seven mice (Extended Data Fig. 4a, b). The data shown in Fig. 3c are the mean \pm s.e.m. for the seven mice. In Fig. 3b, raw $\Delta F/F$ and reconstructed $\Delta F/F$ traces have been smoothed with an 11-bin Gaussian filter and thresholded at 4 robust σ . For the reconstructed $\Delta F/F$ traces, baseline subtraction before smoothing and thresholding was accomplished by subtracting the mean of the reconstructed activity of each cell from the reconstructed activity of each cell.

Measuring how well the activity of individual neurons is predicted by the manifold

This analysis is similar to the one above (see 'Measuring how well the coordinated activity of neurons is predicted by the manifold') but tailored to quantify the predictive power of the manifold on a

single-cell level (Extended Data Fig. 4c). To this end, we removed one test neuron from all the N cells in the data and used MIND to fit a manifold to the remaining $N - 1$ training neurons⁴⁶. We then used GPR to learn a map $g(x)$ from manifold coordinates x to the activity of the test neuron in 80% of the trials. We used a squared exponential kernel function to specify the covariance, making these fits smooth and differentiable, as expected for a response similar to a firing field. In the remaining 20% of trials, we evaluated $g(x)$ and measured the correlation coefficient between the predicted and observed data of the test neuron. This was repeated four more times for fivefold cross-validation and the correlation coefficient over the five folds was averaged. This value was calculated for 10 randomly chosen neurons from the 25 most active neurons in each mouse and averaged to form the decoding index (Extended Data Fig. 4e). In Extended Data Fig. 4d, reconstructed $\Delta F/F$ traces were baseline-subtracted, smoothed and thresholded identically to the procedure mentioned above for Fig. 3b (see ‘Measuring how well the coordinated activity of neurons is predicted by the manifold’).

Comparing MIND with principal component analysis

To compare this nonlinear dimensionality reduction technique with a linear method, we also calculated the decoding index (cross-validated correlation coefficient between predicted and observed data in a held-out trial) using principal component analysis (PCA). To this end, we removed a held-out trial from the data, calculated the principal components of the remaining data and identified the d principal components with greatest coefficients in the training data. We then projected the held-out trial onto these d principal components and used the obtained coefficients to project back into neural state space. The similarity of the observed held-out trial and the reconstruction from PCA was assessed with the correlation coefficient and averaged across 10 random held-out trials. To reach the same mean cross-validated decoding index as MIND for manifolds embedded in $d = 4, 5$ and 6 dimensions, PCA required $d = 29, 40$ and 47 principal components, respectively.

Decoding position and evidence from the manifold and neural activity

We used GPR to learn a function from latent space or neural activity (selecting only the top 10% of cells with highest mutual information for position or evidence to limit overfitting) to position and evidence. Other nonlinear regression methods such as LLE yielded similar results, whereas linear decoding methods generally failed. Figure 3g shows the correlation coefficients between the position and evidence values in the behavioural session of each mouse predicted from the learned regression model (trained on 80% of trials, applied to the test dataset of 20% of trials, and repeated for five folds) and true position and evidence values (averaged over the five folds)—the decoding index. To visualize the position and evidence values (Fig. 3f) as well as the luminance (Extended Data Fig. 5b, c) and view angle (Extended Data Fig. 6a), we smoothed across the 20 nearest neighbours in latent space.

Similar methods were used in analyses shown in Extended Data Fig. 6. To assess whether knowledge of variable X adds to how well variable Y predicts the manifold, we decoded manifold dimensions with GPR using both X and Y as inputs or X and shuffled Y . To assess whether correlated and orthogonal components of X and Y could both be decoded, we used PCA on variables X and Y to linearly regress out Y from X and decoded both principal component (PC)1 and PC2 from the manifold dimensions. To evaluate the accuracy of the decoding for binary variables, such as the upcoming choice, the choice in the previous trial and whether the previous trial was correct, we averaged the prediction from the GPR across the trial to come up with a single value, which was binarized as a single prediction—for example, predicting a left or right choice in the trial—and compared it with the true value—in this case, whether the mouse actually makes a left or right choice.

Hyperalignment procedure

Hyperalignment across two mice was performed as follows. We first fit the neural data of mouse A with MIND to obtain a set of T d -dimensional latents x_A^t . We then perform GPR to learn a map from the d -dimensional latents to a behavioural variable $e_A^t = \text{GPR}(x_A^t)$. Next, we perform MIND on the data of mouse B. This yields a different set of d -dimensional latent vectors x_B^t . From these latents, we predict the behaviour of mouse B using the GPR trained on mouse A and a five-dimensional rotation matrix R with $e_B^t = \text{GPR}(Rx_B^t)$. The rotation matrix was calculated from a five-dimensional representation of the

special orthogonal group of degree 5 (SO(5)) so that $R = \prod_{i=1}^{10} \text{expm}(g_i c_i)$. Here, $\text{expm}()$ indicates the matrix-exponential of g_i , the ten generators of SO(5), multiplied with a scalar angular parameter c_i . These parameters were cross-validated by optimizing on the first half of the data and decoding of position and evidence assessed on the second half. For each mouse, we decoded position and evidence using the hyperaligned five-dimensional manifolds of the six other mice. In Fig. 3i, we show the maximum decoding that can be done across the six other mice for each mouse, compared with the cross-validated GPR (Fig. 3g) for five-dimensional embeddings of the manifold in the same mouse. Means were then calculated across the seven mice. We estimated the contribution of shared geometry for each mouse in terms of fractional variance explained by dividing the r^2 of position and evidence decoding obtained with hyperalignment by the r^2 of the best decoding that could be done with either method.

Task trajectories

To visualize the sequential patterns of the task (Extended Data Fig. 8), we first extracted ‘task trajectories’ as smooth spline interpolations of the specific trajectory through $E \times Y$ space experienced over trials. The task trajectories for single trials in a behavioural session are plotted as thin lines in Extended Data Fig. 8a together with fits across all left or right trials (thick lines). In addition, we also visualized task trajectories as a flow field, for which we binned $E \times Y$ space into 10-cm and 1-tower bins and calculated the trial-averaged gradient in the position and evidence directions for every bin. The resulting gradient matrices were then individually smoothed by convolution with a Gaussian filter with a length of five bins and a σ of one bin. Every other bin was plotted as arrows centred on the respective bin and pointing to the average direction of the gradient (Extended Data Fig. 8b).

Identification and analysis of sequences

A pair of cells was classified as a doublet if the number of trials in which the first cell had a transient event before the second cell was greater than 2σ above the mean of the same value obtained from a shuffled dataset (100 times) in which neural activity was circularly shifted in each trial. Doublets that appeared in fewer than three trials were removed. A transient event was defined as any time $\Delta F/F$ (smoothed with a Gaussian filter with a length of five bins) for that cell was greater than a threshold equal to $11 \times (\text{Ai93} \times \text{EMX1})$ or $5 \times (\text{GP5.3})$ the robust σ across the entire imaging session. Different thresholds for event detection were used for the two mouse strains due to the difference in signal-to-noise ratios. Triplets were constructed by simply combining doublets without allowing the same cell to appear twice, that is, a cell cannot be the first cell and the third cell in the triplet, and tested using the same significance test as was used for doublets.

Even if two place cells had activity that was completely independent, we would still expect, by chance, that they fire in the same trials for a subset of trials. For example, two place cells with fields at 100 cm and 200 cm that are each active in 100 random trials in a session with 200 trials would, on average, show up together in 50 random trials. However, if these two cells appeared in all 100 trials together, it would be unlikely that their activity was independent. To test whether doublets appeared

more often than chance, trial IDs of each cell were independently shuffled, so that relationships between cells were disrupted without affecting the neural activity of each cell (Extended Data Fig. 8d), and then we searched this shuffled dataset for the doublets again to determine the number of instances a doublet would show up if the activity of the two cells were independent ($n = 100$ shuffles).

