Recurrent Connectivity Shapes Spatial Coding in Hippocampal CA3 Subregions

Eunji Kong[1,2*], Erfan Zabeh[2,3*], Zhenrui Liao[1,2*], Tiberiu S. Mihaila[1,2], Caroline Wilson[2], Charan Santhirasegaran[2], Darcy S. Peterka[2], Attila Losonczy[1,2], and Tristan Geiller[4,5]

1. Department of Neuroscience, Columbia University, New York, NY, United States
2. Mortimer B. Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY, United States
3. Department of Biomedical Engineering, Columbia University, New York, NY, USA.
4. Department of Neuroscience, Yale University, New Haven, CT, USA.
5. Wu Tsai Institute, Yale University, New Haven, CT, USA.

Correspondence: ek3324@columbia.edu, al2856@columbia.edu, tristan.geiller@yale.edu

*: These authors contributed equally

Abstract

Stable and flexible neural representations of space in the hippocampus are crucial for navigating complex environments. However, how these distinct representations emerge from the underlying local circuit architecture remains unknown. Using two-photon imaging of CA3 subareas during active behavior, we reveal opposing coding strategies within specific CA3 subregions, with proximal neurons demonstrating stable and generalized representations and distal neurons showing dynamic and context-specific activity. We show in artificial neural network models that varying the recurrence level causes these differences in coding properties to emerge. We confirmed the contribution of recurrent connectivity to functional heterogeneity by characterizing the representational geometry of neural recordings and comparing it with theoretical predictions of neural manifold dimensionality. Our results indicate that local circuit organization, particularly recurrent connectivity among excitatory neurons, plays a key role in shaping complementary spatial representations within the hippocampus.

Introduction

Recurrent connectivity among excitatory neurons is a critical feature of the cortical microcircuit that underlies functions such as sensory processing, working memory, and decision-making[1-6]. In the hippocampus, strong recurrent excitatory connections among pyramidal neurons within the CA3 region are thought to provide a basis for memory-guided behaviors[7-10]. This circuitry is believed to endow CA3 with the ability to rapidly store patterns of integrated sensory information for the initial encoding of contextual representations, and to retrieve them later from partial or degraded inputs[11,12]. However, the density of recurrent collaterals changes on a gradient along the transverse axis, from low recurrence in proximal CA3 (pCA3; close to the dentate gyrus) to higher recurrence in distal CA3 (dCA3; close to CA2)[13,14] , mirroring distinct molecular and physiological gradients recognized along this axis[15,16]. Given these differences, it is natural to ask whether CA3 subregions hold different roles in encoding and retrieving spatial information, and how these roles translate more broadly to learning and memory functions[17-25].

To address this question, we performed a detailed functional characterization of pCA3 and dCA3 neurons in mice performing a spatial navigation task in virtual reality (VR)[26]. We characterized the dynamics of neural representations over time and across contexts, revealing prominent differences between the two subregions. Using recurrent neural network (RNN) and Hopfield network models, we confirmed that varying the level of recurrent connectivity between excitatory neurons is sufficient to explain these differences, demonstrating that distinct connectivity patterns in pCA3 and dCA3 lead to differential mechanisms of information processing.

We conclude that distinct subregions of CA3, shaped by their respective levels of recurrent connectivity, play complementary roles in spatial memory: dCA3 may be more involved in the flexible encoding of new contexts, while pCA3 may contribute to the stable

retrieval of established memories, thus balancing adaptability and consistency in memory-guided behaviors[27-29].

Results

Long-term dynamics of spatial representations

To first characterize the functional dynamics of different recurrent connections within CA3 in the dorsal hippocampus, we performed *in vivo* two-photon calcium imaging in a VR spatial navigation task[30,31] (Fig. 1a-c, Supplementary Fig. 1a-b). We stereotactically injected Cre-dependent adeno-associated virus (AAV) into pCA3 and dCA3 regions of Grik4-Cre mice to selectively express GCaMP8s in CA3 pyramidal neurons, and implanted a chronic imaging window above the hippocampus. Mice were head-fixed and trained to perform a goal-oriented learning (GOL) task[26,31], where a spatially fixed operant reward was given in a 2-m virtual linear environment with a 2-s inter-trial interval between laps. After two weeks of training, mice successfully learned the reward zone (RZ), as reflected by anticipatory changes in animal's velocity and licking (Fig. 1d). On each day, we observed a subset of place cells (0.13 ± 0.01 for pCA3, 0.31 ± 0.02 for dCA3), whose responses reliably tiled the entire virtual track (Supplementary Fig. 1c, g). We quantified individual place cells' specificity, sensitivity, and spatial information and found higher proportions of place cells with lower levels of spatial information and place field specificity in dCA3 compared to pCA3 (Supplementary Fig. 1c-f).

To investigate the long-term spatial coding properties of CA3 PNs, we tracked the same field of view every other day (Fig. 1e) as mice achieved stable behavioral performance in the GOL task. Among cross-registered cells, a subset of neurons was identified as place cells over days (Fig. 1f, Supplementary Fig. 2a, b), and we assessed whether neurons stably encoded their location over the course of the week. We found that dCA3 place cells had a higher recurrence probability (Fig. 1g), defined as the probability that a given neuron displays a significant place field across two paired sessions, irrespective of its location on

the track. This result shows that dCA3 neurons overall remain more spatially selective with time compared to pCA3 neurons. To address whether place fields maintain their selectivity, we quantified the tuning curve correlations between sessions, and found that pCA3 place cells, though less numerous, are more stable over time than dCA3 place cells (Fig. 1h). At the population level, pCA3 population vectors remained highly correlated with each other across days compared to dCA3 population vectors (Fig. 1i, j, Supplementary Fig. 2c, d). In summary, these results imply that the dCA3 code is more dynamic and adaptable over time, while pCA3 implements a more stable, long-term code of spatial memory.

Generalized vs context-discriminative neural representations

Next, to investigate the neural dynamics of proximal and distal CA3 PNs in response to contextual changes, we performed a context-switching task in which animals were exposed to a (pretrained) Familiar (F) and a Novel (N) environment in 15-lap alternating blocks (Fig. 2a, b, Supplementary Fig. 3a, b) while recording from, either pCA3 or dCA3 (Fig. 2c). To quantify how contextual information is differentially processed at the single-cell level, we calculated a discrimination index (DI) of neural responses between familiar and novel contexts (Fig. 2d). We found that a significant proportion of neurons discriminated between contexts in both subregions, though a greater proportion did in dCA3 compared to pCA3 (Fig. 2e, f). Using a linear classifier with matched sample sizes, we also found that the information encoded in population activity in dCA3 discriminated between two different contexts with higher accuracy compared to that in pCA3 (Fig. 2g, h). We also evaluated the decoding performance of randomly chosen subsets of recorded neurons from pCA3 (Supplementary Fig. 3e). This analysis showed a correlation between the number of subsampled neurons and peak decoding accuracy, with the performance of the classifier being similar between all recorded pCA3 neurons and dCA3 neurons (Supplementary Fig. 3f).

In addition to contextual representations, we analyzed place cell dynamics within and between environments. While both pCA3 and dCA3 neurons exhibited spatially tuned responses across context blocks (Fig. 2i, j, Supplementary Fig. 4a), pCA3 place cells showed a higher level of correlation between familiar and novel environments (Fig. 2k). In dCA3, F-N similarity increased during the second switch compared to the first switch, whereas no significant change was observed in pCA3 (Supplementary Fig. 4b, c).

Taken together, these findings suggest subregion-specific complementary roles for encoding dynamic environments, with pCA3 producing a general template for the context and dCA3 providing more specific responses (Supplementary Fig. 4d). Together, these populations might work in concert to balance flexibility and stability for encoding contextual information.

Recurrence level interpolates opposing coding strategies

Our functional imaging approach revealed dynamic and heterogeneous neural representations along the proximodistal axis of CA3 in relation to both time and context. How two subnetworks of the same hippocampal area implement opposing coding strategies remains an unanswered question. We hypothesized that these differences were driven by recurrence level as the most salient difference between these subregions. To test this hypothesis, we employed multi-level computational modeling.

First, we constructed two distinct recurrent neural network (RNN) models of CA3: a sparsely connected RNN and a densely connected RNN (Fig. 3a). The differences in recurrent connectivity of the RNNs were chosen based on earlier studies that reported recurrence level heterogeneity across the proximodistal axis of CA3[13,14] (Supplementary Fig. 5a).

Informed by biological evidence showing projections from speed-selective cells in the entorhinal cortex to the hippocampus[32], we designed our RNN models to receive inputs

of speed and the initial location of the animal in a virtual arena (Fig. 3b). The number of training iterations was chosen to ensure both the sparse and dense RNN reached a plateau in their loss functions, ensuring a fair comparison of their performance (Fig. 3c, Supplementary Fig. 5b). Post-training analysis showed that the denser RNN model exhibited a higher proportion of place cells compared to the sparser model, consistent with our experimental observations (Fig. 1f, Fig. 3d, e).

To further explore the adaptability of place cells in the RNNs to new environments, we simultaneously provided a visual context signal and the instantaneous speed of the animal as the inputs to both network models and trained the networks to identify both location and context (Fig. 3f). Given that animals utilize visual cues to identify the context of the VR track[31], we added noise to the context signal, representing perceptual context evidence from the visual system, to the hippocampal network. We maintained the noise level high enough to avoid a trivial scenario for the network and accordingly reinforced the engagement of a higher proportion of cells in context discrimination. As a result, the RNN model successfully learned the context-switching task with low error rates (Supplementary Fig. 5d); analysis of the hidden layer representations revealed the presence of multiple types of conjunctive and selective cells for context and place (Fig 3.g-h). A greater number of neurons in the dense RNN discriminated between familiar and novel contexts (Fig. 3i), consistent with *in vivo* recording data that showed higher context selectivity in dCA3 (Fig. 2f, h).

Finally, we investigated the effect of varying connectivity on memory generalization and discrimination using a Hopfield network model of CA3 subregions (Fig. 3j). Hopfield networks are widely used as models of content-addressable memory, but the classical Hopfield model requires full connectivity, which is not biologically plausible. However, we find that sparsely connected Hopfield networks, trained with the classic Hebbian rule (Fig. 3k), retain the ability to store and retrieve memories. We also find differences between memory performance in sparsely and densely connected Hopfield networks that mirror differences between pCA3 and dCA3. The dense network exhibited higher memory

capacity, i.e., the ability to accurately discriminate a greater number of patterns compared to the sparse network (Fig. 3l). However, the sparse network was better at recovering patterns from significant distortions (Fig. 3m), suggesting fewer, deeper wells in its energy landscape that could aid generalization. Finally, we find that these abilities interpolate: intermediate levels of connectivity give rise to intermediate levels of discrimination and generalization (Fig. 3l, m; black line).

Dimensionality of neural manifolds altered by recurrency

Our modeling results suggest that the level of recurrence alone can account for the functional heterogeneities between pCA3 and dCA3 observed *in vivo*. Theoretical work has established a direct link between recurrency level and representational geometry of neurons, suggesting higher connectivity between neurons increases the mutual activity of cell populations[33], thereby constraining their representational geometry in population space (Fig. 4a).

We first examined this theoretical intuition in simulations of RNNs during the context-switching task. Our simulations demonstrated that altering connectivity towards sparser networks increases the dimensionality of neural manifolds, as defined by the explained variance of cumulative principal components (Fig. 4b-d). We conducted a similar analysis on the neural activities recorded during the same task to examine the neural representation geometry of proximodistal CA3 neurons (Fig. 4e). As predicted, the manifold dimensionality of representations in dCA3 was lower than that in pCA3 (Fig. 4f, g), further supporting the role of recurrent connectivity in functional heterogeneity within the CA3 circuit.

These findings underscore the critical role of recurrent connectivity in shaping the distinct functional properties of CA3 subregions. The gradient of generalization observed as a function of network sparsity highlights the importance of recurrent connectivity in the encoding and recalling of spatial and contextual information. Our computational modeling,

aligned with experimental data, provides deeper insight into the mechanisms by which CA3 PNs contribute to spatial navigation and context discrimination.

Discussion

Despite extensive evidence of anatomical, molecular, and physiological heterogeneity in the CA3 along the transverse axis[13-16,34,35], the relationship between these differences and neural computations performed by CA3 subregions has remained largely unknown. While both stable and flexible dynamics as well as varying degrees of context generalization have been reported in CA3 using *in vivo* imaging[36-39], these representational features have not been systematically investigated along the proximodistal axis, nor have they been linked to CA3 subregions. In this study, we bridged these gaps by combining *in vivo* two-photon calcium imaging of CA3 subregions with computational modeling at multiple levels of abstraction. Our findings provide significant insights into how distinct hippocampal subcircuits differentially contribute to memory functions[36,40,41].

We found that dCA3 neurons are highly sensitive to contextual changes as well as tuned to different locations over days. This suggests that they are flexible and adaptive, possibly playing a role in learning and responding to new experiences or changes in the environment. In contrast, pCA3 neurons show more stability in their tuning over time, even though fewer neurons are tuned compared to dCA3. This stability suggests that they might be involved in encoding persistent or core aspects of the context that remain relevant across different experiences. These neurons may be essential for maintaining a stable representation of the environment, supporting consistent recall of key contextual features or stable memory traces, and could be rooted in how they integrate inputs at the subcellular level[42-44].

Consistent with recent reports on the rapid formation of place fields in the hippocampal CA1 and CA3[30,37,38,45], we also found that both pCA3 and dCA3 PNs exhibit spatially

tuned neural responses that rapidly formed when the mice are exposed to novel environments. Future studies will explore the synaptic plasticity rules governing the dynamics of place field formation in CA3 along the transverse axis[46]. We further observe that dCA3 PNs are more significantly tuned to novel environments compared to pCA3 PNs. It is well known that hippocampal circuits are under neuromodulatory control related to the behavioral state of animals[47-51]. Correspondingly, the CA3 subregions may receive different levels of catecholaminergic input, such as dopamine and norepinephrine from the locus coeruleus, which facilitates memory formation in a new context[52-54].

