Conserved hippocampal population geometry supports task generalization

Hannah S Wirtshafter[1*], Sara A Solla[1], John F Disterhoft[1]

*Lead contact: hsw@northwestern.edu (HSW)

1. Department of Neuroscience, Northwestern University Feinberg School of Medicine, Chicago, IL, USA.

Abstract

How learning generalizes across contexts is a fundamental question in neuroscience, as successful behavior often requires transferring acquired knowledge to new environments. Conditioning tasks provide a clear example of such generalization, with learned responses rapidly expressed across distinct spatial contexts. A central challenge in understanding the neural basis of this ability is determining how the hippocampus represents task-related information across environments, given that its spatial representations remap with context. Here, we used calcium imaging to record hippocampal population activity as rats performed a conditioning task across multiple spatial contexts. To characterize task-related population structure, we applied dimensionality reduction and alignment methods to construct low-dimensional manifolds of hippocampal activity. We found that task-related population activity occupied a stable geometric structure across contexts, despite pronounced remapping of spatial representations. Strikingly, this task-related geometry was conserved not only across contexts within individual animals but also across animals, revealing a shared organization of task representations in the hippocampus. These findings provide a population-level account of how task-related information is preserved across changing spatial environments and suggest that hippocampal task representations follow shared population-level geometric organization across individuals.

Introduction

How can learning generalize across contexts while remaining sensitive to contextual differences? This question is fundamental to both neuroscience and philosophy[1-4]. Deficits in generalization or inappropriate generalization are hallmarks of many disorders, including autism[5,6], schizophrenia[7,8], and post-traumatic stress disorder[9,10]. In spite of their importance, many questions related to generalization remain to be answered.

The hippocampus (HPC) is important for learning, memory, and navigation, and damage to this region can disrupt contextual learning[11-17]. Many aspects of context, including an animal's spatial location and the presence of local and distal cues, can be represented by 'place cells' in the HPC[18-22]. The activity of many HPC cells therefore changes drastically in different environments (i.e. place cells remap), even when a task can be generalized across these different contexts[23-25]. A major open question is how representations of task-relevant stimuli, which are also found in the HPC[26-31], are organized across environments in which spatial representations remap. A further question is whether different animals solve this problem using the same, or similar, neural strategies.

Until recently, the neural mechanisms behind contextual learning have been challenging to investigate due to the need for tracking large numbers of cells across various environments and learning stages— tasks which were unachievable prior to the development of calcium imaging[32-37]. Calcium imaging enables population-level analyses of hippocampal activity across environments and learning stages, allowing direct comparisons of neural representations beyond single-cell tuning.

This study addresses two population-level questions about hippocampal representations during learning. First, how is task-related information structured across neural populations when spatial representations reorganize across contexts[24,25,38,39]? Second, to what extent is this population structure conserved across individuals? If task-related hippocampal activity follows shared organizational principles across animals, this would suggest that learning is supported by common population-level frameworks rather than idiosyncratic neural solutions.

To answer these questions, we trained animals on an HPC-dependent conditioning task which is rapidly generalized between spatial contexts[40,41]. We examined the same task in disparate environments and quantified how population-level representations of spatial location changed while task-related population activity was preserved. We found that task representations were maintained as animals generalized learning across contexts, despite pronounced changes in spatial representations. Strikingly, task-related population structure was also conserved across animals.

Together, these results demonstrate that hippocampal task representations exhibit a stable population-level organization that persists across spatial contexts and is conserved across individuals. This conserved structure provides a framework for understanding how learning generalizes across environments at the level of neural populations.

84

85 **Results**

86 We trained five freely moving rats in a conditioning task in one of two distinct
87 environments, labeled A and B; the rats had been previously familiarized with the environment
88 by the time the training began (Fig. 1a, Fig. S1). Environment A was an unscented rectangular
89 enclosure with wire floor and walls, and white lighting; environment B was a scented ovular
90 enclosure with white solid floor and walls, and red lighting (Fig. 1b). Both environments were
91 located at the same spot in the room relative to external cues (see Methods); animals could see
92 external cues out of the top of both environments, as well as out of the sides of environment A
93 (animals also reared often, allowing them to see out of the sides of environment B). During
94 training sessions, we recorded cellular activity in the hippocampal CA1 region via miniscopes,
95 using GCaMP8m for calcium imaging (CaImg). To perform data analysis, we used both calcium
96 events and calcium traces (see Methods), as indicated when applicable.

97

98 **Animals easily transfer a conditioning task across environments despite place cell**
99 **remapping**

100 During the initial phase of our study, freely-moving rats (Fig. 1b) underwent training for
101 trace eyeblink conditioning (tEBC) (Fig. S2), a hippocampus-dependent classical conditioning
102 task that serves as a robust model for associative memory formation[42-45]. This paradigm involves
103 presenting a 250ms conditioned stimulus (CS, in the form of a tone) followed by a 500ms trace
104 interval, followed by the 100ms presentation of an unconditioned stimulus (US, an eyelid shock)
105 (Fig. 1c). Both shock and blinking were recorded with wires inserted into the muscle of the
106 eyelid (see Methods). As rats were trained, they exhibited a conditioned blink (CR) to the tone.

Animals were considered to have learned the task after reaching criterion (70% CRs in 50 trials) on three consecutive training sessions (termed ‘criterion sessions’) or when the previous four training sessions averaged over 70% (in this instance, only the final three of those sessions were considered ‘criterion sessions’) (Fig. 1d). There was substantial variability in the number of sessions it took to learn the task, for an average of 20±4.2 training sessions (note: number of sessions always includes criterion sessions). The two rats that learned the fastest reached criterion in 14 sessions, and the rat that learned the slowest reached criterion after 24 sessions (Fig. S2). After reaching criterion, the rats were introduced to environment B, where their ability to perform tEBC was assessed over a two-day period (one session per day). Comparative analysis revealed no significant difference in performance (measured in % CRs) between the criterion sessions in environment A and the testing phase in environment B (mean in environment A criterion sessions was 74.75±6.49, mean in environment B test sessions was 77.70±11.68, two tailed t-test(24)=-0.83, $p>0.05$) (Fig. 1d), indicating the successful transfer of tEBC learning to a new environment.

Consistent with place cell remapping, individual hippocampal neurons exhibited distinct spatial representations in environments A and B (Fig. 1e). Putative place field centers shifted significantly more between environments than between consecutive sessions within the same environment (Fig. 1f), indicating a loss of spatial correspondence across contexts (double-sided t-test, $t(1430) = -2.5$, $p < 0.01$). At the population level, population vector correlations (PVC37,45,46; see Methods) revealed significant similarity between consecutive sessions within environment A but no significant similarity between sessions in environments A and B (Fig. 1g),

129 confirming robust remapping of spatial representations. Additional analyses using shuffled

130 controls and population embedding methods yielded consistent results (Fig. S3–S4).

131

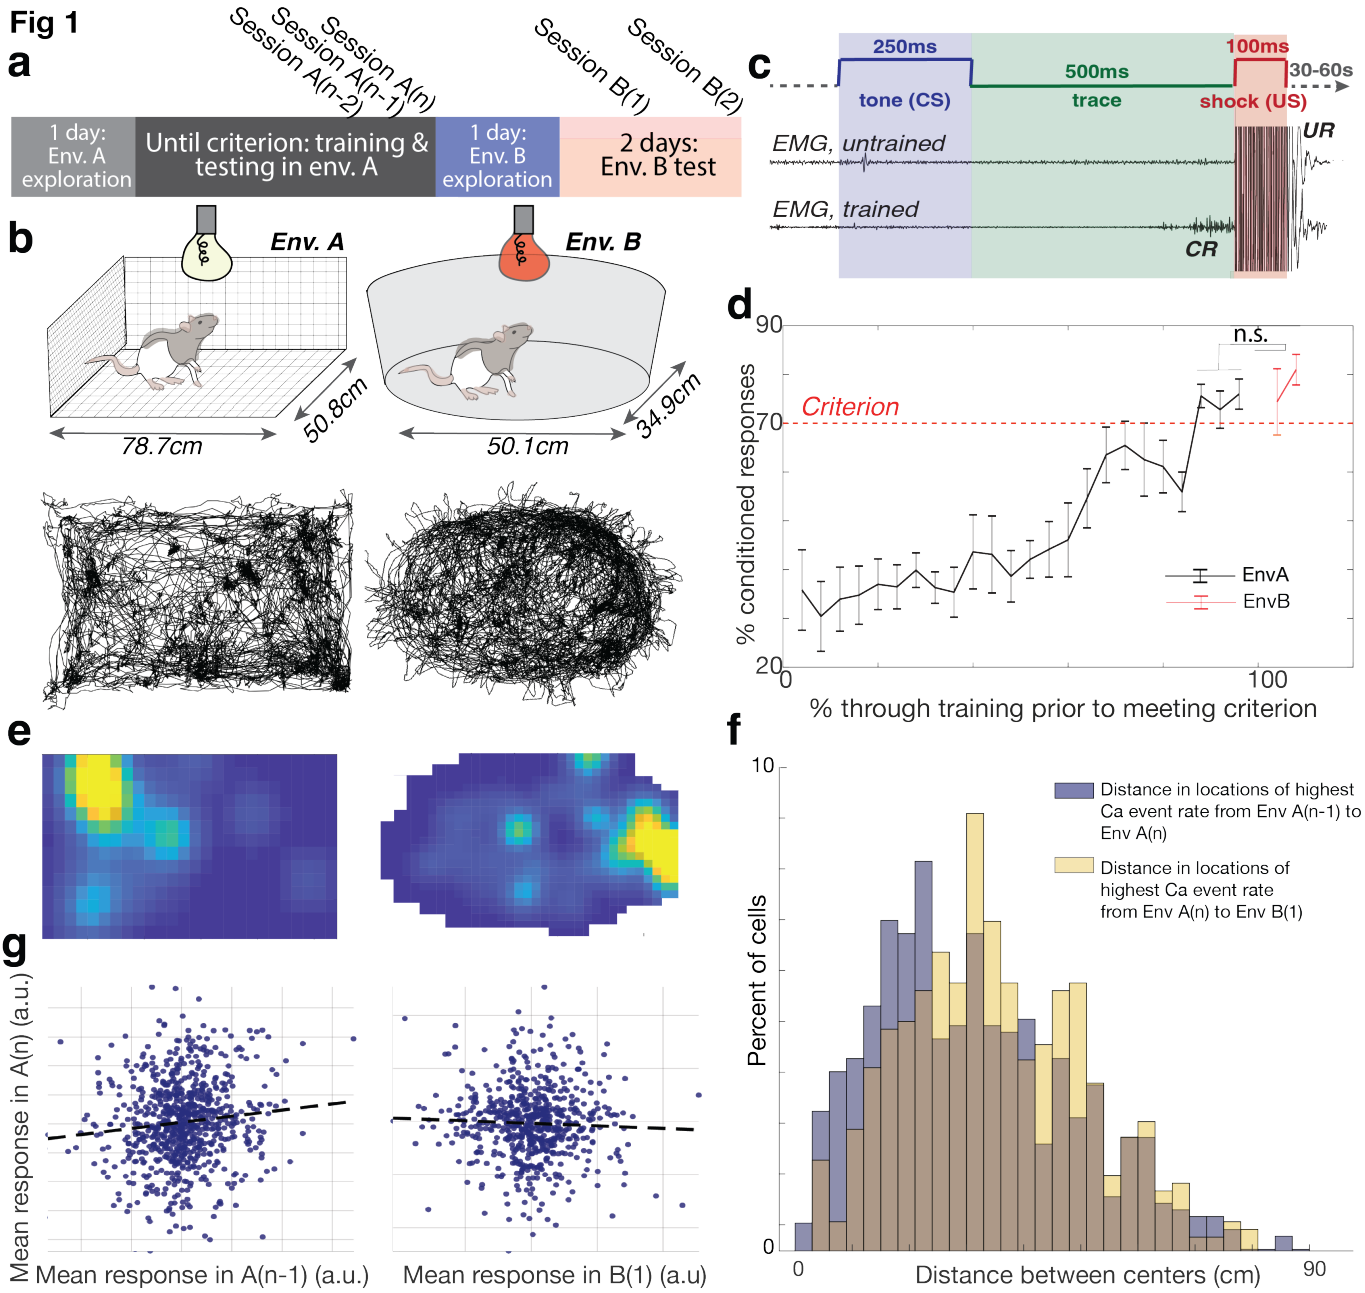

132

Fig 1. Conditioning is maintained while place cells remap between environment A and environment B.

- **a.** Animals explored environment A for one session before undergoing eyeblink conditioning until they reached the learning criterion. Criterion was defined as achieving 70% conditioned responses (CRs) in 50 trials across three consecutive training sessions (A(n-2), A(n-1), and A(n)) or averaging over 70% CRs across the previous four training sessions. After meeting the criterion in environment A, animals were allowed one session of exploration in environment B, followed by two test sessions of trace eyeblink conditioning (B(1) and B(2)) in environment B.
- **b.** (Top) Schematics of environments A and B. Environment A is a rectangular enclosure with wire walls, floor, and ceiling, lit with white light, and unscented. Environment B is oval-shaped with solid white floors and walls, without a ceiling, lit with red light, and scented with clove oil. Both environments provided distal cues visible from the top and sides. (Bottom) Animal trajectories in environments A and B during a single session, extracted using DeepLabCut.
- **c.** Trace eyeblink conditioning (tEBC) paradigm. A 250 ms tone (conditioned stimulus, CS) was followed by a 500 ms trace interval, then a 100 ms eyelid shock (unconditioned stimulus, US). Eyelid activity was recorded using an EMG electrode implanted above the eye. Untrained animals only blinked in response to the US (unconditioned response, UR); trained animals began blinking during the trace interval after the CS and before the US (conditioned response, CR).
- **d.** Performance of animals (n=5) in the tEBC task. Animals learned tEBC while freely moving in environment A and successfully transferred this learning to environment B. The dotted line indicates the performance criterion. No significant difference was found between performance in the criterion sessions in environment A (mean 74.75 ± 6.49%) and the test sessions in environment B (mean 77.70 ± 11.68%; two-tailed t-test, $t(24) = -0.83$, $p > 0.05$). Error bars represent standard error.
- **e.** Example activity map of a specific single cell in environments A and B. Yellow indicates the highest firing rates. This cell exhibited remapping between environments, showing different place fields relative to external cues in the two contexts.
- **f.** Distribution of distances between place field centers (determined by the highest calcium event rate), comparing sessions A(n) and A(n-1) versus sessions A(n) and B(1). Place field centers shifted significantly more when the animal was moved to environment B compared to shifts between two consecutive sessions in environment A (Wilcoxon rank sum test, $p = 0.002$; two-sided t-test, $t(1430) = -2.5$, $p = 0.01$; two-sample Kolmogorov-Smirnov test, $p = 3.6 \times 10$). The brown shaded area represents the overlap between distance histograms.
- **g.** Population vector correlation (PVC) based on calcium events for sessions A(n-1), A(n), and B(1). A significant positive correlation was found between sessions A(n-1) and A(n) for cells present in both sessions ($p = 0.0023$, $r = 0.11$). In contrast, there was no significant correlation between sessions A(n) and B(1) for shared cells ($p > 0.05$, $r = -0.04$). These findings suggest that calcium event patterns are significantly similar between sessions A(n-1) and A(n) but not between session A(n) and B(1). Dashed lines represent lines of best fit. (a.u. = arbitrary units).

Conditioning task representations are consistent across environments

We then asked whether task-related information could be decoded across environments at the population level. Using CEBRA models trained on calcium imaging data labeled with CS and US periods from day A(n) (termed CSUS2 for the division of the task into 2 labeled epochs), we decoded task epochs for day A(n-1) as well as in environment B. Models trained only on data from environment A reliably decoded CS and US periods in both environments, significantly outperforming shuffled controls (Fig. 2a–c, all double sided t-tests $p<0.001$), demonstrating that task-related population activity generalized across contexts.

To determine whether finer temporal structure within the conditioning period was also preserved, we trained models to decode the temporal sequence of task epochs spanning CS presentation, the trace interval, and US delivery (CSUS5). Models trained on session A(n) accurately decoded both session A(n−1) and session B(1), indicating that detailed temporal aspects of task encoding generalized across environments (Fig. 2d–g, all double sided t-tests $p<0.001$). Remarkably, for both CSUS2 and CSUS5, the model trained on session A(n) was no less accurate decoding session B(1) than it was decoding session A(n-1) (double sided t-tests, both $p>0.05$) (Fig. 2h).