A doublet was determined to be choice-predictive if the probability that the mouse was going to turn right in trials in which the given doublet occurred was greater than 2σ above or below the mean probability of a right turn after shuffling the choices for each trial ($n = 1,000$ shuffles). The same assessment was made to determine choice-predictiveness in triplets. Once choice-predictive doublets and triplets were identified, we compared the predictiveness of real doublet events to events obtained from datasets in which trial IDs were shuffled ($n = 100$ shuffles).

Comparison of sequences and the predictions from the manifold

To show that the manifold can predict the presence of sequences, we used the manifold to reconstruct the $\Delta F/F$ of each cell in each imaging session with LLE (described above in 'Manifold inference from neural dynamics'). We then detected doublet events from this reconstructed data and compared the trials in which doublet events were found against the real data to generate the true-positive rate (TPR) and false-positive rate (FPR) for doublet events in each mouse.

More specifically, we constructed a Boolean array B_{data} of size $N_{\text{cells}} \times N_{\text{cells}} \times N_{\text{trials}}$ indicating the presence or absence of a doublet in a specific trial. We populated this array with the doublet-finding algorithm described above (see 'Identification and analysis of sequences') using the observed calcium data. This constitutes ground truth. We then reconstructed all neural activity from the latent dimensions of the five-dimensional embedding of the manifold. This activity was then thresholded at an activity level θ , and we considered only transients that exceed this threshold. By definition, these data are the manifold prediction. We identified doublet events in these surrogate data with the same algorithm to construct a Boolean array $B_{\text{prediction}}$. Comparing this prediction with the ground truth, we can count the number of true positives ('1' in both the ground truth and the surrogate array), false positives ('1' in the surrogate array, '0' in the ground truth), false negatives ('0' in the surrogate array, '1' in the ground truth) and true negatives ('0' in both the ground truth and the surrogate array). TPR was defined as $TP/(TP + FN)$, in which TP is the number of true positives, and FN is the number of false negatives. FPR was defined as $FP/(FP + TN)$, in which FP is the number of false positives, and TN is the number of true negatives. We then scanned across θ (1–100, in intervals of 5) to construct a receiver operating characteristic (ROC) curve (Extended Data Fig. 9d) and calculated the distance d between the point (0, 1) in the top left corner of ROC space and any point on the ROC curve, $d^2 = (1 - TPR)^2 + FPR^2$. We chose the threshold θ such that this distance was minimal to identify a point of best discriminant capacity. The values of TPR and FPR reported in the main text are averages across these points for all seven mice.

We next calculated the predictive power of the manifold for the exact timing of a doublet. For all doublets, we measured the length of the trajectory between the firing of the first cell and the firing of the second cell on the manifold in each trial when the doublet was active. This length, plotted against the time between the sequentially active cells, is shown in Fig. 4e. To test whether the observed correlation of time elapsed and distance on the manifold was significantly greater than the

correlation between time elapsed and any distance on the manifold, we compared the observed correlation to the correlation coefficients obtained from comparing time elapsed in a trial with manifold distances over the same time interval obtained from a different trial. We averaged the correlations across 100 random trajectories obtained from other trials and over all doublets for each mouse and performed a two-tailed Wilcoxon signed-rank test on the average real and random correlation values of the mice ($n = 7$) to test whether real correlation values were significantly greater than the random correlation values.

Statistical tests

All statistical tests were performed with MATLAB (2015b, 2018a, 2018b and 2020a; Mathworks). Bonferroni correction of P values was performed by multiplying the unadjusted P value by the number of multiple comparisons made. In cases in which the corrected P value exceeded 1.0, we reported the value as 1.0.

Reporting summary

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

Data availability

The datasets generated in this study are available from the corresponding authors on reasonable request. Source data are provided with this paper.

Code availability

The code used for all analyses in this study is available on GitHub (https://github.com/BrainCOGS/HPC_manifolds). All other codes are available on reasonable request.

- Gauthier, J. L. & Tank, D. W. A dedicated population for reward coding in the hippocampus. *Neuron* **99**, 179–193.e7 (2018).
- Pinto, L. et al. Task-dependent changes in the large-scale dynamics and necessity of cortical regions. *Neuron* **104**, 810–824.e9 (2019).
- Domnisoru, C., Kinkhabwala, A. A. & Tank, D. W. Membrane potential dynamics of grid cells. *Nature* **495**, 199–204 (2013).
- Rich, P. D., Liaw, H.-P. & Lee, A. K. Large environments reveal the statistical structure governing hippocampal representations. *Science* **345**, 814–817 (2014).
- Tenenbaum, J. B., de Silva, V. & Langford, J. C. A global geometric framework for nonlinear dimensionality reduction. *Science* **290**, 2319–2323 (2000).
- Yu, B. M. et al. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J. Neurophysiol.* **102**, 614–635 (2009).

Acknowledgements We thank A. Song and S. Thiberge for assistance with two-photon imaging, S. Stein and S. Baptista for technical support with animal training, M. L. Ioffe for providing code and E. M. Diamanti and B. Engelhard for discussions. This work was supported by NIH grants U01NS090541, U19NS104648 and F32MH119749, as well as the Simons Collaboration on the Global Brain.

Author contributions E.H.N., C.D.B. and D.W.T. designed the experiments. E.H.N. and N.W.F. performed the experiments. E.H.N., M.S. and N.W.F. analysed the data. S.A.K. wrote the code for the processing of the calcium-imaging data. L.P. provided the code for behavioural analysis. J.L.G. provided the code for mutual information calculations. M.S., R.J.L. and S.L. developed the MIND algorithm. E.H.N., M.S., C.D.B. and D.W.T. wrote the paper.

Competing interests The authors declare no competing interests.

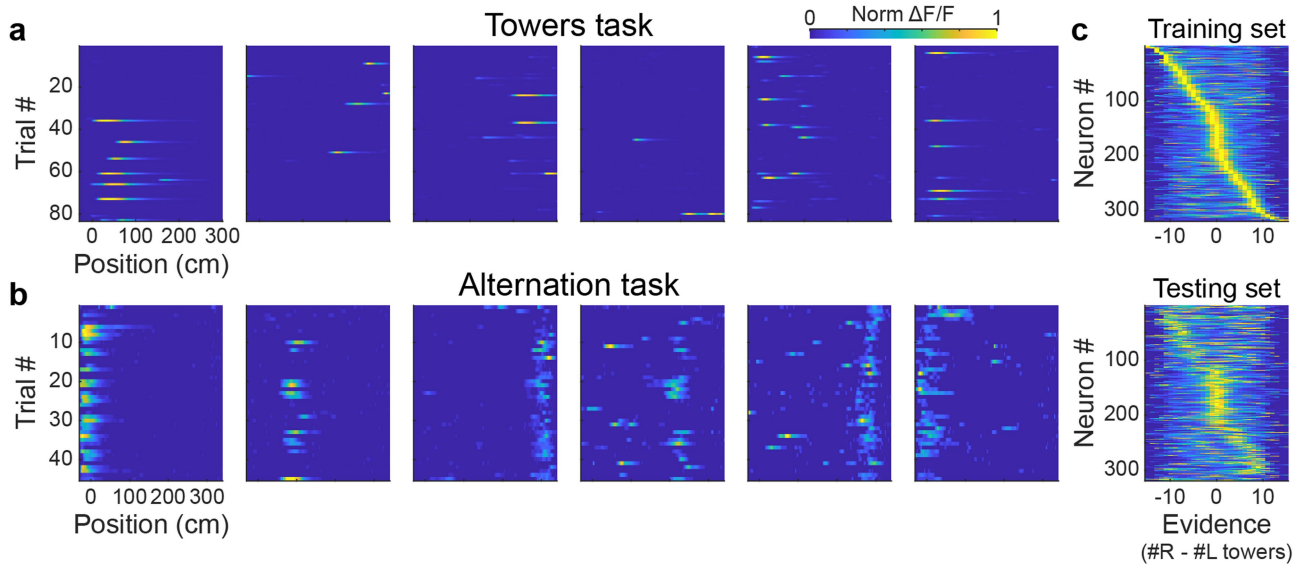
Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41586-021-03652-7>.