It is striking that two subregions of the same hippocampal area, separated by less than a millimeter in the mouse brain and (notwithstanding the important differences we have highlighted) broadly similar in their anatomy, physiology, and development, can implement seemingly opposing computational functions. Our RNN modeling results, in which spatial and contextual representations by artificial cells corresponded closely to the empirical data in both single-cell and population level, underscore the critical role of local connectivity patterns in shaping the functional properties of neural networks. Furthermore, simulation results with dense, intermediate, and sparse Hopfield networks imply a smooth gradient between discrimination and generalization based on the level of recurrence in a subnetwork. Taken together, these findings emphasize the specific importance of recurrent connectivity in determining the computational features of neural circuits, supporting the idea that small variations in structural connectivity can lead to substantial differences in neuronal function and behavior. We further speculate that a gradient of connectivity implementing these intermediate levels exists between pCA3 and dCA3.

Our findings show that recurrent connectivity shapes the dimensionality of neural manifolds in CA3, with sparser networks yielding higher-dimensional manifolds and denser networks producing lower-dimensional ones. While an increasing body of literature associates dimensionality in neural spaces with flexibility and generalization[55-58], dimensionality alone does not fully account for a network's computational capacity. The geometry of neural representations—encompassing factors like curvature,

separation, and topology—plays a critical role in how information is processed[59-63]. To fully understand flexibility and generalization, it is necessary to study the precise geometry of these representations. Previous studies have shown divergent roles in task performance are observed in networks with varying neural trajectory geometries[64,65].

Recent studies have explored how structural connectivity within networks affects their representational space, often by examining constraints such as connectivity rank[66,67]. However, whether biological networks are inherently low-rank remains under debate. In our approach, we focus on binary (present-or-absent) connectivity, providing a clearer structural interpretation of how network architecture shapes representational dimensionality, complementing earlier studies that link structural biases to neuronal variability[68,69]. Theory predicts that low-dimensional activity remains possible in sparse networks through fine-tuning of synaptic weights[70]. Accordingly, the direct measurement of synaptic weights *in vivo* will be an important direction for future experimental work to dissect the synaptic basis for computational capacity.

In addition to recurrent local connectivity, subregion-specific differences in afferent inputs to CA3 can contribute to the observed differences in context-related dynamics between pCA3 and dCA3. Recent reports have highlighted a generalized neural representation by dentate gyrus granule cells, one of the main input sources to CA3[36], which may contribute to the more generalized contextual and spatial dynamics we observe in pCA3 PNs compared to dCA3 PNs. Similarly, subregion-specific differences in afferent inputs from the entorhinal cortex, which is also implicated in context generalization and discrimination in the CA3[71,72], may also shape context processing in CA3. Finally, the role of inhibition and subtype-specific inhibitory control of feature representations and their stability remains unknown[73-75]. Future studies should investigate how external inputs modulate the cognitive flexibility of different subregions within CA3. Understanding local circuit[76] and long-range interactions will provide deeper insights into how the hippocampus supports complex behaviors, integrating contextual and spatial information to facilitate adaptive navigation and memory.

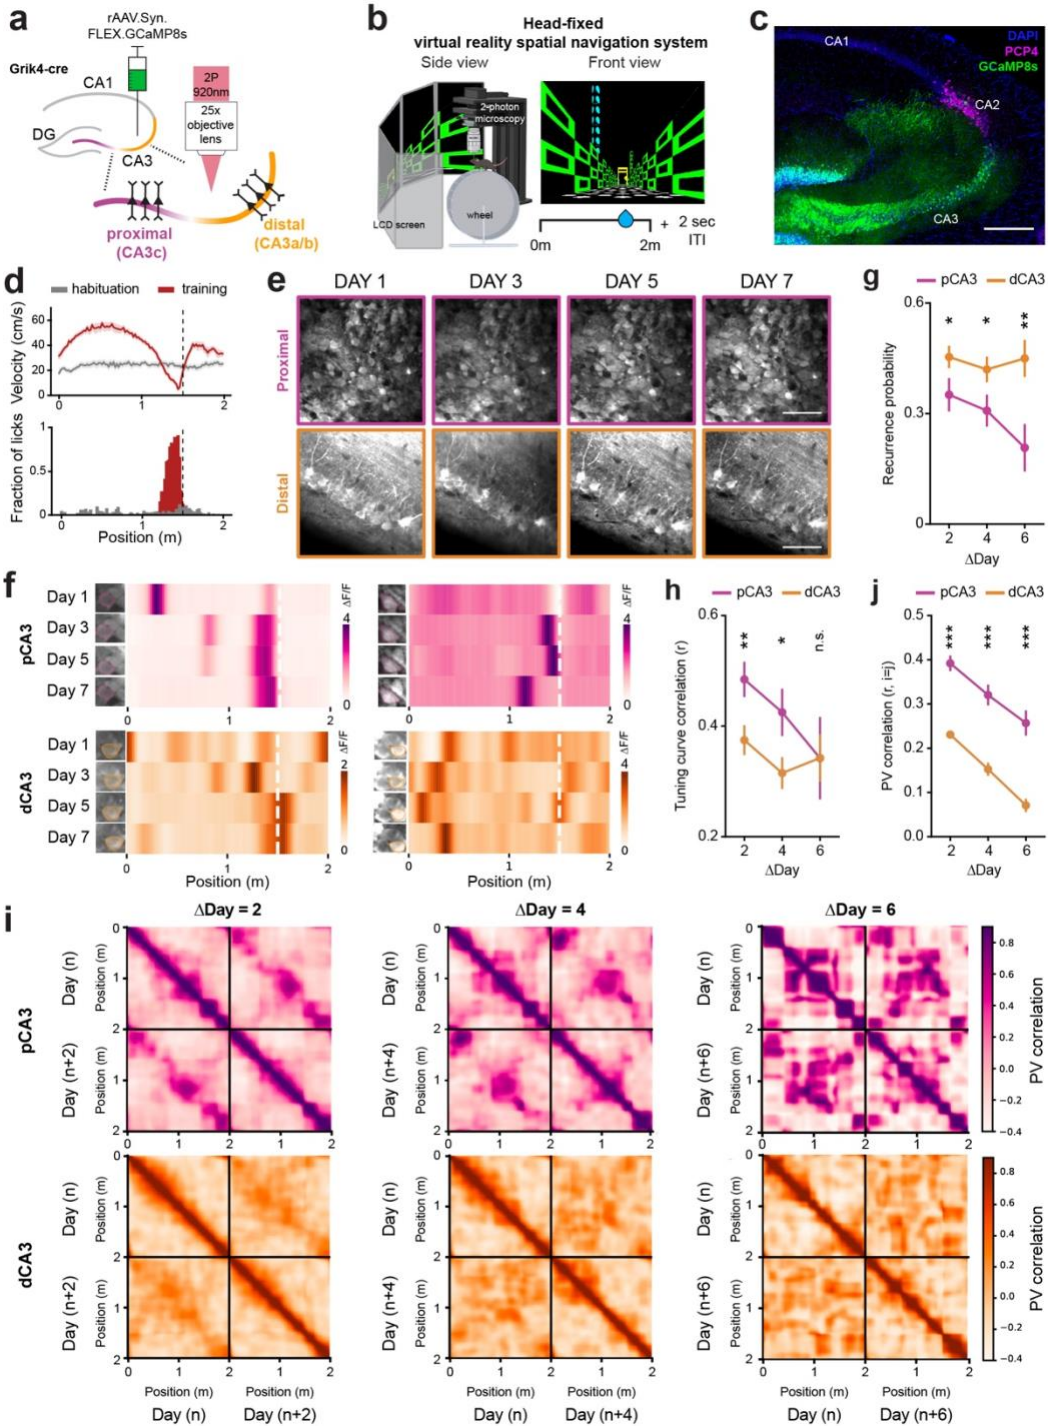
a
rAAV.Syn.
FLEX.GCaMP8s
Grik4-cre
CA1
DG
CA3
2P
920nm
25x
objective
lens
proximal
(CA3c)
distal
(CA3a/b)
b
Head-fixed
virtual reality spatial navigation system
Side view
Front view
2-photon
microscopy
LCD screen
wheel
0m
2m
+ 2 sec
ITI
c
CA1
CA2
CA3
DAPI
PCP4
GCaMP8s
d
habituation
training
Velocity (cm/s)
Fraction of licks
Position (m)
e
DAY 1
DAY 3
DAY 5
DAY 7
Proximal
Distal
g
pCA3
dCA3
Recurrence probability
ΔDay
f
Day 1
Day 3
Day 5
Day 7
pCA3
dCA3
ΔF/F
Position (m)
h
pCA3
dCA3
Tuning curve correlation (r)
n.s.
ΔDay
j
pCA3
dCA3
PV correlation (r, i≠j)
ΔDay
i
ΔDay = 2
ΔDay = 4
ΔDay = 6
pCA3
dCA3
Day (n)
Day (n+2)
Day (n+4)
Day (n+6)
Position (m)
PV correlation

Figure 1. Dynamic and stable spatial representations by distal and proximal CA3. a. Schematic of *in vivo* two-photon imaging of proximodistal CA3 neurons. Grik4-cre transgenic mice were injected with a Cre-dependent GCaMP8s virus to record the neural activities of PNs in dorsal CA3. **b**. Schematic of a virtual reality navigation system combined with *in vivo* two-photon microscopy (*left*) and the task structure of a goal-oriented learning (GOL) task with a fixed reward zone in a 2m virtual environment (*right*). ITI, intertrial interval. **c**. Representative confocal image showing GCaMP8s expression along the transverse axis of the CA3 region in Grik4-cre transgenic mice. Blue: DAPI, Green: GCaMP8s, Magenta: PCP4. Scale bar: 200um. **d**. Examples of velocity (*top*) and licking responses (*bottom*) before and after operant conditioning of the GOL task. **e**. Longitudinal and repetitive imaging of individual cells in proximal (*top*) and distal (*bottom*) CA3 subregions. Scale bar: 50 μm**. f**. Examples of cross-day registered place cells and their heatmaps of spatial tuning curves with stable place fields in pCA3 (*top*) and dCA3 (*bottom*). **g**. Probability of recurrence of place cells across days (mean ± SEM; two-sided unpaired t-test, p = 0.037, 0.032 and 0.0042 for ΔDay 2, 4, and 6; n = 5 and 7 mice for pCA3 and dCA3, respectively). **h**. Average tuning curve correlations of place cells across days (mean ± SEM; two-sided unpaired t-test, p = 0.0092, 0.029, and 0.10 for ΔDay 2, 4, and 6; n = 5 and 7 mice for pCA3 and dCA3, respectively). **i**. Heatmaps of population vector (PV) correlation between all pairs of positions across days for pCA3 (*top*) and dCA3 (*bottom*) place cells. The interval between days was 2 (*left*), 4 (*middle*), and 6 (*right*). **k**. PV correlation between the same position across days for pCA3 and dCA3 place cells (mean ± SEM, two-sided unpaired t-test, p = 1.06 e-14, 7.17 e-09, and 2.80 e-08 for ΔDay 2, 4, and 6, n = 5 and 7 mice for pCA3 and dCA3, respectively). Colors are matched with purple for pCA3 and orange for dCA3. n.s. = non-significant, *p < 0.05, **p < 0.01, *p < 0.001.

