ENVIRONMENT DIVERSITY DRIVES ZERO-SHOT GENERALISATION OF PREDICTIVE SPATIAL REPRESENTATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

The hippocampus rapidly establishes spatial responses in novel environments. Recurrent networks trained to predict their own sensory inputs develop similar spatial representations, with units tuned to place and head direction, but are typically trained and tested in a single environment. How such networks come to transfer to novel environments, and which parts of their representations carry over, is largely unexplored. Here we varied the diversity of training environments to address both questions. A recurrent network was trained only to predict its next view from egocentric views and self-motion, in 1, 5, 20, or 40 of 50 procedurally generated environments that differed in shape, size, texture, and landmarks. The weights were then fixed and the network tested in the remaining, unseen environments. Spatially and directionally tuned units emerged in every network, but generalisation depended on training diversity. The increase in prediction error from seen to unseen environments fell from 89% after training in one environment to 3% after 40 and, with the total amount of training data matched, was 5% after training in 40 environments against 50% in five. In nearly every network, a ring of directionally tuned units rotated together between environments, as head direction cells in the brain do. Their preferred directions were bunched after training in one or five environments but tiled the circle more evenly after training in 20 or 40, and this evenness predicted generalisation across training conditions. Place units tended to rotate with the ring between environments, more closely than with environment geometry. An evenly spread head direction signal, conserved across species and present early in development, may therefore be a spatial primitive that allows maps to form rapidly in novel environments. We predict that, between differently shaped environments, part of the arrangement of place fields is preserved, rotated with head direction cells.

1 INTRODUCTION

The hippocampus is widely held to function as a cognitive map (O'Keefe & Nadel, 1978), with place cells, neurons with spatially stable firing fields, as its neural basis (O'Keefe & Dostrovsky, 1971). Head direction cells, which encode the direction the animal is facing (Taube et al., 1990), grid cells (Hafting et al., 2005), and border-responsive neurons (Barry et al., 2006; Solstad et al., 2008; Lever et al., 2009) belong to a wider suite of spatially modulated neurons that support this map. On first exposure to a novel environment, place fields are established rapidly, often on the first pass through a location and typically within minutes (Karlsson & Frank, 2008; Barry et al., 2012; Dong et al., 2021). These fields are largely unrelated to those in familiar environments, a phenomenon known as remapping (Muller & Kubie, 1987), yet remapping is less complete and less random than often assumed (Whittington et al., 2020; Fenton, 2024). When the shape of an enclosure changes, place fields move in ways that can be predicted from the positions of its boundaries (O'Keefe & Burgess, 1996; Hartley et al., 2000). Head direction cells, in contrast, behave as a ring attractor (Zhang, 1996). Their preferred directions keep fixed offsets from one another, so the population rotates as a whole between environments (Yoganarasimha et al., 2006; Peyrache et al., 2015). This coherence is present in rat pups before eye opening (Bjerknes et al., 2015; Tan et al., 2015; Bassett et al., 2018),

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

and ring-like head direction systems are found across species, from flies (Kim et al., 2017) to bats (Finkelstein et al., 2015), suggesting a conserved scaffold for mapping new environments.

Normative models suggest how such maps could be learned. Following the view of the hippocampus as a predictive map (Stachenfeld et al., 2017), recurrent networks trained to predict their upcoming sensory input develop units resembling place and head direction cells (Recanatesi et al., 2021; Uria et al., 2020; Gornet & Thomson, 2024; Levenstein et al., 2024; Abrate et al., 2026), as do models trained on related spatial objectives (Banino et al., 2018; Cueva & Wei, 2018; Sorscher et al., 2023; Franzius et al., 2007). Most of this prior work trains and analyses networks in a single environment, leaving open whether they learn a map of that environment or a way of rapidly building maps of new ones. The exceptions do not isolate this question: a predictive model trained in three enclosures kept head direction and boundary tuning across them, but had a separate memory for each enclosure and a visual encoder pre-trained on the test enclosures (Uria et al., 2020), and models that generalise structure across many environments do so in discrete graphs rather than from continuous egocentric views (Whittington et al., 2020). In reinforcement learning, agents generalise better to new procedurally generated environments the more environments they are trained in (Cobbe et al., 2020; Kirk et al., 2023), but how training diversity affects transfer in self-supervised spatial representations, and which parts of those representations carry over, remains largely unexplored.

Here we trained recurrent networks (Abrate et al., 2026) to predict the next egocentric view from the current view and self-motion, in 1, 5, 20, or 40 out of 50 procedurally generated polygonal environments that differed in shape, size, wall and floor texture, and landmark placement. Self-motion was given as egocentrically, so no reference frame shared across environments was supplied, and the only training signal was next-view prediction error. With their weights fixed, the networks were then tested in the remaining, unseen environments. Our contributions are as follows.

- **Training diversity drives generalisation.** The increase in prediction error from seen to unseen environments fell from 89% when trained in one environment to 3% when trained in 40, with experience per environment held fixed. Training on 40 environments still generalised better than on 5 when the total amount of training data was matched (5% against 50%).
- **Training diversity spreads a coherent directional ring evenly.** Regardless of how many environments a network was trained in, a subset of directionally tuned units rotated together between environments, as head direction cells in the brain do, but their preferred directions tiled the circle more evenly after training in more environments, and how evenly they did so predicted generalisation across arms (Spearman $\rho = 0.66$, 19 networks). In 40-environment networks, clamping these units to their mean activation at test time raised prediction error 1.7-3.0 times as much as clamping control units matched in tuning class and Rayleigh vector length (mean 2.2, five networks), and more than any of 20 control sets in every case.
- **A procedurally generated testbed for spatial generalisation.** Fifty polygonal environments with randomised geometry, textures, and landmarks, rendered along rodent-like trajectories (George et al., 2024; Abrate, 2026), allow the number of training environments to be varied with the model fixed. We will release the code and environments.
- **A testable prediction for place cells.** Place units tended to rotate with the ring more closely than with environment geometry (66-81% of place units in each of five networks; Wilcoxon $p < 10^{-8}$ in each). This predicts that, between differently shaped environments, part of the arrangement of place fields is preserved, rotated with head direction cells.

2 SETUP

Environments. We generated 50 polygonal environments procedurally (Fig. 1a). Each was a polygon with 4-9 walls, no interior angle below 45°, and a largest width or depth of 0.8-1.2 m, placed at a random orientation so that walls did not share directions across environments, and areas varied up to 4.25-fold. Walls and floor carried textures, and 4-7 walls carried landmarks (large visual patterns), drawn from shared pools, so the same patterns recurred across environments but never at a consistent position (Appendix A). In each environment the agent explored for 2,000 s, sampled at 10 Hz, along smooth, rodent-like trajectories generated with RatInABox (George et al., 2024) (Fig. 1c; parameters in Appendix A). At each step we rendered a $240^\circ \times 120^\circ$ greyscale view of

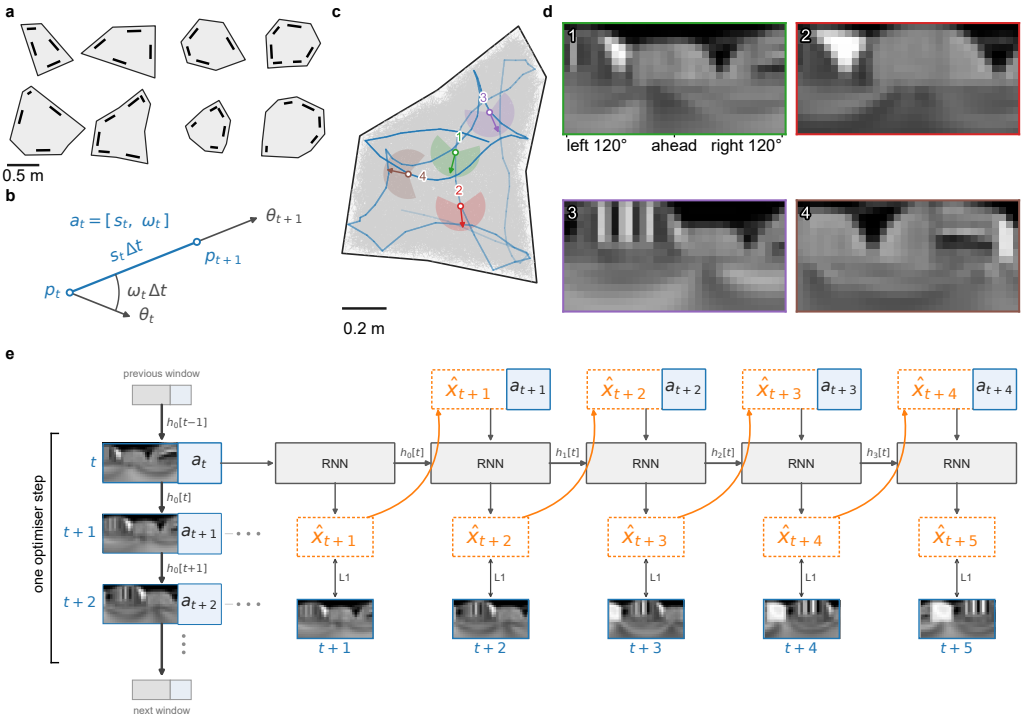

Figure 1: **Networks were trained to predict upcoming views in procedurally generated environments. a** Outlines of eight of the 50 environments; bars show the number of landmarks in each (positions schematic). **b** Self-motion $\mathbf{a}_t = [s_t, \omega_t]$: speed and angular velocity between consecutive positions $\mathbf{p}_t$ and headings θ_t. **c** One environment with its full trajectory (grey), a 10-s excerpt (blue), and four numbered positions with their fields of view. **d** Views from the four positions in **c** (colours as in **c**), spanning 120° either side of the heading, at the network's 32×16 input resolution. **e** Training: each 10-step window was unrolled from the hidden state carried over from the previous window; from every step the network predicted five views ahead autoregressively, receiving its own predictions ($\hat{\mathbf{x}}$, orange) and the true self-motion ($\mathbf{a}$, blue).