Correspondence and requests for materials should be addressed to C.D.B. or D.W.T.

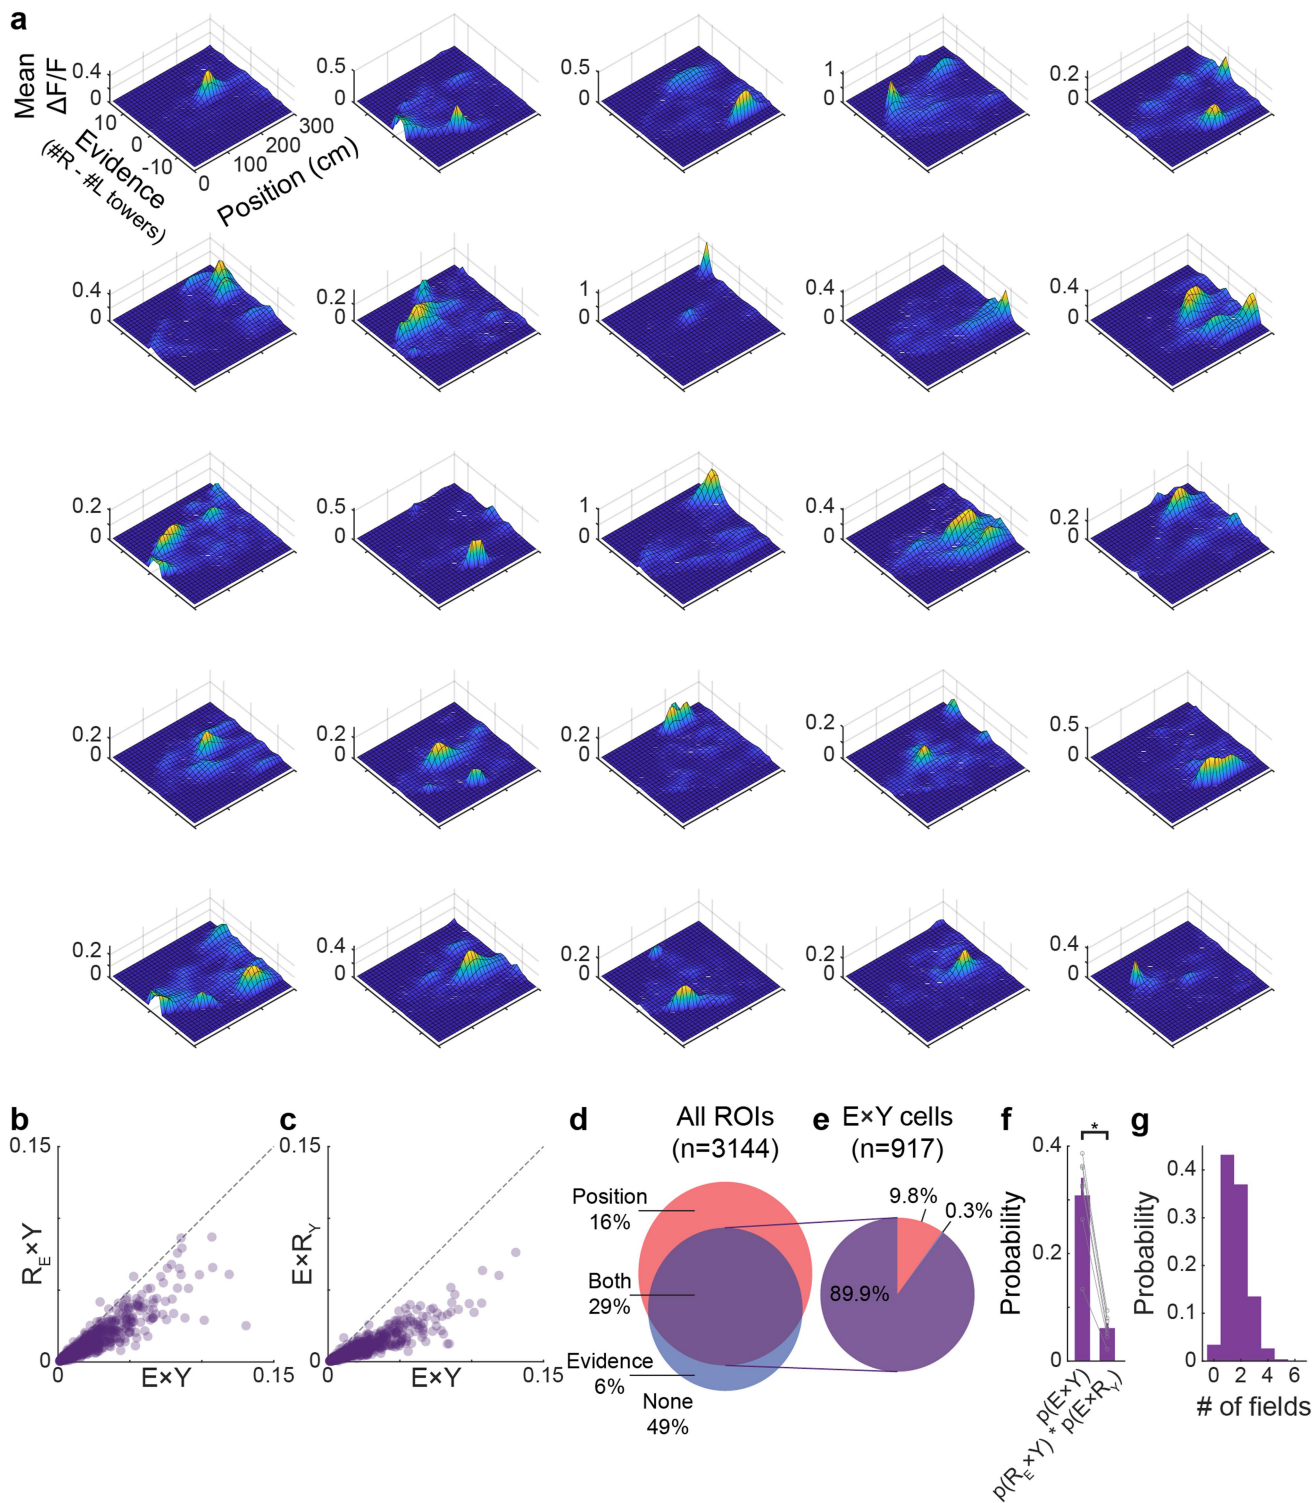
Peer review information Nature thanks the anonymous reviewers for their contribution to the peer review of this work.