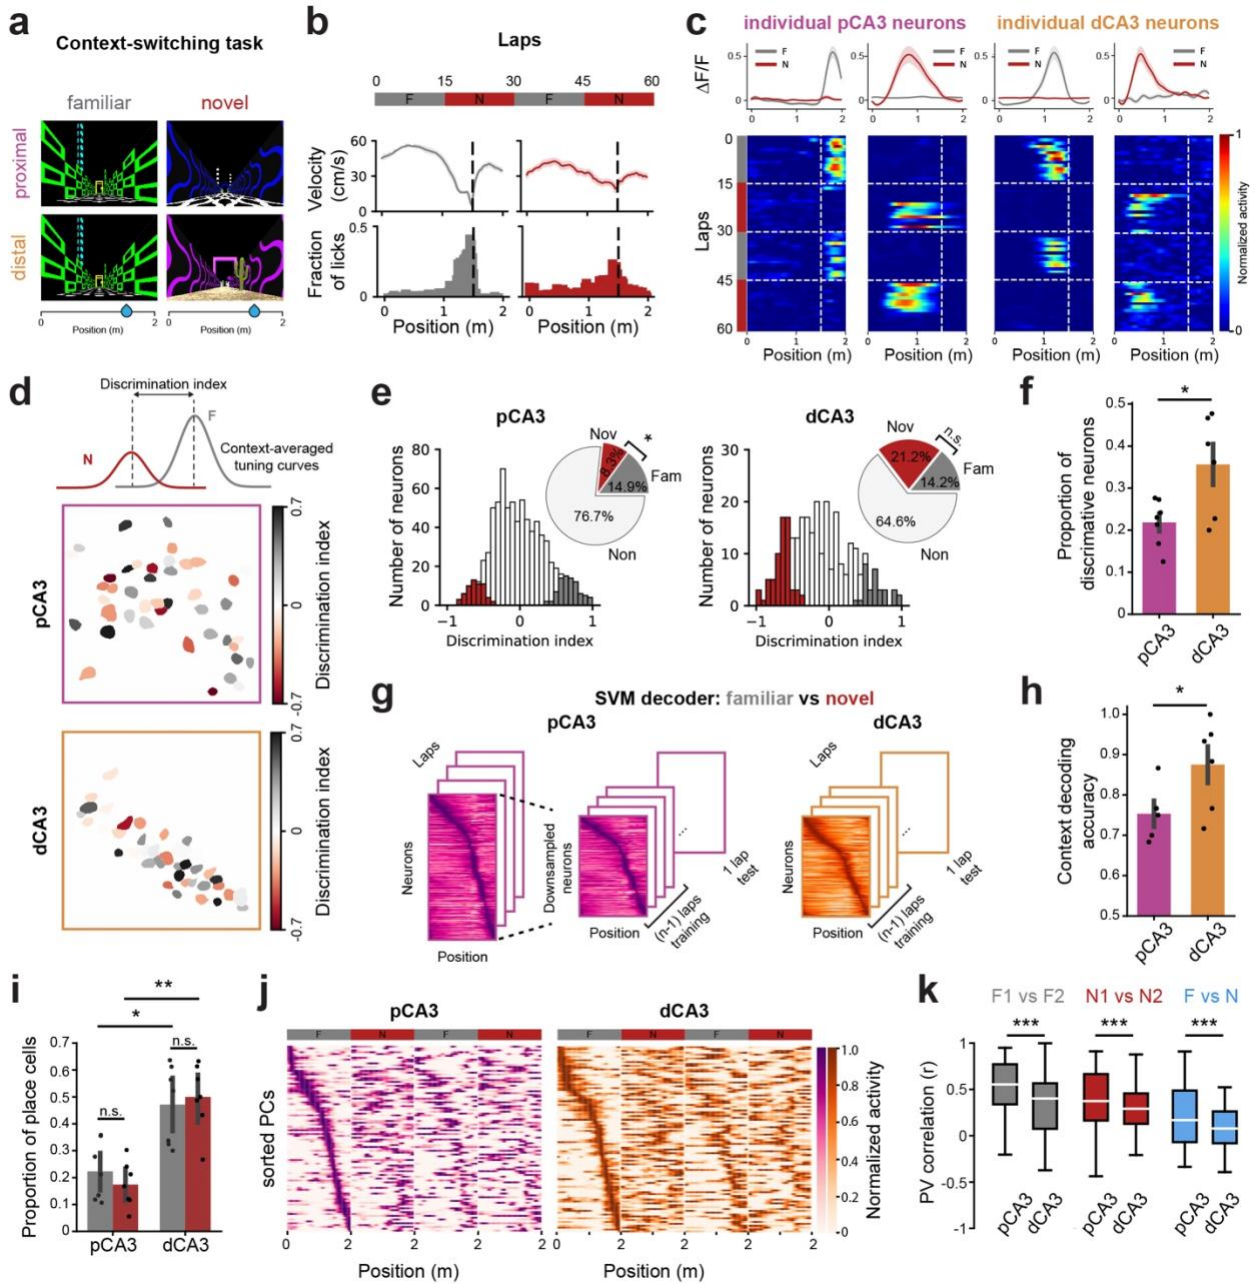

Figure 2. A generalized contextual- and spatial representation by proximal CA3 neurons. a-b. Schematic of the context-switching task (**a**) and behavioral responses from an example mouse (**b**). Each session consisted of 60 laps with alternating blocks between familiar (gray) and novel (red) contexts. Both environments were 2m in length with the same location of operant reward. Mice showed different velocity and licking responses in familiar and novel contexts. **c**, Mean calcium traces (*top*) and heatmaps of lap-by-lap activity (*bottom*) for example neurons from pCA3 (**b**) and dCA3 (**c**). **d**. Calculation of discrimination index (DI) from context-averaged tuning

curves (*top*) with the spatial distribution of context-discriminative PNs in pCA3 (*middle*) and dCA3 (*bottom*). **e**. Histogram of DI and significant DI values (colored bars, shuffling test, two-sided, $p < 0.05$). Inset: The fraction of neurons with significant DI in familiar (gray) and novel (red) environments. pCA3 neurons are more significantly tuned for familiar environments ($p = 0.035$; 7 mice), while dCA3 neurons are not significantly different between environments ($p = 0.27$; 7 mice). **f**. Proportion of neurons with significant DI value (mean ± SEM). dCA3 showed a higher level of context-discriminative neurons compared to pCA3 ($p = 0.047$; $n = 7$ mice for both pCA3 and dCA3). **g.** Schematic of the support vector machine (SVM) used to decode contextual information. Separate linear decoders were trained for subsampled pCA3 and dCA3 place matrices. **h**. SVM linear decoder performance on context discrimination (mean ± SEM). The peak decoding accuracies of dCA3 PNs in discriminating F-N contexts were higher compared to pCA3 PNs (0.61 ± 0.050 for pCA3, 0.85 ± 0.046 for dCA3; $p = 0.0011$; $n = 7$ mice for both pCA3 and dCA3). **i**. Proportion of place cells in familiar and novel environments. dCA3 showed a higher level of place cells in either familiar or novel environments compared to pCA3 (Kruskal-Wallis H-test, $H = 16.40$, $p = 0.00093$; post-hoc Dunn's multiple comparison with Bonferroni correction, $p = 0.015$ and 0.0012 for familiar and novel context, respectively; $n = 7$ mice for both pCA3 and dCA3). **j**. Heatmaps of spatial tuning curves for pCA3 (*left*) and dCA3 (*right*) PNs during context switching. Rows are average and normalized responses for all CA3 PNs, sorted by the position of peak activity during the F1 block. **k**. Spatial tuning curve similarities within (F1 vs F2, N1 vs N2) and between contexts (F vs N). The population of place cells in pCA3 showed a higher level of PV correlation across conditions ($p = 6.80$ e-29, 4.35 e-13, 6.25 e-19 for F1 vs F2, N1 vs N2, and F vs N, respectively, $n = 7$ mice for both pCA3 and dCA3). Mann-Whitney U tests were used to determine statistical significance unless otherwise stated. Colors are matched as follows: purple for pCA3, orange for dCA3, gray for familiar context, and red for novel context. n.s. = non-significant, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Recurrent neural network model

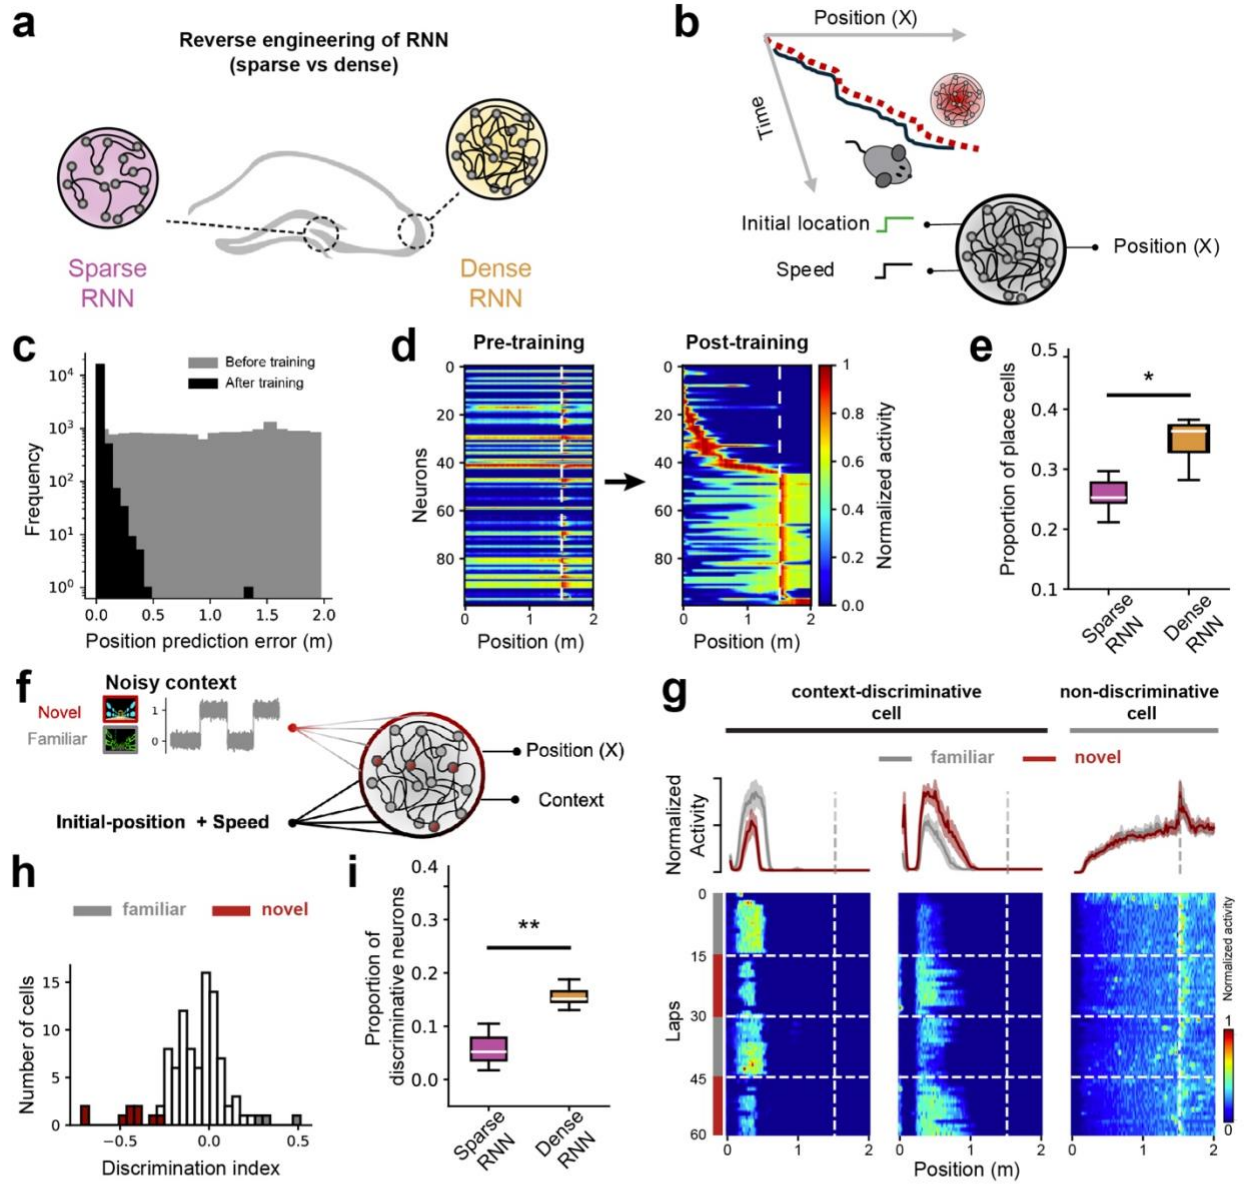

Hopfield network model

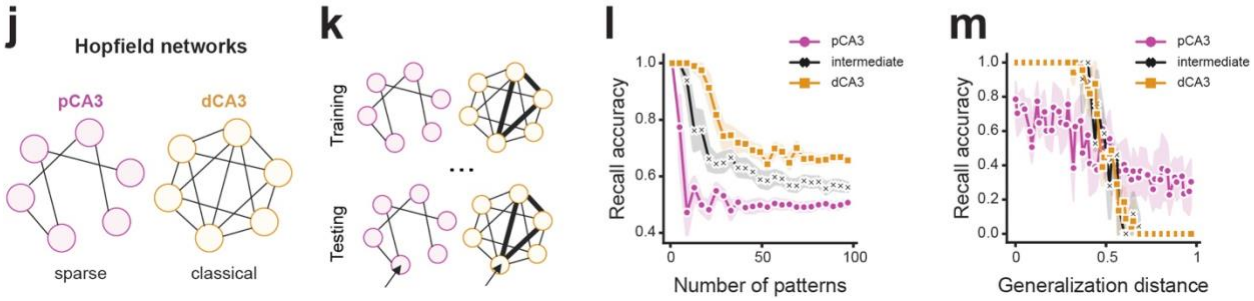

Figure 3. Distinct spatial and memory properties in sparse and dense neural networks. a. Schematic representation of modeling proximal and distal CA3 with sparse and dense RNNs, respectively. **b.** Training of RNNs for a linear track navigation task using inputs of instantaneous speed and lap initiation to mimic mouse trajectory. Network connectivity and activation functions are updated during training, and post-training, activation patterns are analyzed for reverse engineering. **c.** Decoding accuracy of the RNNs before (gray) and after (black) training, shown by the position prediction error. **d.** Tuning curve evolution of RNNs with 100 cells pre- and post-training, sorted by activity levels. The tuning curves cover the spatial scale of the virtual arena. **e.** Proportion of spatial-tuned artificial cells in RNNs with sparse versus dense recurrence. **f.** Training of RNNs for context-switching tasks, where task-relevant inputs are fed to RNNs. **g.** Examples of artificial cell activation post-training in the context-switching task. The left shows two cells with context modulation (familiar vs. novel), while the right shows a spatially tuned cell without context-dependent activity. **h.** Histogram of DI and significant DI values (colored bars, shuffling test, two-sided, $p<0.05$) for artificial cells. **i.** Dependence of context-discriminative cells on recurrence level. Box plots represent the fraction of context-tuned cells for RNNs with sparse versus dense recurrence. **j.** Sparse (*left*) and classical (*right*) Hopfield model of associative memory, corresponding to areas pCA3 and dCA3, respectively. **k.** Schematic of training (storage) and testing (recall). During storage, binary patterns are sequentially presented to the network and pairwise synaptic weights are updated according to a Hebbian plasticity rule. During retrieval, corrupted patterns are presented to the network, whose dynamics recover the closest previously stored pattern. **l.** Dense network (dCA3) has greater discrimination capacity than the sparse (10% connectivity) network (pCA3). In repeated trials, N patterns are stored in each network and then recalled (y axis: mean bit-recall accuracy; chance level 50%). Both networks achieve 100% accuracy for small N with accuracy decreasing as N increases, but the number of patterns the sparse network is able to discriminate is much smaller than the dense network. **m.** Sparse network (pCA3) has greater generalization capacity than the dense network (dCA3). Each network was trained to its maximum capacity while still achieving 100% accuracy in (**l**). To test generalization performance, increasingly different patterns were then presented to each network (number of corrupted bits). The sparse network was able to generalize more accurately from patterns with a larger distance to the original compared to the dense network. Colors are matched as follows: purple for pCA3, orange for dCA3, gray for familiar context, and red for novel context. n.s. = non-significant, $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