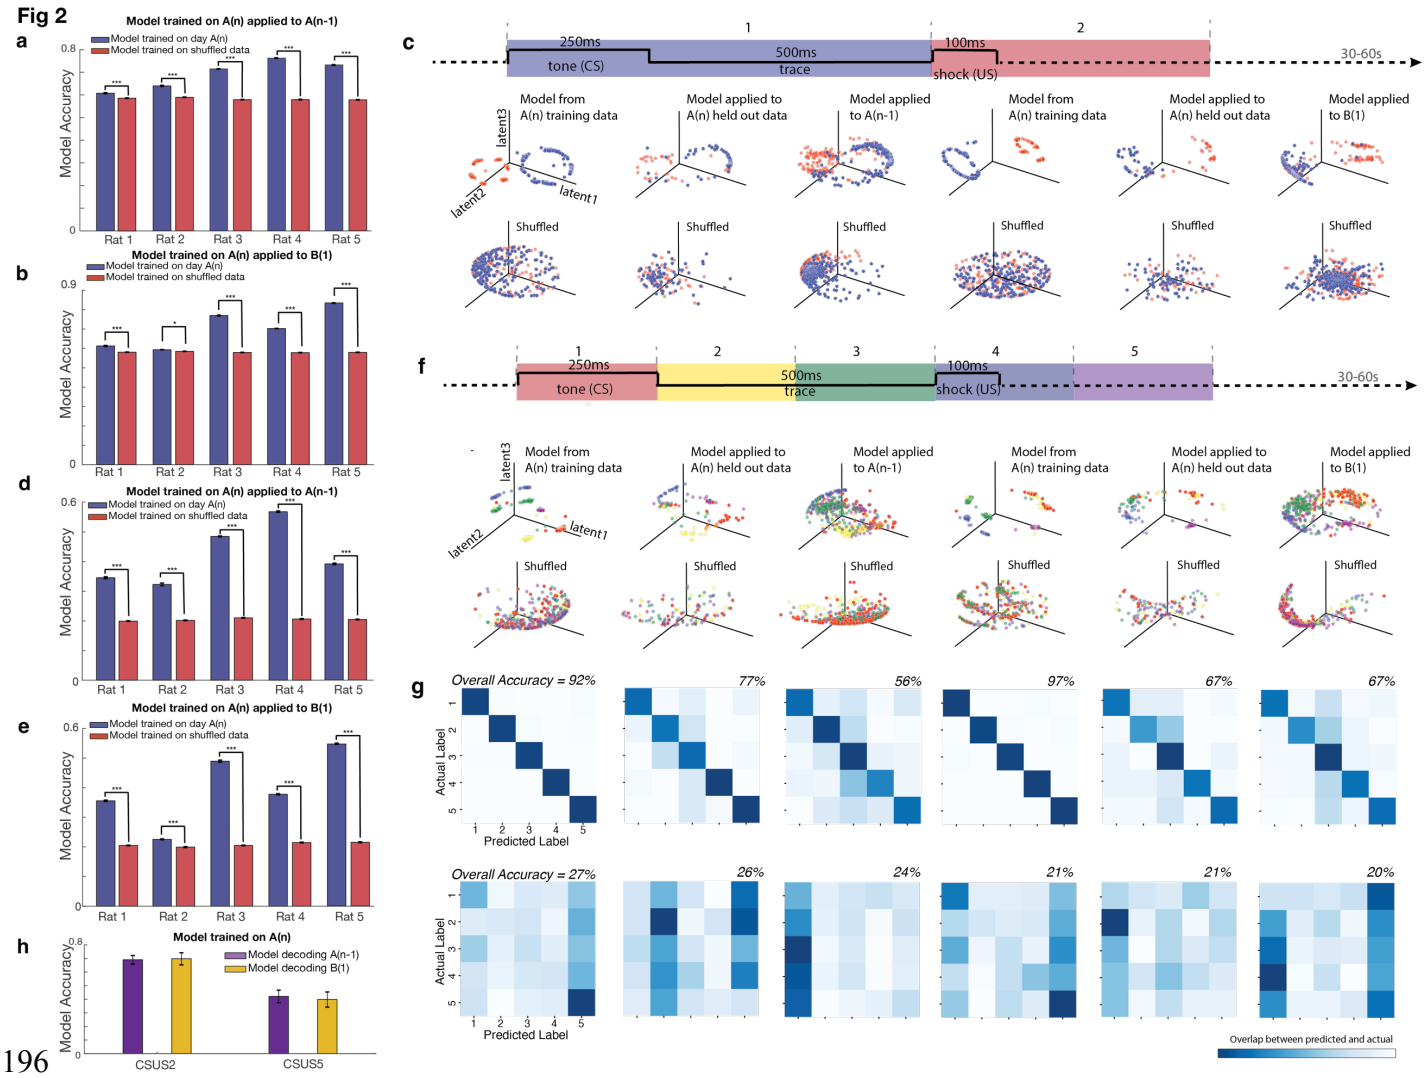

Fig. 2. A model trained in environment A can decode CS and US periods in both environment A and B at above chance levels, including fine-grained temporal decoding.

a. A CEBRA model was trained using calcium imaging data and time-stamped CS/US periods from environment A, using only cells that were recorded in both environments A and B. The model was used to decode whether the animal was in a CS or US period in another session in environment A; all models successfully decoded these periods compared to shuffled data (all double-sided t-tests: Rat 1: $t(998) = 6.7$, $p = 2.5 \times 10^{-11}$; Rat 2: $t(998) = 16.1$, $p = 7.4 \times 10^{-52}$; Rat 3: $t(998) = 83.5$, $p = 0$; Rat 4: $t(998) = 80.1$, $p = 0$; Rat 5: $t(998) = 61.0$, $p = 0$). Bars indicate standard error; significance denoted as ***$p < 10^{-10}$.

b. The same models from panel a were now applied to environment B; they significantly outperformed chance level (all double-sided t-tests: Rat 1: $t(998) = 10.4$, $p = 2.6 \times 10^{-24}$; Rat 2: $t(998) = 2.7$, $p = 7.6 \times 10^{-3}$; Rat 3: $t(998) = 75.3$, $p = 0$; Rat 4: $t(998) = 63.7$, $p = 0$; Rat 5: $t(998) = 106.3$, $p = 0$). Error bars represent standard error, *$p < 10^{-3}$, ***$p < 10^{-10}$.

c. (Top row) A CEBRA model trained on CS/US periods, divided into two time bins (data from Rat 4). For the top first three graphs, the model was trained using position and calcium trace data from session A(n), using cells present in both A(n) and A(n-1). First graph: The trained model applied to training data from session A(n). Second graph: The model applied to decode CS/US periods from held-out data (25%) from session A(n). Third graph: The model applied to session A(n-1). For the top last three graphs, the model was trained using position and calcium trace data from session A(n), using cells present in both A(n) and B(1). Fourth graph: The trained model applied to training data from session A(n). Fifth graph: The model applied to decode CS/US periods from held-out data (25%) from session A(n). Sixth graph: The model applied to session B(1). (Bottom row) The same model trained on shuffled data. The model trained on data from session A(n) significantly outperformed the model trained on shuffled data in decoding both session A(n-1) and session B(1) (double-sided t-tests, $t(998) = 83.5$, $p = 0$, and $t(998) = 75.3$, $p = 0$, respectively).

d. A CEBRA model was trained on data from environment A to decode temporal order within the CS, trace, and US periods split into five divisions, applied to an alternate session in environment A (session A(n-1)). The model significantly outperformed chance (all double-sided t-tests for accuracy: Rat 1: $t(998) = 41.5$, $p = 3.3 \times 10^{-220}$; Rat 2: $t(998) = 28.7$, $p = 7.8 \times 10^{-133}$; Rat 3: $t(998) = 122.6$, $p = 0$; Rat 4: $t(998) = 118.5$, $p = 0$; Rat 5: $t(998) = 62.6$, $p = 0$). Bars indicate standard error; significance denoted as ***$p < 10^{-10}$.

e. The same five models from panel d were now applied to environment B and outperformed models trained on shuffled data in decoding the temporal aspects of the CS/US periods, indicating that fine-grained temporal encoding is stable across environments (all double-sided t-tests for accuracy: Rat 1: $t(998) = 55.1$, $p = 4.9 \times 10^{-305}$; Rat 2: $t(998) = 9.3$, $p = 1.1 \times 10^{-19}$; Rat 3: $t(998) = 71.4$, $p = 0$; Rat 4: $t(998) = 62.6$, $p = 0$; Rat 5: $t(998) = 106.2$, $p = 0$). Bars indicate standard error; significance denoted as ***$p < 10^{-10}$.

f. Same analysis as in panel c, but for a CEBRA model trained on CS/US periods split into five time bins. Top: The five divisions of the conditioning period are shown. Data from Rat 3 show that the model trained on session A(n) outperformed models trained on shuffled data when decoding both session A(n-1) and session B(1) (double-sided t-tests, $t(998) = 88.9$, $p = 0$, and $t(998) = 71.4$, $p = 0$, respectively).

g. Confusion matrices displaying CEBRA decoding of five CS/US time bins, as shown in panel f. Top row: Models trained on data from session A(n). Bottom row: Models trained on shuffled data. Darker colors along the diagonal indicate higher model accuracy.

h. Model accuracy for decoding sessions A(n-1) and B(1) using a model trained on data from session A(n) with CSUS2 and CSUS5 divisions. For both CSUS2 and CSUS5, a model trained in session A(n) decoded session B(1) with accuracy similar to that achieved when decoding session A(n-1) (CSUS2: double-sided t-test, $t(8) = -0.13$, $p > 0.05$; CSUS5: $t(8) = 0.32$, $p > 0.05$). Bars represent standard error

Finally, to directly test whether task representations shared a conserved population geometry across environments, we compared the embedding structure of task-related activity using CEBRA. Embedding geometries for task epochs were highly similar across sessions in environments A and B and were significantly more similar than shuffled controls. This similarity was robust across multiple embedding dimensionalities and was preserved when task periods were divided into either coarse (CSUS2) or fine (CSUS5) temporal segments (Fig. 3a–f).

Together, these results indicate that task-related hippocampal representations maintain a stable population-level geometry across environments, despite changes in spatial representations.

Fig 3

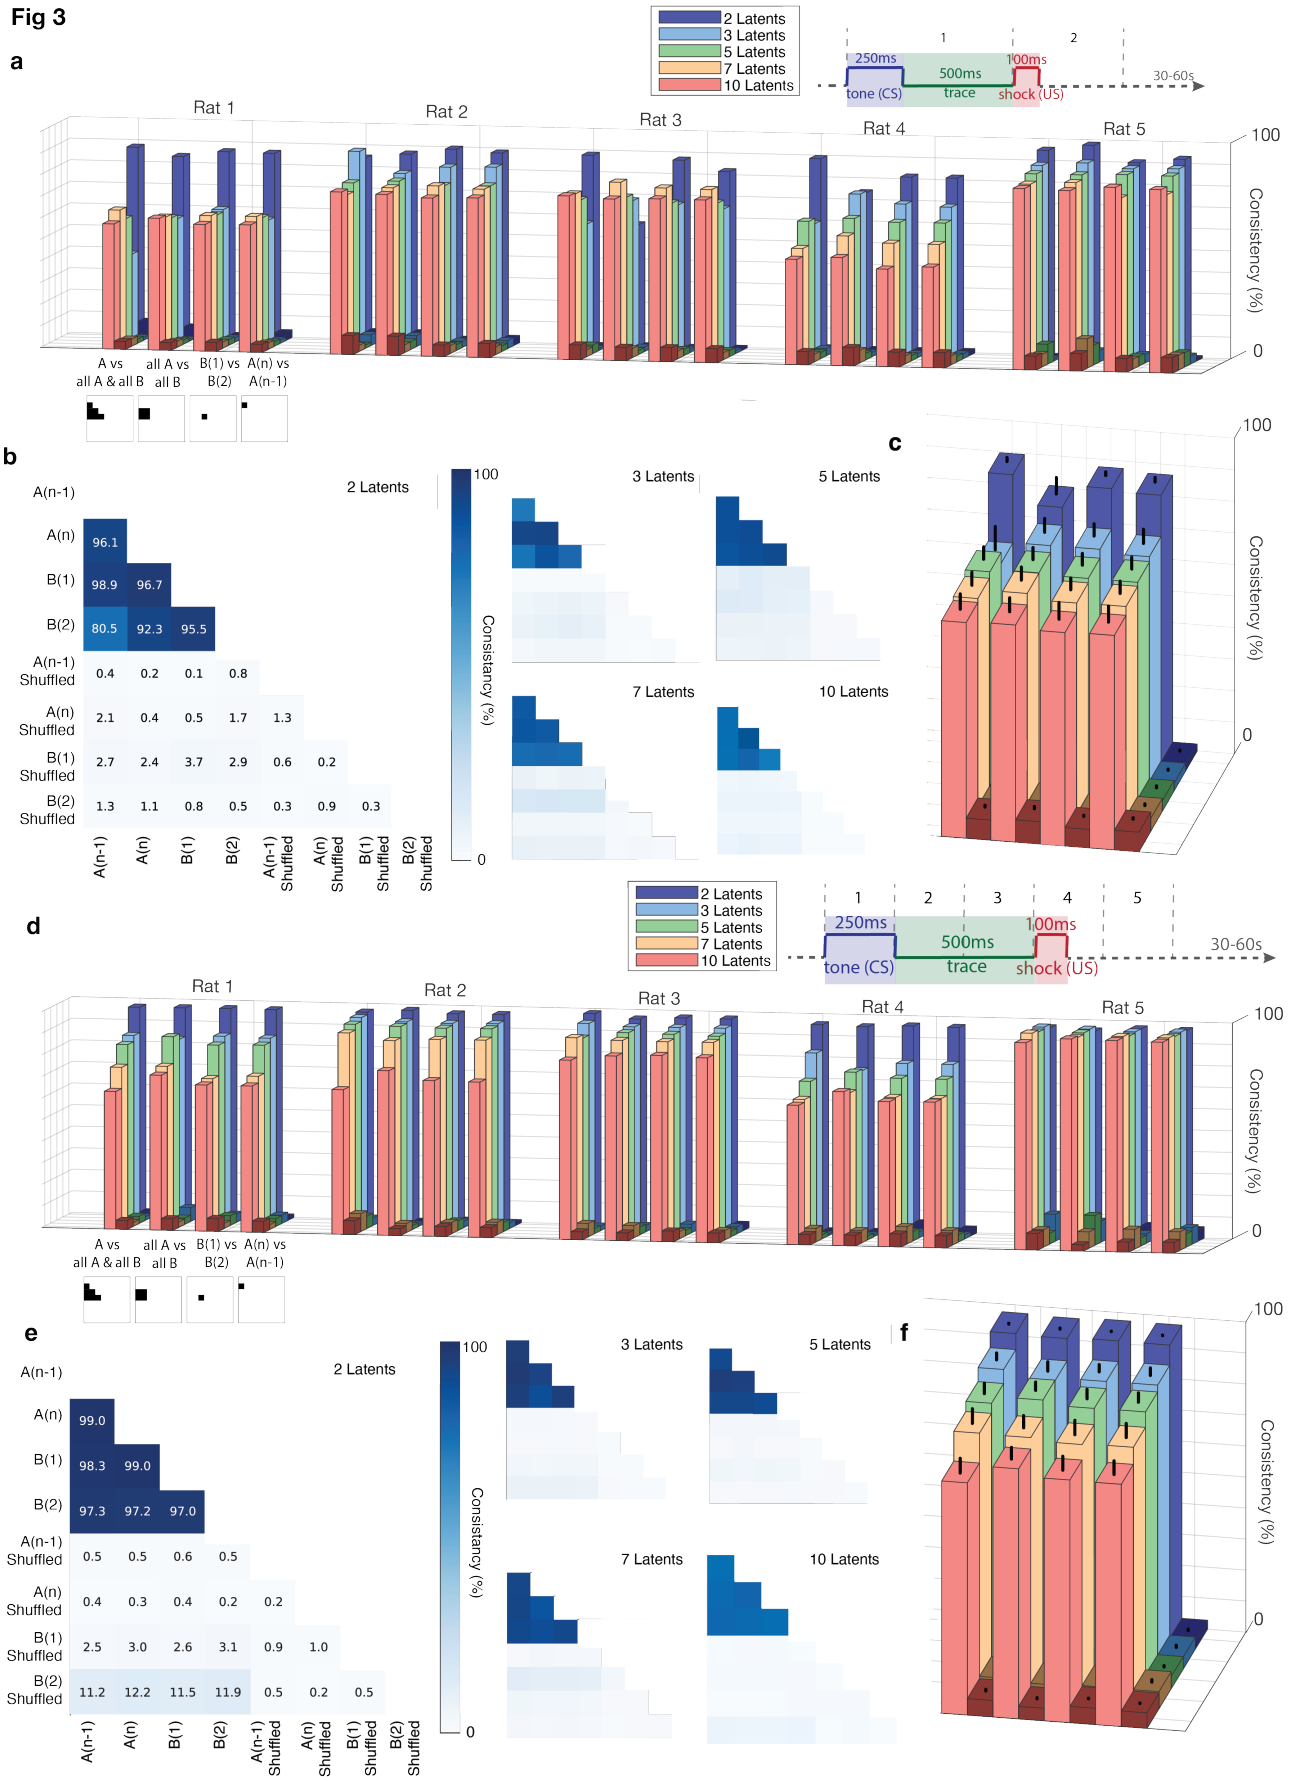
a
2 Latents
3 Latents
5 Latents
7 Latents
10 Latents
250ms
500ms
100ms
tone (CS)
trace
shock (US)
30-60s
Rat 1
Rat 2
Rat 3
Rat 4
Rat 5
Consistency (%)
A vs all A & all B
all A vs all B
B(1) vs B(2)
A(n) vs A(n-1)
b
A(n-1)
A(n) 96.1
B(1) 98.9 96.7
B(2) 80.5 92.3 95.5
A(n-1) Shuffled 0.4 0.2 0.1 0.8
A(n) Shuffled 2.1 0.4 0.5 1.7 1.3
B(1) Shuffled 2.7 2.4 3.7 2.9 0.6 0.2
B(2) Shuffled 1.3 1.1 0.8 0.5 0.3 0.9 0.3
A(n-1) A(n) B(1) B(2) A(n-1) Shuffled A(n) Shuffled B(1) Shuffled B(2) Shuffled
2 Latents
3 Latents
5 Latents
7 Latents
10 Latents
Consistancy (%)
c
d
e
A(n-1)
A(n) 99.0
B(1) 98.3 99.0
B(2) 97.3 97.2 97.0
A(n-1) Shuffled 0.5 0.5 0.6 0.5
A(n) Shuffled 0.4 0.3 0.4 0.2 0.2
B(1) Shuffled 2.5 3.0 2.6 3.1 0.9 1.0
B(2) Shuffled 11.2 12.2 11.5 11.9 0.5 0.2 0.5
f