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



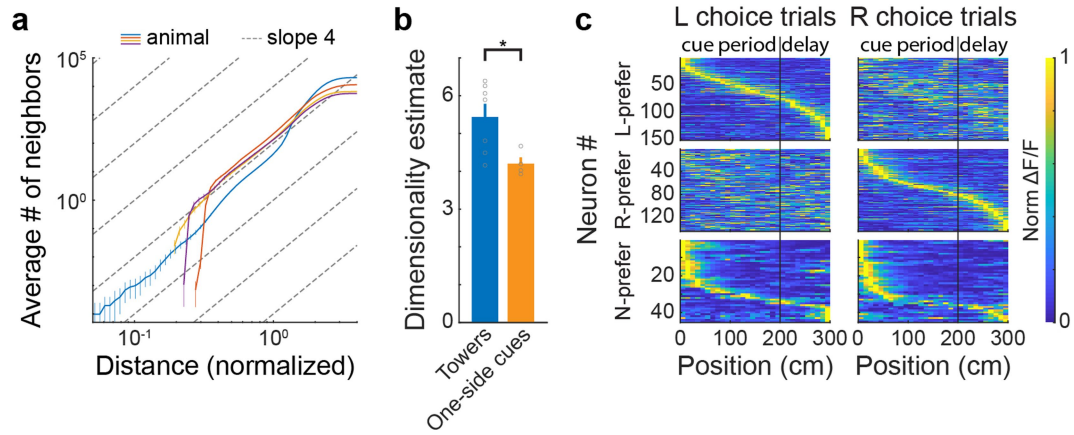
Extended Data Fig. 1 | Characterization of CA1 neural variability in the accumulating towers task. a, Each heat map represents one neuron and the trial-by-trial activity of that neuron in the towers task for left-choice trials. Each row in each heat map is the $\Delta F/F$ (normalized within each neuron) of the neuron in that trial. **b**, Same as in **a**, but for the alternation task. Note that the single-trial activity appears more variable in the towers task and more reliable in the

alternation task, consistent with the results that evidence is also being represented by neurons in the towers task. **c**, Neural activity ($\Delta F/F$ normalized within each neuron) of cells significantly encoding evidence, sorted by activity in half the trials (top) and plotted using the same sorting in the other half of the trials (bottom).



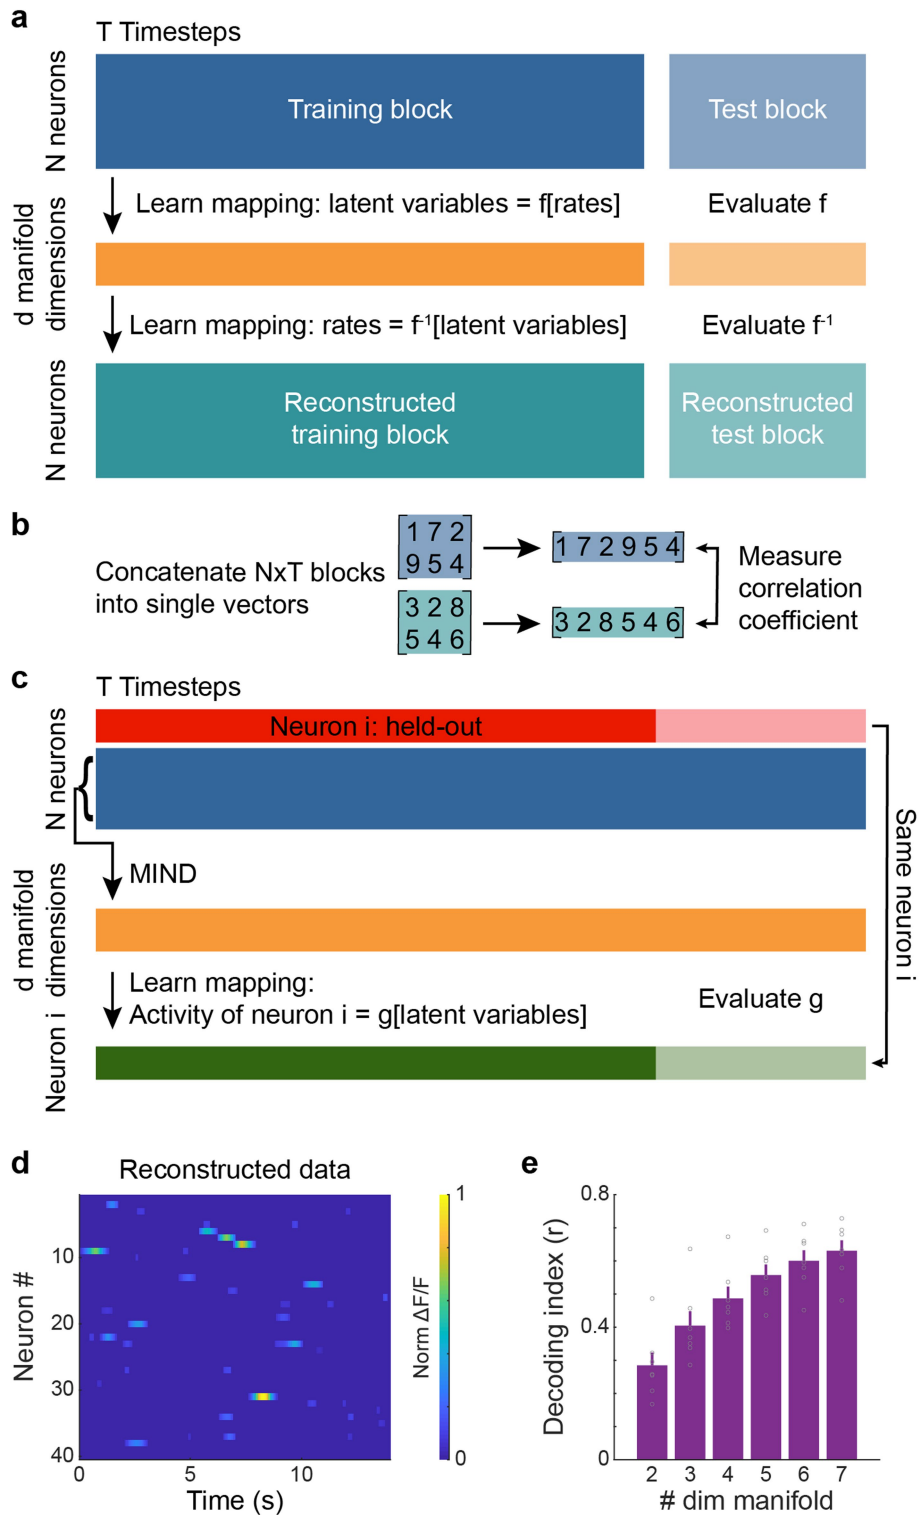
Extended Data Fig. 2 | Place fields in evidence-by-position space. a, Each heat map shows the mean $\Delta F/F$ of a neuron with significant mutual information in $E \times Y$ space. **b**, Scatterplot of the mutual information in $R_e \times Y$ space versus $E \times Y$ space for each cell with significant information in $E \times Y$ space ($n = 917$ neurons). R_e is randomized evidence. **c**, Same as in **b**, but for $E \times R_y$ space versus $E \times Y$ space. R_y is randomized position. **d**, In total, 29% of imaged neurons had significant mutual information in $E \times Y$ space, whereas 16% had significant mutual information only for position and 6% had significant mutual information only for evidence. **e**, Of the cells with significant mutual

information in $E \times Y$ space, 89.9% had significantly more information in $E \times Y$ space than just place or evidence information alone, whereas 9.8% could not be differentiated from place cells and 0.3% could not be differentiated from evidence cells (Methods). **f**, The probability of a cell having significant mutual information in $E \times Y$ space is significantly greater than the joint probability of a cell being a place cell and a cell being an evidence cell. Two-tailed Wilcoxon signed-rank test, $*P = 0.016$; $n = 7$ mice; data are mean \pm s.e.m. **g**, Cells with significant mutual information in $E \times Y$ space had 1.7 ± 0.03 (mean \pm s.e.m.) firing fields ($n = 917$ cells).