128×64 pixels from 3.5 cm above the floor with the GPU ray caster of RatVision (Abrate, 2026). Views were blurred and downsampled to 32×16 pixels to reduce input dimensionality (Fig. 1d; Appendix A).

Model. The network was a single-layer recurrent network of 500 sigmoid units with a linear readout (Fig. 1e), similar to that of Abrate et al. (2026). At each 0.1 s step it received the current view $\mathbf{x}_t$ (512 pixels) and self-motion $\mathbf{a}_t = [s_t, \omega_t]$ (speed and angular velocity; Fig. 1b), and predicted the next view:

$$\mathbf{h}_{t+1} = \sigma\big(\mathbf{W}_{\mathrm{ih}}\,[\,\mathbf{x}_t;\, g(\mathbf{a}_t)\,] + \mathbf{W}_{\mathrm{hh}}\,\mathbf{h}_t + \mathbf{b}_{\mathrm{h}}\big), \qquad \hat{\mathbf{x}}_{t+1} = \mathbf{W}_{\mathrm{ho}}\,\mathbf{h}_{t+1} + \mathbf{b}_{\mathrm{o}}, \tag{1}$$

$$g(\mathbf{a}_t) = \tanh(\mathbf{W}_a\,\mathbf{a}_t + \mathbf{b}_a), \tag{2}$$

where σ is the sigmoid function and g a learned 16-unit embedding that prevented the two self-motion inputs being swamped by the 512 pixels. Because self-motion was egocentric and each environment's orientation and start position were random, any reference frame shared across environments had to be inferred from vision. Training used backpropagation through time on 10-step windows, and from every step the network predicted five steps ahead autoregressively, given the true self-motion, with an L_1 loss weighted equally across the five steps (Fig. 1e). The hidden state was carried across consecutive windows. No allocentric information, such as position, heading, or environment variables, was given to the model. Optimisation details are in Appendix B.

Training arms and splits. In each arm the network was trained on a different number of environments, our measure of training diversity (1, 5, 20, or 40 at 3, 3, 3, and 5 seeds, details in Appendix B),

drawn at random from the pool of 50 for each seed, and the remaining environments were unseen and used only for testing. In each seen environment, the final 20% of the trajectory was held out from training as one contiguous block, to test prediction in seen environments. Each environment's trajectory was cut into 100 s segments, each starting from a zero hidden state, and the segments of all environments were stacked along the batch axis, so the four main arms had the same number of optimiser steps and training windows per environment, but different batch sizes. Arms with more seen environments therefore saw more data in total. The matched-data arm (three seeds) controlled for this: networks were trained in 40 of a separately generated set of 50 environments with 255 s rather than 2,000 s trajectories, supplying about as many training windows as the five-environment arm (80 segments per batch against 75; 300,000 optimiser steps in every arm; Fig. A2). We defined the generalisation gap as the relative increase in prediction error from seen to unseen environments, $G = (\mathcal{L}_{\text{unseen}} - \mathcal{L}_{\text{seen}})/\mathcal{L}_{\text{seen}}$, expressed as a percentage, where $\mathcal{L}_{\text{seen}}$ was the mean one-step L_1 loss over the held-out blocks of the seen environments and $\mathcal{L}_{\text{unseen}}$ the mean over the trajectories of the unseen environments, both also evaluated in 100 s segments. As a reference, using the current view as the prediction gave $\mathcal{L}_{\text{unseen}} = 0.120 \pm 0.004$ (mean $\pm$ SD over the five 40-environment networks), a baseline that differed due to environment shuffle.

Analysis. After training, each network was run over the full trajectory in all 50 environments, including, for seen environments, the training portion (which favours them below). For each unit and environment we built a rate map of mean activation per spatial bin (25×25 bins of 4.8 cm over a fixed 1.2 m square; Gaussian smoothing, $\sigma = 0.75$ bins, normalised by smoothed occupancy) and a 60-bin head direction tuning curve, whose Rayleigh vector (the activation-weighted mean of unit vectors in each direction) gave the preferred direction (its angle) and tuning strength (its length, 0-1). Fields (Fig. 3c) were detected with an iterative threshold on each rate map (Appendix D). Units were classified in each environment following Abrate et al. (2026): place units had spatial information (Skaggs et al., 1993) above 0.3; head direction units had directional information above 0.2 or a Rayleigh vector length above 0.3; conjunctive units met both. Place units that met the criterion in at least half of the environments were classed as consistent. Map stability was the Pearson correlation, across spatial bins, between rate maps from the two halves of each trajectory (Muller & Kubie, 1987). Linear decoders of position and heading (ordinary least squares from all hidden units; heading error was the absolute angular difference, wrapped to at most $180°$) were fitted on the first half of each trajectory and tested on the second.

Directional units were those whose mean Rayleigh vector length across all 50 environments exceeded 0.15, a less stringent criterion than for head direction units so that conjunctive and unclassified units could qualify. The rotation $\Delta\theta$ of a unit between two environments was the change in preferred direction. The common rotation φ of a set of units was the mode of their $\Delta\theta$, and the residual of a unit was $\Delta\theta - \varphi$, near zero for a unit that rotated with the rest. Ring units were the 25 directional units with the highest co-rotation score, the consistency of their residuals from the common rotation of all directional units across pairs of 25 selection environments, and all ring statistics were computed in the other 25, the evaluation environments (Appendix E). Ring co-rotation was the resultant length of the $\Delta\theta$ of ring units (1 if all rotated by the same angle), averaged over pairs of environments. Ring tiling was the Kullback-Leibler (KL) divergence from uniform of the histogram of preferred directions of ring units (lower when more even), compared with random sets of 25 directions. To test whether prediction depended on the ring, we clamped ring units to their mean activation at test time and compared the increase in loss with that for 20 control sets matched in tuning class and Rayleigh vector length.

To test whether place units rotated with the ring, we took the bearing of the main field (the field containing the peak) of each consistent place unit outside the ring, from the centre of the environment, averaged over maps built separately for eight heading sectors to remove heading dependence (Appendix D). Like all angles, bearings were measured from a fixed axis of the simulator coordinates, in which each environment was generated at a random orientation, and, unlike field positions, can be compared between environments of different shapes. For example, a field at $30°$ from the centre in one environment and $120°$ in another has rotated by $90°$, and if the ring also rotated by $90°$, the unit rotated with the ring. A unit tied to environment shape would instead follow the geometric rotation, the rotation (in $15°$ steps) that maximised the overlap of the two floor areas. Agreement with each rotation was the resultant length, across pairs of evaluation environments, of the difference between the change in bearing and that rotation, compared with shuffled unit identities.