Fig 3. High consistency in neural representations between environments A and B.

a. Consistency scores for each rat, calculated with 2, 3, 5, 7, and 10 latents. Lighter bars represent consistency percentage for actual data, while adjacent darker bars represent consistency for shuffled data. The schematic below the x-axis illustrates the data sets included in each comparison (for position and labelling of data sets see figure 6b). For the CSUS2 division of conditioning periods (split into CS and US components), all five rats show significantly higher consistency between environments A and B in the actual data compared to shuffled data. This consistency remains significant with up to 10 latents.

b. Example consistency measurements from Rat 5 for CSUS2 with 2, 3, 5, 7, and 10 latents.

c. Average consistency scores across all rats for CSUS2. The figure legend follows the format of panel a. Bars represent standard error. Consistency scores for actual data versus shuffled data were significantly different across all latent dimensions (all double-sided t-tests: 2 latents, $p < 1 \times 10^{-5}$; 3 latents, $p < 1 \times 10^{-4}$; 5 latents, $p < 1 \times 10^{-5}$; 7 latents, $p < 1 \times 10^{-5}$; 10 latents, $p < 1 \times 10^{-4}$).

d. Consistency scores for each rat, as in panel a, but for CSUS5, where the conditioning period is divided into five time bins. All five rats show significantly higher consistency between environments A and B in the actual data compared to shuffled data; this is maintained across latents for up to 10.

e. Example consistency measurements from Rat 3 for CSUS5 with 2, 3, 5, 7, and 10 latents.

f. Average consistency scores across all rats for CSUS5. The figure legend matches that of panel c. Bars represent standard error. Significant differences between actual and shuffled data were observed for all latent dimensions (all double-sided t-tests: 2 latents, $p < 1 \times 10^{-12}$; 3 latents, $p < 1 \times 10^{-7}$; 5 latents, $p < 1 \times 10^{-6}$; 7 latents, $p < 1 \times 10^{-5}$; 10 latents, $p < 1 \times 10^{-5}$).

Task representations are conserved across animals

Given the similarity of task-related representations across environments within individual animals, we next asked whether task representations also exhibit a shared population-level organization across animals. To address this question, we trained animal-specific models based on calcium activity and task structure and quantified the similarity between models derived from different animals.

Task-related population representations showed a high degree of similarity across animals compared to shuffled controls. This similarity was evident both for models distinguishing coarse

task epochs (CSUS2) and for models capturing finer temporal structure across five task segments (CSUS5) (Fig. 4a–b).

Notably, the similarity of task-related population geometry across individuals was comparable in magnitude to the similarity observed across sessions within the same animal. In fact, across-animal similarity was not significantly different from within-animal similarity for any tested embedding dimensionality, for either coarse (CSUS2) or fine-grained (CSUS5) task segmentation) (Fig. 4c–d) (all paired t-tests, $p>0.05$). This indicates that task-related population geometry is conserved across animals to a comparable extent as it is across sessions within an individual.

Together, these results demonstrate that hippocampal task representations follow a shared population-level organization across animals, consistent with conserved geometric structure rather than idiosyncratic task encoding.

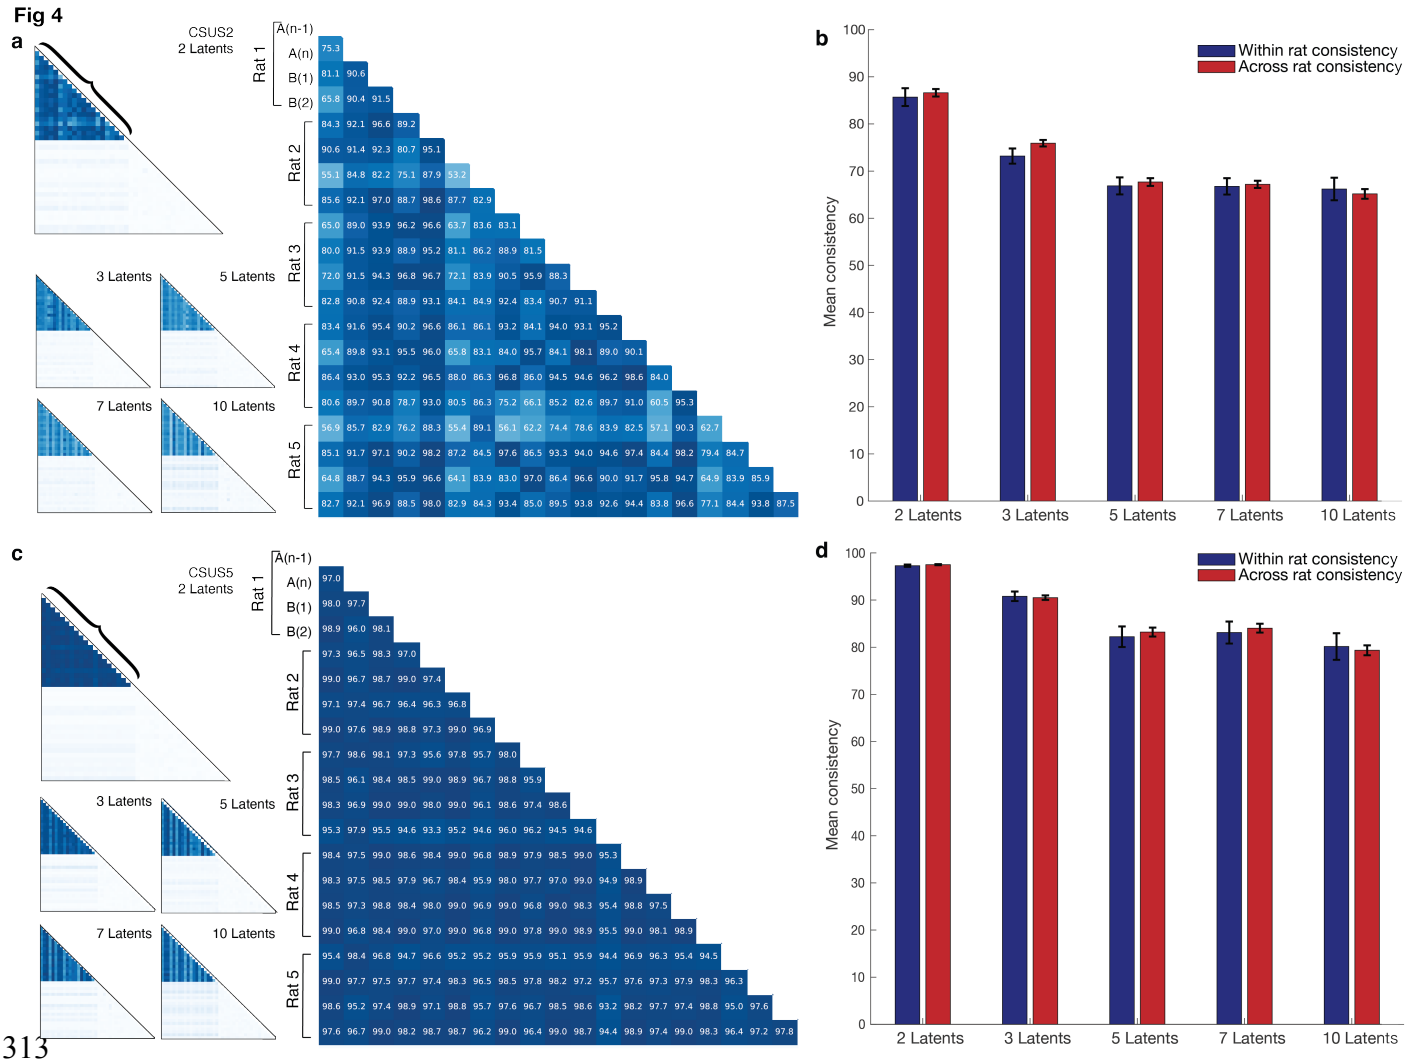

Figure 4. High degree of consistency between the conditioning representations across animals

a. Consistency across animals for CSUS2 with 2, 3, 5, 7, and 10 latents. The larger graph highlights the bracketed area from the graph with two latents.

b. Conditioning period is divided into 2 time bins (CSUS2): the similarity between models across animals is not significantly different from the similarity between models for an individual animal, regardless of the number of latents (2 latents: $t(188) = -0.45$, $p > 0.05$; 3 latents: $t(188) = -1.57$, $p > 0.05$; 5 latents: $t(188) = -0.40$, $p > 0.05$; 7 latents: $t(188) = -0.22$, $p > 0.05$; 10 latents: $t(188) = 0.40$, $p > 0.05$). Error bars represent standard error.

c. Consistency across animals for CSUS5 with 2, 3, 5, 7, and 10 latents. The larger graph highlights the bracketed area from the graph with two latents.

d. Conditioning period is divided into 5 time bins (CSUS5); the similarity across animal models remains comparable to that for within-animal models, regardless of the number of latents (2 latents: $t(188) = -0.79$, $p > 0.05$; 3 latents: $t(188) = 0.25$, $p > 0.05$; 5 latents: $t(188) = -0.42$, $p > 0.05$; 7 latents: $t(188) = 0.70$, $p > 0.05$; 10 latents: $t(188) = 0.30$, $p > 0.05$). Error bars represent standard error.

Discussion

Our results show that hippocampal representations of a conditioning task are organized in a manner that supports generalization across spatial contexts and consistency across individuals. Although spatial representations in CA1 remap robustly between environments, task-related population activity preserves a stable structure across contexts and animals. These findings indicate that task-relevant information is maintained at the level of population organization, providing a mechanism by which learned behaviors can be flexibly expressed despite changes in spatial context.

Task abstraction across environments

A central finding of this study is that hippocampal representations of a conditioning task remain stable across distinct spatial environments, even as place cell representations reorganize. This dissociation between spatial remapping and task stability suggests that task-related information is not simply inherited from spatial coding, but is supported by an additional representational structure at the population level. Such organization allows the hippocampus to support behavioral generalization when task demands remain constant across environments.

These results align with theoretical frameworks proposing that hippocampal representations extend beyond physical space to include abstract task structure and relational information[46]. In such models, spatial and task-related variables are jointly embedded within a higher-dimensional representational space, allowing learned behaviors to be applied in novel environments that share cognitive demands but differ in sensory or contextual features. Our findings provide empirical

support for this view by showing that hippocampal population activity preserves task-related structure across environments when task contingencies are unchanged.

Importantly, the conditioning task used here was explicitly designed to be identical across contexts, enabling us to isolate generalization of task representations from contextual discrimination. Under these conditions, hippocampal population activity maintained task-related structure across environments, suggesting that the hippocampus can abstract task-relevant information from surrounding sensory details when such abstraction is behaviorally appropriate[47-49,55,56].

Pattern separation and completion reflect task demands

The dissociation between spatial remapping and task stability in our data speaks to longstanding debates regarding hippocampal pattern separation and pattern completion[46-48,50-56]. Robust remapping of place cell representations across environments is consistent with pattern separation, supporting the formation of distinct contextual representations and reducing interference between memories. At the same time, the preservation of task-related representations across environments reflects pattern completion, allowing a learned response to be expressed despite changes in contextual cues.

Our results suggest that the balance between pattern separation and pattern completion depends on whether task demands change across contexts, rather than on contextual change alone. When task demands are preserved, as in our study, hippocampal representations favor generalization of task-related information. In contrast, studies in which task contingencies differ across environments have reported strong representational divergence, even when behavioral demands

appear superficially similar[54,57-60]. Together, these findings indicate that hippocampal representations flexibly balance separation and completion in a manner that reflects task structure.

Population-level organization supports non-spatial representations

There remains active debate regarding the extent to which hippocampal pyramidal neurons encode non-spatial information[31,61-63]. Importantly, the stability of task-related population structure observed here cannot be explained by the engagement of a distinct, task-specific subpopulation. Instead, prior analyses demonstrate substantial overlap between neurons encoding spatial and task-related information[31], indicating that task generalization emerges from shared population resources rather than segregated cell classes.

Specifically, task-related activity formed stable low-dimensional population structures across environments, despite changes in spatial coding. Such population-level organization offers a solution to the problem of mixed selectivity at the single-cell level, allowing different task and contextual variables to be represented along distinct dimensions within a shared population space[64]. Similar organizational principles have been observed in other brain regions, including prefrontal and cingulate cortex[65-67], and in artificial neural networks trained on context-dependent tasks[68]. Our findings extend this population-level framework to hippocampal representations of learned tasks across contexts.

Task representations are conserved across animals

Beyond stability across environments, we found that task-related population organization was conserved across animals. Notably, the similarity of task-related population geometry across

individuals was comparable to the similarity observed across sessions within the same animal. This result suggests that hippocampal task representations follow shared organizational principles across animals, rather than being entirely idiosyncratic to individual experience.

Such conservation is striking given the hippocampus's established role in episodic memory and its sensitivity to individual experience[69-73]. Prior studies reporting conserved neural dynamics across animals have largely focused on systems associated with motor control or basic motivational states

[74-82], which are often linked to stereotyped behaviors [79,81,83,84]. In contrast, hippocampal-dependent tasks involve flexible cognitive processes shaped by learning and context[85-88]. The observation that hippocampal task representations exhibit consistent population-level organization across animals suggests that common representational strategies may underlie flexible cognitive functions.

Conclusion

Together, our findings indicate that the hippocampus supports behavioral generalization by organizing task-related activity into stable population-level structures that persist across contexts and individuals. This organization allows learned behaviors to be flexibly expressed despite changes in spatial representations, providing a neural basis for cognitive flexibility. Rather than encoding task information solely at the level of individual neurons, the hippocampus appears to rely on population-level geometry to preserve task structure across varying conditions.

These results add to a growing body of work emphasizing population-level organization as a fundamental principle of neural computation and suggest that conserved representational geometry may support flexible cognition. By demonstrating how task representations are maintained across environments and animals, this study advances our understanding of how the hippocampus integrates learning, memory, and generalization.

Citations

1 Bouton, M. E. A learning theory perspective on lapse, relapse, and the maintenance of behavior change. *Health psychology* **19**, 57 (2000).

2 Herszage, J. & Censor, N. Modulation of learning and memory: a shared framework for interference and generalization. *Neuroscience* **392**, 270-280 (2018).

3 Linda, Q. Y., Wilson, R. C. & Nassar, M. R. Adaptive learning is structure learning in time. *Neuroscience & Biobehavioral Reviews* **128**, 270-281 (2021).

4 Censor, N. Generalization of perceptual and motor learning: a causal link with memory encoding and consolidation? *Neuroscience* **250**, 201-207 (2013).

5 Brown, S. & Bebko, J. Generalization, overselectivity, and discrimination in the autism phenotype: A review. *Research in Autism Spectrum Disorders* **6**, 733-740 (2012).

6 Church, B. A. *et al.* Learning, plasticity, and atypical generalization in children with autism. *Psychonomic Bulletin & Review* **22**, 1342-1348 (2015).

7 Weiler, J. A., Bellebaum, C., Brüne, M., Juckel, G. & Daum, I. Impairment of probabilistic reward-based learning in schizophrenia. *Neuropsychology* **23**, 571 (2009).

8 Shohamy, D. *et al.* Learning and generalization in schizophrenia: effects of disease and antipsychotic drug treatment. *Biological psychiatry* **67**, 926-932 (2010).