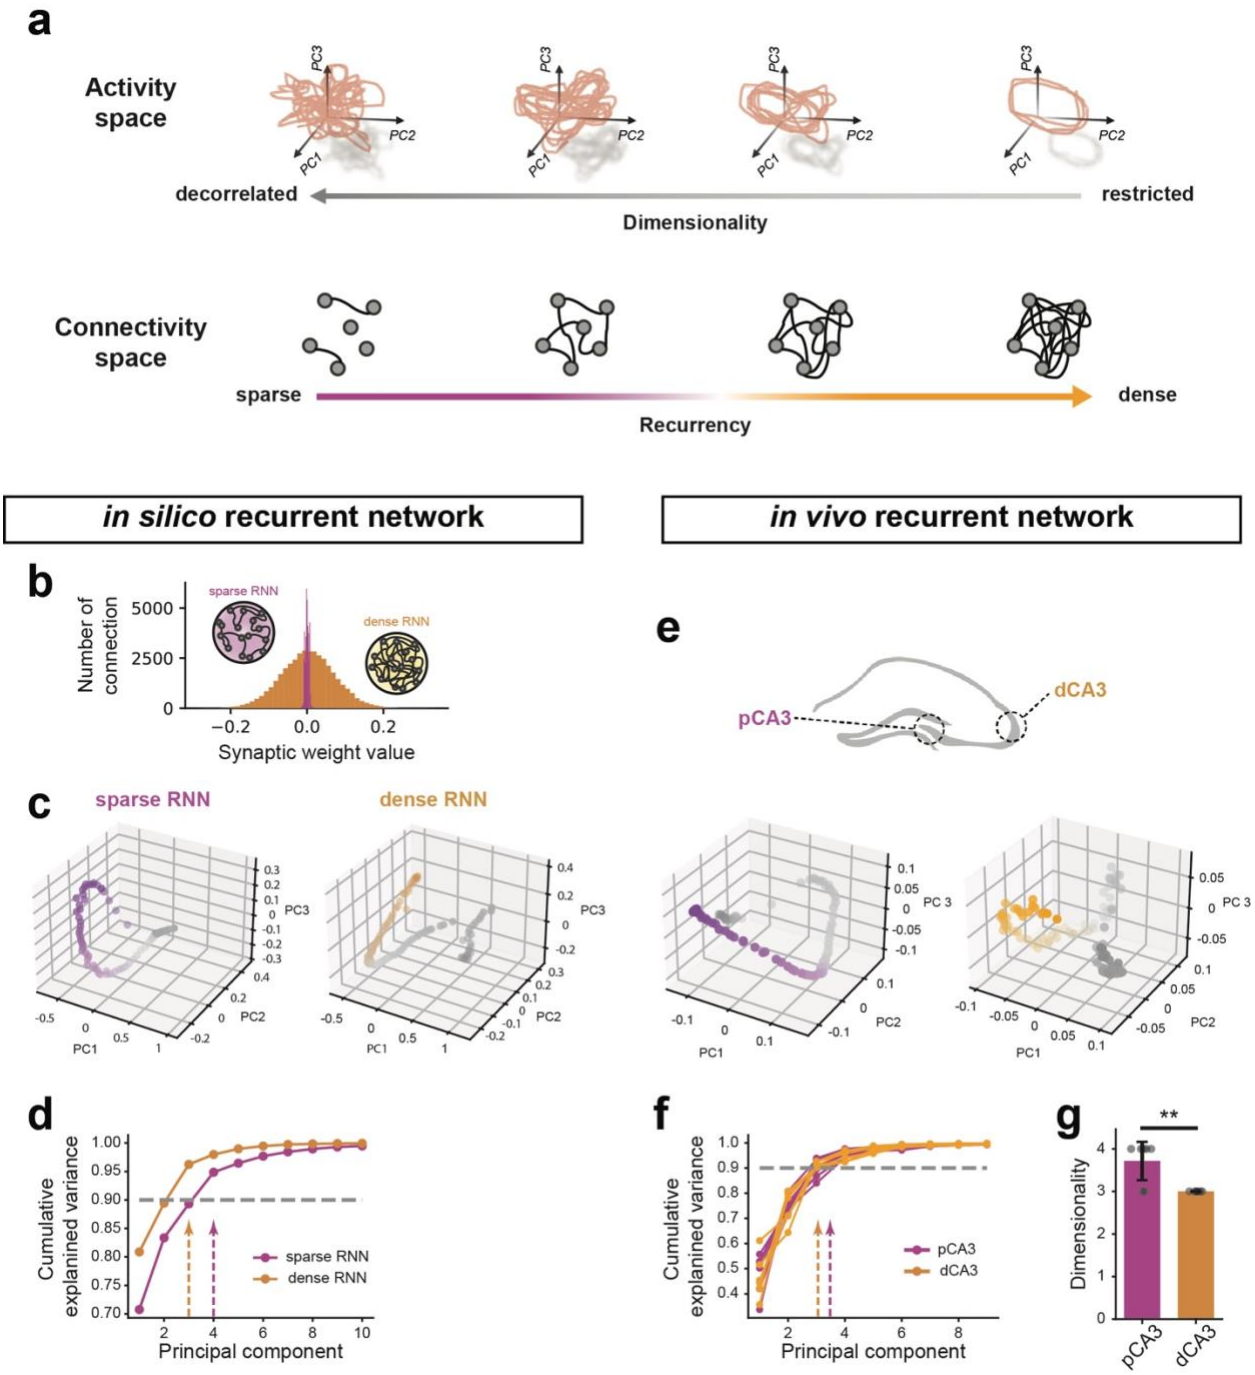

Figure 4. Neural manifold dimensionality in *in vivo* and *in silico* recurrent network. a. Illustration of a direct link between recurrency and neuronal manifold dimensionality. **b.** Distribution of recurrent weights in sparse and dense RNN models. **c.** Principal Component Analysis (PCA) of population activities from simulated sparse (*left*) and dense (*right*) RNNs, with projections onto the top three principal components shown. **d.** Cumulative explained variance as

a function of the number of principal components extracted from PCA with *in silico* simulated datasets. Dashed lines indicate the principal component at which 90% cumulative explained variance is reached. **e**. PCA of population activities in pCA3 and dCA3 during spatial navigation tasks, with projections onto the top three principal components shown. **f.** Cumulative explained variance as a function of the number of principal components extracted from PCA for *in vivo* neural activities. Dashed lines indicate the principal component at which 90% cumulative explained variance is reached. **g.** The average number of principal components required to reach 90% cumulative explained variance, as indicated by the dashed lines in **f** ($p = 0.009$). Mann-Whitney U tests were used to determine statistical significance unless otherwise stated. Colors are matched as follows: purple for pCA3 or sparse RNN, orange for dCA3 or dense RNN. n.s. = non-significant, *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$.

Supplementary Figures

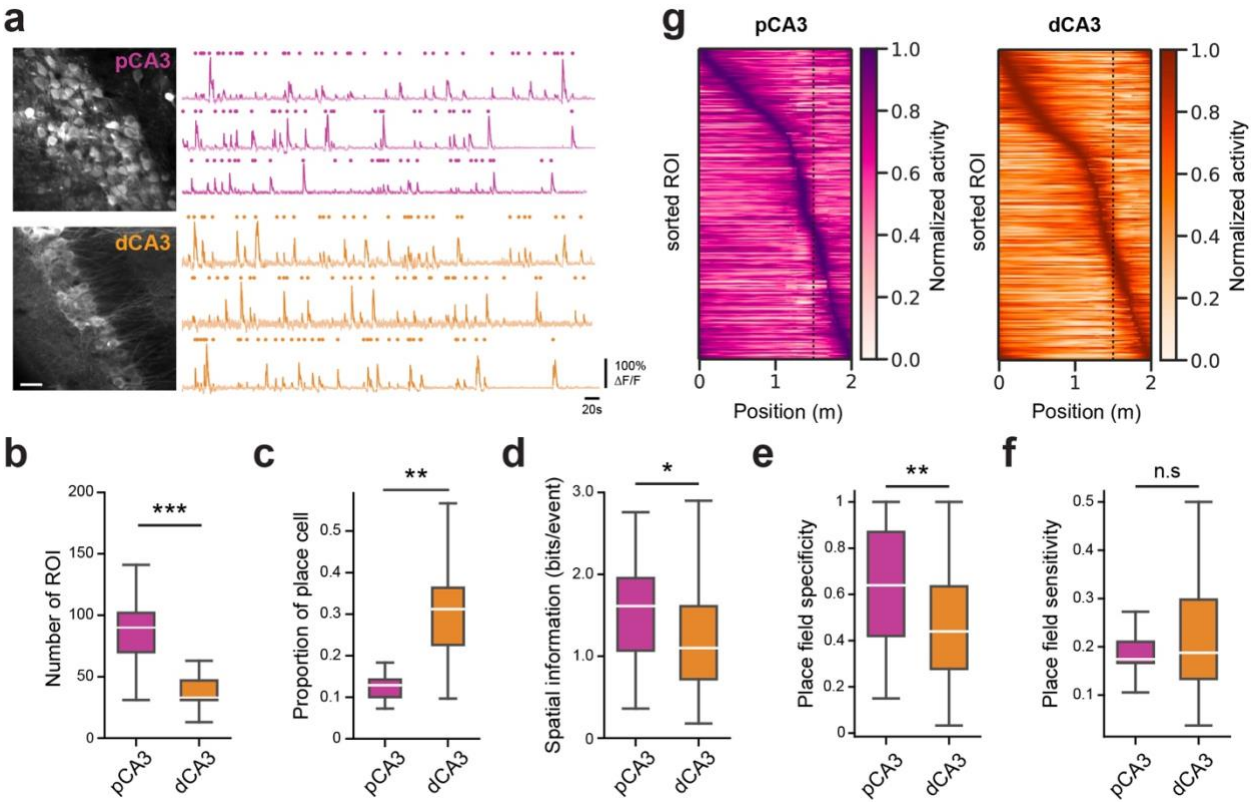

Supplementary Figure 1. (Related to Figure 1) **Spatial tuning of proximal and distal CA3 during spatial navigation. a**. Representative time average image (*left*) and ΔF/F traces (*right*) of proximal (*top*) and distal (*bottom*) CA3 neurons. Transient events in bold and start times indicated by circles above traces. **b**. Average number of detected ROI in proximal and distal CA3 ($p = 3.25$ e-5). **c**. Proportion of place cells in each CA3 subregion (mean ± SEM; $p = 0.0025$; n = 5 and 7 mice for pCA3 and dCA3, respectively). **d**. Spatial information (bits/event) of place cells in each subregion ($p = 0.037$) **e.** Place field specificity in each subregion ($p = 0.0076$). **f**. Place field sensitivity in each subregion ($p = 0.58$). **g**. Heatmaps of spatial tuning curves for pCA3 (*left*) and dCA3 (*right*) cells. Rows show average and normalized responses for all cells, sorted by the position of the peak activity. Mann-Whitney U tests were used to determine statistical significance unless otherwise stated. Colors are matched with purple for pCA3 and orange for dCA3. n.s. = non-significant, *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$.

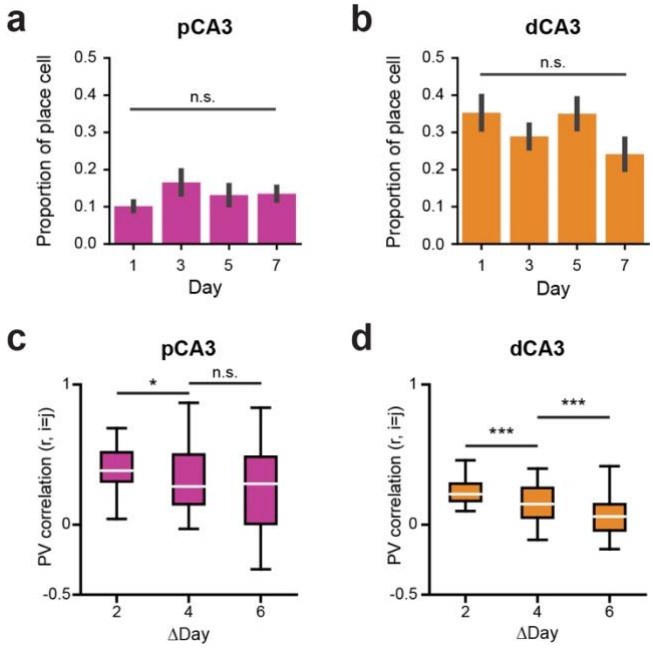

Supplementary Figure 2. (Related to Figure 1) **Long-term stability of place cells in proximodistal CA3. a-b**. Proportion of place cells across days in pCA3 (**a**) and dCA3 (**b**). No significant difference was observed across days in either region (Kruskal-Wallis *H*-test, pCA3: H = 2.69, p = 0.91; dCA3: H = 2.70, p = 0.72; post-hoc Dunn's multiple comparison with Bonferroni correction). **c**-**d**. PV correlation between the same position across days of place cells from pCA3 (**c**) and dCA3 (**d**). Correlation values decreased as interval between days increased. (**c**) p = 0.015 between intervals 2 and 4, p= 0.088 between intervals 4 and 6. (**d**) p = 6.81 e-06 between intervals 2 and 4, p = 0.00014 between intervals 4 and 6. Two-sided unpaired t-tests were used to determine statistical significance. Colors are matched with purple for pCA3 and orange for dCA3. n.s. = non-significant, *p < 0.05, **p < 0.01, ***p < 0.001.

