

A neuronal code for space in hippocampal coactivity dynamics independent of place fields

Elliott R.J. Levy¹, Simon Carrillo-Segura^{1,2}, Eun Hye Park¹, William T. Redman³, José Hurtado¹, André A. Fenton^{1,4*}

¹Center for Neural Science, New York University, New York, NY 10003, USA

²Graduate Program in Mechanical and Aerospace Engineering, Tandon School of Engineering, New York University, Brooklyn, NY 11201

³Interdepartmental Graduate Program in Dynamical Neuroscience, University of California, Santa Barbara, Santa Barbara, CA 93106, USA.

⁴Neuroscience Institute at the NYU Langone Medical Center, New York, NY 10016, USA

**Correspondence:*

André Fenton, Neurobiology of Cognition Laboratory, Center for Neural Science, New York University,

4 Washington Place, Room 980. Email: afenton@nyu.edu

Abstract

Hippocampus is comprised of ~20% place cells, discharging in cell-specific locations (“place fields”), standardly interpreted as a dedicated neuronal code for space. However, place cell discharge is temporally unreliable across seconds and days, and spatially multimodal, suggesting an alternative “ensemble cofiring” spatial code with manifold dynamics that does not require reliable spatial tuning. We evaluated these hypotheses using GCaMP6f and miniature microscopes to image mouse CA1 ensemble activity in two environments, across 3 weeks. Both place fields and ensemble coactivity relationships “remap,” being distinct between and (weakly) similar within environments. We find that decoding the environment from cell-pair coactivity relationships is effective and improves with experience, even after removing place tuning. Decoding from 1-s ensemble activity is indifferent to firing fields, relying crucially on the most anti-coactive cells because ensemble activity is internally-organized on a low-dimensional manifold of non-linear cofiring relationships that intermittently reregisters to environments according to anti-cofiring subpopulation activity.

Keywords:

Hippocampus, place cell, cell assembly, neural synchrony, synapseensemble, environment, context, remapping, reregistering

Introduction

The hippocampus is crucial for spatial navigation as well as memory, and although hippocampal place cells are tuned to discharge in “place fields,” cell-specific locations within an environment, their relationship to spatial cognition is unclear (Muller et al., 1987; O'Keefe, 1976, 1979). When animals change environments, the location-specific tuning of each place cell changes uniquely, a phenomenon called “remapping” (Fig. 1A). Remapping means that the ensemble pattern of cofiring activity is unique and environment-specific, constituting a spatial code that relates neural activity to the environment (Alme et al., 2014; Bostock et al., 1991; Colgin et al., 2008; Kubie et al., 2020; Muller and Kubie, 1987; Wills et al., 2005).

Although the concept of remapping means a rearrangement of temporal cofiring relationships, the vast majority of studies have inferred that a rearrangement of cofiring relationships has occurred by only measuring that the session-averaged place fields of individual cells has changed. It is, however, possible that the place fields of individual cells change without substantially changing their temporal cofiring relationships. It is especially easy (but not necessary) to imagine how this can happen if one accepts that most place cells have multiple non-periodic place fields, as is observed in environments with linear dimensions greater than 1 m (Fenton et al., 2008; Harland et al., 2021). As Fig 1A illustrates, remapping implies that place fields are determined by environmental features (O'Keefe and Burgess, 1996), whereas it is possible to observe changed place fields, the traditional evidence of remapping, when only the registration changes between the environment and a network of invariant cofiring relationships that are intrinsically-organized (Fig. 1B). Whether remapping is actually reregistration has important implications not only for understanding the hippocampal space code and navigation, but also for understanding memory.

Because both memory and environmental context representations rely on the hippocampus, the dominant hypothesis asserts that remapping and memory are intimately linked; changes in place tuning are assumed to correspond to changes in memory representations (Guzowski et al., 2004; Leutgeb et al., 2005b; Leutgeb et al., 2005c; Lever et al., 2002; Wills et al., 2005), in particular, episodic memory representations that include information about environments (Kentros, 2006; Mizumori, 2006). We note however, that it is not straightforward how the arrangement of place fields can represent a particular environment at the millisecond-to-second timescale of neural computations without initial, extensive spatial exploration (McHugh and Tonegawa, 2007) and for a number of fundamental reasons: 1) place cells are multimodal, having multiple place fields in environments bigger than 1m² (Fenton et al., 2008; Harland et al., 2021); 2) discharge in firing fields is extremely variable during the 1-to-5 seconds it takes to cross the place field (Fenton et al., 2010; Fenton and Muller, 1998; Jackson and Redish, 2007); 3) the firing rate in a place field varies systematically across behavioral episodes, called “rate remapping” (Leutgeb et al., 2005b); and 4) only a minority of place cells have stable place fields across days in familiar environments (Lee et al., 2020; Ziv et al., 2013). The consequence of cells with multiple place fields demonstrates that this alone can degrade the ability to decode environments from ensemble discharge (Fig. S1). We also note that studies that tested the hypothesis that remapping is related to memory have been remarkably unsupportive (Duvell et al., 2019; Jeffery et al., 2003; Leutgeb et al., 2005a; Moita et al., 2004; van Dijk and Fenton, 2018).