9 Anastasides, N. *et al.* Increased generalization of learned associations is related to re-experiencing symptoms in veterans with symptoms of post-traumatic stress. *Stress* **18**, 484-489 (2015).

10 Lis, S. *et al.* Generalization of fear in post-traumatic stress disorder. *Psychophysiology* **57**, e13422 (2020).

11 Calandreau, L., Desgranges, B., Jaffard, R. & Desmedt, A. Switching from contextual to tone fear conditioning and vice versa: the key role of the glutamatergic hippocampal-lateral septal neurotransmission. *Learn Mem* **17**, 440-443 (2010). https://doi.org:10.1101/lm.1859810

12 Desmedt, A., Garcia, R. & Jaffard, R. Vasopressin in the lateral septum promotes elemental conditioning to the detriment of contextual fear conditioning in mice. *Eur J Neurosci* **11**, 3913-3921 (1999).

13 Vouimba, R. M., Garcia, R. & Jaffard, R. Opposite effects of lateral septal LTP and lateral septal lesions on contextual fear conditioning in mice. *Behav Neurosci* **112**, 875-884 (1998).

14 Anagnostaras, S. G., Gale, G. D. & Fanselow, M. S. Hippocampus and contextual fear conditioning: recent controversies and advances. *Hippocampus* **11**, 8-17 (2001). https://doi.org:10.1002/1098-1063(2001)11:1<8::AID-HIPO1015>3.0.CO;2-7

15 Honey, R. C. & Good, M. Selective hippocampal lesions abolish the contextual specificity of latent inhibition and conditioning. *Behav Neurosci* **107**, 23-33 (1993). https://doi.org:10.1037//0735-7044.107.1.23

16 Plitt, M. H. & Giocomo, L. M. Experience-dependent contextual codes in the hippocampus. *Nat Neurosci* **24**, 705-714 (2021). https://doi.org:10.1038/s41593-021-00816-6

17 Devito, L. M., Kanter, B. R. & Eichenbaum, H. The hippocampus contributes to memory expression during transitive inference in mice. *Hippocampus* **20**, 208-217 (2010). https://doi.org:10.1002/hipo.20610

18 Schimanski, L. A., Lipa, P. & Barnes, C. A. Tracking the course of hippocampal representations during learning: when is the map required? *J Neurosci* **33**, 3094-3106 (2013). https://doi.org:10.1523/JNEUROSCI.1348-12.2013

19 Gauthier, J. L. & Tank, D. W. A Dedicated Population for Reward Coding in the Hippocampus. *Neuron* **99**, 179-193 e177 (2018). https://doi.org:10.1016/j.neuron.2018.06.008

20 Mamad, O. *et al.* Place field assembly distribution encodes preferred locations. *PLoS Biol* **15**, e2002365 (2017). https://doi.org:10.1371/journal.pbio.2002365

21 Hok, V. *et al.* Goal-related activity in hippocampal place cells. *J Neurosci* **27**, 472-482 (2007). https://doi.org:27/3/472 [pii] 10.1523/JNEUROSCI.2864-06.2007

22 Knierim, J. J. Dynamic interactions between local surface cues, distal landmarks, and intrinsic circuitry in hippocampal place cells. *J Neurosci* **22**, 6254-6264 (2002). https://doi.org:20026608

23 Allen, K., Rawlins, J. N., Bannerman, D. M. & Csicsvari, J. Hippocampal place cells can encode multiple trial-dependent features through rate remapping. *J Neurosci* **32**, 14752-14766 (2012). https://doi.org:10.1523/JNEUROSCI.6175-11.2012

24 Jeffery, K. J., Gilbert, A., Burton, S. & Strudwick, A. Preserved performance in a hippocampal-dependent spatial task despite complete place cell remapping. *Hippocampus* **13**, 175-189 (2003). https://doi.org:10.1002/hipo.10047

25 Colgin, L. L., Moser, E.I., May-Britt Moser. Understanding memory through hippocampal remapping. *Cell* (2008).

26 Christian, K. M. & Thompson, R. F. Neural substrates of eyeblink conditioning: acquisition and retention. *Learn Mem* **10**, 427-455 (2003). https://doi.org:10.1101/lm.59603

27 McEchron, M. D., Tseng, W. & Disterhoft, J. F. Single neurons in CA1 hippocampus encode trace interval duration during trace heart rate (fear) conditioning in rabbit. *J Neurosci* **23**, 1535-1547 (2003).

28 Wirtshafter, H. S. & Wilson, M. A. Locomotor and Hippocampal Processing Converge in the Lateral Septum. *Current biology : CB* **29**, 3177-3192 (2019). https://doi.org:10.1016/j.cub.2019.07.089

29 Berger, T. W., Rinaldi, P. C., Weisz, D. J. & Thompson, R. F. Single-unit analysis of different hippocampal cell types during classical conditioning of rabbit nictitating membrane response. *J Neurophysiol* **50**, 1197-1219 (1983). https://doi.org:10.1152/jn.1983.50.5.1197

30 Wirtshafter, H. S. & Wilson, M. A. Differences in reward biased spatial representations in the lateral septum and hippocampus. *eLife* **9** (2020). https://doi.org:10.7554/eLife.55252

31 Wirtshafter, H. S., Mehta, M. R., Solla, S. A. & Disterhoft, J. F. Hippocampal trace coding dominates and disrupts place coding. *bioRxiv*, 2026.2002.2017.706430 (2026). https://doi.org:10.64898/2026.02.17.706430

32 Aharoni, D. & Hoogland, T. M. Circuit Investigations With Open-Source Miniaturized Microscopes: Past, Present and Future. *Frontiers in cellular neuroscience* **13**, 141 (2019). https://doi.org:10.3389/fncel.2019.00141

33 Aharoni, D., Khakh, B. S., Silva, A. J. & Golshani, P. All the light that we can see: a new era in miniaturized microscopy. *Nat Methods* **16**, 11-13 (2019). https://doi.org:10.1038/s41592-018-0266-x

34 Silva, A. J. Miniaturized two-photon microscope: seeing clearer and deeper into the brain. *Light: Science & Applications* **6**, e17104-e17104 (2017).

35 Wirtshafter, H. S. & Disterhoft, J. F. In Vivo Multi-Day Calcium Imaging of CA1 Hippocampus in Freely Moving Rats Reveals a High Preponderance of Place Cells with Consistent Place Fields. *J Neurosci* (2022). https://doi.org:doi: 10.1523/JNEUROSCI.1750-21.2022

36 Cunningham, J. P. & Yu, B. Dimensionality reduction for large-scale neural recordings. *Nature neuroscience* **17**, 1500-1509 (2014).

37 Bassett, D. S. & Sporns, O. Network neuroscience. *Nature neuroscience* **20**, 353-364 (2017).

38 Fenton, A. A. Remapping revisited: how the hippocampus represents different spaces. *Nat Rev Neurosci* **25**, 428-448 (2024). https://doi.org:10.1038/s41583-024-00817-x

39 Sanders, H., Wilson, M. A. & Gershman, S. J. Hippocampal remapping as hidden state inference. *eLife* **9** (2020). https://doi.org:10.7554/eLife.51140

40 Bouton, M. E., Nelson, J. B. & Rosas, J. M. Stimulus generalization, context change, and forgetting. *Psychological bulletin* **125**, 171-186 (1999). https://doi.org:10.1037/0033-2909.125.2.171

41 Lovibond, P. F., Preston, G. & Mackintosh, N. Context specificity of conditioning, extinction, and latent inhibition. *Journal of Experimental Psychology: Animal Behavior Processes* **10**, 360 (1984).

42 Moyer, J. R., Jr., Deyo, R. A. & Disterhoft, J. F. Hippocampectomy disrupts trace eye-blink conditioning in rabbits. *Behav Neurosci* **104**, 243-252 (1990).

43 McEchron, M. D. & Disterhoft, J. F. Hippocampal encoding of non-spatial trace conditioning. *Hippocampus* **9**, 385-396 (1999). https://doi.org:10.1002/(SICI)1098-1063(1999)9:4<385::AID-HIPO5>3.0.CO;2-K

44 Woodruff-Pak, D. S. & Disterhoft, J. F. Where is the trace in trace conditioning? *Trends Neurosci* **31**, 105-112 (2008). https://doi.org:10.1016/j.tins.2007.11.006

45 Weiss, C. & Disterhoft, J. F. The impact of hippocampal lesions on trace-eyeblink conditioning and forebrain-cerebellar interactions. *Behav Neurosci* **129**, 512-522 (2015). https://doi.org:10.1037/bne0000061

46 Whittington, J. C. R. *et al.* The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. *Cell* **183**, 1249-1263 e1223 (2020). https://doi.org:10.1016/j.cell.2020.10.024

47 Wirtshafter, H. S. & Wilson, M. A. Artificial intelligence insights into hippocampal processing. *Front Comput Neurosci* **16**, 1044659 (2022). https://doi.org:10.3389/fncom.2022.1044659

48 Lehnert, L., Littman, M. L. & Frank, M. J. Reward-predictive representations generalize across tasks in reinforcement learning. *PLoS Comput Biol* **16**, e1008317 (2020). https://doi.org:10.1371/journal.pcbi.1008317

49 Collins, A. G. E. & Frank, M. J. Neural signature of hierarchically structured expectations predicts clustering and transfer of rule sets in reinforcement learning. *Cognition* **152**, 160-169 (2016).

50 Stachenfeld, K. L., Botvinick, M. M. & Gershman, S. J. The hippocampus as a predictive map. *Nat Neurosci* **20**, 1643-1653 (2017). https://doi.org:10.1038/nn.4650

51 Eichenbaum, H. Hippocampus: cognitive processes and neural representations that underlie declarative memory. *Neuron* **44**, 109-120 (2004).

52 Wills, T. J., Lever, C., Cacucci, F., Burgess, N. & O'Keefe, J. Attractor dynamics in the hippocampal representation of the local environment. *Science* **308**, 873-876 (2005). https://doi.org:10.1126/science.1108905

53 Colgin, L. L. *et al.* Attractor-map versus autoassociation based attractor dynamics in the hippocampal network. *J Neurophysiol* **104**, 35-50 (2010). https://doi.org:10.1152/jn.00202.2010

54 Smith, D. M. & Mizumori, S. J. Learning-related development of context-specific neuronal responses to places and events: the hippocampal role in context processing. *J Neurosci* **26**, 3154-3163 (2006). https://doi.org:10.1523/JNEUROSCI.3234-05.2006

55 Colgin, L. L. *et al.* Attractor-map versus autoassociation based attractor dynamics in the hippocampal network. *Journal of neurophysiology* **104**, 35-50 (2010).

56 Lee, I., Yoganarasimha, D., Rao, G. & Knierim, J. J. Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. *Nature* **430**, 456-459 (2004).

57 Bahar, A. S., Shirvalkar, P. R. & Shapiro, M. L. Memory-guided learning: CA1 and CA3 neuronal ensembles differentially encode the commonalities and differences between situations. *J Neurosci* **31**, 12270-12281 (2011). https://doi.org:10.1523/JNEUROSCI.1671-11.2011

58 Markus, E. J. *et al.* Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J Neurosci* **15**, 7079-7094 (1995).

59 Singer, A. C., Karlsson, M. P., Nathe, A. R., Carr, M. F. & Frank, L. M. Experience-dependent development of coordinated hippocampal spatial activity representing the similarity of related locations. *J Neurosci* **30**, 11586-11604 (2010). https://doi.org:10.1523/JNEUROSCI.0926-10.2010

60 McKenzie, S. *et al.* Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. *Neuron* **83**, 202-215 (2014). https://doi.org:10.1016/j.neuron.2014.05.019

61 Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M. & Tanila, H. The hippocampus, memory, and place cells: is it spatial memory or a memory space? *Neuron* **23**, 209-226 (1999). https://doi.org:S0896-6273(00)80773-4 [pii]

62 O'Keefe, J. & Krupic, J. Do hippocampal pyramidal cells respond to nonspatial stimuli? *Physiol Rev* **101**, 1427-1456 (2021). https://doi.org:10.1152/physrev.00014.2020

63 Knierim, J. J. From the GPS to HM: Place cells, grid cells, and memory. *Hippocampus* **25**, 719-725 (2015). https://doi.org:10.1002/hipo.22453

64 Nieh, E. H. *et al.* Geometry of abstract learned knowledge in the hippocampus. *Nature* **595**, 80-84 (2021).
65 Johnston, W. J., Fine, J. M., Yoo, S. B. M., Ebitz, R. B. & Hayden, B. Y. Semi-orthogonal subspaces for value mediate a binding and generalization trade-off. *Nature Neuroscience*, 1-13 (2024).
66 Machens, C. K., Romo, R. & Brody, C. D. Functional, but not anatomical, separation of "what" and "when" in prefrontal cortex. *Journal of Neuroscience* **30**, 350-360 (2010).
67 Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *nature* **503**, 78-84 (2013).
68 Flesch, T., Juechems, K., Dumbalska, T., Saxe, A. & Summerfield, C. Orthogonal representations for robust context-dependent task performance in brains and neural networks. *Neuron* **110**, 1258-1270. e1211 (2022).
69 Poulter, S., Hartley, T. & Lever, C. The neurobiology of mammalian navigation. *Current Biology* **28**, R1023-R1042 (2018).
70 Tucker, M. A., Ord, T. J. & Rogers, T. L. Evolutionary predictors of mammalian home range size: body mass, diet and the environment. *Global Ecology and Biogeography* **23**, 1105-1114 (2014).
71 Schilder, B. M., Petry, H. M. & Hof, P. R. Evolutionary shifts dramatically reorganized the human hippocampal complex. *Journal of Comparative Neurology* **528**, 3143-3170 (2020).
72 Sherry, D. F., Jacobs, L. F. & Gaulin, S. J. Spatial memory and adaptive specialization of the hippocampus. *Trends in neurosciences* **15**, 298-303 (1992).
73 Pravosudov, V. V. & Roth II, T. C. Cognitive ecology of food hoarding: the evolution of spatial memory and the hippocampus. *Annual Review of Ecology, Evolution, and Systematics* **44**, 173-193 (2013).
74 Safaie, M. *et al.* Preserved neural dynamics across animals performing similar behaviour. *Nature* **623**, 765-771 (2023). https://doi.org:10.1038/s41586-023-06714-0
75 Gallego, J. A. *et al.* Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nature communications* **9**, 1-13 (2018).
76 Talpir, I. & Livneh, Y. Stereotyped goal-directed manifold dynamics in the insular cortex. *Cell reports* **43** (2024).
77 Harris, K. D. & Shepherd, G. M. The neocortical circuit: themes and variations. *Nature neuroscience* **18**, 170-181 (2015).
78 Shmuelof, L. & Krakauer, J. W. Are we ready for a natural history of motor learning? *Neuron* **72**, 469-476 (2011).
79 Manoli, D. S., Meissner, G. W. & Baker, B. S. Blueprints for behavior: genetic specification of neural circuitry for innate behaviors. *Trends in neurosciences* **29**, 444-451 (2006).
80 Tierney, A. Evolutionary implications of neural circuit structure and function. *Behavioural processes* **35**, 173-182 (1995).
81 LeDoux, J. E. As soon as there was life, there was danger: the deep history of survival behaviours and the shallower history of consciousness. *Philosophical Transactions of the Royal Society B* **377**, 20210292 (2022).

82 Martínez-García, F. & Lanuza, E. Evolution of vertebrate survival circuits. *Current Opinion in Behavioral Sciences* **24**, 113-123 (2018).
83 Kogan, E., Lu, J. & Zuo, Y. Cortical circuit dynamics underlying motor skill learning: from rodents to humans. *Frontiers in Molecular Neuroscience* **16**, 1292685 (2023).
84 LeDoux, J. & Daw, N. D. Surviving threats: neural circuit and computational implications of a new taxonomy of defensive behaviour. *Nature Reviews Neuroscience* **19**, 269-282 (2018).
85 Buzsáki, G. & Moser, E. I. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nature neuroscience* **16**, 130-138 (2013).
86 Sweis, B. M., Mau, W., Rabinowitz, S. & Cai, D. J. Dynamic and heterogeneous neural ensembles contribute to a memory engram. *Curr Opin Neurobiol* **67**, 199-206 (2021). https://doi.org:10.1016/j.conb.2020.11.017
87 Davidson, T. J., Kloosterman, F. & Wilson, M. A. Hippocampal replay of extended experience. *Neuron* **63**, 497-507 (2009). https://doi.org:10.1016/j.neuron.2009.07.027 S0896-6273(09)00582-0 [pii]
88 Rubin, R. D., Watson, P. D., Duff, M. C. & Cohen, N. J. The role of the hippocampus in flexible cognition and social behavior. *Frontiers in human neuroscience* **8**, 742 (2014).

651
652 **Methods**

653 **LEAD CONTACT AND MATERIALS AVAILABILITY**

654 Questions and requests for information should be directed to and will be fulfilled by the
655 Lead Contact, Hannah Wirtshafter (hsw@northwestern.edu). This study did not generate new
656 unique reagents. The data that support the findings of this study are available from the
657 corresponding author.