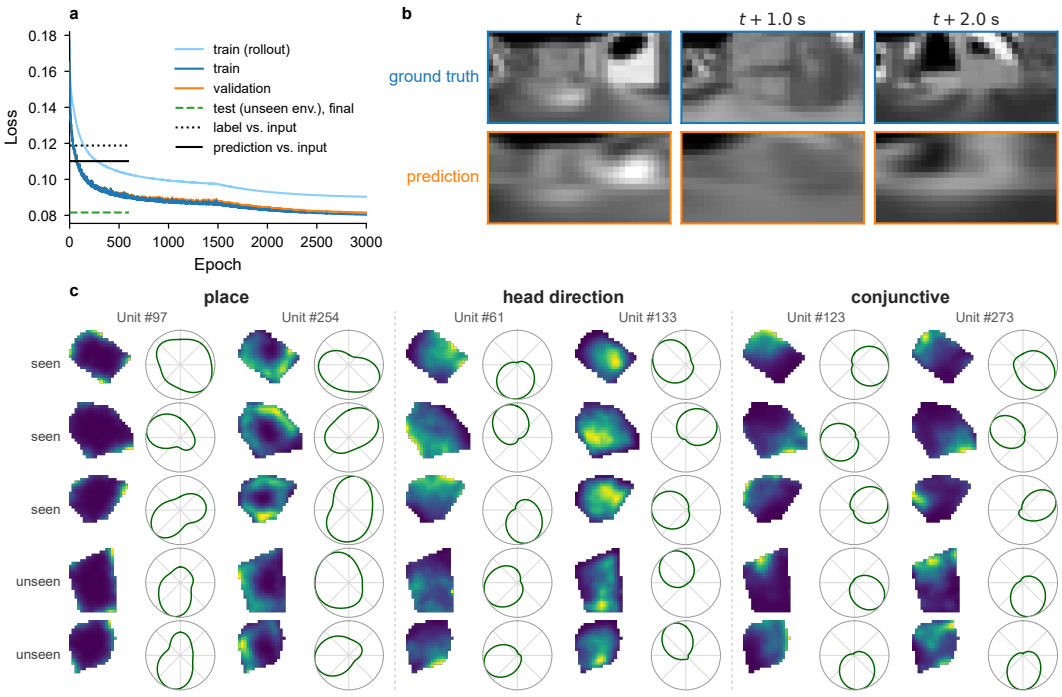

Figure 2: **The network learned to predict upcoming views and developed spatially and directionally tuned units.** Results from the network trained in 40 environments (seed 12). **a** L_1 loss during training. Only the five-step rollout loss on the training data (*train (rollout)*) was optimised; the one-step loss on the same data (*train*) is shown for comparison with the one-step loss on the held-out final 20% of each trajectory in seen environments (*validation*). Dashed green line, final loss in unseen environments (*test*). Dotted black line, loss when the current view is used as the prediction (*label vs. input*, unseen environments). Solid black line, distance between prediction and current view (*prediction vs. input*), close to the dotted line, so the network did not simply copy the input. **b** Views at t, $t + 1$ s, and $t + 2$ s from a held-out portion of a seen environment (top, blue) and the network's one-step predictions (bottom, orange). **c** Rate maps, each scaled to its own range (dark blue, low; yellow, high) and polar tuning curves, scaled to each curve's maximum, of two place, two head direction, and two conjunctive units (columns) consistently classified in three seen and two unseen environments (rows).

3 Results

3.1 Generalisation improved with the number of training environments

Every network learned to predict upcoming views. In the 40-environment network, loss fell to 0.08 in seen and unseen environments alike (generalisation gap $+0.1\%$; $+3.3 \pm 3.7\%$ across all five), below the 0.120 ± 0.004 (five seeds) obtained by using the current view as the prediction (Fig. 2a), and units with spatial, directional, or conjunctive tuning emerged in every arm (Table A1; Fig. 2c). The generalisation gap, the relative increase in prediction error from seen to unseen environments, fell from 89% when training in one environment to 50%, 22%, and 3% when training in 5, 20, and 40 (mean $\pm$ SD over 3, 3, and 5 networks; Fig. 3a; Spearman $\rho = -0.96$ across 14 networks, $p = 3 \times 10^{-8}$). Networks trained in 40 environments with the five-environment data budget generalised like the other 40-environment networks, not like the five-environment ones ($5 \pm 2\%$, three networks; against five environments, Welch $t = -24.3$, $p = 3 \times 10^{-5}$; against 40 environments, $p = 0.48$), as did networks with 128 rather than 500 hidden units ($2 \pm 4\%$, three networks; $p = 0.78$). After training in one or five environments, the loss in unseen environments (0.185 and 0.137) exceeded that of using the current view as the prediction (0.124). With data matched, the 40-environment networks had lower loss in unseen environments than the five-environment ones (0.116 against 0.137) despite higher loss in seen ones (0.111 against 0.091), with unseen loss below the 0.135 ± 0.001 obtained

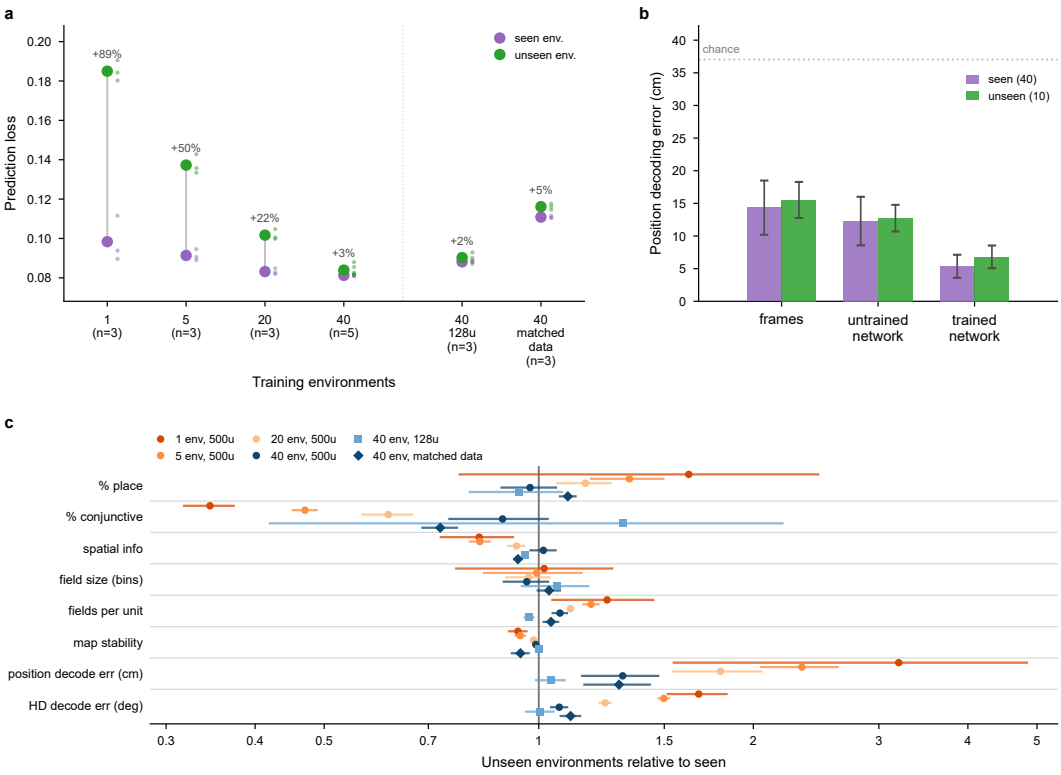

Figure 3: **Generalisation to unseen environments improved with the number of training environments. a** Prediction loss in seen (purple; held-out final 20%) and unseen (green) environments after training in 1, 5, 20, or 40 environments (500 hidden units) and, right of the vertical line, in 40 with 128 units (128u) or with total training data matched to the five-environment arm (matched data). Large dots, mean; small dots, seeds; percentages, generalisation gap. **b** Error of a linear position decoder, fitted separately in each environment, applied to the input views, an untrained network, and the 40-environment network, in seen and unseen environments (colours as in **a**; mean ± 1 SD over environments, *trained network*: seed 12); dotted line, chance. **c** Each per-environment metric in unseen environments relative to seen (1, no change; log scale), for every arm; markers, mean over seeds ($n = 3$; 5 for 40 environments, 500 units); bars, ± 1 SD.

by using the current view as the prediction there. Training in more environments therefore improved prediction in unseen environments, an advantage that did not depend on seeing more data.

3.2 Position decoding and tuning statistics in unseen environments

In the 40-environment network, a linear decoder fitted in each environment recovered position in unseen environments with an error of 6.8 cm, 27% higher than in seen environments (5.4 cm) but well below one fitted on the input views alone (15.5 cm), the activity of an untrained network (12.7 cm), or chance (37 cm; Fig. 3b). In the same network, none of the other statistics in Fig. 3c differed detectably between seen and unseen environments (largest difference 5%; bootstrap over environments, Holm-corrected, all $p > 0.5$). Across five seeds, these statistics stayed within 11% of their seen-environment values (Fig. 3c), With more training environments, these statistics in unseen environments moved closer to those in seen ones.