One possible reason for the inability to relate remapping to memory could be incorrect assumptions about how neural information is encoded; indeed, what is the nature of the neural code? Perhaps the code for environments is different from the code for locations (Tanaka et al.,

2018). Studies standardly average the activity of individual cells over minutes, discarding discharge fluctuations in time as well as any cofiring relationships with other cells to extract the cell's relationship to a certain variable, such as its location-specific tuning. Such analyses are performed under the (incorrect) assumption that place-discharge relationships are in a steady state, although they are not (Fenton et al., 2010; Fenton and Muller, 1998; Jackson and Redish, 2007). This standard approach, a dedicated-rate place field hypothesis assumes the cell's momentary firing rate independently carries information that can be adequately extracted by analysis of each cell's place tuning. The firing relations to other neurons are assumed to be uninformative, or at least secondary, which is why they are explicitly ignored by data representations like the firing rate maps that define place fields. Indeed, it is intuitive to imagine that place cells with overlapping fields might be expected to cofire because time and space are confounded during navigation. However, whereas some place cell pairs with overlapping place fields reliably cofire, other pairs do not, and yet other pairs discharge independently on the ms-to-s timescales of crossing firing fields and neural computation (Fig. S3; Fenton, 2015a; Harris et al., 2003; Kelemen and Fenton, 2013). This motivates an alternative cofiring (ensemble or cell assembly) coding hypothesis that asserts information is encoded in the momentary cofiring patterns of activity that are expressed by large groups of neurons -- each cofiring pattern encoding different information (Harris et al., 2003; Meshulam et al., 2017; Stefanini et al., 2020). Such cofiring neural codes are explicit in the increasingly popular recognition that high-dimensional neural population activity (each dimension being a neuron) can be constrained to a manifold, a low-dimensional sub-space defined by recurring across-cell patterns of coactivity (Chaudhuri et al., 2019; Ebitz and Hayden, 2021; Gardner et al., 2022; Jazayeri and Afraz, 2017; McNaughton et al., 2006; Park et al., 2019; Peyrache et al., 2015; Rubin et al., 2019; Umakantha et al., 2021). Note that higher-order firing patterns constrained on a manifold carry information, and that in principle, without changing shape, such manifolds can describe an

infinity of individual cell discharge patterns and even an infinity of lower-order cofiring patterns (Heeger and Mackey, 2019; Jazayeri and Afraz, 2017), as recently demonstrated for MEC grid cells (Gardner et al., 2022). Such discharge patterns are constrained by “on-manifold” cofiring relationships and because they need not ever repeat, they need not depend on tuning to external variables like locations. Here we repeated a standard place cell remapping procedure to evaluate the dedicated place field and ensemble cofiring hypotheses for representing distinct environments in the hippocampus.