658

659 **EXPERIMENTAL MODEL AND SUBJECT DETAILS**

660 All procedures were performed within Northwestern Institutional Animal Care and Use
661 Committee and NIH guidelines. Five male Long Evans rats (275–325 g) were sourced from
662 Charles River Laboratories, injected with AAV9-GCaMP8m, implanted with a 2-mm GRIN lens,
663 and trained and tested on eyeblink conditioning in two apparatuses (Fig. 1). Animals were
664 individually housed in an animal facility with a 12/12 h light/dark cycle.

665

666 **METHOD DETAILS**

667 **GCaMP7c injection, lens implantation, EMG implantation**

668 GCaMP8 injection and lens implantation were completed as reported in Wirtshafter and
669 Disterhoft, 2022 and Wirtshafter and Disterhoft, 2023[1,2]. Briefly, rats were anesthetized with
670 isoflurane (induction 4%, maintenance 1-2%) and a craniotomy was performed at stereotaxic
671 coordinates Bregma AP −4.00mm, ML 3.00mm. 0.06uL of GCaMP8m (obtained from
672 AddGene, packaged AAV9 of pGP-AAV-syn-jGCaMP8m-WPRE, lot v175525, titer 1.3E+13
673 GC/mL) was injected over 12 minutes (approximate coordinates Bregma AP −4.00mm, ML
674 3mm, DV 2.95mm relative to skull); then the syringe was raised 0.2mm and an additional 0.6ul

675 of GCaMP8 was injected. We repeated this process once more and at slightly different
676 coordinates in the craniotomy hole, resulting in 4 total injections.
677 We then aspirated tissue from the craniotomy site using a vacuum pump and 25 gauge
678 needle. Tissue was aspirated up to and including the horizontal striations of the corpus collosum.
679 A 2mm GRIN lens (obtained from Go!Foton, CLH lens, 2.00mm diameter, 0.448 pitch, working
680 distance 0.30mm, 550nm wavelength) was then inserted into the craniotomy hole and cemented
681 in place using dental acrylic. Animals were given buprenorphine (0.05mg/kg) and 20mL saline,
682 taken off anesthesia, and allowed to recover in a clean cage placed upon a heat pad.
683 Six to eight weeks after surgery, animals were again anesthetized with isoflurane and
684 checked for GCaMP expression. If expression was seen, baseplates were attached using UV-
685 curing epoxy and dental acrylic. Electrode implantation to record obicularis oculi
686 electromyographic (EMG) activity occurred in the same surgery as baseplate attachment, as
687 described previously[3,4]. Briefly, a connector containing 5 wires was cemented on the front of the
688 animal's head: 4 wires were implanted directly above the eye in the surrounding muscle (2 for
689 recording, 2 for electrical stimulation). An additional wire was attached to a connector attached
690 to a ground screw located above the cerebellum; this screw was implanted during lens
691 implantation surgery.
692
693 **Behavioral environment and training**
694 Two behavioral apparatuses were used in these experiments: Environment A was a
695 78.7cm x 50.8cm unscented rectangular enclosure with wire floor and walls and white lighting.
696 Environment B was a 50.1cm x 34.9cm scented (with two dabs of clove essential oil on opposite
697 walls) ovular enclosure with white solid floor and walls, and red lighting. Both environments
698 were located at the same spot in the room relative to external cues (see Figures 1b and S1).

A tether containing a plug to relay the EMG activity and to deliver a shock to the rat's eye was attached to a the eyeblink connector on the rat's head. The miniscope was plugged into the cemented baseplate. The miniscope and EMG cords were all attached to a commutator for ease of animal movement.

The CS was a 250ms, 85dB free-field tone (5ms rise-fall time). The US was a 100ms shock directed to the left eye. Shock amount varied per session per animal and was calibrated, if needed, at the end of a training session for the next session's training. Shock level was deemed appropriate when a shock was met with a firm shake of the animal's head.

The trace interval was 500ms and the intertrial interval (ITI) was randomized between 30s and 60s, with a 45s average. EMG signal output was amplified (5000×) and filtered (100 Hz to 5 kHz), then digitized at 3 kHz and stored by computer.

A conditioned response (CR) was identified as an increase in integrated EMG activity that exceeded the baseline mean amplitude by more than four standard deviations, sustained for a minimum duration of 15ms. Baseline mean amplitude was calculated during the 500ms preceding CS onset. Additionally, the response had to commence at least 50ms after the conditioned stimulus (CS) onset and before the unconditioned stimulus (US) onset.

The animal's first exposure to each environment was a 38min exploration session, in which the animal was able to freely move and explore the environment without any conditioning (Figure 1a). Animals were then trained in one environment per session, with no more than one session per day, and were considered to have learned the task after reaching criterion (70% CRs in 50 trials) on three consecutive training sessions (termed 'criterion sessions') or when the previous four training sessions averaged over 70% (in this instance, only the final three of those sessions were considered 'criterion sessions'). Following the last session in environment A, the animal was given an exploratory session in environment B. The session after that, the animal was

723 tested on eye blink conditioning in environment B, using the same parameters as used in
724 environment A.
725
726 **Calcium imaging**
727 Calcium imaging was completed as reported in Wirtshafter and Disterhoft, 2022 and
728 Wirtshafter and Disterhoft, 2023[1,2]. Briefly, calcium imaging was done using UCLA V4
729 Miniscopes[5,6], assembled with two 3mm diameter, 6mm FL achromat lens used in the objective
730 module and one 4mm diameter, 10mm FL achromat lens used in the emission module.
731
732 **QUANTIFICATION AND STATISTICAL ANALYSIS**
733 Means are presented as mean+-standard deviation. All analysis code is available
734 at https://github.com/hsw28/ca_imaging and https://github.com/hsw28/Hannahs-CEBRAs. Code
735 to create specific figures is also available at the former github repository.
736
737 **Position and speed analysis**
738 Position was sampled by an overhead camera at 30Hz. Position tracking was done post-
739 recording using DeepLabCut[7]. Position was then converted from pixels to cm. Position was
740 smoothed using a Gaussian filter with standard deviation of 2cm. Speed was calculated by taking
741 the hypotenuse of the coordinates one before and after the time of interest.
742
743 **Video pre-processing and cell identification**
744 Video pre-processing and cell identification were performed as reported in Wirtshafter
745 and Disterhoft, 2022 and Wirtshafter and Disterhoft, 2023[1,2]. In brief, videos were recorded with
746 Miniscope software at 15frames/second. Video processing was done using CIATAH software[8].

747 Videos were down sampled in space and normalized by subtracting the mean value of each
748 frame from the frame. Each frame was then normalized using a bandpass FFT filter (70-
749 100cycles/pixel) and motion corrected to a using TurboReg[9]. Videos were then converted to
750 relative florescence (dF/F_0); F_0 was the mean over the entire video.

751 Cells were automatically identified using CIATAH[8] using CNMF-E[10]. Images were
752 filtered with a gaussian kernel of width 2 pixels and neuron diameter was set at a pixel size of 8.
753 The threshold for merging neurons was set at a calcium trace correlation of 0.65; neurons were
754 merged if their distances were smaller than 4 pixels and they had highly correlated spatial shapes
755 (correlation>0.8) and small temporal correlations (correlation <0.4).

756 In vivo calcium imaging involves detecting changes in intracellular calcium levels, which
757 serve as proxies for neuronal activity. Calcium events refer to transient increases in calcium
758 concentration above a threshold level; these crossings putatively correspond to spikes in neuronal
759 firing. These events typically appear as peaks in the data and indicate an active response from the
760 neuron. Calcium traces are continuous recordings of calcium levels over time. Thus, calcium
761 events highlight specific neuronal activations, while calcium traces provide a full temporal
762 picture of these activations together with baseline activity.

763 All cells identified using CNMF-E were then scored as neurons or not by a human scorer.
764 Scoring was also done within CIATAH software in a Matlab GUI. Scoring was done while
765 visualizing and considering a calcium activity trace, average waveform, a montage of the
766 candidate cell's Ca2+ events, and a maximum projection of all cells on which the candidate cell
767 was highlighted. The relative fluorescence ($\Delta F/F_0$) local maxima of each identified cell were
768 considered calcium event times.

769

770 **Cell cross registration across sessions and within session**

771 Validation and registration were completed as documented in Wirtshafter and Disterhoft[2].
772 Briefly, videos underwent five rounds of registration using Turboreg image rotation[9] with the
773 CIATAH software[8,11]. Background noise, axons, and dendrites were removed using an image
774 binarization threshold of 40% of the images' maximum value. Cells were matched across
775 sessions using a distance threshold of a maximum of five pixels, with a minimum 2-D correlation
776 coefficient of 0.5. Sessions were aligned to session A(n), the last session in environment A.
777 Calcium imaging (CaImg) enabled the longitudinal monitoring of the same hippocampal
778 cells over multiple sessions in both environments. For criterion and testing sessions, we observed
779 an average of 459.85±265.31 cells per session per animal, with no significant difference between
780 the number of cells recorded in environment A and environment B (two-tailed t-test(24)= -0.56,
781 $p>0.05$) On average, 132±95 cells were present in both the last criterion session in A and the first
782 testing session in B. This was not a significantly different numbers of cells that were present, on
783 average, in both the semi-final session in A, session A(n-1), and the final session in A, session
784 A(n) (155±115 cells).
785
786 **Place cell identification and computing spatial mutual information**
787 Place cells were identified using mutual information computed when the animals were
788 running at speeds greater than or equal to 4cm/s. MI was computed for all cells; there was no
789 calcium event rate criterion for included cells. To be considered significant, the computed mutual
790 information (MI) must be greater than 95% of MI scores computed 500 times from shuffled
791 positions[12]. To compute the MI for each cell, the training environments were divided into 2.5cm
792 x 2.5cm bins. The calcium event rate of each cell and the occupancy of the animal were found
793 for each bin. Rate and occupancy were smoothed with a Gaussian kernel with filter width of 3cm

794 and Sigma of 0.5cm. Mutual information was computed during periods of movement as

795 follows[2,12,13]:

796 $$p = \frac{P_s}{P_o}$$

797 $$M_s = \sum P_s$$

798 $$M_o = \sum P_o$$

799 $$MI = \sum p * \log_2\left(\frac{p}{M_s * M_o}\right)$$

800 where:

801 P_s = calcium event probability in each bin

802 P_o = occupancy probability at each bin

803

804 Mutual information using calcium traces was computed as above, except instead of P_s

805 being calcium event probability per bin, the value of P_s was the average value of calcium trace in

806 the bin.

807 We computed MI using both calcium events and calcium trace data. There was no

808 significant difference between the number of place cells detected using calcium event data and

809 calcium trace data, (paired t-test $t(24)=1.01$, $p>0.05$). Note that all place cell and place field

810 measurements are presented with conditioning periods included, as the animal was frequently

811 moving during conditioning periods. We also computed results while excluding conditioning

812 periods and found no significant differences.

813

814 **Computing CSUS mutual information**

815 The computation of CSUS mutual information was very similarly to that for spatial

816 mutual information. A 1.3 second period beginning at the start of the CS tone was either divided

817 into 2 bins (CSUS-MI2) or 5 bins (CSUS-MI5) (Fig. 3c-3d). Mutual information was then

818 computed using the following:

$$p = \frac{P_s}{P_o}$$

$$M_s = \sum P_s$$

$$M_o = \sum P_o$$

$$MI = \sum p * \log_2\left(\frac{p}{M_s * M_o}\right)$$

Where:

$$P_s = \text{calcium event probability in each CSUS bin}$$

$$P_o = \text{probability of individual CSUS occuring out of all CSUS bins}$$

Mutual information using calcium traces was computed as above, except that P_s did not represent the calcium event probability per bin, but the average value of calcium trace within the bin.

Remapping quantification

The place cell center was defined as the occupancy-normalized location with maximum number of calcium events while the animal was moving at 5cm/s or faster. Position was binned into 2.5cm square bins. The place cell centers at environments A and B, as well as at environment A across days and at environment B across days, were used to align each environment across days, as well as to align environment A to environment B.

Population vector correlation was calculated between two environments using calcium event data. Neurons present in both datasets (such as sessions A(n) and A(n-1), or A(n) and B(1)) were identified and their calcium event times were converted to rates using 0.75 second binning. These firing rates were then normalized using z-score normalization across each neuron's activity across time. The mean calcium event rate for each neuron in each environment was then computed. The correlation among population vectors of firing rates in the two environments was cpmputed:

$$population\ vector\ correlation = \frac{\sum_{i=1}^{N}\left(r_{\{i,A\}} - \overline{\{r\}}_A\right)\left(r_{\{i,B\}} - \overline{\{r\}}_B\right)}{\sqrt{\sum_{i=1}^{N}\left(r_{\{i,A\}} - \overline{\{r\}}_A\right)^2}\sqrt{\sum_{i=1}^{N}\left(r_{\{i,B\}} - \overline{\{r\}}_B\right)^2}}$$

Where:

$r_{\{i,A\}}$ and $r_{\{i,B\}}$ = firing rates of neuron *i* in environments A and B

$\overline{\{r\}}_A$ and $\overline{\{r\}}_B$ = mean firing rates across neurons in environments A and B

The factors in the denominator correspond to the standard deviation of the components of each population vector relative to their mean, computed in each environment.

Use of CEBRA versus alternative methods

We explored multiple different methods before settling on the use of CEBRA for this study. A short summary of each tested method can be found below:

- **Principal Component Analysis (PCA)**[14]: Principal component analysis (PCA) is a statistical method used to reduce the dimensionality of data while retaining as much variability as possible. This linear technique identifies the axes (principal components) in the dataset that maximize variance. The first principal component explains the most variance, the second explains the second most, and so on. Principal components are combinations of original features and may not always have clear or intuitive meanings. In agreement with previous hippocampal data[15], PCA required upwards of 15-25 components to capture 95% of the variance of the data. In addition, across and within all sessions and representations (spatial and task representations), the manifolds spanned by the largest PCs remained highly similar, with small principal angles in pairwise comparisons. This similarity in the orientation of the leading subspaces suggested that PCA did not distinguish between spatial or behavioral components of the task (Figure S5).

- **Independent Component Analysis (ICA)**[16]: Independent Component Analysis (ICA) is a computational technique used to separate a multivariate signal into additive, independent components. ICA operates under the assumption that observed data are linear mixtures of underlying, independent sources. It aims to find a linear transformation that maximizes the statistical independence of the estimated components. We found that ICA embeddings were unstable throughout the length of the recordings, and also did not clearly map onto behavioral states (Figure S6).
- **Isomap**[17]: Isomap is a manifold learning technique that seeks to capture the intrinsic geometric structure of data. Isomap is useful when linear methods like PCA cannot capture the intrinsic structure of the data, as it preserves the geodesic (curved) distances in the reduced dimensionality space. Unlike linear methods such as PCA, Isomap can capture nonlinear relationships in the data. Interestingly, using Isomap, only about 5 neural modes were required to achieve a residual variance of 5-10%. However, the embedding shape did not relate to any discernable property of neural data or behavior (Figure S7). Dimensionality reduction was achieved, but the resulting representations were not interpretable (Figure S7).
- **MIND**[15,18]: MIND is a decoding method designed for integrating multiple data modalities to predict various features, particularly sensory and motor functions. MIND uses recurrent neural networks whose hidden variables provide a memory mechanism for remembering previous inputs; this approach is particularly apt for the analysis of time series data such as neural recordings. While MIND was very robust at distinguishing the different environments, it was not equipped to handle relatively short signals separated in time, such as the conditioning trials separated by intertrial intervals. Our analyses using MIND resulted in poor and unstable embeddings that could not be analyzed (Figure S8).

896 CEBRA[19] was chosen for this project for its ability to capture nonlinear relationships in the data
897 and to create stable embeddings over short and long time periods. Additionally, spatial
898 separations of components were well isolated and correlated well with observed behaviors.
899
900 **Use of CEBRA for position decoding**
901 Optimal parameters for decoding the position of each animal from neural activity were
902 determined using an extensive grid search across learning rate, temperature, and number of
903 iterations (Figure S9). Models created to compare different sessions, such as a model trained on
904 data from session A(n) used to decode session B(1), were only trained on cells that occurred in
905 both sessions. Models were trained on spike traces of these cells, labeled with the animal's (X,Y)
906 position. In all cases, 75% of data was used to train the model while 25% of data was held out for
907 verification. All models were run 500 times. Optimal embeddings were determined based on the
908 minimum median absolute error between the predicted and true positions. The optimal
909 parameters for each rat are as follows:

	Rat 1	Rat 2	Rat 3	Rat 4	Rat 5
Model Architecture	'Offset10-model'				
Batch size	512				
Learning rate	$5.5*10^{-5}$	$6.625*10^{-4}$	$5.5*10^{-4}$	$1.0*10^{-3}$	$1.0*10^{-3}$
Temperature mode	'Auto'				
Minimum temperature	No minimum	1.5	0.95	$1.0*10^{-9}$	No minimum
Output dimensions (# of latents)	3				
Max iterations	25000	8000	26500	30000	18000
Distance	'Cosine'				
Conditional	'Time delta'				
Number of hidden units	32				

Time offsets	1

The number of output dimensions was chosen based on the fewest number of dimensions under which all 5 models consistently outperformed shuffled data for both position and conditioning decoding (Figure S10-12).