3.3 Diverse training spread a coherent directional population into a ring

Regardless of how many environments a network was trained in, a subset of directional units rotated together between environments, similar to head direction cells in the brain (Figs 4a and A5). These ring units, the 25 directional units that rotated most consistently with the directional population

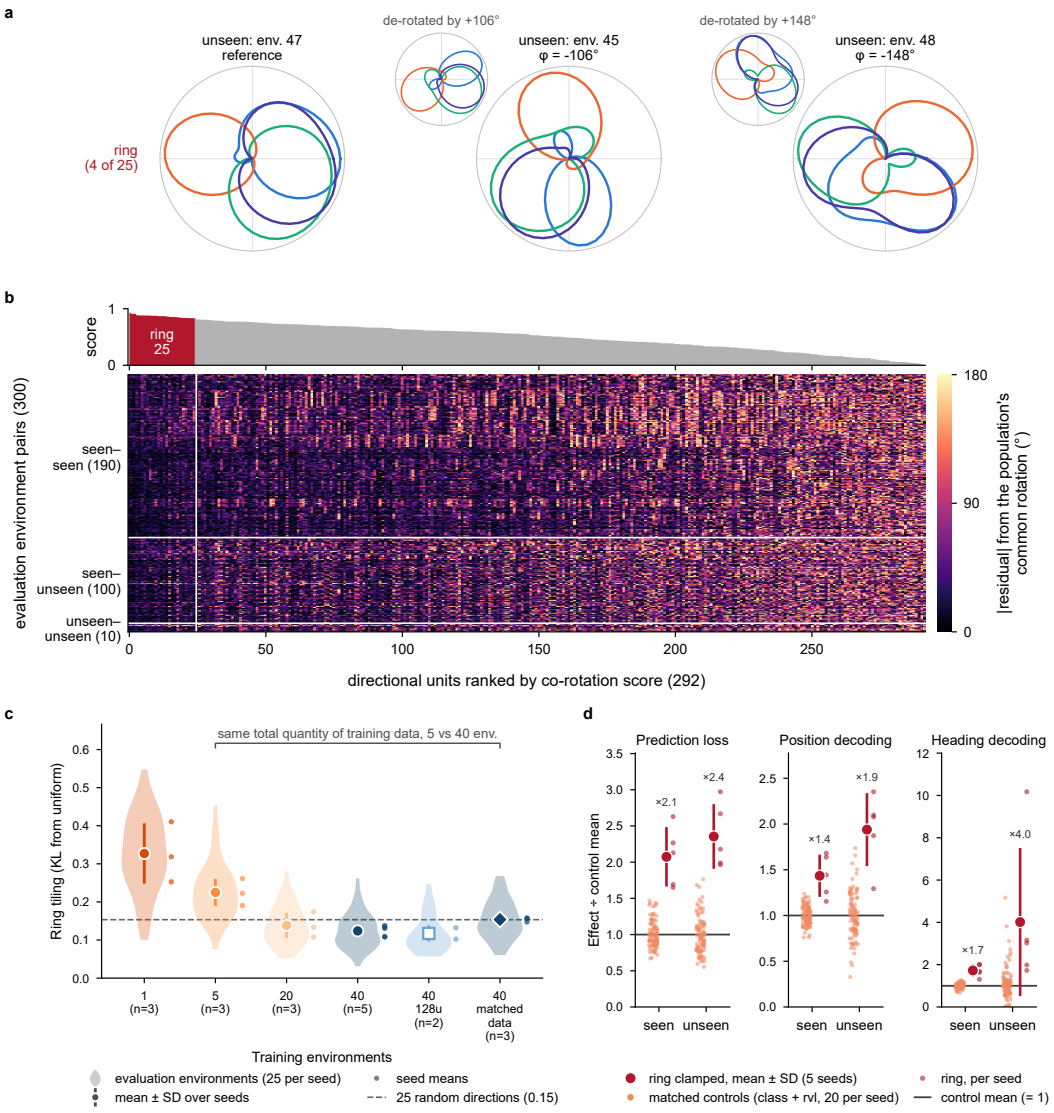

Figure 4: **A subset of directional units rotated together between environments, diverse training spread their preferred directions, and clamping these units impaired prediction.** The 40-environment network (seed 12) unless stated. **a** Head direction tuning curves (each scaled to its own range) of four of the 25 ring units in three unseen environments; φ, common rotation of the ring relative to the first; insets, curves rotated back by φ. **b** Absolute residual of each of the 292 directional units (columns, ranked by co-rotation score in the selection environments; bar above) from their common rotation, in 300 pairs of evaluation environments (rows; darker, smaller residual). The 25 highest-scoring units (left) formed the ring and had the smallest residuals. **c** Ring tiling (KL divergence from uniform, 8 bins) in the evaluation environments of every seed (violins; 3 seeds per arm, 5 for 40 environments, and 2 for 128 units, whose third seed had only 24 directional units); dots, seed means; markers and bars, mean ± SD; dashed line, mean for random sets of 25 directions (0.15; 5th-95th percentile 0.07-0.27). Ring units were bunched after training in 1 or 5 environments (mean KL 0.33 and 0.23) and near the random mean after 20 or 40 (0.12-0.14), as in the 128-unit (square; hollow, as nearly all directional units were ring units) and matched-data (diamond, 0.15) networks. **d** Increase in loss, and in position and heading decoding error from the remaining units, when the 25 ring units (red) or each of 20 control sets matched in modal tuning class and Rayleigh vector length (orange; line, mean) were clamped to their mean activation within the recurrence; numbers, ratio to control mean.

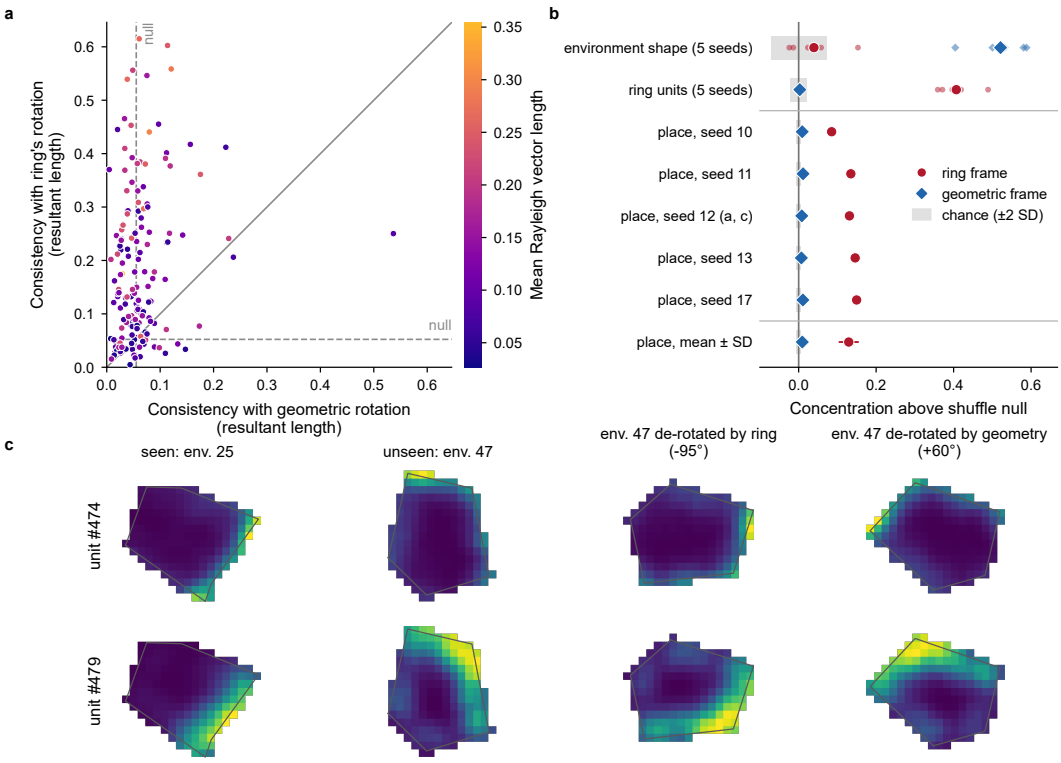

Figure 5: **Place fields rotated with the ring more closely than with environment geometry. a** Agreement of the field bearing of each consistent place unit with ring rotation (y) and the geometric rotation (x) across pairs of evaluation environments (seed 12); colour, mean directional tuning strength; dashed lines, shuffle nulls. **b** Agreement above the shuffle null with ring rotation (circles) and the geometric rotation (diamonds) in five seeds; grey bars, chance (± 2 SD). Positive controls: *environment shape*, the long axis of each outline in place of a field bearing, followed the geometric rotation, so the test detects geometric anchoring where it exists; *ring units*, the fields of the 25 ring units, followed the ring. **c** Two example units in environments 25 (seen) and 47 (unseen); undoing ring rotation, not geometric rotation, aligned the map in 47 with that in 25.

(Fig. 4b), were selected regardless of tuning class and typically included roughly 3 head direction, 14 conjunctive, 6 place, and 2 unclassified units (ranges 0-8, 8-19, 3-13, and 1-6 across the 17 networks with 500 units; each unit classified by its most common class across environments). In the 40-environment network their co-rotation in the evaluation environments was 0.86, against a ceiling of 0.99 from split halves of the data and a chance level of 0.18 from shuffled unit identities. In every arm it was 0.76-0.89, including after training in a single environment (Fig. A5b). Of the residuals of ring units from their common rotation, $43 \pm 2\%$ were within $15°$, against $27 \pm 1\%$ for matched control units and 15% after shuffling (Fig. A5c).

Training diversity instead changed how the preferred directions of these units were spread around the circle. After training in one or five environments these were bunched (mean KL 0.33 and 0.23, against 0.15 for random directions; 36-96% and 4-40% of evaluation environments per network above the 95th percentile for random directions, against 5% expected), whereas after 20 or 40 they were as even as random directions (0.12-0.14; Fig. 4c; Mann–Whitney $p = 3 \times 10^{-4}$). With data matched, 40 environments still gave a more evenly spread ring than five (KL 0.15 against 0.23).