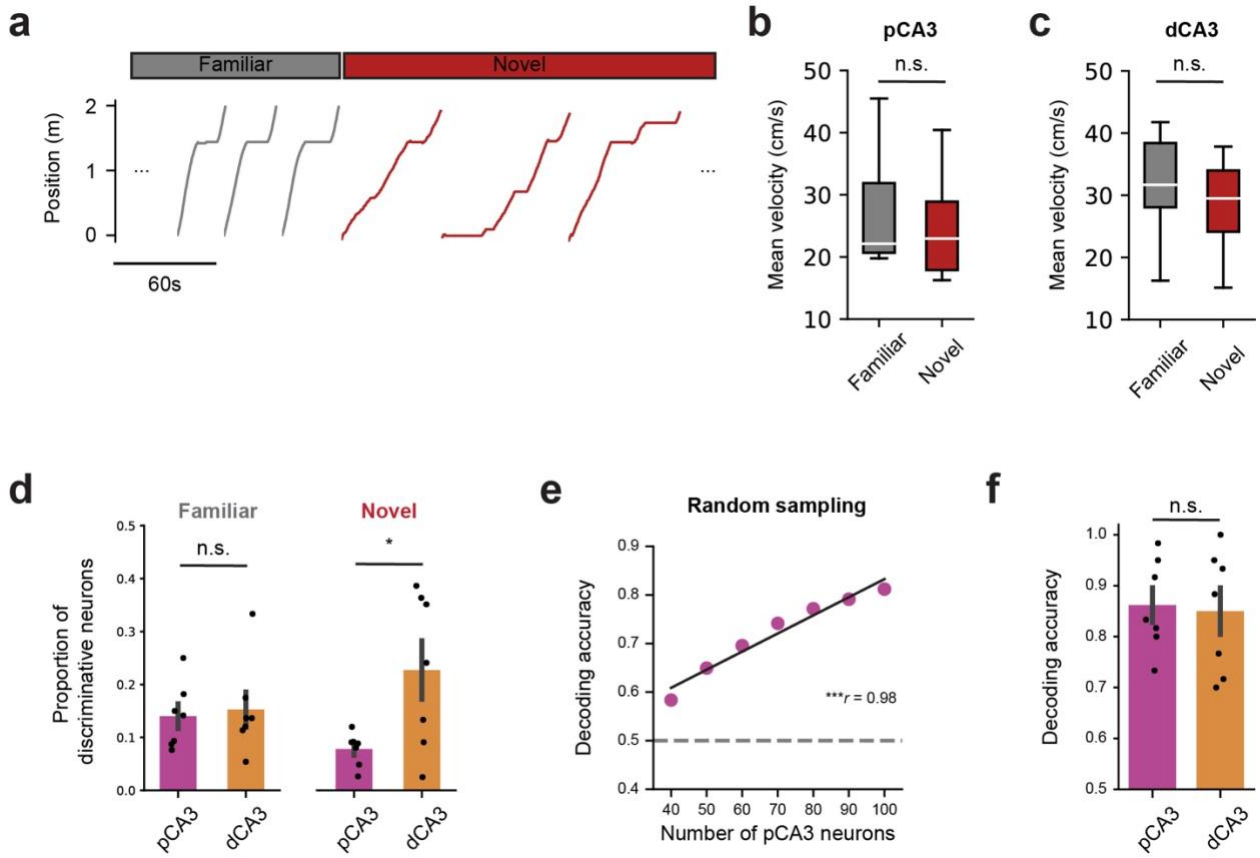

Supplementary Figure 3. (Related to Figure 2) **Context-discriminative neural activity in proximodistal CA3. a**. Position of the animal in the familiar and novel environments from an example session. **b-c**. Comparison of mean velocity in the familiar and novel environments for pCA3 (b) ($p = 0.53$; n = 7mice) and dCA3 (c) ($p = 0.32$; n = 7mice) recording sessions. **d**. Proportion of neurons with significant DI value in familiar and novel environments (7 mice for both pCA3 and dCA3). The proportion of PNs significantly tuned to the familiar environment was similar between pCA3 and dCA3 (0.14 ± 0.23 for pCA3, 0.15 ± 0.03 for dCA3; $p = 1.0$), but dCA3 showed a higher level of novelty-tuned neural responses compared to pCA3 (0.078 ± 0.012 for pCA3, 0.23 ± 0.055 for dCA3; $p = 0.021$). **e**. Peak decoding accuracy correlated with the number of randomly subsampled pCA3 neurons (Pearson's $r = 0.98$, $p = 0.00015$, 7 mice). The black line indicates the regression line. **f**. Performance of the linear decoder for context in pCA3 and dCA3 without subsampling. No significant differences were found in peak decoding accuracies between pCA3 and dCA3 (0.86 ± 0.034 for pCA3, 0.85 ± 0.046 for dCA3, $p = 0.85$, 7 mice for both pCA3 and dCA3). Mann-Whitney U tests were used to determine statistical significance unless otherwise stated. Colors are matched as follows: purple for pCA3, orange for dCA3, gray for familiar context, and red for novel context. n.s. = non-significant, *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$.

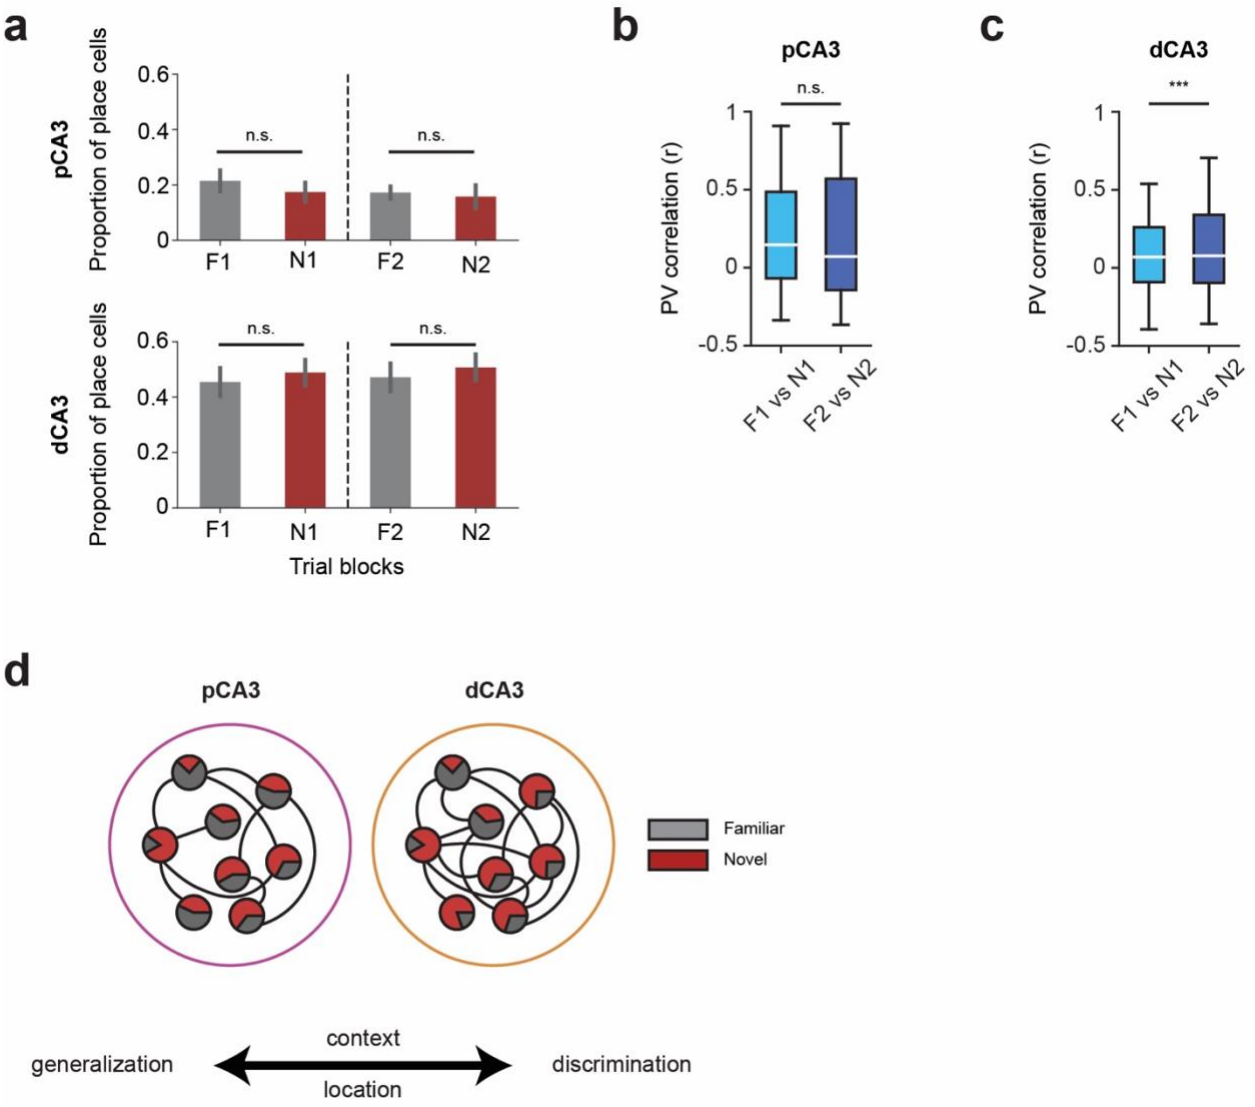

Supplementary Figure 4. (Related to Figure 2) **Place cell responses across blocks in a context-switching task. a.** Proportion of place cells in each trial block in pCA3 (top) and dCA3 (bottom). Familiar and novel blocks showed similar levels of place cells in both the first and second exposures (Kruskal-Wallis H-test, H = 1.43, p = 0.69 for pCA3, H = 0.73, p = 0.87 for dCA3). **c**-**d**. PV correlation between familiar and novel contexts for the first and second blocks. No significant difference was observed in pCA3 between the first and second blocks (p = 0.71), whereas a significant difference was found in dCA3 (p = 0.0003). **d**. Schematics showing gradients of generalization and discrimination of context and location by proximodistal CA3 with different levels of recurrent connection. Colors are matched as follows: purple for pCA3, orange for dCA3, gray for familiar context, and red for novel context. n.s. = non-significant, *p < 0.05, **p < 0.01, ***p < 0.001.

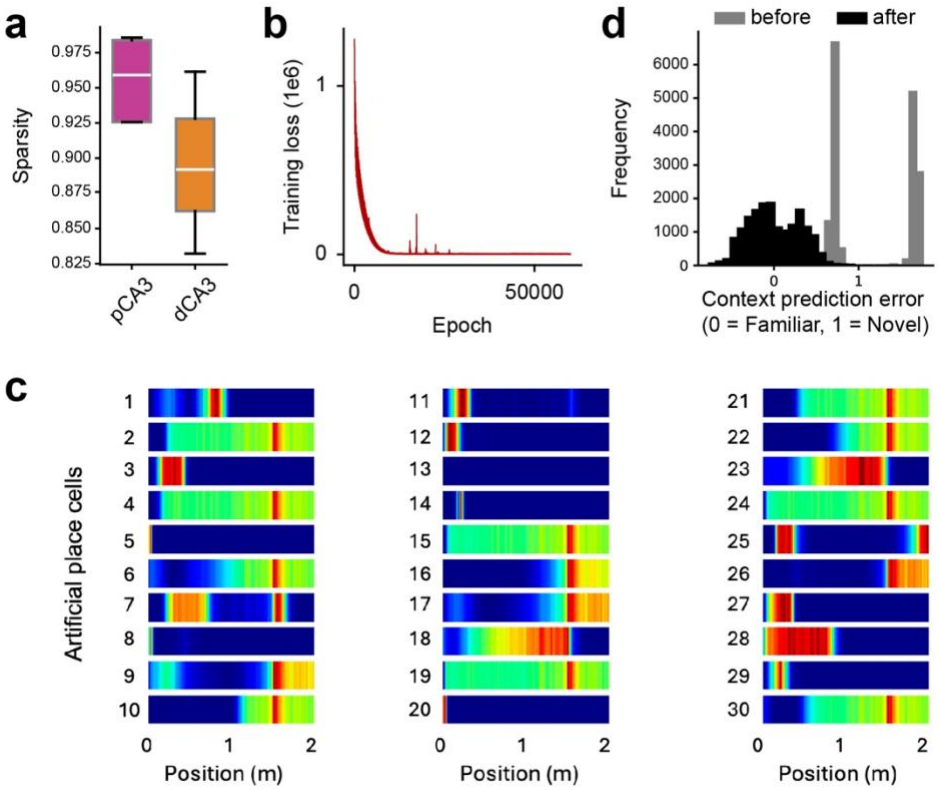

Supplementary Figure 5. (Related to Figure 3) **Spatial tuning properties of *in silico* place cells in RNN. a**. Comparing sparsity level via functional connectivity matrices: sparsity is evaluated based on the functional connectivity matrix of neurons in each subregion, using a 20-neuron window for both distal and proximal recordings. Results averaged across a recording session align with physiological studies, indicating greater recurrence in the dCA3 subregion. **b**. Training loss of multi-task RNN: the training loss curve for the multi-task RNN (referenced in Fig. 3f) demonstrates the stability of the model's convergence. **c**. spatial RNN tuning curve with different neuron number: Example of tuning curves for a spatial RNN network, modeled with varying neuron counts (30 neurons), highlighting the network's spatial encoding capabilities. **d**. Distribution of prediction error across all laps before (gray) and after (black) training RNN.