Extended Data Fig. 3 | Dimensionality of an earlier training stage. During the training of the towers task, mice proceed through various stages of training. In one of these training stages, mice perform a task nearly identical to the towers task, except that visual cues only show up on one side of the maze. **a**, The intrinsic dimensionality of the one-side cues task is approximately 4.2 (4.0–4.5; 95% bootstrapped confidence interval). $n = 4$ mice; data are mean \pm 95% bootstrapped confidence intervals. Grey dashed lines illustrate the slope expected for a four-dimensional manifold. **b**, Intrinsic dimensionality of

the one-side cues task is significantly lower than the dimensionality of the towers task. Two-tailed Wilcoxon rank-sum test, $*P = 0.042$; $n = 7$ mice (towers task) and $n = 4$ mice (one-side cues task); data are mean \pm s.e.m. **c**, Choice-specific place cell sequences in the one-side cues task, similar to Fig. 2a. Sequences are divided into left-choice-prefering (top), right-choice-prefering (middle) and non-prefering (bottom) cells. Data are split between left-choice trials (left) and right-choice trials (right). Cells are shown in the same order within each row group. $\Delta F/F$ was normalized within each neuron.



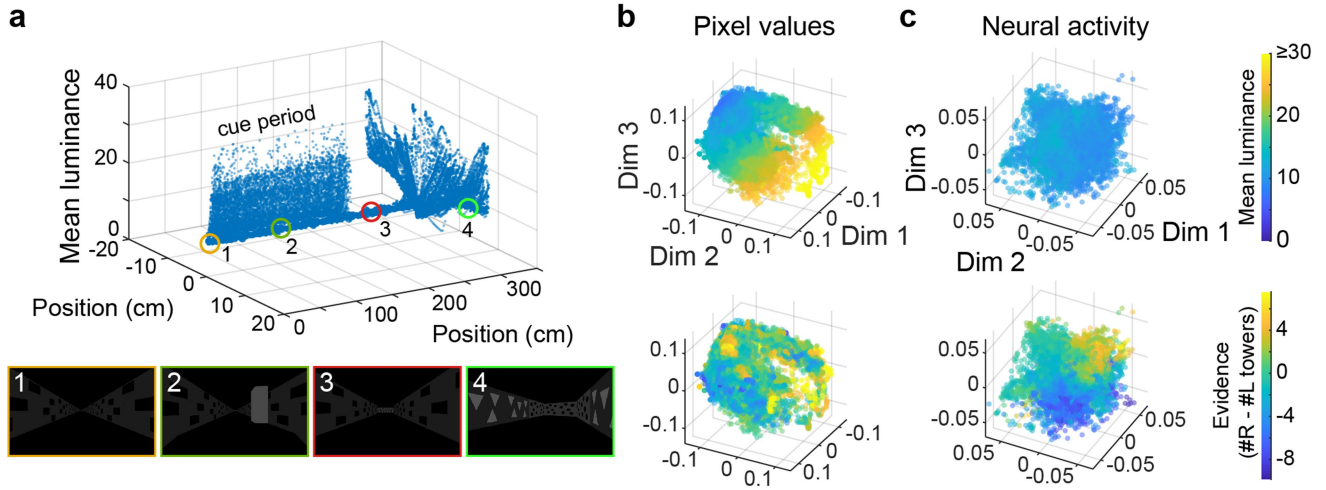
Extended Data Fig. 4 | See next page for caption.

Article

Extended Data Fig. 4 | Cross-validation methods and results demonstrating how neural activity from single neurons is captured by coordinated population activity.

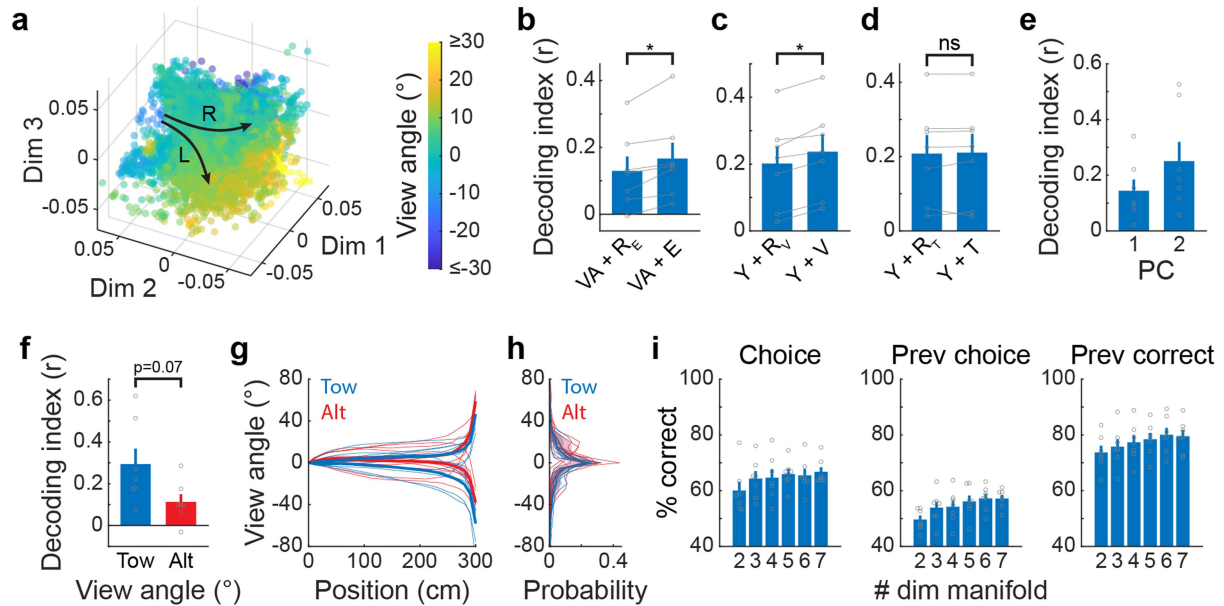
a, Illustration of the cross-validation method to calculate the decoding index in Fig. 3c. Data are split for training (solid colours) and testing (shaded colours). With the training data, a map is obtained from $\Delta F/F$ to latent dimensions and back. This map is then evaluated on the test data. **b**, To assess the performance of the map, we concatenate the neuron \times time data in the test block and reconstructed test block into two vectors and calculate the correlation coefficient from the elementwise pairwise comparison of the vectors. The correlation coefficient was averaged across 10 individually held-out trials to yield the decoding index. **c**, Illustration of a similar analysis in which the activity of a single cell is decoded from a manifold fit to the rest of the neural population. One neuron (red) is removed before using MIND to obtain a set of latents. Next, in the training data (solid green), a map is calculated from

the manifold to the activity of the held-out neuron. The map is then used to predict the test data (shaded green). The correlation coefficient is calculated as in **b** and averaged over five folds across 10 individually held-out neurons as the decoding index. **d**, Example of neural activity from 40 individually reconstructed neurons, in which the activity of each neuron was decoded from the five-dimensional manifold fit to the other cells following procedures in **c** (comparable to Fig. 3b, for which the method in **a** and **b** was used). $\Delta F/F$ is normalized to the maximum $\Delta F/F$ in the window shown. **e**, Cross-validated correlation coefficients between the activity of individual neurons in the real and reconstructed data, in which the reconstruction was accomplished with d -dimensional embeddings of the neural manifold. The decoding index is the correlation coefficient between the predicted and real $\Delta F/F$ of the held-out ROIs. $n = 7$ mice; data are mean \pm s.e.m.