3.4 Ring tiling predicted generalisation; place units followed the ring

Across the networks with enough directional units to form a 25-unit ring (19/20), the more evenly the preferred directions of ring units were spread, the smaller the generalisation gap (ring tiling against gap, Spearman $\rho = 0.66$, $p = 0.002$). The relation held between arms but not within

them ($\rho = 0.09$, $p = 0.73$) and did not reflect the amount of training data, as the matched-data networks had an even ring and a small gap. In the 40-environment networks, clamping ring units raised prediction error 2.1 ± 0.4 and 2.4 ± 0.5 times as much in seen and unseen environments, respectively, as clamping matched controls (mean $\pm$ SD over five networks; Fig. 4d).

We next asked whether consistent place units rotated with the ring. Between pairs of evaluation environments, their field bearings rotated with the ring more closely than after shuffling unit identities ($z = 30$-53 against 200 shuffles in each of five networks), and for 66-81% of place units in each of five networks they followed the ring more closely than the geometric rotation (Wilcoxon $p < 10^{-8}$ in each; Fig. 5a,b). Undoing the rotation of the ring aligned the map of a unit in an unseen environment with that in a seen one, whereas undoing the geometric rotation did not (Fig. 5c). The rotation of the ring was itself unrelated to the geometric rotation (the two agreed within 30° in 11-20% of evaluation pairs, against 17% expected by chance), so this agreement did not reflect environment geometry. Thus spatial maps rotated far more closely with the ring than with environment shape, so the ring could provide a shared reference that aligns maps across environments.

4 Discussion

We explored how exposure to diverse training environments allows networks trained to predict their next view to generalise zero-shot to unseen environments. We found that transfer depended on the diversity of the training data rather than the amount. In the hidden layer, all networks developed place, head direction, and conjunctive units, including a ring of directional units that rotated together between environments, as head direction cells do. Only after training in diverse environments did the preferred directions of these units spread evenly around the circle, and this evenness predicted generalisation. Place units followed the ring more closely than the geometry of the environments, suggesting that the ring anchors spatial firing across environments.

The two properties of the ring seem to have different origins. Ring units rotated together even after training in a single environment, perhaps because of the angular velocity input, as recurrent networks integrating angular velocity develop head direction representations (Cueva et al., 2020; Uria et al., 2020). Even tiling, in contrast, required diverse training, and the correlation between tiling and generalisation held across arms but not within them. Because place units rotated with the ring, an evenly spread ring could orient spatial maps in any environment, whereas a bunched ring would represent some headings poorly.

The rodent head direction system is coherent before eye opening (Bjerknes et al., 2015; Tan et al., 2015; Bassett et al., 2018), matures before other spatial cell types (Langston et al., 2010; Wills et al., 2010), and is needed for stable place and grid cells (Calton et al., 2003; Winter et al., 2015), suggesting a spatial primitive on which other representations build. Our results are consistent with this: place units rotated with the ring more closely than with environment geometry, and the evenness of the ring predicted generalisation to unseen environments. It is therefore plausible that an intact head direction signal is a prerequisite for place fields to form rapidly in novel environments (Karlsson & Frank, 2008; Dong et al., 2021). Equally, given the propensity of place cells to rotate with head direction cells when cues rotate within an environment (Knierim et al., 1995), the model predicts that, even between environments of different shape, part of the arrangement of place fields should be preserved, rotated by the same angle as head direction cells. This preserved arrangement would be revealed once maps are realigned by that rotation.

Limitations. Our setting was minimal: a single-layer network with a linear readout, low-resolution greyscale views, and small, flat environments from a single generator, so transfer to out-of-distribution environments is untested. Each arm had three to five seeds, generalisation was measured on one-step prediction, and whether place units depend on the ring was not tested.

Conclusion. Networks trained to predict their next view formed a coherent directional population, but only after diverse training did the preferred directions of these units cover all directions evenly, and this evenness predicted generalisation to unseen environments. For predictive models, transfer may depend on the breadth rather than the amount of training. For the brain, an early heading signal covering every direction may scaffold the rapid mapping of novel environments.

AI USE STATEMENT

In this work, we used generative AI tools (Claude, Anthropic) to give feedback on methodology and on the interpretation of results, and to help write code, which the authors checked by hand. We did not use them to generate data or to develop the models or hypotheses; the other tasks requiring disclosure are not applicable to this work. Additionally, we used them to search and summarise the literature, check citations, suggest a structure for the paper, and draft and edit text for readability. We have reviewed all AI-assisted work and take responsibility for the final content of this paper, including text, claims, or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

All authors adhere to the ICLR Code of Ethics. This work uses only simulated agents and involves no human or animal subjects or personal data.

REPRODUCIBILITY STATEMENT

Section 2 describes the environments, network, training, and analyses; Appendix A gives the parameters of the environments, trajectories, and rendering, and Appendix B the optimisation, seeds, checkpoint selection, and compute. Code and instructions for reproducibility will be attached along with this submission. Afterwards, the full code will be made public on GitHub at [anonymised URL].

REFERENCES

Marco P Abrate. RatVision, 2026. Code at https://github.com/marcoabrate/ratvision.

Marco P. Abrate, Laurenz Muessig, Joshua P. Bassett, Hui Min Tan, Francesca Cacucci, Thomas J. Wills, and Caswell Barry. From movement to cognitive maps: recurrent neural networks reveal how locomotor development shapes hippocampal spatial coding. In *International Conference on Learning Representations*, 2026.

Andrea Banino, Caswell Barry, Benigno Uria, Charles Blundell, Timothy Lillicrap, Piotr Mirowski, Alexander Pritzel, Martin J. Chadwick, Thomas Degris, Joseph Modayil, et al. Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557(7705):429–433, 2018.

Caswell Barry, Colin Lever, Robin Hayman, Tom Hartley, Stephen Burton, John O'Keefe, Kate Jeffery, and Neil Burgess. The boundary vector cell model of place cell firing and spatial memory. *Reviews in the Neurosciences*, 17(1–2):71–97, 2006.

Caswell Barry, Lin Lin Ginzberg, John O'Keefe, and Neil Burgess. Grid cell firing patterns signal environmental novelty by expansion. *Proceedings of the National Academy of Sciences*, 109(43): 17687–17692, 2012. doi: 10.1073/pnas.1209918109.

Joshua P. Bassett, Thomas J. Wills, and Francesca Cacucci. Self-organized attractor dynamics in the developing head direction circuit. *Current Biology*, 28(4):609–615, 2018.

Tale L. Bjerknes, Rosamund F. Langston, Ingvild U. Kruge, Edvard I. Moser, and May-Britt Moser. Coherence among head direction cells before eye opening in rat pups. *Current Biology*, 25(1): 103–108, 2015.

Jeffrey L. Calton, Robert W. Stackman, Jeremy P. Goodridge, William B. Archey, Paul A. Dudchenko, and Jeffrey S. Taube. Hippocampal place cell instability after lesions of the head direction cell network. *Journal of Neuroscience*, 23(30):9719–9731, 2003. doi: 10.1523/JNEUROSCI.23-30-09719.2003.

Karl Cobbe, Christopher Hesse, Jacob Hilton, and John Schulman. Leveraging procedural generation to benchmark reinforcement learning. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 2020.

Christopher J. Cueva and Xue-Xin Wei. Emergence of grid-like representations by training recurrent neural networks to perform spatial localization. In *International Conference on Learning Representations*, 2018.

Christopher J. Cueva, Peter Y. Wang, Matthew Chin, and Xue-Xin Wei. Emergence of functional and structural properties of the head direction system by optimization of recurrent neural networks. In *International Conference on Learning Representations*, 2020.

Can Dong, Antoine D. Madar, and Mark E. J. Sheffield. Distinct place cell dynamics in CA1 and CA3 encode experience in new environments. *Nature Communications*, 12:2977, 2021.

André A. Fenton. Remapping revisited: how the hippocampus represents different spaces. *Nature Reviews Neuroscience*, 25:428–448, 2024.

Arseny Finkelstein, Dori Derdikman, Alon Rubin, Jakob N. Foerster, Liora Las, and Nachum Ulanovsky. Three-dimensional head-direction coding in the bat brain. *Nature*, 517(7533):159–164, 2015.

Mathias Franzius, Henning Sprekeler, and Laurenz Wiskott. Slowness and sparseness lead to place, head-direction, and spatial-view cells. *PLoS Computational Biology*, 3(8):e166, 2007.

Tom M. George, Mehul Rastogi, William de Cothi, Claudia Clopath, Kimberly Stachenfeld, and Caswell Barry. RatInABox, a toolkit for modelling locomotion and neuronal activity in continuous environments. *eLife*, 13:e85274, 2024.

James Gornet and Matt Thomson. Automated construction of cognitive maps with visual predictive coding. *Nature Machine Intelligence*, 6:820–833, 2024.

Torkel Hafting, Marianne Fyhn, Sturla Molden, May-Britt Moser, and Edvard I. Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806, 2005.

Tom Hartley, Neil Burgess, Colin Lever, Francesca Cacucci, and John O'Keefe. Modeling place fields in terms of the cortical inputs to the hippocampus. *Hippocampus*, 10:369–379, 2000.