Use of CEBRA for conditioning decoding

As in position decoding, the optimal parameters for decoding conditioning were determined for each animal using an extensive grid search across learning rate, temperature, and number of iterations (Figure S11). Models created to compare different sessions of neural activity, such as a model trained on data from session A(n) used to decode session B(1), were only trained on cells that occurred in both sessions. Models were trained on spike traces of these cells, with labels corresponding to the CSUS bin during which the signal occurred (either one out of 2 bins or out of 5 bins, see Figures 3c-d). In all cases, 75% of data was used to train the model while 25% of data was held out for verification. All models were run 500 times. Optimal embeddings were determined based on the percent of correctly binned time points. The optimal parameters for each rat are as follows:

	Rat 1	Rat 2	Rat 3	Rat 4	Rat 5
Model Architecture	‘Offset10-model’				
Batch size	512				
Learning rate	$3.5*10^{-3}$	$7.0*10^{-3}$	$3.5*10^{-3}$	$7.5*10^{-3}$	$9.5*10^{-3}$
Temperature mode	‘Constant’	‘Constant’	‘Auto’	‘Constant’	‘Constant’
Minimum temperature	2.33	1.75	1.67	1.67	2.66
Output dimensions (# of latents)	3				
Max iterations	50000	7500	20000	18000	25000

Distance	'Euclidian'	'Cosine'	'Cosine'	'Euclidian'	'Cosine'
Conditional	'Time delta'				
Number of hidden units	32				
Time offsets	1				

The number of output dimension was chosen based on the fewest number of dimensions under which all 5 models consistently outperformed shuffled data for both position and conditioning decoding (Figure S10-12). The parameters listed above were used for decoding into 2 or 5 bins, including the use of 3 output dimensions (# of latents).

The accuracy of results was computed from the entries in the confusion matrix:

$$Accuracy = \frac{Sum\ of\ number\ of\ diagonal\ (correct)\ elements}{Sum\ of\ number\ of\ total\ elements}$$

Precision was calculated for each class i, $1 \leq i \leq n$, where n is the number of classes:

$$Precision_i = \frac{(TP_i)}{TP_i + FP_i}$$

where:

$$TP = true\ positives$$
$$FP = false\ positives$$

The global precision is given by the average:

$$Precision = \frac{1}{n}\sum_{i=1}^{n} Precision_i$$

Recall, also known as sensitivity, was calculated for each class i, $1 \leq i \leq n$, where n is the number of classes:

$$Recall_i = \frac{(TP_i)}{TP_i + FN_i}$$

where:

$$TP = true\ positives$$
$$FN = false\ negatives$$

The global recall is given by the average:

$$Recall = \frac{1}{n}\sum_{i=1}^{n} Recall_i$$

The F1 score was calculated for each class i, $1 \leq i \leq n$, where n is the number of classes:

$$F1_i = 2 * \frac{Precision_i * Recall_i}{Precision_i + Recall_i}$$

The global F1 score is given by the average:

$$F1 = \frac{1}{n}\sum_{i=1}^{n} F1_i$$

The area under the receiver operating characteristic (ROC) curve was calculated for each class i, $1 \leq i \leq n$, where n is the number of classes:

$$AUC_i = roc_auc_score(y_{true-bin}[:, i], y_{pred-prob}\,[:, i])$$

The global value is given by the average:

$$ROC\ AUC = \frac{1}{n}\sum_{i=1}^{n} AUC_i$$

where:

$$y_{true-bin} = binarized\ true\ labels$$
$$y_{pred-prob} = predicted\ probabilities\ for\ each\ class$$

Model consistency

Model consistency was computed using a built-in CEBRA function which relies on the function 'sklearn.metrics.consistency_score'[20]. The function compares the embeddings from different models by calculating pairwise consistency scores. This comparison involves measuring the similarity of the embeddings using statistical metrics; i.e. this metric calculates how similar a model's labels are for similar instances in the data set.

To determine consistency between environments and across animals, the data was fit to each model 20 times. The model with the lowest loss was selected and compared to other models with the lowest loss. Models were created using each individual animal's optimal parameters (see above).

Acknowledgements: This work was supported by an NIA T32 (T32-AG020506/AG/NIA), an NIA R37 (R37-AG008796/AG/NIA), an NINDS R01 (R01 NS113804/NS/NINDS), and a K99 award (K99 MH135062).

This research was supported in part through the computational resources and staff assistance provided by the Quest high performance computing facility at Northwestern University, which is jointly supported by the Office of the Provost, the Office for Research, and Northwestern University Information Technology.

We would like to thank all members of the Disterhoft lab, especially Mackenzie Kneisly. We would like to thank the following individuals for their assistance with CEBRA: Mackenzie Mathis and Steffen Schneider. And additional thank you to David Wirtshafter for his feedback.

Author contributions: Investigation, H.S.W.; Formal Analysis, H.S.W.; Writing –Original Draft, H.S.W.; Writing – Review & Editing, H.S.W., S.A.S., J.F.D.; Funding: H.S.W, J.F.D., Supervision, S.A.S., J.F.D.

Declaration of interests: The authors declare no competing interests.

Additional Information

Supplementary Information is available for this paper.

Correspondence and requests for materials should be addressed to Hannah S Wirtshafter, hsw@northwestern.edu .

Reprints and permissions information is available at www.nature.com/reprints.

Supplementary figures

Figure S1

Figure S1. Photos of testing chambers.

Photos of testing chambers. Top: Environment A, an unscented rectangular enclosure with wire floor and walls, and white lighting. Bottom: Environment B, a scented ovular enclosure with white solid floor and walls, and red lighting. Both environments were located at the same spot in the room relative to external cues. The door to the chamber was closed during testing, to accentuate the distinction between the white and red lighting.

Figure S2

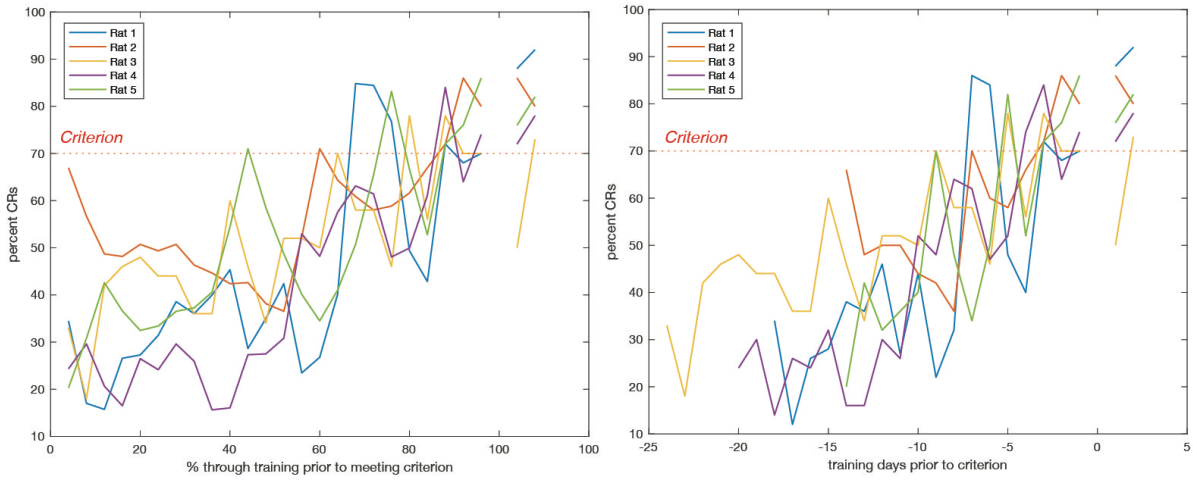

Figure S2. Learning curves for the five rats.

Substantial variability in the number of sessions required to learn the task; average number of sessions was 20 ± 4.2 (including the criterion sessions). The fastest learners (2 rats) reached criterion after 14 sessions, while the slowest rat required 24 sessions to reach criterion.

Figure S3

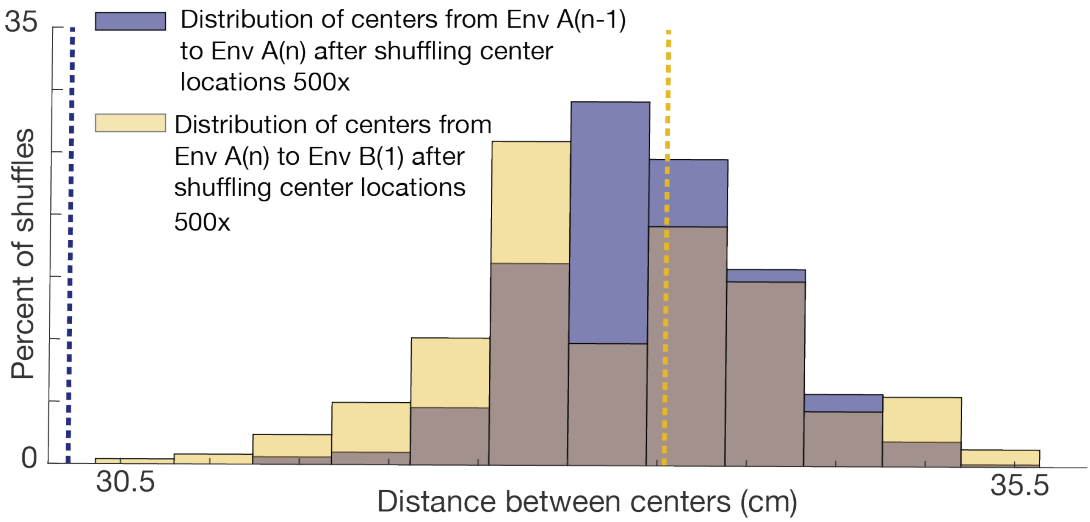

Figure S3. Place cells remap based on comparison to shuffled data.

Distribution of the median distance between place field centers after shuffling center locations 500 times. The actual median value for session A(n) to A(n-1) was smaller than all shuffled medians (p = 0, dashed blue line), while the median for session A(n) to B(1) was greater than or equal to 56% of shuffled medians (p = 0.56, dashed yellow line). The brown shaded area represents the overlap between shuffled distributions.

Figure S4

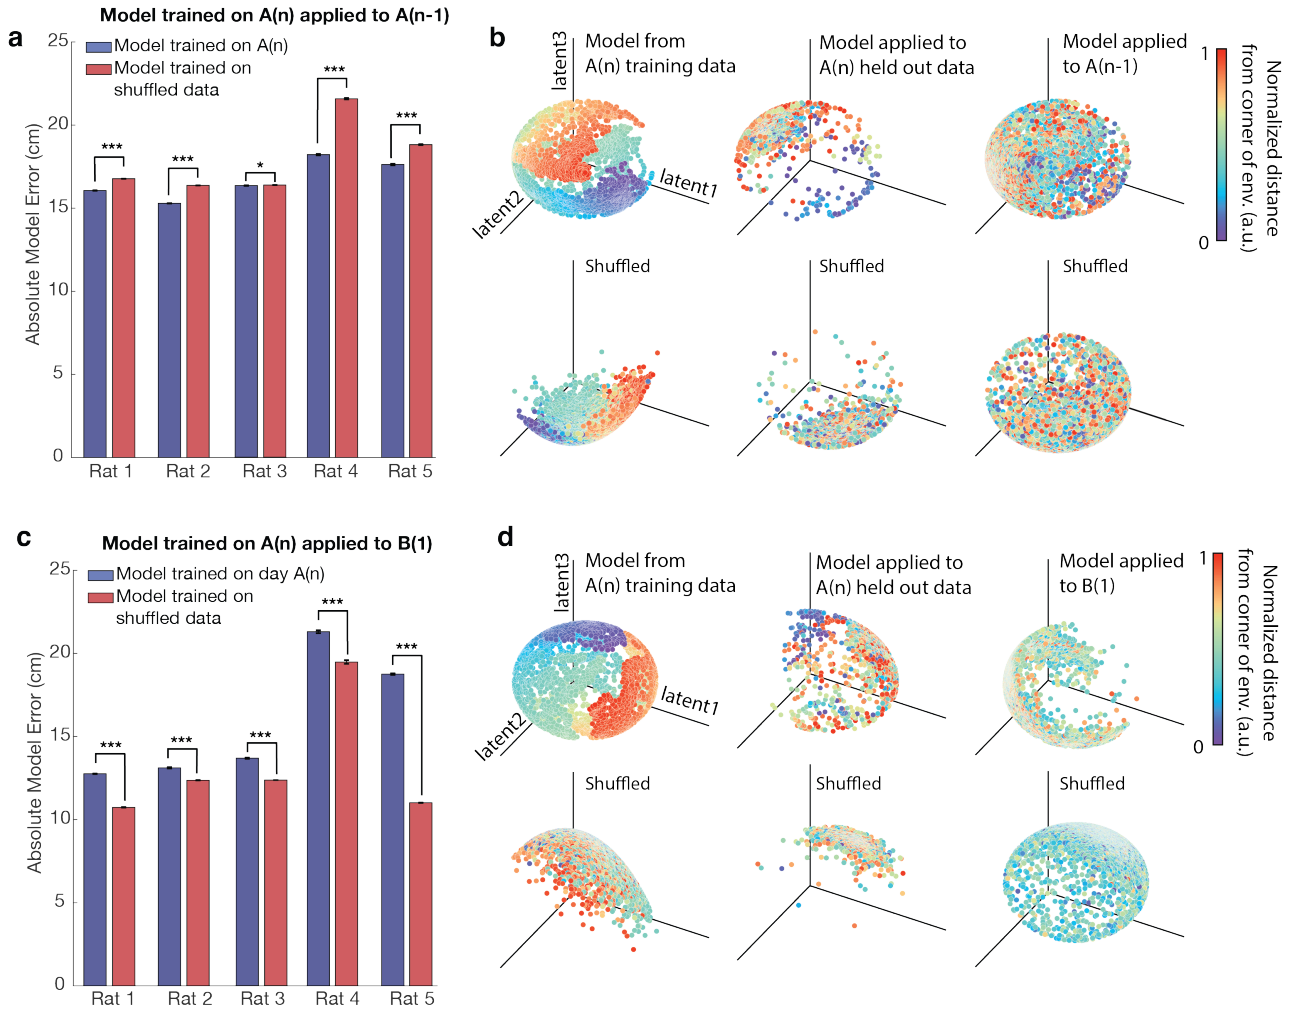

Figure S4. A model trained in environment A can decode positions within environment A but not in environment B.

a. A model trained on calcium trace and position data from session A(n), using cells present in both A(n) and A(n-1), predicted positions in environment A(n-1) with significantly greater accuracy than that of a model trained on shuffled position data. The model was run 500 times, and the accuracy of predictions was assessed. The model trained on actual

data significantly outperformed the shuffled model across rats (double-sided t-tests: Rat 1: $t(998) = -34.3$, $p = 5.0 \times 10^{-171}$; Rat 2: $t(998) = -72.5$, $p = 0$; Rat 3: $t(998) = -1.96$, $p = 0.05$; Rat 4: $t(998) = -53.7$, $p = 2.6 \times 10^{-299}$; Rat 5: $t(998) = -20.0$, $p = 4.1 \times 10^{-75}$). Error bars represent standard error. A single asterisk (*) indicates $p \leq 0.05$, and three asterisks (***) indicate $p < 10^{-74}$.

b. (Top row) Visualization of model performance for decoding animal position in environment A(n). Left: The trained model demonstrated on the training data from session A(n). Middle: The same model applied to held-out trace data (25%) from session A(n). Right: The model applied to predict animal position in session A(n-1). (Bottom row) The same models trained on shuffled position data. Shown here is the model for Rat 4, where the model trained on real data significantly outperformed model trained on shuffled data for decoding position in session A(n-1) (500 simulations, $t(998) = -53.7$, $p = 2.6 \times 10^{-299}$). For visualization purposes, distance from the corner of the environment is plotted using normalized values in arbitrary units [a.u.].

c. A model trained on data from session A(n), using cells present in both A(n) and B(1), was applied to environment B(1). The model's predictions were significantly less accurate than those of a model trained on shuffled position data (double-sided t-tests: Rat 1: $t(998) = 84.1$, $p = 0$; Rat 2: $t(998) = 18.8$, $p = 7.4 \times 10^{-68}$; Rat 3: $t(998) = 74.7$, $p = 0$; Rat 4: $t(998) = 13.1$, $p = 2.0 \times 10^{-36}$; Rat 5: $t(998) = 154.4$, $p = 0$). Error bars represent standard error. Three asterisks (***) indicate $p < 10^{-35}$.

d. (Top row) Model trained on position and calcium trace data from session A(n), using cells present in both A(n) and B(1). Left: The model demonstrated on the training data from session A(n). Middle: The same model applied to held-out trace data (25%) from session A(n). Right: The model applied to predict the animal position in session B(1).