Extended Data Fig. 5 | Task manifold and neural manifold encode different variables. **a**, The visual space of the accumulating towers task across a representative session. Top, the mean luminance of the virtual-reality visual field as a function of position in the T-maze. Bottom, four example frames. Note the high variability of luminance during the cue period, where bright towers are randomly presented on the left and right walls. **b**, Performing dimensionality reduction on the time series of the pixel values in the raw video stream using

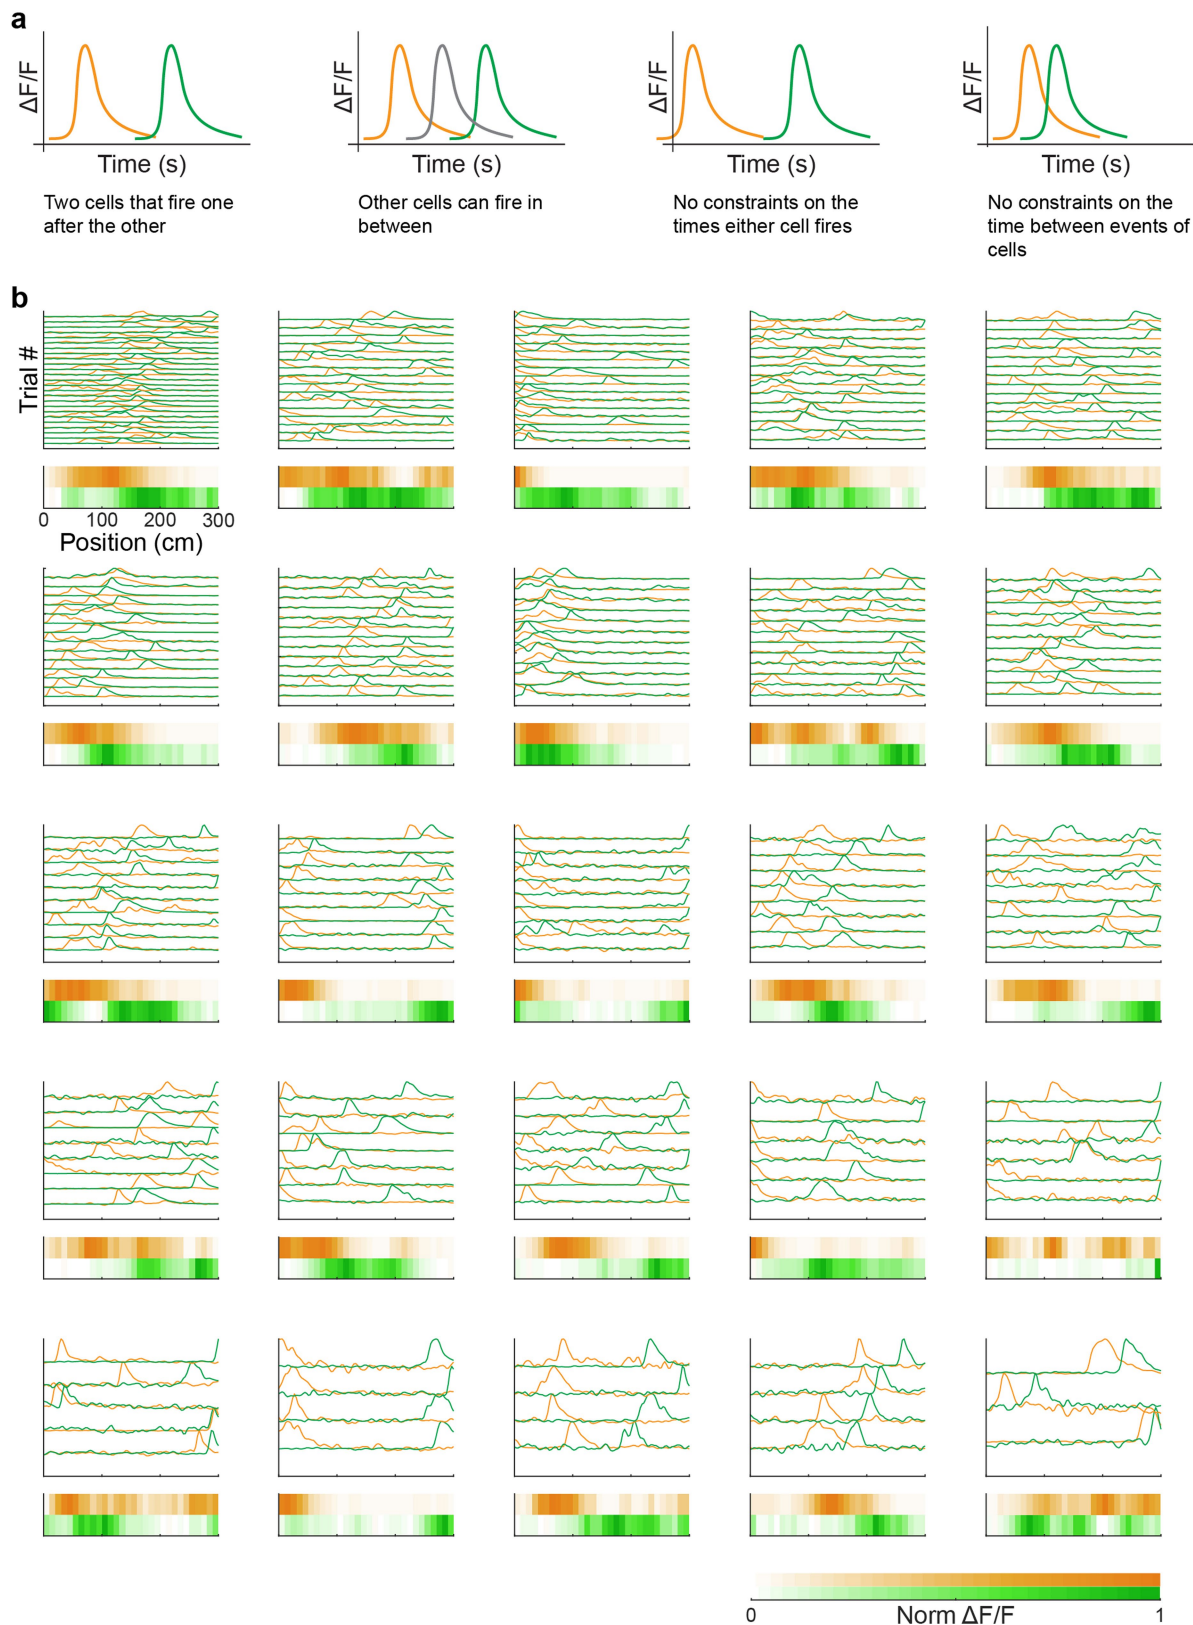
MIND reveals a low-dimensional manifold, reflecting the visual sensory structure of the accumulating towers task. Plotting luminance (top) and evidence (bottom) on the manifold reveals that luminance is represented as a smooth gradient, whereas evidence requires memory and is thus absent on the task manifold. **c**, Same as in **b**, but showing the neural manifold obtained from the mouse that ran the session (Fig. 3f). Note the absence of a luminance representation, but the emergence of evidence.



Extended Data Fig. 6 | Decoding other variables from the neural manifold.

a, Similar to Fig. 3f, the view angle is plotted as colour on the three-dimensional embedding of the manifold. **b**, The five latent variables of the neural manifold embedded in a five-dimensional Euclidean space are better predicted by GPR from view angle and evidence values than from view angle and shuffled evidence values. Two-tailed Wilcoxon signed-rank test, $*P=0.016$; $n=7$ mice; data are mean \pm s.e.m. Decoding index is the correlation coefficient between the predicted manifold values and true manifold values, averaged over the five dimensions of the manifold. **c**, Same as in **b**, but for decoding manifold values using position and velocity. The addition of velocity to position information significantly improves the decoding of manifold values. Two-tailed Wilcoxon signed-rank test, $*P=0.016$; $n=7$ mice; data are mean \pm s.e.m. **d**, Same as in **b**, but for decoding using position and time. The addition of time information does not significantly increase how well manifold values are decoded. Two-tailed Wilcoxon signed-rank test, $P=0.30$ (ns, not significant); $n=7$ mice; data are mean \pm s.e.m. **e**, We used PCA to separate the correlated and orthogonal dimensions between evidence and view angle and decoded both PC1 (correlated) and PC2 (orthogonal) from the five-dimensional embedding