Mattias P. Karlsson and Loren M. Frank. Network dynamics underlying the formation of sparse, informative representations in the hippocampus. *Journal of Neuroscience*, 28(52):14271–14281, 2008.

Sung Soo Kim, Hervé Rouault, Shaul Druckmann, and Vivek Jayaraman. Ring attractor dynamics in the Drosophila central brain. *Science*, 356(6340):849–853, 2017.

Robert Kirk, Amy Zhang, Edward Grefenstette, and Tim Rocktäschel. A survey of zero-shot generalisation in deep reinforcement learning. *Journal of Artificial Intelligence Research*, 76:201–264, 2023.

James J. Knierim, Hemant S. Kudrimoti, and Bruce L. McNaughton. Place cells, head direction cells, and the learning of landmark stability. *Journal of Neuroscience*, 15(3):1648–1659, 1995.

Rosamund F. Langston, James A. Ainge, Jonathan J. Couey, Cathrin B. Canto, Tale L. Bjerknes, Menno P. Witter, Edvard I. Moser, and May-Britt Moser. Development of the spatial representation system in the rat. *Science*, 328(5985):1576–1580, 2010. doi: 10.1126/science.1188210. URL https://www.science.org/doi/abs/10.1126/science.1188210.

Daniel Levenstein, Aleksei Efremov, Roy Henha Eyono, Adrien Peyrache, and Blake Richards. Sequential predictive learning is a unifying theory for hippocampal representation and replay. *bioRxiv*, pp. 2024.04.28.591528, 2024.

Colin Lever, Stephen Burton, Ali Jeewajee, John O'Keefe, and Neil Burgess. Boundary vector cells in the subiculum of the hippocampal formation. *Journal of Neuroscience*, 29(31):9771–9777, 2009.

Robert U. Muller and John L. Kubie. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, 7(7):1951–1968, 1987.

John O'Keefe and Neil Burgess. Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381:425–428, 1996.

John O'Keefe and Jonathan Dostrovsky. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34(1):171–175, 1971.

John O'Keefe and Lynn Nadel. *The Hippocampus as a Cognitive Map*. Clarendon Press, Oxford, 1978.

Adrien Peyrache, Marie M. Lacroix, Peter C. Petersen, and György Buzsáki. Internally organized mechanisms of the head direction sense. *Nature Neuroscience*, 18(4):569–575, 2015.

Stefano Recanatesi, Matthew Farrell, Guillaume Lajoie, Sophie Deneve, Mattia Rigotti, and Eric Shea-Brown. Predictive learning as a network mechanism for extracting low-dimensional latent space representations. *Nature Communications*, 12:1417, 2021.

William E. Skaggs, Bruce L. McNaughton, Katalin M. Gothard, and Etan J. Markus. An information-theoretic approach to deciphering the hippocampal code. In *Advances in Neural Information Processing Systems 5*, pp. 1030–1037, 1993.

Trygve Solstad, Charlotte N. Boccara, Emilio Kropff, May-Britt Moser, and Edvard I. Moser. Representation of geometric borders in the entorhinal cortex. *Science*, 322(5909):1865–1868, 2008.

Ben Sorscher, Gabriel C. Mel, Samuel A. Ocko, Lisa M. Giocomo, and Surya Ganguli. A unified theory for the computational and mechanistic origins of grid cells. *Neuron*, 111(1):121–137, 2023.

Kimberly L. Stachenfeld, Matthew M. Botvinick, and Samuel J. Gershman. The hippocampus as a predictive map. *Nature Neuroscience*, 20(11):1643–1653, 2017.

Hui Min Tan, Joshua P. Bassett, John O'Keefe, Francesca Cacucci, and Thomas J. Wills. The development of the head direction system before eye opening in the rat. *Current Biology*, 25(4): 479–483, 2015.

Jeffrey S. Taube, Robert U. Muller, and James B. Ranck, Jr. Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience*, 10(2):420–435, 1990.

Benigno Uria, Borja Ibarz, Andrea Banino, Vinicius Zambaldi, Dharshan Kumaran, Demis Hassabis, Caswell Barry, and Charles Blundell. The Spatial Memory Pipeline: a model of egocentric to allocentric understanding in mammalian brains. *bioRxiv*, pp. 2020.11.11.378141, 2020.

James C. R. Whittington, Timothy H. Muller, Shirley Mark, Guifen Chen, Caswell Barry, Neil Burgess, and Timothy E. J. Behrens. The Tolman-Eichenbaum machine: Unifying space and relational memory through generalization in the hippocampal formation. *Cell*, 183(5):1249–1263, 2020.

Tom J. Wills, Francesca Cacucci, Neil Burgess, and John O'Keefe. Development of the hippocampal cognitive map in preweanling rats. *Science*, 328(5985):1573–1576, 2010. doi: 10.1126/science.1188224. URL https://www.science.org/doi/abs/10.1126/science.1188224.

Shawn S. Winter, Benjamin J. Clark, and Jeffrey S. Taube. Disruption of the head direction cell network impairs the parahippocampal grid cell signal. *Science*, 347(6224):870–874, 2015. doi: 10.1126/science.1259591.

D. Yoganarasimha, Xintian Yu, and James J. Knierim. Head direction cell representations maintain internal coherence during conflicting proximal and distal cue rotations: comparison with hippocampal place cells. *Journal of Neuroscience*, 26(2):622–631, 2006.

Kechen Zhang. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *Journal of Neuroscience*, 16(6):2112–2126, 1996.

A Environments and data

Floor plans. Each environment was a polygon with 4-9 walls (uniformly random), generated by placing vertices at evenly spaced angles around a centre, each jittered by up to $\pm 40\%$ of the spacing and at a radius drawn uniformly from 0.7-1.0. With probability 0.5 the polygon was biased to be concave by pulling between one and $\lfloor n/3 \rfloor$ vertices 30-60% of the way towards the centroid. Candidates were rejected if they self-intersected or had any angle between adjacent walls below 45°. Each polygon was rotated by a uniformly random angle, so that environments shared no orientation, and scaled so its larger bounding-box extent was 0.8-1.2 m (uniform). Of the 50 environments, 14 were concave.

Textures and landmarks. Each wall carried a texture drawn from a set of ten, without replacement within an environment, and the floor one from a set of four; the ceiling was black. Each environment had 4-7 landmarks (uniform, capped at the number of walls), large visual patterns drawn without replacement from a pool of ten (striped and flipped gratings, horizontal and vertical stripes, a ring, a filled circle, and upright and inverted triangles in black or white), each placed on a different randomly chosen wall and at most 0.4×0.4 m. Walls were 0.5 m high.

Trajectories. Trajectories were simulated with RatInABox (George et al., 2024) at 10 Hz for 2,000 s per environment (20,000 steps), (speed mean = 1, thigmotaxis = 0.8, speed coherence time = 0.8, rotational velocity coherence time = 0.8). Across environments the agent's speed actually averaged 0.66 m s^{-1} due to wall-bounce and environment size, and its absolute angular velocity 256° s^{-1}. Heading was taken from RatInABox's head direction, and the action given to the network was speed and angular velocity. The matched-data dataset contained 50 further environments drawn from the same distributions, each with a 255 s trajectory.

Rendering. Views were rendered by ray casting in RatVision (Abrate, 2026) from a camera 3.5 cm above the floor facing the head direction, with a $240^\circ \times 120^\circ$ field of view, at 128×64 greyscale pixels.

Preprocessing. Each view was blurred with a Gaussian of radius 2 pixels, downsampled to 32×16 pixels (Lanczos), and scaled to $[0, 1]$.

B Network and training

Architecture and parameters. Each network is a single-layer recurrent network with sigmoid units and biases throughout. The input at each step is a 32×16 greyscale view (512 pixels) concatenated with a 16-dimensional encoding of the action (speed and angular velocity, through one $\tanh$ layer), and a linear decoder maps the hidden state to the predicted next view. The 500-unit network has 771,560 parameters (input weights 264,000, recurrent weights 250,000, decoder 256,000, action encoder 48, biases 1,512); the 128-unit network has 150,320. Views and actions were min-max normalised with statistics from the training windows only. There was no dropout, weight decay, or activity regularisation.

Objective. Trajectories were cut into windows of 10 steps. From every position in a window the network predicted the next view, then continued autoregressively for 4 further steps, receiving its own previous prediction (clamped to $[0, 1]$) and the true action; the loss was the L1 error averaged over the 5 steps with equal weight. Gradients were truncated at window boundaries, while the hidden state was carried across consecutive windows: each environment's training stream was cut into chains of 100 windows (100 s), each starting from a zero state. Reported test losses are one-step L1 errors.

Optimisation. RMSprop with learning rate 5×10^{-4}, held for 1,500 epochs and then decayed exponentially to 5×10^{-5} at epoch 3,000, with gradient norms clipped at 1. Each epoch advances every chain by one window per step, so every network received 100 updates per epoch and 300,000 in total, whatever the number of training environments.