(Bottom row) The same models trained on shuffled position data. Shown here is the model for Rat 4, where the shuffled model performed significantly better than the model trained on session A(n) for decoding position in session B(1) (500 simulations, $t(998) = 13.1$, $p = 2.0 \times 10^{-36}$). For visualization purposes, distance from the corner of the environment is plotted using normalized values in arbitrary units [a.u.].

Figure S5

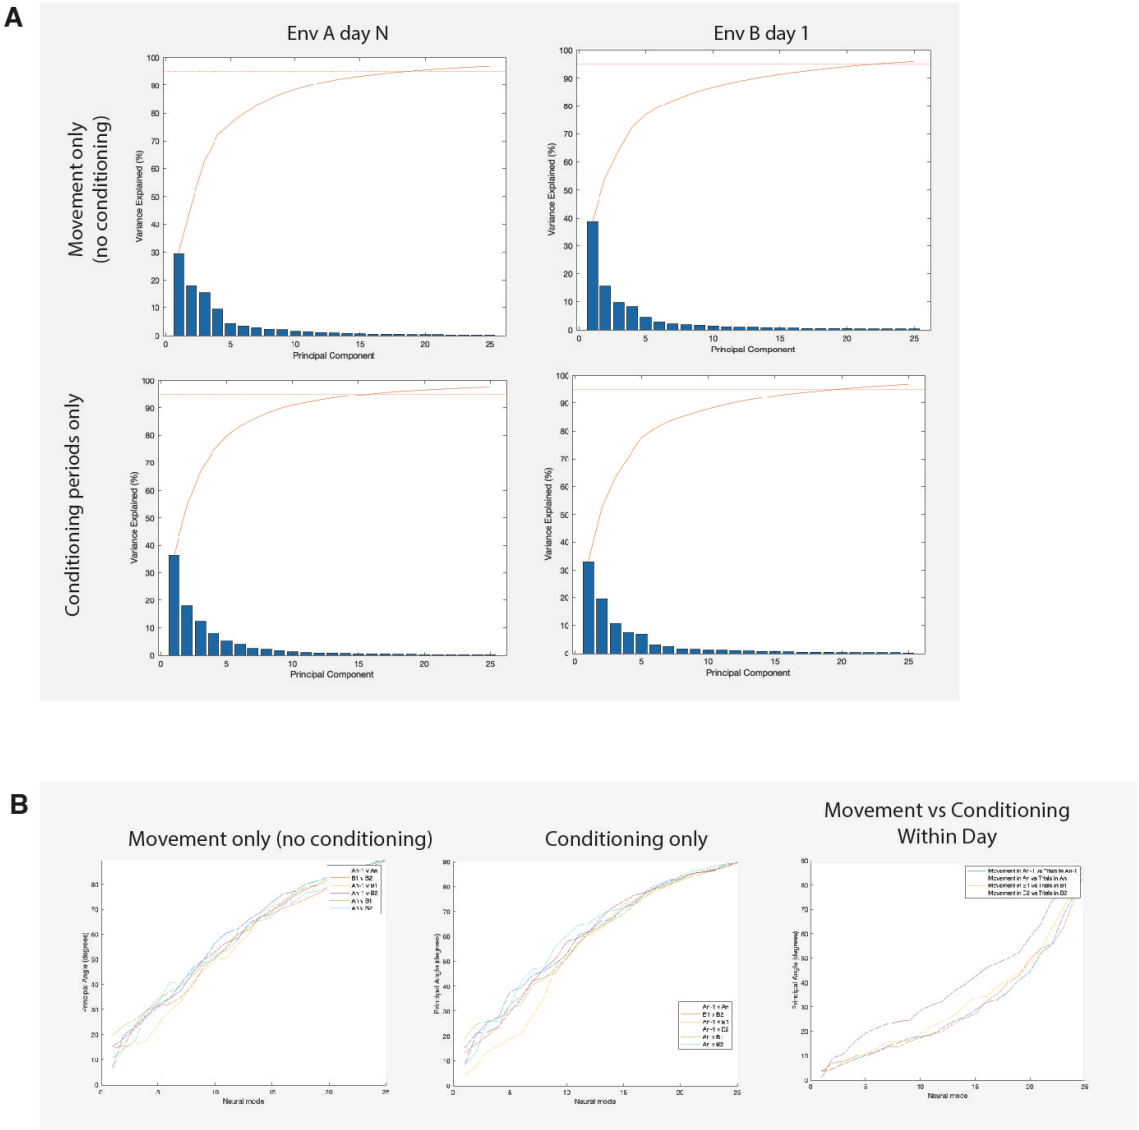

Figure S5. PCA computations for session A(n) and session B(1).

Only cells present in both sessions were used. Principal component analysis (PCA) revealed that approximately 15-25 principal components (PCs) are needed to account for 95% of the variance in the data. When using the complete cell population (not shown), more than 25 PCs are required to achieve the same variance. Across and within all sessions and representations (spatial and task), the principal angles between manifolds remain highly similar.

Figure S6

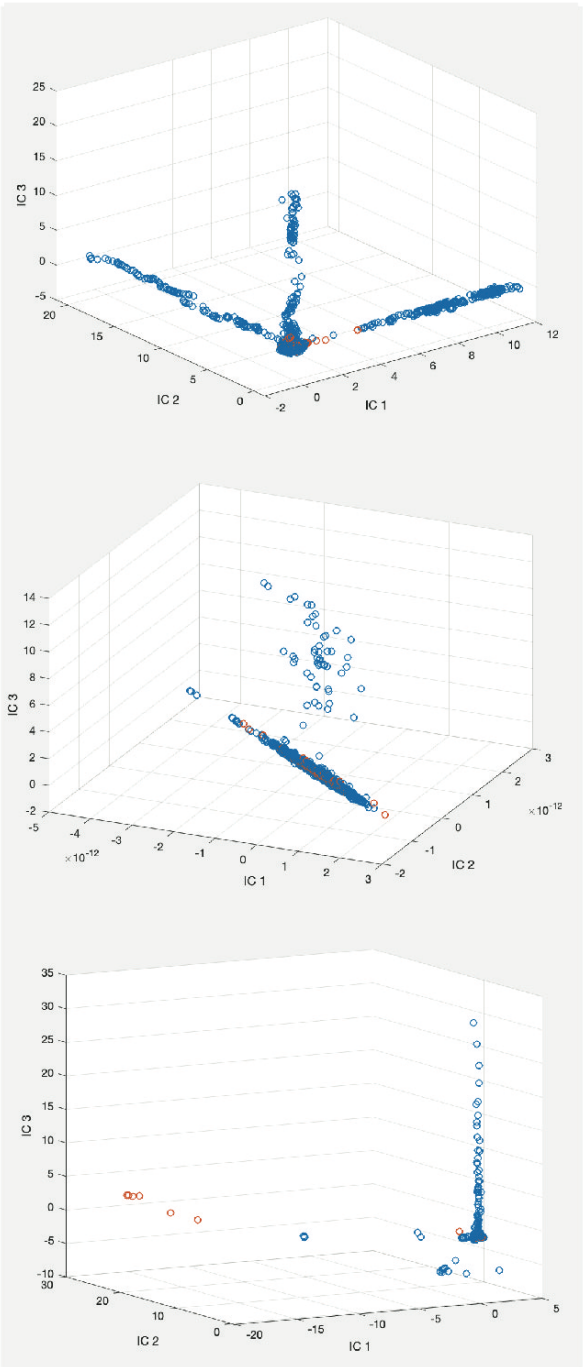

Figure S6. ICA computations across different segments of a session.

Top: Independent component analysis (ICA) was for data over an entire session, using three independent components (ICs). Blue dots represent non-trial times, while red dots represent trial times. Middle: ICA computed over the last two-thirds of the same session shows variability in

ICs across segments within a session. Bottom: ICA computed over the second half of the session shows additional variability in the results of the analysis depending on how the session is divided. These results indicate that independent components are highly variable over the course of one session and are sensitive to how the session is partitioned.

Figure S7

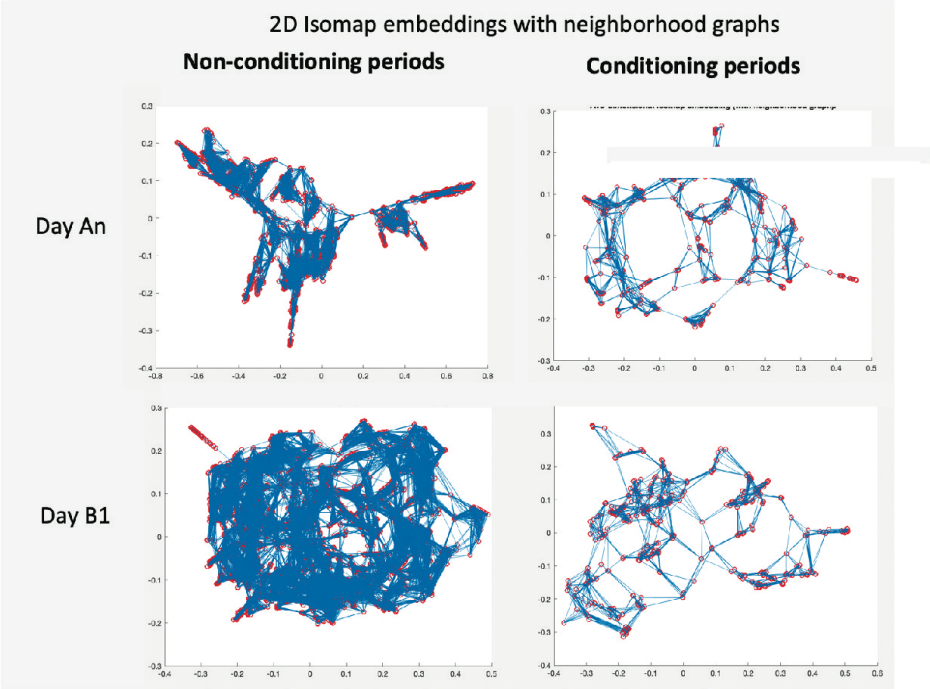

Figure S7. Isomap computations for session A(n) and session B(1).

Isomap computations suggest that approximately five neural modes are sufficient to achieve a residual variance of 5-10%. However, the shape of the Isomap embedding does not correlate with any discernable properties of neural activity or behavior, suggesting limited interpretability of the embedding structure.

Figure S8

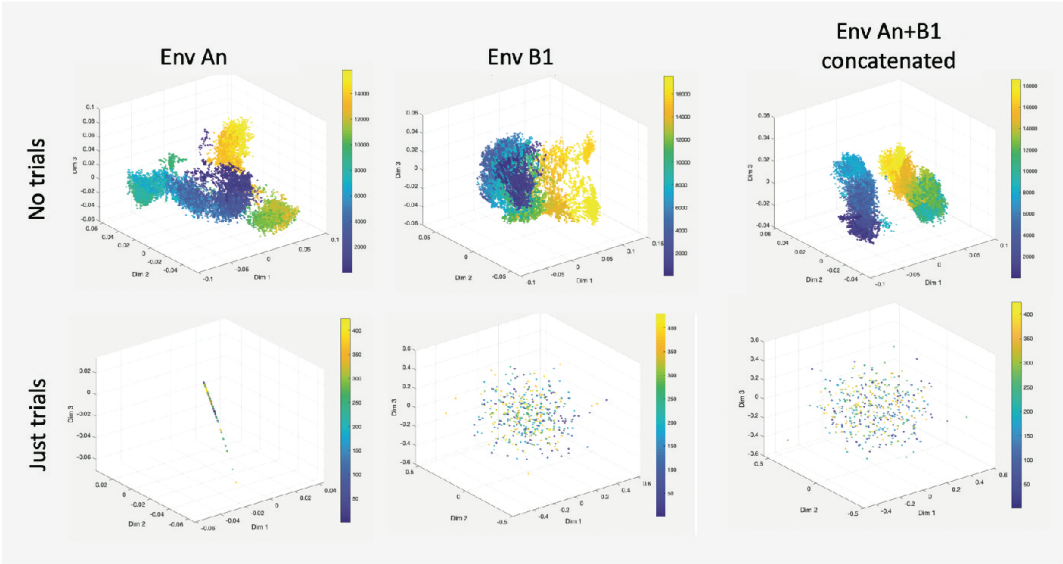

Figure S8. MIND outputs for sessions A(n), B(1), and concatenated sessions. Top row: MIND embeddings during movement, excluding trial periods, with color bars representing frames. The temporal structure of the data is well captured, with clear separation between A(n) and B(1). Bottom row: MIND embeddings during conditioning periods are highly unstable. Small changes in parameters result in substantial shifts in the embedding structure, transitioning from a linear structure (left) to an undefined, unstable cloud (middle and right).

Figure S9

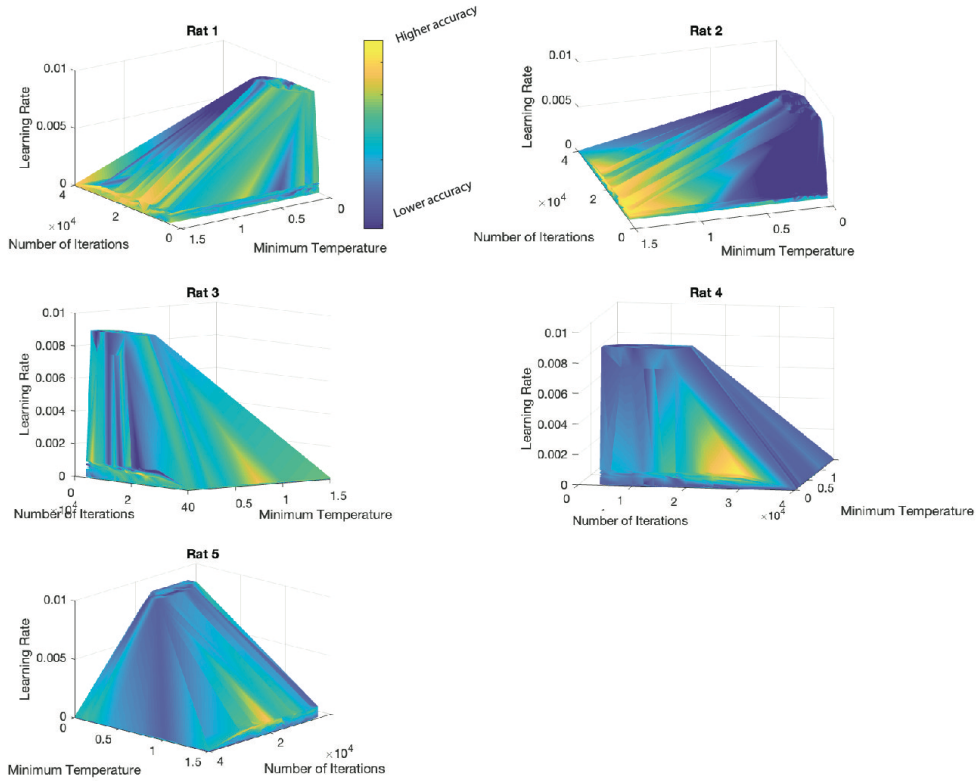

Figure S9. Grid search over decoding parameters for position.

A grid search was performed over three parameters: minimum temperature, learning rate, and number of iterations, for decoding position. Models were trained using cells from session A(n) that also appeared in session A(n-1). The figure shows decoding accuracy for session A(n-1) using the models trained on data from session A(n). Yellow areas indicate higher decoding accuracy.