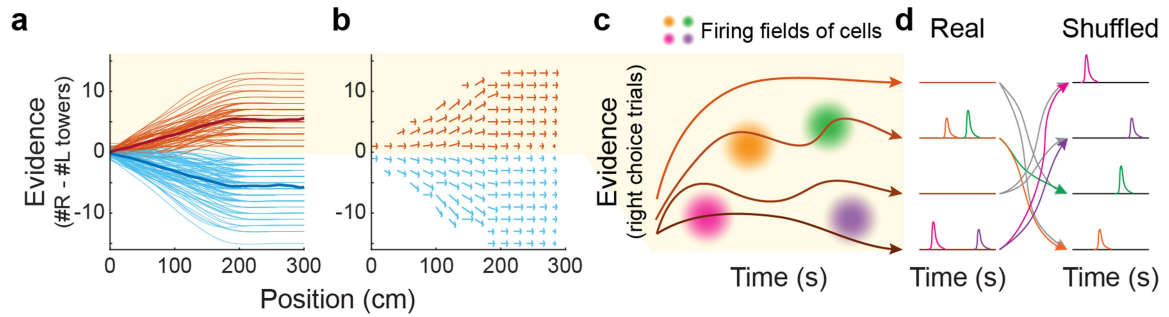
of the manifold. $n=7$ mice; data are mean \pm s.e.m. The decoding index is the correlation coefficient between the predicted and true principal component values. **f**, The view angle is better decoded from the neural manifold (five-dimensional embedding) in the towers task ('Tow'), when evidence is also present, than in the alternation task ('Alt') when evidence is not present. Two-tailed Wilcoxon rank-sum test, $P=0.07$; $n=7$ mice (towers task) and $n=7$ mice (alternation task); data are mean \pm s.e.m. The decoding index is the correlation coefficient between the predicted and true view angle values. **g**, Average view angle trajectories, separated between left- and right-choice trials, for the mice in the towers task ($n=7$ mice; blue, thin lines) and the alternation task ($n=7$ mice; red, thin lines). Thick lines represent averages across mice. **h**, Average view angle values in the towers task ($n=7$ mice; blue, thin lines) and the alternation task ($n=7$ mice; red, thin lines) over all trials. Thick lines and shaded area are mean \pm 95% bootstrapped confidence interval. **i**, Accuracy in predicting the upcoming choice (left), the choice of the mouse in the previous trial (centre) and whether the previous trial was rewarded (right) from d -dimensional embeddings of the neural manifold. $n=7$ mice; data are mean \pm s.e.m.



Extended Data Fig. 7 | Examples of sequences in CA1 neural activity.

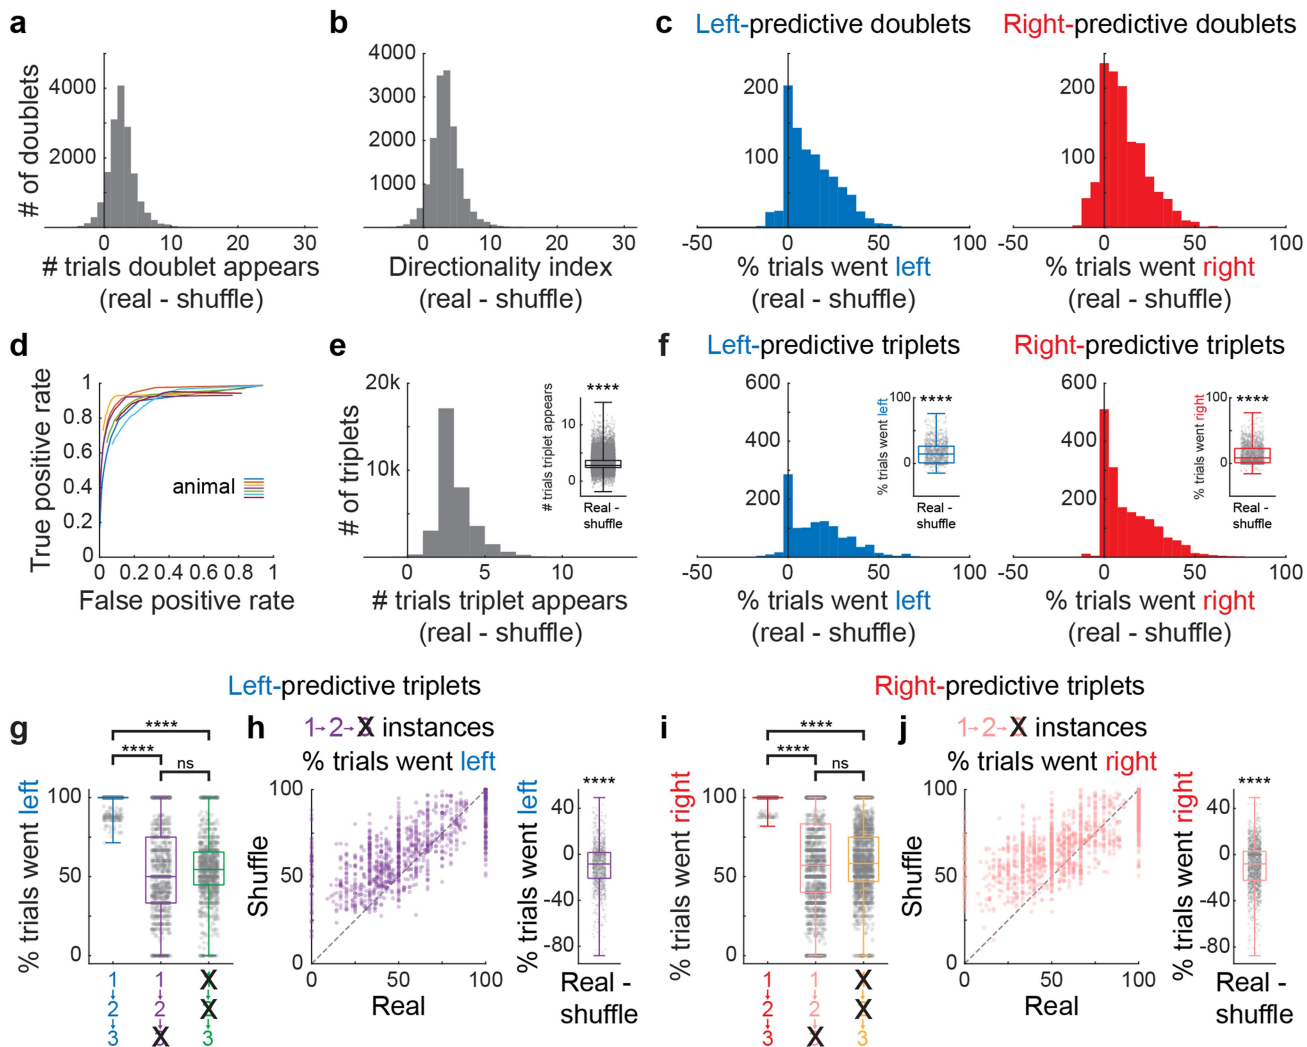
a, Schematic to describe how doublets were defined. Orange and green are calcium traces of the cells under consideration. Grey is the calcium trace of a third cell. **b**, Twenty-five examples of doublets in a single session from one mouse. Each panel shows traces for trials in which the doublet was present.

Orange traces are the neural activity from the first cell in the doublet; green traces are the neural activity from the second cell in the doublet. Heat maps represent the normalized neural activity of each cell across all trials in the session.



Extended Data Fig. 8 | Neural activity generated by trajectories through the task. **a**, Trajectories through evidence and position in one session of the task. Each thin line represents a fit with a cubic spline to a single trial; thick lines represent fits over all trials in which the mouse was supposed to turn left or right. **b**, The average change in position and evidence over time across trials in a single session for a set of representative states in evidence and position space.

c, Conceptual diagram showing four trajectories through the neural manifold in right-choice trials. Two different doublets are activated because the trajectories pass through their firing fields. **d**, Shuffling trial IDs within right-choice trials will disrupt doublet activity while maintaining trial-averaged place and choice preferences of each cell.