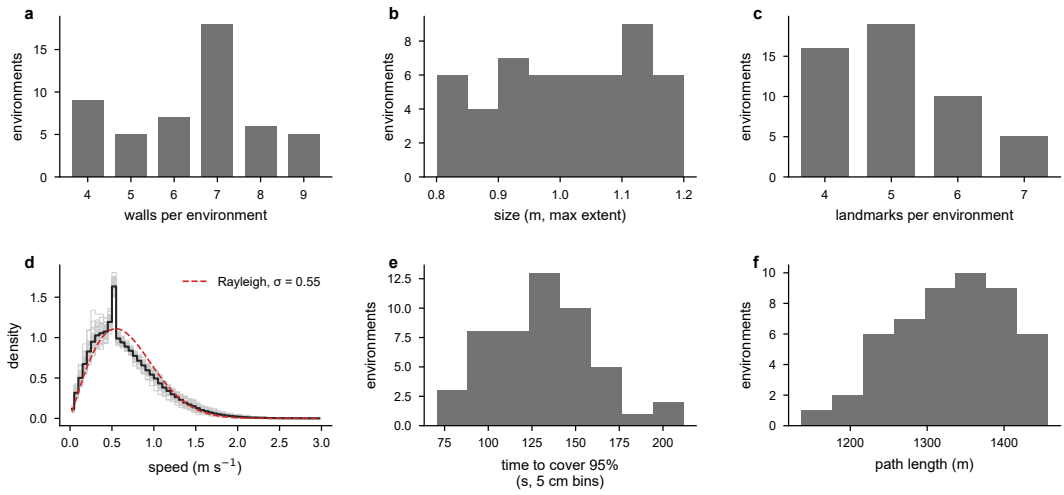

Figure A1: **Statistics of the 50 environments and of the agent's trajectories through them.** Each environment was a procedurally generated polygon (14 of 50 concave, the rest convex) with a random global rotation. **a** Number of walls (polygon vertices) per environment. **b** Environment size (maximum extent of the polygon). **c** Number of landmarks per environment, each drawn without replacement from a shared pool of 10 and placed on a different, randomly chosen wall. **d** Distribution of the agent's speed in each environment (grey) and pooled (black; $0.05\,\mathrm{m\,s^{-1}}$ bins); dashed line, Rayleigh distribution fitted to the pooled data by the method of moments ($\sigma = 0.55\,\mathrm{m\,s^{-1}}$). The spike at $0.5\,\mathrm{m\,s^{-1}}$ is RatInABox's wall-bounce speed: when a step would cross a wall, RatInABox reflects the velocity and sets the speed to half the mean-speed parameter. **e** Time taken for the trajectory to visit 95% of the environment's interior (5 cm bins). **f** Total path length of each trajectory (20,000 steps at $\Delta t = 0.1$ s, 2,000 s in total).

Data and batch composition. Each environment's trajectory was 2,000 s at 10 Hz (20,000 views, 1,995 windows); the last 20% of windows in each training environment were held out as the seen-environment validation set, leaving 1,596 for training (15 chains). The matched-data networks used 255 s trajectories (250 windows; 200 for training, 2 chains). A batch holds one window from every chain:

Training environments	Chains per environment	Batch size
1	15	15
5	15	75
20	15	300
40 (500 and 128 units)	15	600
40, matched-data	2	80

Seeds. The seed set both the initial weights and which environments were held out as unseen. Seeds: 1 environment, 14, 18, 19; 5 environments, 15, 20, 21; 20 environments, 22, 23, 24; 40 environments, 10, 11, 12, 13, 17; 40 environments with 128 units, 16, 25, 26; matched data, three seeds on its own dataset.

Checkpoint selection. All analyses use the final checkpoint (epoch 3,000); no checkpoint was selected on test performance.

Hardware and run time. Networks were trained on an NVIDIA GeForce RTX 4070 SUPER (AMD Ryzen 5 7600X, 32 GB RAM; PyTorch 2.6). A network took about 3.5-6.5 h to train on the main dataset and about 2 h on the matched-data dataset.

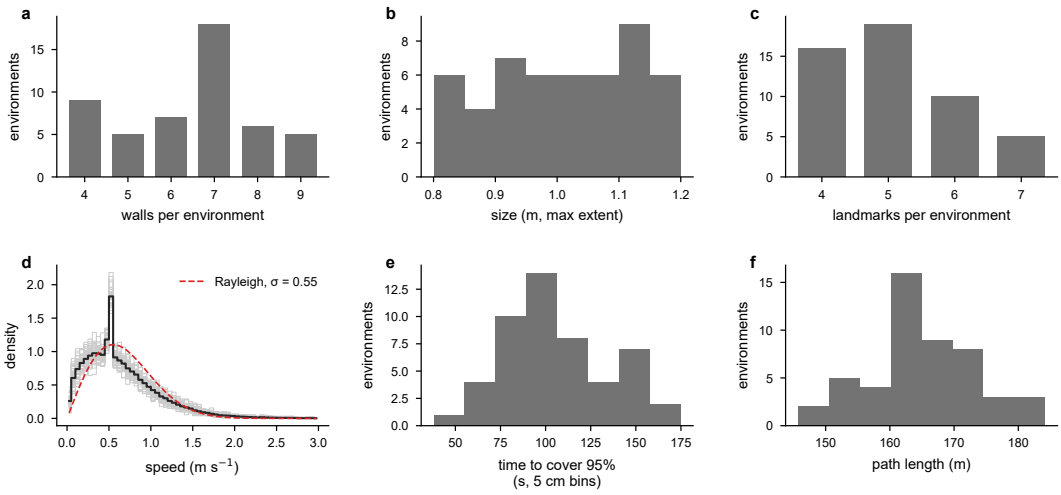

Figure A2: **Environment and trajectory statistics for the matched-data arm.** As Figure A1, for the 50 environments in which the matched-data networks were trained and tested. Each trajectory was 2,550 steps at $\Delta t = 0.1\,\mathrm{s}$ (255 s), against 20,000 steps (2,000 s) in the main dataset, so 40 training environments supplied about the same total amount of training data as five full-length ones (8,000 against 7,980 training windows). Each environment's number of walls, size, and number of landmarks were drawn with the same random seed as its counterpart in the main dataset, so **a**-**c** match Figure A1; its shape, rotation, landmark placement, and trajectory were drawn anew. **d** The pooled speed distribution matched the main dataset's (Rayleigh scale $\sigma = 0.55\,\mathrm{m\,s^{-1}}$), with the wall-bounce spike at $0.5\,\mathrm{m\,s^{-1}}$. **e** Every environment was 95% covered within its 255 s trajectory. **f** Total path lengths were about $1/8$ of those in the main dataset, in proportion to the shorter trajectories.

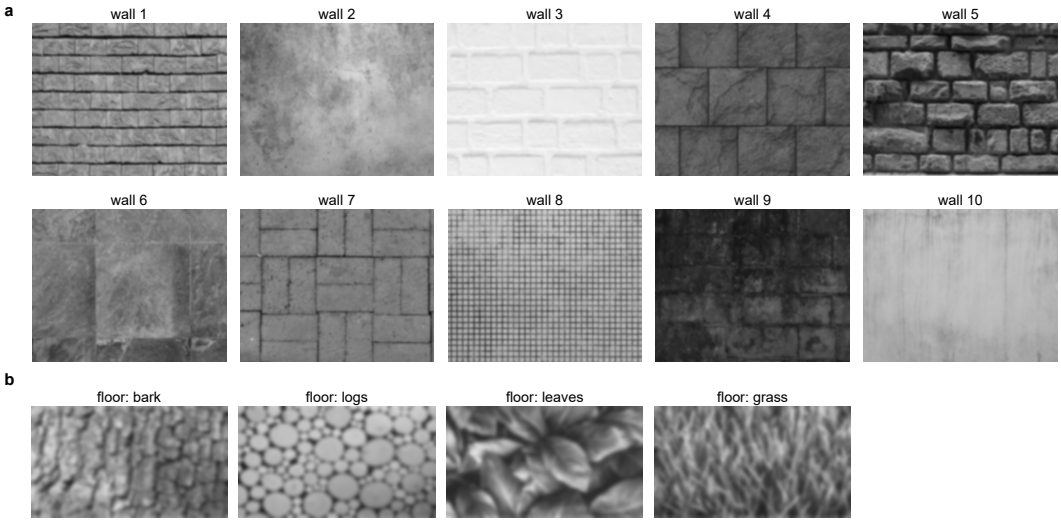

Figure A3: **Texture pools. a** The ten wall textures; each wall of an environment carried a different one, drawn at random. **b** The four floor textures; each environment used one. Textures are shown as rendered: greyscale, blurred with a Gaussian of $\sigma = 3$ pixels, before rendering and before the blur and downsampling applied to each view.

C NETWORK SIZE

A network with 128 rather than 500 hidden units, trained in 40 environments (three seeds), also generalised to unseen environments (generalisation gap $+2.5 \pm 3.7\%$, against $+3.3 \pm 3.7\%$ for 500 units; mean $\pm$ SD over seeds), at a cost of about 8% in prediction error.