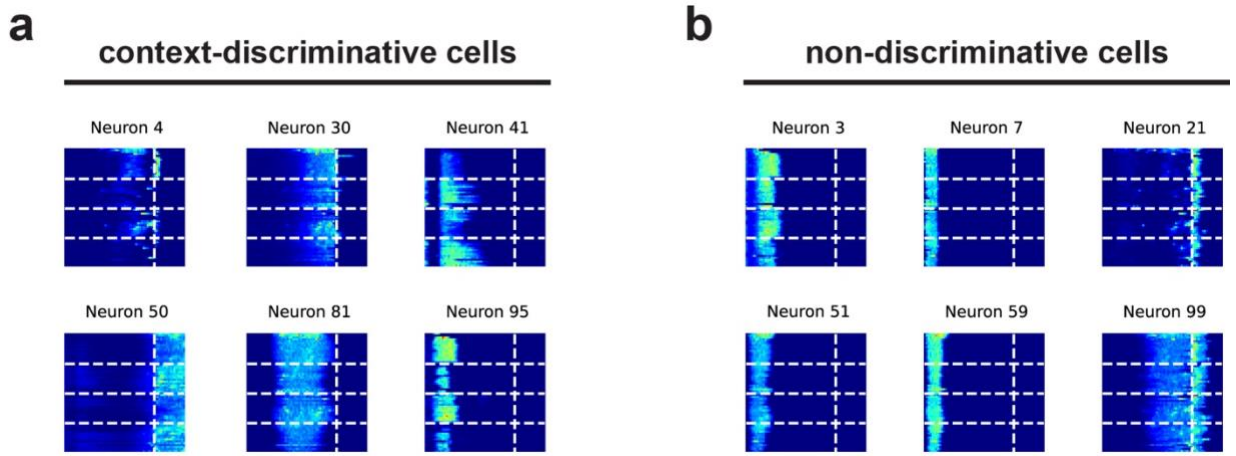

Supplementary Figure 6. (Related to Figure 3) **RNN artificial cell responses during context-switching task.** Example heatmaps of lap-by-lap activity for RNN artificial cells are shown. Context-discriminative cells (**a**) and non-discriminative cells (**b**). Horizonal dashed lines indicate context changes, and vertical dashed lines indicate the reward location.

Methods

All experiments were conducted in accordance with NIH guidelines and with the approval of the Columbia University Institutional Animal Care and Use Committee. No statistical methods were used to predetermine sample sizes. The experiments were not randomized, and the investigators were not blinded to allocation during experiments and outcome assessment.

Animals

Imaging experiments were performed with healthy, 2- to 4-month-old, heterozygous adult male and female Grik4-cre mice (Jackson Laboratory, 006474) on a C57BL/6J background. Mice were group-housed under normal lighting conditions in a 12-hour light/dark cycle. Ad libitum water was provided until the beginning of training for the spatial navigation task.

Viruses

CA3 pyramidal cell imaging experiments were performed by injecting a Cre-dependent rAAV expressing GCaMP8s under the control of the synapsin promoter (rAAV9-Syn-FLEX-jGCaMP8s-WPRE-SV40; Addgene, 162377; titer, 2.3×10^{13} viral genomes per ml).

Surgery

All surgical procedures were performed on Grik4-cre mice (2-4 months old) under isoflurane anesthesia (4% induction, 1.5% maintenance in 95% oxygen). Mice were placed on a stereotaxic surgery instrument (Kopf Instruments), and their body temperature was maintained using a heating pad. Meloxicam and bupivacaine were administered subcutaneously to reduce discomfort. Following a skin incision, a craniotomy was performed over the right hippocampus using a drill. A sterile glass capillary loaded with rAAV was attached to a Nanoject syringe (Drummond Scientific) and slowly lowered into the right hippocampus. The proximal and distal CA3 were targeted with two x–y coordinates, each consisting of two different injection sites separated in z:

AP –1.46, ML –1.75, DV –2.1, –1.9 and AP –1.46, ML –1.3, DV –2.3, –2.1 relative to bregma, with 60 nl of virus injected at each DV location. The pipette was held for 5-10 minutes after the last injection and then slowly retracted from the brain. The skin was sutured, and the mice were allowed to recover for 4 days before the window/headpost implant.

For hippocampal window and headpost implantation, as described previously[23,30,31], a 3-mm craniotomy was performed over the right anterior hippocampus, centered between the two injection coordinates. We then slowly aspirated the cortex overlying the right dorsal hippocampus and implanted a 3-mm glass-bottomed stainless-steel cannula for optical access. Subsequently, a titanium headpost with layers of dental cement was attached to the skull. The mice received 1.0 ml of saline subcutaneously and recovered in their homecage on the heating pad. All mice were monitored for 3 days of post-operative care until behavior training began.

Behavioral paradigm

After a 14-day recovery period from implant surgery, mice were water-restricted to 85-90% of their original body weight and habituated to handling and head fixation. They were then exposed to a 2-meter-long linear virtual reality (VR) corridor that remained consistent during training and recording[26,31]. At the end of the environment, an inter-trial interval of 2 seconds with a blank screen was included, and the mouse was simply teleported to the next lap.

For the next 10-14 days, mice were trained to run through the virtual environment and lick for a water reward. In the habituation phase, a drop of water was given in a non-operant manner when the subject mice entered the reward zone (RZ). Additional water droplets were provided as they continued licking within a 30 cm area. In the training phase, the reward was given in an operant manner where the correct lick within the RZ could trigger the first drop of reward and subsequent droplets with additional licking. Mice were trained

to run at least 30 laps in the environment. After recording during the GOL task, the context-switching task was performed the next day.

For context-switch experiments, we used a completely distinct set of visual cues along the 200 cm track as a novel environment. Each recording session consisted of four blocks: 15 laps in the Familiar context, 15 laps in the Novel context, 15 laps in the Familiar context, and 15 laps in the Novel context. The reward location remained fixed at 150 cm for both contexts. For subject mice used in both pCA3 and dCA3 imaging experiments, we used different visual cues for the novel environments, while the familiar context remained the same for both sequences.

Histology

After the completion of imaging experiments, mice were transcardially perfused with 40 ml of ice-cold PBS (Thermo Fisher), followed by 40 ml of ice-cold 4% paraformaldehyde (PFA; Electron Microscopy Sciences). The brain was then kept submerged in 4% PFA overnight at 4 °C. On the next day, coronal hippocampal sections (thickness: 100 μm) were obtained using a vibratome (VT1200-S, Leica, Germany). After brief washing with 0.3% Triton X-100 dissolved in phosphate-buffered saline (PBST), brain slices were incubated in a blocking solution (5% normal donkey serum, 0.3% PBST) for two hours at room temperature. Subsequently, slices were incubated with the primary antibody (rabbit anti-PCP4 diluted 1:500; HPA005792; Sigma-Aldrich) in a blocking solution overnight at 4 °C. After the sections were washed three times for 15 minutes with 0.3% PBST at room temperature, they were incubated with the secondary antibody (Alexa Fluor 647-conjugated donkey anti-rabbit IgG antibody diluted 1:800; 711-605-152; Jackson ImmunoResearch) in the blocking solution for 2 hours. Slices were rinsed three times for 15 minutes with 0.3% PBST, stained with DAPI (1:1000 dilution in 0.3% PBST), and then mounted on microscope slides. Fluorescence imaging was performed using a Nikon Ti-E A1R laser-scanning confocal microscope with ×10/0.45-NA Plan Apo (Nikon) or ×20/0.75-

NA Plan Apo (Nikon) objective lens. 3 µm Z-stack images were processed using NIS Elements software (Nikon) or Fiji (ImageJ).

In vivo two-photon imaging

All imaging was conducted using a 2-photon 8 kHz resonant scanner (Bruker) and high-NA multiphoton objective (XLPLN25XWMP2, 25x/1-NA, Olympus). For excitation, we used a 920 nm femtosecond-pulsed laser (Chameleon Ultra II, Coherent). Pockels cells were used to regulate the power of the laser reaching the tissue. Green (GcaMP8s) fluorescence was collected through an emission cube filter set (HQ525/70 m-2p) to a GaAsP photomultiplier tube detector (Hamamatsu, 7422P-40). A custom dual-stage preamp (1.4 x 105 dB, Bruker) was used to amplify signals prior to digitization. All images were acquired with 512 x 512 pixel resolution at 30 Hz.

Calcium imaging data preprocessing

The preprocessing steps for the raw fluorescence signal have been described elsewhere[51,77]. Briefly, the imaging data was motion corrected using the SIMA software package[78]. The time average of each imaged cell was manually inspected, and an ROI was hand-drawn over each cell using a data visualization server program developed in the lab. The same CA3 PNs were transferred across sessions wherever possible, and identified with a unique ID, so that their activity across sessions could be tracked. Fluorescence was extracted from each ROI using the FISSA software[79] package to correct for neuropil contamination using six patches of 50% the size of the original ROI. For each resulting raw fluorescence trace, a baseline F was calculated by taking the first percentile in a rolling window of 30 s, and a $\Delta F/F$ trace was calculated. The $\Delta F/F$ trace for each cell was smoothed using an exponential filter, and all further analyses were performed on the resulting $\Delta F/F$ traces. We detected statistically significant transients as described previously[77] to use for place field calculations. All further analyses were implemented using Python using custom-written scripts.

Spatial tuning analysis

The virtual environment was divided into 100 evenly spaced bins (2 cm), which were then utilized to bin a histogram of each cell's neuronal activity. Neuronal activity was filtered to include activity from when the animal was running above 3 cm/s and to exclude activity during the 2-sec teleportation at the end of the 2-m track. The spatial tuning curves were normalized for the animal's occupancy and then smoothed with a Gaussian kernel ($\alpha = 6$ cm) to obtain a smoothed activity estimate.

To detect place fields, we generated a baseline distribution of spatial tuning curves for each neuron by randomly and independently shifting the activity on each lap a random distance in a circular manner. Following this, we recalculated the smoothed, lap-averaged tuning curve as detailed above. We repeated this process a thousand times, determining the 95th percentile of baseline tuning values at each spatial bin, which served as the threshold for significant spatial tuning ($p < 0.05$). Areas of space where the true spatial tuning curve surpassed this baseline threshold were identified as potential PFs. In order to be classified as a place cell, ROIs needed to have at least 3 consecutive significant bins (6cm), but less than 25 consecutive bins (50cm). Moreover, in order to avoid spurious detection of significant bins, we set an additional criterion that there must be activity within the PF boundaries for a minimum of 20% of laps. For place cell sensitivity, representing the proportion of laps that had active events within the significant field, the number of laps that had at least 1 binarized event within the significantly detected fields was divided by the total number of laps. For place cell specificity, the total number of events within the significant field was divided by the total number of events observed within the lap. Then, it was averaged across all laps to have a single value for each ROI. If multiple fields were detected, they were computed separately and averaged across fields to have a single value for each ROI. For spatial information, it was calculated as described previously[80].

Tuning-curve correlation

To determine the level of similarity between spatial representations across different sessions for each place cell, we calculated the tuning curve correlation (Pearson

correlation) for each pair of sessions imaged on different days. Only tuning curves of cells with place fields on both days were included in the analysis.

Population vector correlation

To assess the stability of spatial representations over days, we calculated the population vector (PV) correlation (Pearson correlation) between two sessions imaged on different days. Only cells defined as place cells on both days were included in the analysis. The diagonal values of the correlation matrix represent the correlation between corresponding spatial bins across different sessions. Specifically, each diagonal value indicates the correlation between the PV in a given spatial bin in one session and the PV in the same bin in another session. We averaged these diagonal correlation values across all spatial bins. This process was repeated for all pairs of sessions, and the correlation coefficient was quantified as a function of the day interval between sessions.

We also used PV correlations to quantify the similarity of the spatial code between and across contexts – F1 vs F2, N1 vs N2, and F vs N. Cells that were defined as place cells in at least one of the two blocks were included in this analysis. The same quantification process for correlation coefficients was applied

Discrimination index

As described previously[81], the discrimination index (DI) was calculated for each cell to estimate the response preference toward a given context. The DI was defined by the following formula:

$$Discrimination\ index = \frac{\mu_1 - \mu_2}{\sqrt{\frac{1}{2}({\sigma_1}^2 + {\sigma_2}^2)}}$$

where μ and σ represent mean and standard deviations of calcium response amplitudes across laps of each environment, respectively[82].

To test the significance of the DI value, we generated 1,000 surrogate data sets in which laps were randomly permuted and calculated the 95th percentile of the null value per cell. If the experimentally observed DI value surpassed the top or bottom 2.5% of the distribution of surrogate DI values, we considered it context-selective. Therefore, 5% of neurons were expected to have a significant DI value in a given session by chance.

Support vector machine (SVM) analysis

To assess whether neuronal population activity patterns convey contextual information, we utilized a linear classifier support vector machine (SVM) to classify neuronal activity patterns into familiar or novel environments. For each imaging session, independent SVMs were trained and tested at each time point using a leave-one-trial-out cross-validation approach. Specifically, each decoder was trained with the neural population activity pattern from all laps except one withheld lap (n-1). We then evaluated whether the trained decoder accurately classified the held-out lap into the correct context. This process was repeated for each lap in the session, ensuring that every lap was utilized as a test lap at least once. Decoding accuracy was measured as the proportion of correct classifications. To compare the decoding efficacy between CA3 subregions, cells from proximal CA3 were subsampled to match the average detected ROI from distal CA3. The same SVM analysis procedures were then applied separately to the pCA3 and dCA3 datasets.

Neural manifold analysis

We employed Principal Component Analysis (PCA) to reduce the dimensionality of neural data and visualize neural trajectories in a low-dimensional subspace. Neural activity was represented as an N × T matrix, where N denotes the number of neurons, and T is the number of time points. To ensure comparability between pCA3 and dCA3 regions, we matched the number of recorded neurons across sessions. The data were mean-centered before PCA. The covariance matrix of the centered data, C, was calculated as:

$$C = \frac{1}{T-1} X_{centered} X_{centered}^{T},$$

where $X_{centered}$ is the mean-centered data matrix. Eigenvalue decomposition was performed on C, yielding eigenvalues λ_i and corresponding eigenvectors (principal components). The neural data were projected onto the subspace spanned by the leading eigenvectors.