Extended Data Fig. 9 | Choice-predictive sequences in CA1 neural activity.
a, Distribution of the values in Fig. 4b. **b**, Distribution of the values in Fig. 4c.
c, Distribution of the values in Fig. 4f. **d**, ROC curves for sequential activity predicted from the five-dimensional embedding of the manifold compared to sequential activity in real data. $n = 7$ mice. **e**, Similar to **a**, but for triplets. Inset shows that triplets are significantly more likely to appear in the real data than in the shuffled dataset in which trial IDs were shuffled. Two-tailed paired Student's t -test, real versus shuffled data, **** $P < 0.0001$; $n = 34,737$ triplets.
f, Similar to **c**, but for triplets, showing that left- and right-choice-predictive triplets from real data are more predictive than triplets obtained from the shuffled dataset in which trial IDs were shuffled. Left inset, left-predictive, $n = 1,135$ triplets, two-tailed paired Student's t -test, real versus shuffled data, **** $P < 0.0001$; right inset, right-predictive, $n = 1,755$ triplets, two-tailed paired Student's t -test, real versus shuffled data, **** $P < 0.0001$. **g**, Left-choice-predictive triplets are significantly more predictive than instances in which the

first two cells in the triplet fire, but the third does not, or when the third cell fires alone. $n = 1,135$ triplets; two-tailed paired Student's t -tests, Bonferroni corrected; $1 \rightarrow 2 \rightarrow 3$ versus $1 \rightarrow 2 \rightarrow \text{not } 3$, **** $P < 0.0001$; $1 \rightarrow 2 \rightarrow 3$ versus not $1 \rightarrow \text{not } 2 \rightarrow 3$, **** $P < 0.0001$; $1 \rightarrow 2 \rightarrow \text{not } 3$ versus not $1 \rightarrow \text{not } 2 \rightarrow 3$, $P = 0.78$. **h**, Notably, for left-choice-predictive triplets, in trials in which cells 1 and 2 fire, but cell 3 does not, significantly more trials end with the mouse turning right than the same instances in the shuffled dataset. $n = 1,135$ triplets, two-tailed paired Student's t -test, real versus shuffled data, **** $P < 0.0001$. **i**, Same as in **g**, but for right-choice-predictive triplets. $n = 1,755$ triplets; two-tailed paired Student's t -tests, Bonferroni corrected; $1 \rightarrow 2 \rightarrow 3$ versus $1 \rightarrow 2 \rightarrow \text{not } 3$, **** $P < 0.0001$; $1 \rightarrow 2 \rightarrow 3$ versus not $1 \rightarrow \text{not } 2 \rightarrow 3$, **** $P < 0.0001$; $1 \rightarrow 2 \rightarrow \text{not } 3$ versus not $1 \rightarrow \text{not } 2 \rightarrow 3$, $P = 1.0$. **j**, Same as in **h**, but for right-choice-predictive triplets. $n = 1,755$ triplets; two-tailed paired Student's t -test, real versus shuffled data, **** $P < 0.0001$. For box plots, boundaries: 25–75th percentiles; midline, median; whiskers, minimum–maximum.

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see [Authors & Referees](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection

Imaging data was collected using ScanImage 2015. Behavioral data was collected using Matlab code (Matlab 2015b, Mathworks Inc) based on the ViRMEn package (<https://pni.princeton.edu/pni-software-tools/virmen-virtual-reality-matlab-engine>). The CNMF algorithm was used as published by Pnevmatikakis et al., Neuron 2016 and mildly adapted for manual post-hoc curation. This process is documented in Koay et al., eLife 2020. MIND is made publicly available, together with all code to reproduce our analyses in the following Github repository: https://github.com/BrainCOGS/HPC_manifolds. ViRMEn was used as published by Aronov et al., Neuron 2014, together with a custom wrapper that allowed for communication with more recent hardware. This is documented in Pinto*, Koay* et al., Frontiers 2018 and publicly available in <https://github.com/sakoay/AccumTowersTools>. For brevity, refer to these papers for details.

Data analysis

Analyses were done using custom Matlab code (Matlab 2018a, Matlab 2018b, Matlab 2020a, Mathworks Inc).

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

The datasets from this study are available from the corresponding author on reasonable request.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

Life sciences Behavioural & social sciences Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/documents/nr-reporting-summary-flat.pdf](https://www.nature.com/documents/nr-reporting-summary-flat.pdf)

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	Sample sizes were chosen to allow for sufficient statistical power for our experimental procedures, while reducing the number of animals to minimize pain and distress. The number of animals we used is typical for a study of this type. As further confirmation, statistical tests were used in all analyses to demonstrate that the sample sizes used had sufficient statistical power.
Data exclusions	Only neurons with at least one calcium transient are able to be identified by constrained non-negative matrix factorization (CNMF). Therefore, cells in the hippocampus that were silent for the entire imaging session were not included.
Replication	We trained a total of 15 animals on three behavioral tasks which we analyzed across animals (n=7 accumulating towers task, n=7 across two types of alternation tasks, n=4 one-side cues task). Some animals were used for multiple tasks, i.e. the one-side cues task was a training stage in the shaping procedure for the accumulating towers task. Statistical tests were used for all analyses to ensure that the findings were significant and unlikely to occur by chance on standard significance levels, as indicated in the figures (* for p<0.05, ** for p<0.01, *** for p<0.001, and **** for p<0.0001).
Randomization	There was only one experimental group, so no randomization was needed.
Blinding	There was only one group, so no group allocation blinding was needed.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> Antibodies
<input checked="" type="checkbox"/>	<input type="checkbox"/> Eukaryotic cell lines
<input checked="" type="checkbox"/>	<input type="checkbox"/> Palaeontology
<input type="checkbox"/>	<input checked="" type="checkbox"/> Animals and other organisms
<input checked="" type="checkbox"/>	<input type="checkbox"/> Human research participants
<input checked="" type="checkbox"/>	<input type="checkbox"/> Clinical data

Methods

n/a	Involved in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> ChIP-seq
<input checked="" type="checkbox"/>	<input type="checkbox"/> Flow cytometry
<input checked="" type="checkbox"/>	<input type="checkbox"/> MRI-based neuroimaging

Animals and other organisms

Policy information about [studies involving animals](#); [ARRIVE guidelines](#) recommended for reporting animal research

Laboratory animals	Male and female mice aged 2 - 18 months expressing GCaMP6f were used in this study. n=5 triple transgenic crosses expressing GCaMP6f under the CaMK11 α promoter from Ai93-D;CaMK11 α -tTA [lgS7tm93.1(tetO-GCaMP6f)Hze Tg(Camk2a-tTA)1Mmay/J, Jackson Laboratories, stock# 024108] and Emx1-IRES-Cre [B6.129S2-Emx1tm1(cre)Krj/J, Jackson Laboratories, stock# 005628] and n=10 Thy1-GCaMP6f [C57BL/6J-Tg(Thy1-GCaMP6f)GP5.3Dkim/J, Jackson Laboratories, stock# 028280]. Animals were housed in reverse light cycle conditions (light cycle: 8PM - 8AM), temperature was 70 \pm 2 $^{\circ}$ F, and humidity was 50 \pm 10%.
Wild animals	No wild animals were used.
Field-collected samples	No field-collected samples were used.
Ethics oversight	All procedures performed in this study were approved by the Institutional Animal Care and Use Committee at Princeton University and were performed in accordance with the Guide for the Care and Use of Laboratory Animals (National Research Council, 2011).

Note that full information on the approval of the study protocol must also be provided in the manuscript.