Figure S10

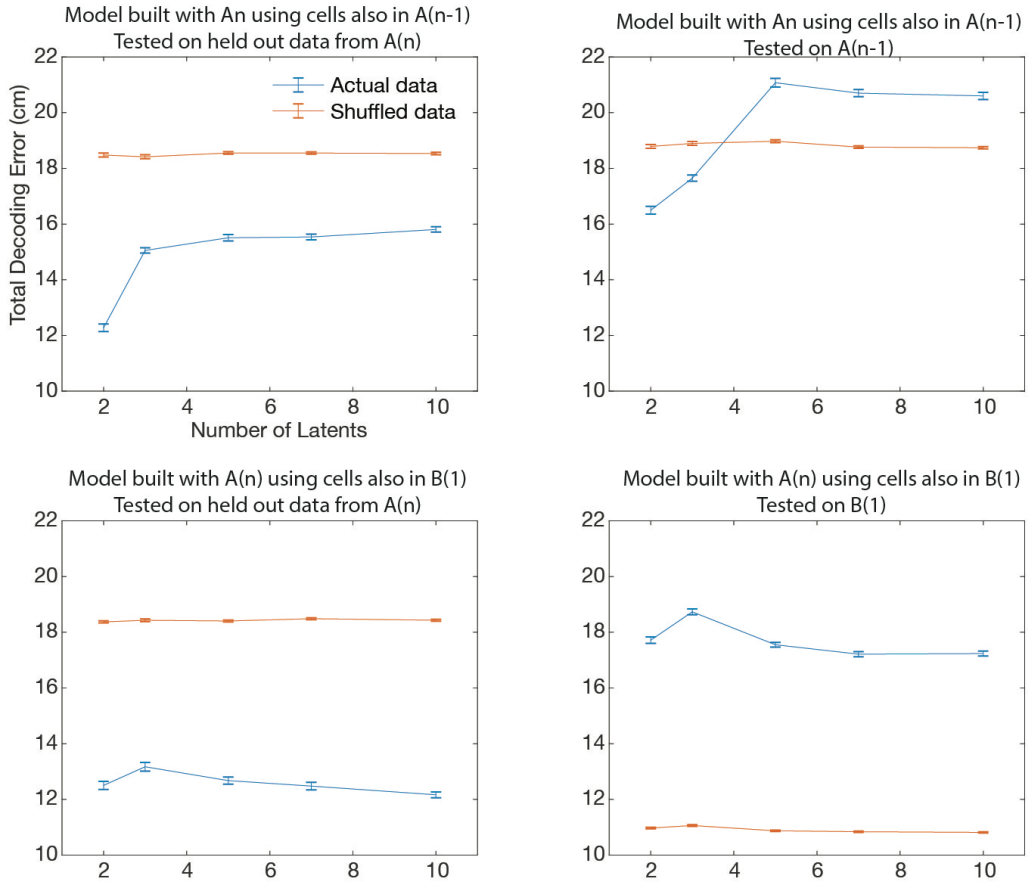

Figure S10. Position decoding error as a function of the number of latents (Rat 5). This figure shows the decoding error for position as the number of latents increases. Left panels: model trained on data from A(n), tested on held out data from A(n). Upper left panel: As the number of latents increases, the model's ability to decode a different session within the same environment decreases. This effect is not consistent across all rats; see lower left panel. Right panels: model trained on data from A(n), tested on A(n-1) (top) or B(1) (bottom). Performance is particularly bad when the model is tested on a different environment (lower right panel).

Figure S11

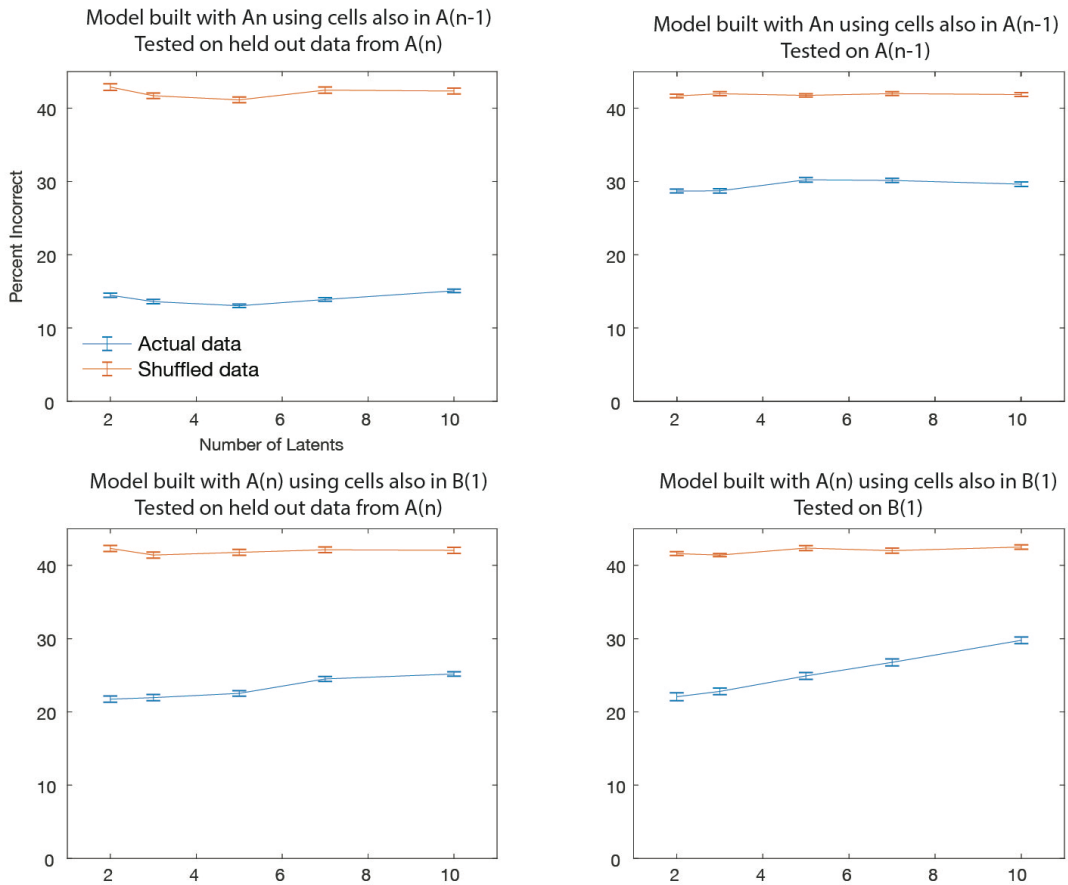

Figure S11. CSUS2 decoding accuracy with increasing number of latents (Rat 3). Percent of incorrect decoding for the CSUS2 model as the number of latents increases. Left panels: model trained on data from A(n), tested on held out data from A(n). As the number of latents increases, the model's ability to decode a different session within the same environment remains stable or slightly decreases. Right panels: model trained on data from A(n), tested on A(n-1) (top) or B(1) (bottom). Performance remains stable for a different session in the same environment (upper right panel) but deteriorates with increasing number of latents when the model is tested in a different environment (lower right panel). Each model was run 100 times.

Figure S12

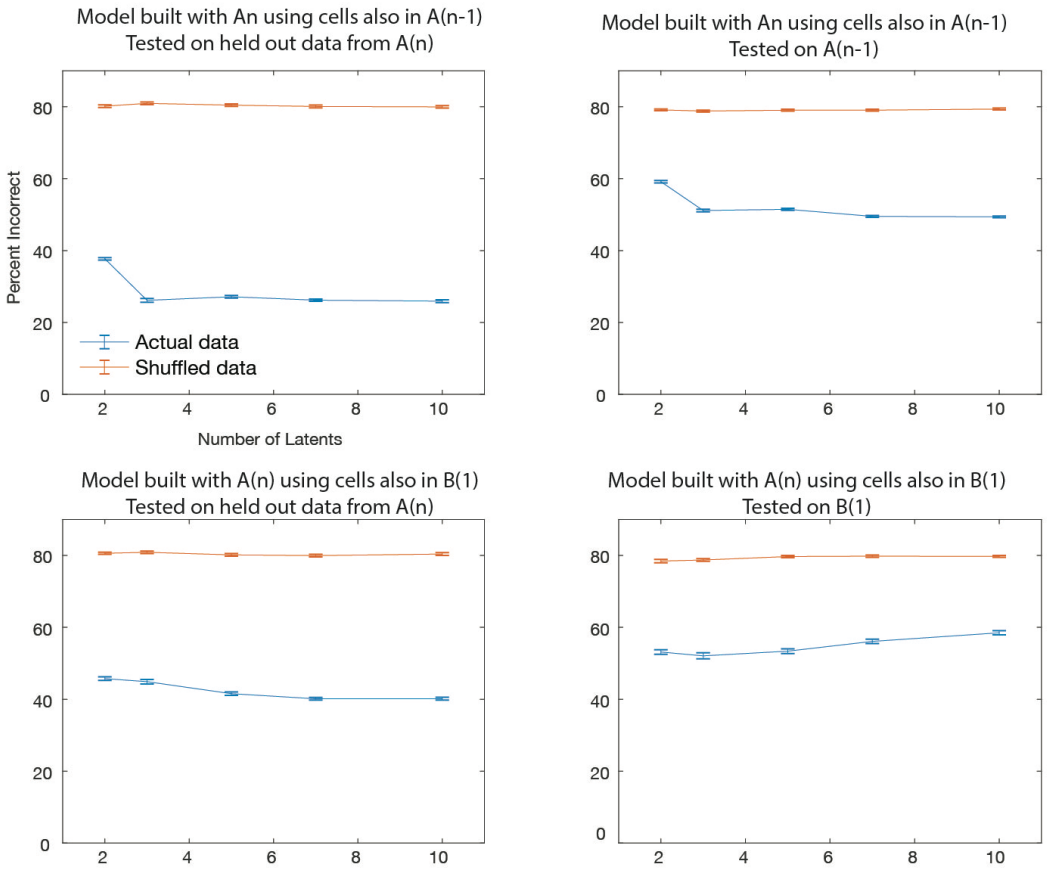

Figure S12. CSUS5 decoding accuracy with increasing number of latents (Rat 5). Same as Figure S11, but for CSUS5; the conditioning period has been divided into five time bins instead of two. The percent of incorrect decoding is shown as the number of latents increases. Left panels: model trained on data from A(n), tested on held out data from A(n). As the number of latents increases, the model's ability to decode a different session within the same environment remains stable or slightly increases. Right panels: model trained on data from A(n), tested on A(n-1) (top) or B(1) (bottom). Performance remains stable for a different session in the same environment (upper right panel) but deteriorates with increasing number of latents when the model is tested in a different environment (lower right panel). Each model was run 100 times.

Figure S13

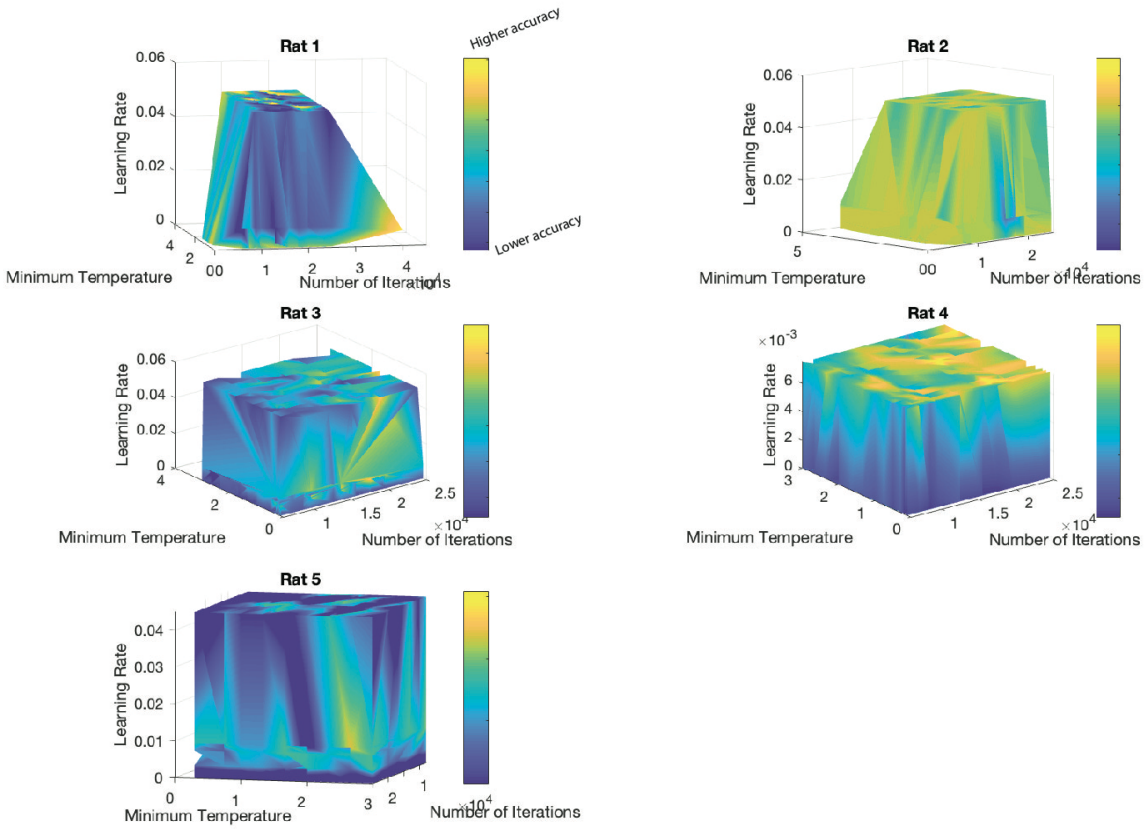

Figure S13. Grid search over decoding parameters for conditioning.

A grid search was performed over three parameters: minimum temperature, learning rate, and number of iterations for conditioning decoding. Models were trained using cells from session A(n) that also appeared in session B(1). The figure shows decoding accuracy for CSUS2 session B(1) using the models trained on data from session A(n). Yellow areas indicate higher decoding accuracy. For Rats 1 and 4, the 'Euclidean' distance with 'constant' temperature mode was used. For Rats 2 and 5, 'cosine' distance with 'constant' temperature mode was used. For Rat 3, 'cosine' distance with 'auto' temperature mode was used.

Methods References

1 Wirtshafter, H. S. & Disterhoft, J. F. Place cells are nonrandomly clustered by field location in CA1 hippocampus. *Hippocampus* **33**, 65-84 (2023). https://doi.org:10.1002/hipo.23489

2 Wirtshafter, H. S. & Disterhoft, J. F. In Vivo Multi-Day Calcium Imaging of CA1 Hippocampus in Freely Moving Rats Reveals a High Preponderance of Place Cells with Consistent Place Fields. *J Neurosci* (2022). https://doi.org:doi: 10.1523/JNEUROSCI.1750-21.2022

3 Weiss, C. & Thompson, R. F. The effects of age on eyeblink conditioning in the freely moving Fischer-344 rat. *Neurobiol Aging* **12**, 249-254 (1991). https://doi.org:10.1016/0197-4580(91)90105-s

4 Skelton, R. W. Bilateral cerebellar lesions disrupt conditioned eyelid responses in unrestrained rats. *Behav Neurosci* **102**, 586-590 (1988). https://doi.org:10.1037//0735-7044.102.4.586

5 Aharoni, D., Khakh, B. S., Silva, A. J. & Golshani, P. All the light that we can see: a new era in miniaturized microscopy. *Nat Methods* **16**, 11-13 (2019). https://doi.org:10.1038/s41592-018-0266-x

6 Silva, A. J. Miniaturized two-photon microscope: seeing clearer and deeper into the brain. *Light: Science & Applications* **6**, e17104-e17104 (2017).

7 Mathis, A. *et al.* DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature neuroscience* **21**, 1281-1289 (2018).

8 Corder, G. *et al.* An amygdalar neural ensemble that encodes the unpleasantness of pain. *Science* **363**, 276-281 (2019). https://doi.org:10.1126/science.aap8586

9 Thevenaz, P., Ruttimann, U. E. & Unser, M. A pyramid approach to subpixel registration based on intensity. *IEEE Trans Image Process* **7**, 27-41 (1998). https://doi.org:10.1109/83.650848

10 Zhou, P. *et al.* Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data. *eLife* **7** (2018). https://doi.org:10.7554/eLife.28728

11 CIAtah: a software package for analyzing one- and two-photon calcium imaging datasets. v. v1.0.0 (Zenodo, 2018).

12 Kinsky, N. R., Sullivan, D. W., Mau, W., Hasselmo, M. E. & Eichenbaum, H. B. Hippocampal Place Fields Maintain a Coherent and Flexible Map across Long Timescales. *Current biology : CB* **28**, 3578-3588 e3576 (2018). https://doi.org:10.1016/j.cub.2018.09.037

13 Olypher, A. V., Lansky, P., Muller, R. U. & Fenton, A. A. Quantifying location-specific information in the discharge of rat hippocampal place cells. *J Neurosci Methods* **127**, 123-135 (2003). https://doi.org:10.1016/s0165-0270(03)00123-7

14 Pearson, K. LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science* **2**, 559-572 (1901).

15 Nieh, E. H. *et al.* Geometry of abstract learned knowledge in the hippocampus. *Nature* **595**, 80-84 (2021).

16 Ans, B., Hérault, J. & Jutten, C. Architectures neuromimétiques adaptatives: Détection de primitives. *Proceedings of Cognitiva* **85**, 593-597 (1985).

17 Tenenbaum, J. B., Silva, V. d. & Langford, J. C. A global geometric framework for nonlinear dimensionality reduction. *science* **290**, 2319-2323 (2000).

18 Low, R. J., Lewallen, S., Aronov, D., Nevers, R. & Tank, D. W. Probing variability in a cognitive map using manifold inference from neural dynamics. *BioRxiv*, 418939 (2018).

19 Schneider, S., Lee, J. H. & Mathis, M. W. Learnable latent embeddings for joint behavioural and neural analysis. *Nature* **617**, 360-368 (2023). https://doi.org:10.1038/s41586-023-06031-6

20 Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *the Journal of machine Learning research* **12**, 2825-2830 (2011).