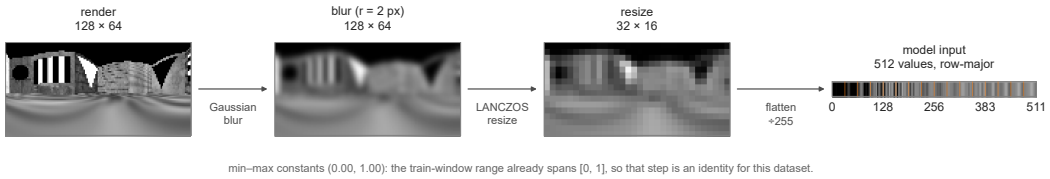

Figure A4: **Each rendered view was blurred, downsampled, and flattened into the network's visual input.** Views were rendered as 128×64 greyscale images covering $240° \times 120°$, blurred with a Gaussian filter (radius 2 px), downsampled to 32×16 with Lanczos resampling, and flattened row by row into a vector scaled to $[0, 1]$ (strip; faint orange lines mark the boundaries between its 16 rows, for visualisation only). Min-max normalisation with constants computed from the training windows (0 and 1 for this dataset) left the values unchanged. The self-motion input (speed and angular velocity), after its learned 16-unit embedding g, was concatenated to this vector.

D FIELDS AND FIELD BEARINGS

Field detection (Fig. 3c). Each rate map (25×25 bins of 4.8 cm over the fixed 1.2 m square) was min-max normalised, and fields were grown with a threshold rising from 0.1 in steps of 0.05. A field was kept if it had at least 10 bins and a peak of at least 0.5, covered no more than half the visited area, and had a within-field split-half stability of at least 0.25 (Pearson r). Field size and number were averaged over the place and conjunctive units with at least one field in each environment.

Main field and bearing (Fig. 5). The bearing test used its own maps and a simpler field definition, independent of the detection above. For each consistent place unit outside the ring, a 25×25 rate map was built in the environment's own reference frame, centred on its centroid and scaled to unit area, so that bin size varied between environments. The main field was the connected region (including diagonal neighbours) of visited bins at or above the midpoint of the map's range that contained its peak bin, with no size, area, or stability filter. Its centre was the mean position of its bins weighted by their rate above that threshold, and its bearing was the angle of this centre from the environment's centre, measured from the same fixed axis as all other angles. Fields whose centre lay within 0.2 of the environment's centre (in units of the square root of its area) were excluded because their bearing was ill defined. So that bearings did not depend on heading, the map was rebuilt for each of eight heading sectors, the main field's bearing was found in each, and these bearings were averaged as angles, each weighted by the range (maximum minus minimum) of its sector's map; a unit needed at least four sectors whose maps covered at least 40% of the environment's bins. Because the main field is whichever field holds the peak, and units had about 2.8 fields per environment, a different field could be selected in different environments; this adds noise that lowers agreement with every candidate rotation equally, making the test conservative rather than biased. The two example units in Fig. 5c were chosen to have exactly one field by the detection above in both environments.

E RING ANALYSIS

The common rotation φ of a set of units between two environments was the circular mode of their $\Delta\theta$, the centre of the $\pm 30°$ window containing the most values, refined by the circular mean of the values in that window. The resultant length of a set of angles, the length of the mean of their unit vectors, is the same measure as the Rayleigh vector length applied to angles rather than to a tuning curve: 1 if all are equal and near 0 if they are evenly spread. The 50 environments were split into two halves of 25, one for selection and one for evaluation, each with the same mix of seen and unseen (20 and 5 after training in 40). A unit's co-rotation score was the resultant length, over pairs of selection environments, of its residuals from the common rotation of all directional units; the 25 highest-scoring units in each network formed the ring. Ring co-rotation, the resultant length of the $\Delta\theta$ of ring units in each pair of evaluation environments averaged over pairs, was compared with shuffled unit identities and with a ceiling from split halves of each trajectory. Ring tiling used a histogram of the preferred directions of ring units in 8 bins, averaged over bin offsets, and was compared with the distribution of the same measure for sets of 25 directions drawn independently and uniformly at random. For clamping, each unit of a set was held at its own mean activation at

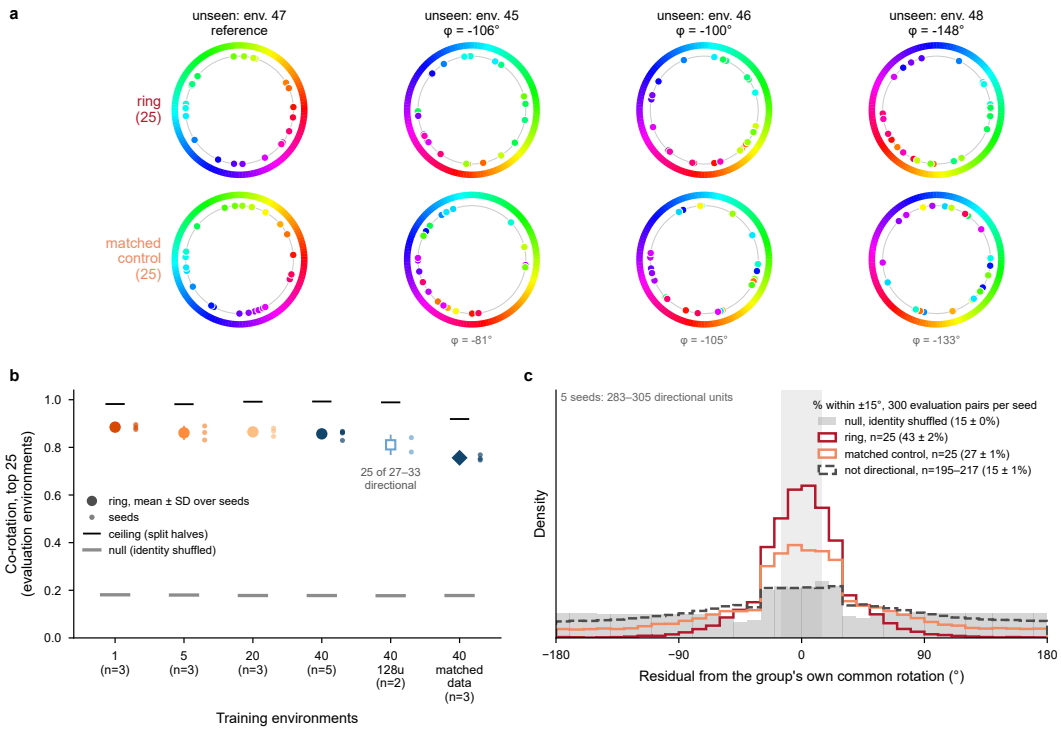

Figure A5: **Ring units rotated together between environments in every arm.** The 40-environment network (seed 12) in **a**; all seeds of each arm in **b** and **c**. **a** Preferred directions of the 25 ring units (top) and 25 matched control units (bottom) in four unseen environments, each unit coloured by its preferred direction in the reference (left) environment; the colour wheel is the reference's turned by φ, each set's own common rotation relative to the reference. Ring units kept their order round the circle; control units did not. **b** Co-rotation of the 25 ring units in the evaluation environments for each arm: large markers, mean $\pm$ SD over seeds ($n = 3$; 5 for 40 environments; 2 for 128 units, whose third seed has too few directional units to form a ring); small dots, individual seeds. Black, ceiling from split halves of the data within environments; grey, shuffle of directional units' identities within each environment (both means over seeds). The 128-unit rings are hollow: they used 25 of only 27-33 directional units. **c** Density of each unit's residual from its group's common rotation over 300 pairs of evaluation environments per seed, pooled over the five 40-environment seeds, for ring units ($n = 25$ per seed), matched control units ($n = 25$), non-directional units ($n = 195$-217), and a shuffle of unit identities; percentages, fraction within $\pm 15°$ (mean $\pm$ SD over seeds).

every time step, within the recurrence, and the increases in loss and in position and heading decoding error were compared between the ring units and 20 control sets of 25 units matched in modal tuning class (each unit's most frequent class across environments) and, within class, in Rayleigh vector length.

Table A1: Percentage of units classified as place, head direction, or conjunctive, averaged over each network's 50 environments; mean $\pm$ SD across seeds.

Arm	Seeds	Place	Head direction	Conjunctive
1 environment	3	46.1 ± 1.9	7.6 ± 1.5	15.4 ± 1.2
5 environments	3	40.2 ± 2.6	9.2 ± 1.6	23.4 ± 1.3
20 environments	3	37.1 ± 0.7	9.4 ± 0.5	24.3 ± 2.0
40 environments	5	38.3 ± 0.4	8.8 ± 0.2	17.1 ± 0.9
40 environments, 128 units	3	43.3 ± 1.3	3.2 ± 0.9	1.8 ± 0.4
40 environments, matched data	3	47.9 ± 0.3	5.6 ± 0.3	26.1 ± 0.8