To determine the dimensionality of the data, we identified the minimum number of principal components k needed to explain at least 90% of the cumulative variance, calculated as:

$$\sum_{i=1}^{k} \frac{\lambda_i}{\sum_{j=1}^{N} \lambda_j} \geq 0.90.$$

The same manifold dimensionality quantification was applied to both neuronal recordings and sparse and dense recurrent neural networks (RNNs) from simulation.

Modeling

Recurrent neural network modeling

Each subregion's recurrent neural network (RNN) model consisted of a recurrent layer, alongside input and output layers. The network architecture and training method of our artificial neural networks was imitated from previous papers on path integrator artificial

networks[83,84]. The recurrent layer comprised 100 rate neurons with tanh activation functions and was implemented using the TensorFlow LSTM package[85]. The output layer is a single (path integration) / two (multi-tasking network) neurons with dense connection from the recurrent layer with a linear activation function. The inputs to both the spatial imitator network and the multitasking network were the instantaneous speed and the initial position of the mice on the linear track. The initial position inputs were held constant, serving as the trial initiation signal. The network's objective was to replicate the mice's momentary position during each trial by integrating momentary speed and internally generated temporal signals within the network.

To stabilize network convergence, we introduced Gaussian noise as input to the network, a method shown to enhance convergence[86]. In the multitasking RNN (context switching + path integration), we included a noisy input signal simulating the context visual signal encoded as a binary code, combined with Gaussian noise centered around zero with a standard deviation of 0.3. This perceptual noise was added to mimic the context-dependent visual cue noise in the virtual arena. With a noiseless context signal, we anticipated a smaller neuronal population to be involved in context identification computations. The goal of the multitasking RNN was to minimize the Euclidean distance in a two-dimensional subspace between predicted location-context pairs and actual location-context pairs.

To implement recurrent sparsity, an *L1* regularization term was incorporated into the loss function, with different coefficients applied to promote either sparse (λ_{sparse} = 0.1) or dense λ_{dense} = 0.001) connectivity. The total loss L is defined as:

$$L = \sum (X - X_{Predict})^2 + \lambda \sum |W_{recurrent}|$$

Where X denotes the vector of momentary location and context, $X_{Predict}$ is the network's output, $W_{recurrent}$ represents the recurrent weights across hidden layer neurons, and λ is the regularization coefficient.

Training was conducted using the Adam optimizer from the Keras package[85], with a learning rate of 1e-4 and a minimum batch size of 8. We ensured the spatiality of training by verifying the saturation of the training loss function.

To reverse-engineer network activity, we analyzed the response functions of individual neurons in the hidden layer during a feedforward epoch, where instantaneous speed and context signals were provided as inputs. Analysis of the artificial neurons, including cell selectivity and population activity, followed the same criteria used for biological neurons in this study. Additionally, we replicated the formation of place-selective cells using a network with a smaller hidden layer size (30 neurons). Population-level results were based on simulations involving 100 hidden neurons.

Hopfield networks

Hopfield networks were implemented in Python 3.8 in line with previous studies[87]. In brief, a network was initialized with *N* binary (±1) valued neurons. In the training phase, *M* patterns ξ are stored in the weight matrix according to the Hebbian rule

$$W = \sum \xi_j \, \xi_j^T$$

In the test phase, the trained network is initialized to a test pattern ξ', constructed by flipping a predetermined fraction of the activities of a parent pattern ξ chosen from the *M* patterns the network was previously trained on. The network is then allowed to evolved according to the classical Hopfield dynamics

$$x(t) = sgn\big(Wx(t-1)\big)$$

where $x(0) = \xi$ until either convergence is reached (i.e., $x(t) = x(t-1)$) or a maximum number of iterations T_{max} is reached.

To introduce sparsity into the Hopfield network, we choose a connection probability *p* and sample a structural connectivity matrix *A* where a connection exists between each pair of neurons i.i.d. with probability *p* (equivalently, a graph is drawn from the Erdős-Rényi process parameterized by *p*). The case *p*=1 reduces to the classical Hopfield network with full connectivity, while smaller *p* represents higher levels of sparsity. In sparse networks, weights between non-structurally connected neurons are not updated during training or used in the dynamics (which are otherwise unchanged), such that the effective weight matrix becomes

$$W_{eff} = A \odot W$$

where $\odot$ is the elementwise (Hadamard) product.

To model discrimination, we investigated how many distinct patterns networks at different levels of sparsity could store and accurately recall by drawing *M* random binary vectors for a range of *M* and storing them in the network. In the test phase, a fixed fraction of neurons were flipped and accuracy was read out as the similarity between the recovered pattern and the ground truth pattern.

To model generalization, we investigated perturbations to a base pattern that could still be identified with the base pattern by networks at different levels of sparsity by fixing the number of patterns *M* and varying the fraction of flipped neurons. We again used similarity between the recovered pattern and the ground truth pattern as our accuracy metric.

Statistics

Statistical differences between means were determined by two-sided unpaired t-test and Mann-Whitney U test, Kruskal–Wallis tests with post hoc Dunn's multiple comparison tests as mentioned in the text or figure legends. A p-value<0.05 was used as the criterion for statistical significance. Boxplots show the 25th, 50th (median), and 75th quartile

ranges with the whiskers extending to 1.5 interquartile ranges below or above the 25th or 75th quartiles, respectively. All data analysis and visualization were done with custom software built on Python version 2.7.15, and 3.8.11. All data were expressed as mean ± SEM.

Acknowledgments

We thank Bovey Y. Rao and Jingcheng Shi for comments on the early version of the manuscript, George Zakka for technical assistance, and Leo Rodriguez for the experimental schematic in Fig. 1b. Confocal imaging was performed with support from the Zuckerman Institute's Cellular Imaging platform. T.G is supported by an R00MH129565 from the National Institute of Mental Health (NIMH). A.L. is supported by National Institute of Mental Health (NIMH) R01MH124047 and R01MH124867; National Institute on Aging (NIA) RF1AG080818; National Institute of Neurological Disorders and Stroke (NINDS) Brain Initiative U01NS115530; NINDS R01NS121106, NINDS R01NS131728, and NINDS Brain Initiative R01NS133381.

Author Contributions

E.K., T.G., and A.L. conceived the project. E.K. and T.G designed experiments. E.K collected data, and performed analysis with assistance from T.G., T.S.M., and D.S.P. E.Z. and Z.L. performed the computational modeling. All authors wrote the paper.

Conflicts of Interest

The authors declare no competing interests.

References

1 Oldenburg, I. A. *et al.* The logic of recurrent circuits in the primary visual cortex. *Nat Neurosci* **27**, 137-147 (2024). https://doi.org:10.1038/s41593-023-01510-5

2 Peron, S. *et al.* Recurrent interactions in local cortical circuits. *Nature* **579**, 256-259 (2020). https://doi.org:10.1038/s41586-020-2062-x

3 Khona, M. & Fiete, I. R. Attractor and integrator networks in the brain. *Nat Rev Neurosci* **23**, 744-766 (2022). https://doi.org:10.1038/s41583-022-00642-0

4 Hopfield, J. J. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America* **79**, 2554-2558 (1982). https://doi.org:10.1073/pnas.79.8.2554

5 Marr, D. Simple memory: a theory for archicortex. *Philos Trans R Soc Lond B Biol Sci* **262**, 23-81 (1971). https://doi.org:10.1098/rstb.1971.0078

6 Wang, X. J. Decision making in recurrent neuronal circuits. *Neuron* **60**, 215-234 (2008). https://doi.org:10.1016/j.neuron.2008.09.034

7 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

8 Kali, S. & Dayan, P. The involvement of recurrent connections in area CA3 in establishing the properties of place fields: a model. *J Neurosci* **20**, 7463-7477 (2000). https://doi.org:10.1523/JNEUROSCI.20-19-07463.2000

9 Kesner, R. P. & Rolls, E. T. A computational theory of hippocampal function, and tests of the theory: new developments. *Neurosci Biobehav Rev* **48**, 92-147 (2015). https://doi.org:10.1016/j.neubiorev.2014.11.009

10 Nakazawa, K., McHugh, T. J., Wilson, M. A. & Tonegawa, S. NMDA receptors, place cells and hippocampal spatial memory. *Nat Rev Neurosci* **5**, 361-372 (2004). https://doi.org:10.1038/nrn1385

11 Knierim, J. J. & Neunuebel, J. P. Tracking the flow of hippocampal computation: Pattern separation, pattern completion, and attractor dynamics. *Neurobiol Learn Mem* **129**, 38-49 (2016). https://doi.org:10.1016/j.nlm.2015.10.008

12 McClelland, J. L. & Goddard, N. H. Considerations arising from a complementary learning systems perspective on hippocampus and neocortex. *Hippocampus* **6**, 654-665 (1996). https://doi.org:10.1002/(SICI)1098-1063(1996)6:6<654::AID-HIPO8>3.0.CO;2-G

13 Ishizuka, N., Weber, J. & Amaral, D. G. Organization of intrahippocampal projections originating from CA3 pyramidal cells in the rat. *J Comp Neurol* **295**, 580-623 (1990). https://doi.org:10.1002/cne.902950407

14 Li, X. G., Somogyi, P., Ylinen, A. & Buzsáki, G. The hippocampal CA3 network: an in vivo intracellular labeling study. *J Comp Neurol* **339**, 181-208 (1994). https://doi.org:10.1002/cne.903390204

15 Thompson, C. L. *et al.* Genomic Anatomy of the Hippocampus. *Neuron* **60**, 1010-1021 (2008). https://doi.org:https://doi.org/10.1016/j.neuron.2008.12.008

16 Sun, Q. *et al.* Proximodistal Heterogeneity of Hippocampal CA3 Pyramidal Neuron Intrinsic Properties, Connectivity, and Reactivation during Memory Recall. *Neuron* **95**, 656-672 e653 (2017). https://doi.org:10.1016/j.neuron.2017.07.012
17 Lee, H., Wang, C., Deshmukh, S. S. & Knierim, J. J. Neural Population Evidence of Functional Heterogeneity along the CA3 Transverse Axis: Pattern Completion versus Pattern Separation. *Neuron* **87**, 1093-1105 (2015). https://doi.org:10.1016/j.neuron.2015.07.012
18 Lu, L., Igarashi, K. M., Witter, M. P., Moser, E. I. & Moser, M. B. Topography of Place Maps along the CA3-to-CA2 Axis of the Hippocampus. *Neuron* **87**, 1078-1092 (2015). https://doi.org:10.1016/j.neuron.2015.07.007
19 Hunsaker, M. R., Rosenberg, J. S. & Kesner, R. P. The role of the dentate gyrus, CA3a,b, and CA3c for detecting spatial and environmental novelty. *Hippocampus* **18**, 1064-1073 (2008). https://doi.org:10.1002/hipo.20464
20 Nakamura, N. H., Flasbeck, V., Maingret, N., Kitsukawa, T. & Sauvage, M. M. Proximodistal segregation of nonspatial information in CA3: preferential recruitment of a proximal CA3-distal CA1 network in nonspatial recognition memory. *J Neurosci* **33**, 11506-11514 (2013). https://doi.org:10.1523/JNEUROSCI.4480-12.2013
21 Nakazawa, Y., Pevzner, A., Tanaka, K. Z. & Wiltgen, B. J. Memory retrieval along the proximodistal axis of CA1. *Hippocampus* **26**, 1140-1148 (2016). https://doi.org:10.1002/hipo.22596
22 Vancura, B., Geiller, T., Grosmark, A., Zhao, V. & Losonczy, A. Inhibitory control of sharp-wave ripple duration during learning in hippocampal recurrent networks. *Nat Neurosci* **26**, 788-797 (2023). https://doi.org:10.1038/s41593-023-01306-7
23 Geiller, T. *et al.* Local circuit amplification of spatial selectivity in the hippocampus. *Nature* **601**, 105-109 (2022). https://doi.org:10.1038/s41586-021-04169-9
24 Terada, S. *et al.* Adaptive stimulus selection for consolidation in the hippocampus. *Nature* **601**, 240-244 (2022). https://doi.org:10.1038/s41586-021-04118-6
25 Li, Y., Briguglio, J. J., Romani, S. & Magee, J. C. Mechanisms of memory storage and retrieval in hippocampal area CA3. *bioRxiv*, 2023.2005.2030.542781 (2023). https://doi.org:10.1101/2023.05.30.542781
26 Bowler, J. C. *et al.* (eLife Sciences Publications, Ltd, 2024).
27 Geiller, T., Fattahi, M., Choi, J. S. & Royer, S. Place cells are more strongly tied to landmarks in deep than in superficial CA1. *Nat Commun* **8**, 14531 (2017). https://doi.org:10.1038/ncomms14531
28 Lee, I., Rao, G. & Knierim, J. J. A double dissociation between hippocampal subfields: differential time course of CA3 and CA1 place cells for processing changed environments. *Neuron* **42**, 803-815 (2004). https://doi.org:10.1016/j.neuron.2004.05.010

29 Nakazawa, K. *et al.* Requirement for hippocampal CA3 NMDA receptors in associative memory recall. *Science* **297**, 211-218 (2002). https://doi.org:10.1126/science.1071795
30 Priestley, J. B., Bowler, J. C., Rolotti, S. V., Fusi, S. & Losonczy, A. Signatures of rapid plasticity in hippocampal CA1 representations during novel experiences. *Neuron* **110**, 1978-+ (2022). https://doi.org:10.1016/j.neuron.2022.03.026
31 Bowler, J. C. & Losonczy, A. Direct cortical inputs to hippocampal area CA1 transmit complementary signals for goal-directed navigation. *Neuron* **111**, 4071-4085.e4076 (2023). https://doi.org:https://doi.org/10.1016/j.neuron.2023.09.013
32 Ye, J., Witter, M. P., Moser, M.-B. & Moser, E. I. Entorhinal fast-spiking speed cells project to the hippocampus. *Proceedings of the National Academy of Sciences* **115**, E1627-E1636 (2018). https://doi.org:doi:10.1073/pnas.1720855115
33 Cayco-Gajic, N. A., Clopath, C. & Silver, R. A. Sparse synaptic connectivity is required for decorrelation and pattern separation in feedforward networks. *Nat Commun* **8**, 1116 (2017). https://doi.org:10.1038/s41467-017-01109-y
34 Ishizuka, N., Cowan, W. M. & Amaral, D. G. A quantitative analysis of the dendritic organization of pyramidal cells in the rat hippocampus. *J Comp Neurol* **362**, 17-45 (1995). https://doi.org:10.1002/cne.903620103
35 Turner, D. A., Li, X. G., Pyapali, G. K., Ylinen, A. & Buzsaki, G. Morphometric and electrical properties of reconstructed hippocampal CA3 neurons recorded in vivo. *J Comp Neurol* **356**, 580-594 (1995). https://doi.org:10.1002/cne.903560408
36 Hainmueller, T. & Bartos, M. Parallel emergence of stable and dynamic memory engrams in the hippocampus. *Nature* **558**, 292-296 (2018). https://doi.org:10.1038/s41586-018-0191-2
37 Dong, C., Madar, A. D. & Sheffield, M. E. J. Distinct place cell dynamics in CA1 and CA3 encode experience in new environments. *Nat Commun* **12**, 2977 (2021). https://doi.org:10.1038/s41467-021-23260-3
38 Madar, A. D., Dong, C. & Sheffield, M. E. J. BTSP, not STDP, Drives Shifts in Hippocampal Representations During Familiarization. *bioRxiv*, 2023.2010.2017.562791 (2023). https://doi.org:10.1101/2023.10.17.562791
39 Sheintuch, L., Geva, N., Deitch, D., Rubin, A. & Ziv, Y. Organization of hippocampal CA3 into correlated cell assemblies supports a stable spatial code. *Cell Rep* **42**, 112119 (2023). https://doi.org:10.1016/j.celrep.2023.112119
40 Soltesz, I. & Losonczy, A. CA1 pyramidal cell diversity enabling parallel information processing in the hippocampus. *Nat Neurosci* **21**, 484-493 (2018). https://doi.org:10.1038/s41593-018-0118-0
41 Geiller, T., Royer, S. & Choi, J. S. Segregated Cell Populations Enable Distinct Parallel Encoding within the Radial Axis of the CA1 Pyramidal Layer. *Exp Neurobiol* **26**, 1-10 (2017). https://doi.org:10.5607/en.2017.26.1.1
42 Makara, J. K. & Magee, J. C. Variable dendritic integration in hippocampal CA3 pyramidal neurons. *Neuron* **80**, 1438-1450 (2013). https://doi.org:10.1016/j.neuron.2013.10.033

43 O'Hare, J. K. *et al.* Compartment-specific tuning of dendritic feature selectivity by intracellular Ca(2+) release. *Science* **375**, eabm1670 (2022). https://doi.org:10.1126/science.abm1670
44 Kim, S., Guzman, S. J., Hu, H. & Jonas, P. Active dendrites support efficient initiation of dendritic spikes in hippocampal CA3 pyramidal neurons. *Nat Neurosci* **15**, 600-606 (2012). https://doi.org:10.1038/nn.3060
45 Leutgeb, S., Leutgeb, J. K., Treves, A., Moser, M. B. & Moser, E. I. Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science* **305**, 1295-1298 (2004). https://doi.org:10.1126/science.1100265
46 Li, Y., Briguglio, J. J., Romani, S. & Magee, J. C. Mechanisms of memory-supporting neuronal dynamics in hippocampal area CA3. *Cell* https://doi.org:10.1016/j.cell.2024.09.041
47 Hasselmo, M. E. The role of acetylcholine in learning and memory. *Curr Opin Neurobiol* **16**, 710-715 (2006). https://doi.org:10.1016/j.conb.2006.09.002
48 Brandon, M. P., Koenig, J., Leutgeb, J. K. & Leutgeb, S. New and distinct hippocampal place codes are generated in a new environment during septal inactivation. *Neuron* **82**, 789-796 (2014). https://doi.org:10.1016/j.neuron.2014.04.013
49 Kentros, C. G., Agnihotri, N. T., Streater, S., Hawkins, R. D. & Kandel, E. R. Increased attention to spatial context increases both place field stability and spatial memory. *Neuron* **42**, 283-295 (2004). https://doi.org:10.1016/s0896-6273(04)00192-8
50 McNamara, C. G., Tejero-Cantero, A., Trouche, S., Campo-Urriza, N. & Dupret, D. Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nat Neurosci* **17**, 1658-1660 (2014). https://doi.org:10.1038/nn.3843
51 Kaufman, A. M., Geiller, T. & Losonczy, A. A Role for the Locus Coeruleus in Hippocampal CA1 Place Cell Reorganization during Spatial Reward Learning. *Neuron* **105**, 1018-1026 e1014 (2020). https://doi.org:10.1016/j.neuron.2019.12.029
52 Kempadoo, K. A., Mosharov, E. V., Choi, S. J., Sulzer, D. & Kandel, E. R. Dopamine release from the locus coeruleus to the dorsal hippocampus promotes spatial learning and memory. *Proceedings of the National Academy of Sciences* **113**, 14835-14840 (2016). https://doi.org:doi:10.1073/pnas.1616515114
53 Takeuchi, T. *et al.* Locus coeruleus and dopaminergic consolidation of everyday memory. *Nature* **537**, 357-362 (2016). https://doi.org:10.1038/nature19325
54 Wagatsuma, A. *et al.* Locus coeruleus input to hippocampal CA3 drives single-trial learning of a novel context. *Proceedings of the National Academy of Sciences* **115**, E310-E316 (2018). https://doi.org:doi:10.1073/pnas.1714082115
55 Tye, K. M. *et al.* Mixed selectivity: Cellular computations for complexity. *Neuron* **112**, 2289-2303 (2024). https://doi.org:10.1016/j.neuron.2024.04.017
56 Barak, O., Rigotti, M. & Fusi, S. The Sparseness of Mixed Selectivity Neurons Controls the Generalization–Discrimination Trade-Off. *The Journal of Neuroscience* **33**, 3844-3856 (2013). https://doi.org:10.1523/jneurosci.2753-12.2013

57 Fusi, S., Miller, E. K. & Rigotti, M. Why neurons mix: high dimensionality for higher cognition. *Current Opinion in Neurobiology* **37**, 66-74 (2016). https://doi.org:https://doi.org/10.1016/j.conb.2016.01.010
58 Badre, D., Bhandari, A., Keglovits, H. & Kikumoto, A. The dimensionality of neural representations for control. *Current Opinion in Behavioral Sciences* **38**, 20-28 (2021). https://doi.org:https://doi.org/10.1016/j.cobeha.2020.07.002
59 Cohen, U., Chung, S., Lee, D. D. & Sompolinsky, H. Separability and geometry of object manifolds in deep neural networks. *Nature Communications* **11**, 746 (2020). https://doi.org:10.1038/s41467-020-14578-5
60 Chung, S. & Abbott, L. F. Neural population geometry: An approach for understanding biological and artificial neural networks. *Current Opinion in Neurobiology* **70**, 137-144 (2021). https://doi.org:https://doi.org/10.1016/j.conb.2021.10.010
61 Esparza, J., Sebastián, E. R. & de la Prida, L. M. From cell types to population dynamics: Making hippocampal manifolds physiologically interpretable. *Current Opinion in Neurobiology* **83**, 102800 (2023). https://doi.org:https://doi.org/10.1016/j.conb.2023.102800
62 Sebastian, E. R., Esparza, J. & M. de la Prida, L. Quantifying the distribution of feature values over data represented in arbitrary dimensional spaces. *PLOS Computational Biology* **20**, e1011768 (2024). https://doi.org:10.1371/journal.pcbi.1011768
63 Okazawa, G., Hatch, C. E., Mancoo, A., Machens, C. K. & Kiani, R. Representational geometry of perceptual decisions in the monkey parietal cortex. *Cell* **184**, 3748-3761.e3718 (2021). https://doi.org:10.1016/j.cell.2021.05.022
64 Russo, A. A. *et al.* Neural Trajectories in the Supplementary Motor Area and Motor Cortex Exhibit Distinct Geometries, Compatible with Different Classes of Computation. *Neuron* **107**, 745-758.e746 (2020). https://doi.org:10.1016/j.neuron.2020.05.020
65 Nakai, S., Kitanishi, T. & Mizuseki, K. Distinct manifold encoding of navigational information in the subiculum and hippocampus. *Science Advances* **10**, eadi4471 (2024). https://doi.org:doi:10.1126/sciadv.adi4471
66 Ostojic, S. & Fusi, S. Computational role of structure in neural activity and connectivity. *Trends in Cognitive Sciences* **28**, 677-690 (2024). https://doi.org:10.1016/j.tics.2024.03.003
67 Dubreuil, A., Valente, A., Beiran, M., Mastrogiuseppe, F. & Ostojic, S. The role of population structure in computations through neural dynamics. *Nature Neuroscience* **25**, 783-794 (2022). https://doi.org:10.1038/s41593-022-01088-4
68 Litwin-Kumar, A., Harris, K. D., Axel, R., Sompolinsky, H. & Abbott, L. F. Optimal Degrees of Synaptic Connectivity. *Neuron* **93**, 1153-1164.e1157 (2017). https://doi.org:https://doi.org/10.1016/j.neuron.2017.01.030
69 Litwin-Kumar, A. & Doiron, B. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience* **15**, 1498-1505 (2012). https://doi.org:10.1038/nn.3220

70 Herbert, E. & Ostojic, S. The impact of sparsity in low-rank recurrent neural networks. *PLOS Computational Biology* **18**, e1010426 (2022). https://doi.org:10.1371/journal.pcbi.1010426
71 Robert, V. *et al.* Entorhinal cortex glutamatergic and GABAergic projections bidirectionally control discrimination and generalization of hippocampal representations. *bioRxiv*, 2023.2011.2008.566107 (2023). https://doi.org:10.1101/2023.11.08.566107
72 Guzman, S. J. *et al.* How connectivity rules and synaptic properties shape the efficacy of pattern separation in the entorhinal cortex–dentate gyrus–CA3 network. *Nature Computational Science* **1**, 830-842 (2021). https://doi.org:10.1038/s43588-021-00157-1
73 Dudok, B. *et al.* Recruitment and inhibitory action of hippocampal axo-axonic cells during behavior. *Neuron* **109**, 3838-3850 e3838 (2021). https://doi.org:10.1016/j.neuron.2021.09.033
74 Geiller, T. *et al.* Large-Scale 3D Two-Photon Imaging of Molecularly Identified CA1 Interneuron Dynamics in Behaving Mice. *Neuron* **108**, 968-983 e969 (2020). https://doi.org:10.1016/j.neuron.2020.09.013
75 Hainmueller, T., Cazala, A., Huang, L. W. & Bartos, M. Subfield-specific interneuron circuits govern the hippocampal response to novelty in male mice. *Nat Commun* **15**, 714 (2024). https://doi.org:10.1038/s41467-024-44882-3
76 Rolotti, S. V. *et al.* Local feedback inhibition tightly controls rapid formation of hippocampal place fields. *Neuron* **110**, 783-794 e786 (2022). https://doi.org:10.1016/j.neuron.2021.12.003
77 Danielson, N. B. *et al.* Sublayer-Specific Coding Dynamics during Spatial Navigation and Learning in Hippocampal Area CA1. *Neuron* **91**, 652-665 (2016). https://doi.org:10.1016/j.neuron.2016.06.020
78 Kaifosh, P., Zaremba, J. D., Danielson, N. B. & Losonczy, A. SIMA: Python software for analysis of dynamic fluorescence imaging data. *Front Neuroinform* **8**, 80 (2014). https://doi.org:10.3389/fninf.2014.00080
79 Keemink, S. W. *et al.* FISSA: A neuropil decontamination toolbox for calcium imaging signals. *Sci Rep* **8**, 3493 (2018). https://doi.org:10.1038/s41598-018-21640-2
80 Skaggs, W. E. & McNaughton, B. L. Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. *J Neurosci* **18**, 8455-8466 (1998). https://doi.org:10.1523/JNEUROSCI.18-20-08455.1998
81 Kong, E., Lee, K. H., Do, J., Kim, P. & Lee, D. Dynamic and stable hippocampal representations of social identity and reward expectation support associative social memory in male mice. *Nat Commun* **14**, 2597 (2023). https://doi.org:10.1038/s41467-023-38338-3
82 Macmillan, N. A. & Creelman, C. D. *Detection theory: A user's guide, 2nd ed*. (Lawrence Erlbaum Associates Publishers, 2005).

83 Banino, A. *et al.* Vector-based navigation using grid-like representations in artificial agents. *Nature* **557**, 429-433 (2018). https://doi.org:10.1038/s41586-018-0102-6

84 Christopher J. Cueva, X.-X. W. Emergence of grid-like representations by training recurrent neural networks to perform spatial localization. *arXiv* (2018). https://doi.org:10.48550/arXiv.1803.07770

85 Hochreiter, S. & Schmidhuber, J. Long short-term memory. *Neural Comput* **9**, 1735-1780 (1997). https://doi.org:10.1162/neco.1997.9.8.1735

86 Rungratsameetaweemana, N., Kim, R., Chotibut, T. & Sejnowski, T. J. Random noise promotes slow heterogeneous synaptic dynamics important for robust working memory computation. *bioRxiv*, 2022.2010.2014.512301 (2023). https://doi.org:10.1101/2022.10.14.512301

87 Hopfield, J. J. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences of the United States of America* **81**, 3088-3092 (1984). https://doi.org:10.1073/pnas.81.10.3088