Results

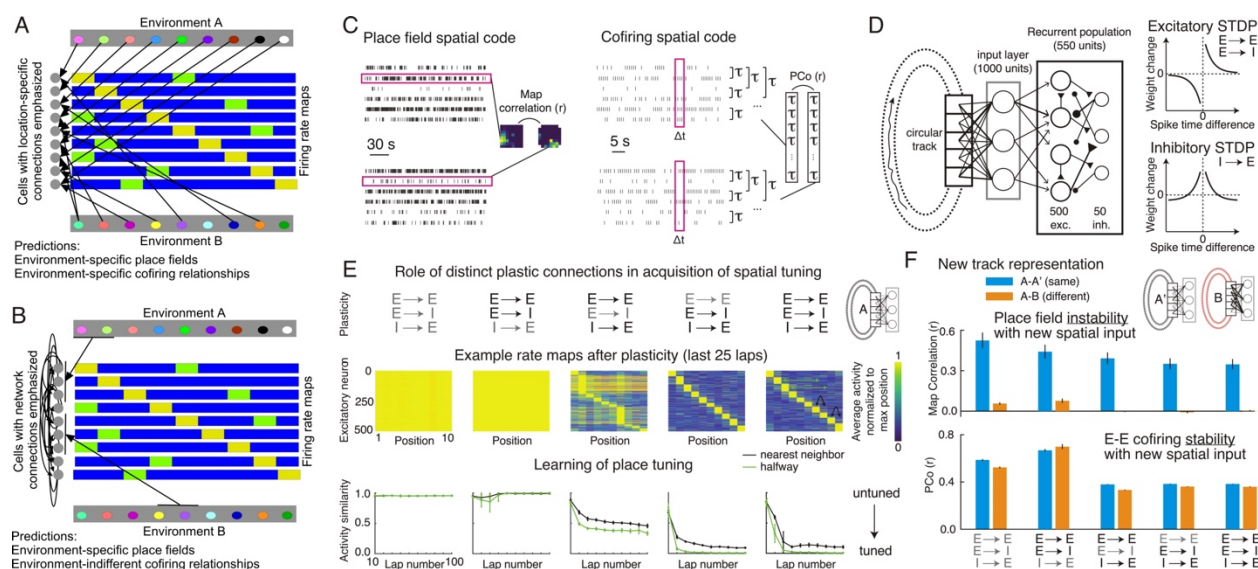


Figure 1. Neural coding hypotheses for representing environments. (A) Schematic illustrating the standard concept of remapping that emphasizes that cell-specific location-specific inputs change both place fields and ensemble cofiring relationship across environments, and (B) the concept of reregistration that emphasizes intrinsic hippocampal inputs that do not change between environments despite a change in the registration of environment information onto the invariant network cofiring relationships. The firing rate maps illustrate environment-specific multi-peaked place fields (yellow and green). (C) Illustration of hypothesized (left) dedicated place field and (right) ensemble cofiring codes for environments, and how they can be quantified, by single cell spatial firing stabilities, or by computing the stability (PCo) of the vector of cell pair correlations. PCo computes Pearson correlation (r) to measure the similarity between the set of $n(n-1)/2$ Kendall correlations (τ) that describe the pairwise activity correlations within an n -cell ensemble (Neymotin et al., 2017). (D) Schematic of neural network encoding of a circular environment with STDP learning rules. (E, top) Learned network representation of the circular environment after (black) enabling or (grey) disabling STDP learning at cell-type specific connections, as measured on (middle) time-averaged rate

maps and by (bottom) activity similarity. The nearest neighbor metric measures the difference between the peak activity and the activity at the closest positions to the location of the peak; the halfway metric measures the difference between the peak activity and the activity at the position halfway across the track. (F) Consequence of remapping the external inputs to the network (track B) compared to an unchanged mapping (null remapping, track A'). Evaluation of (top) dedicated place field and (bottom) ensemble cofiring codes in the track A and B environments as a function of (black) enabled or (grey) disabled STDP learning at cell-type specific connections. Discrimination between environments is robust with the place field code (top) and is most sensitive to inhibitory plasticity. This is especially prominent in the ensemble cofiring code (bottom) that in contrast to place firing, is relatively similar across experiences of the same and different environments.

Distinguishing place tuning and cofiring contributions to representing environments: a neural network model

We formalized the place field and cofiring hypotheses by examining a heuristic neural network model designed to evaluate the contributions of place tuning (Fig. 1C, left) and cofiring (Fig. 1C, right) to representing environments. The recurrently connected excitatory-inhibitory network receives excitatory input from units tuned to randomly sampled positions of a circular track (Fig. 1D, left). The recurrent network weights change according to standard spike-timing dependent plasticity (STDP) rules (Fig. 1D, right) that are sufficient to generate place-selective activity (place fields, Fig. 1E, middle) in the recurrent network after ~30 laps on the track (Fig. 1E, bottom). Among other features, simply enlarging the track generates multiple firing fields once there has been sufficient time for exploration and the learning rules to operate, as is observed in real environments (Fenton et al., 2008; Harland et al., 2021). Location-specificity could be assessed by the similarity of the place fields (Fig. 1E), as well as the ability to decode current

position from momentary network activity (data not shown). Remarkably, $E \rightarrow I$ and $I \rightarrow E$ plasticity are necessary for developing the place tuning, but $E \rightarrow E$ plasticity is not (Fig. 1E), reproducing observations after blocking NMDA receptor-dependent LTP (Kao et al., 2017; Kentros et al., 1998; Lamsa et al., 2005), which dominates at $E \rightarrow E$ connections but not at $E \rightarrow I$ or $I \rightarrow E$ connections in hippocampus (Lamsa et al., 2005). This illustrates that tunable excitation-inhibition coordination can be important for network and memory function, as observed (Caroni, 2015; Chung et al., 2021; Dvorak et al., 2021; Fenton, 2015b; Lamsa et al., 2010; Mongillo et al., 2018; van Dijk and Fenton, 2018); see papers dedicated to the topic (Kullmann and Lamsa, 2011). To simulate remapping after learning a second environment, we remapped the position-to-input relationships and allowed the recurrent weights to change according to the same STDP rules during another 30 laps. As expected, the place fields of the network units changed; they remapped the new track (Fig. 1F, top). The population of pairwise cofiring relationships behaves differently, demonstrating it is possible to change firing fields (the conventional measure of remapping) without substantially changing cofiring relationships (Fig. S2). Co-firing relationships are marginally more stable when spatial inputs are unchanged compared to when they are changed (i.e. remap). Co-firing relationships tend to change with experience, regardless of whether the environmental inputs change, and these changes are most strongly dependent on $I \rightarrow E$ plasticity and they are required for location-specific activity, in this model. This network simulation is a heuristic demonstration that several key ideas are feasible: 1) that a place field code and a cofiring code could be distinct and operate in parallel; 2) that place field remapping between environments may occur without changing the majority of $E \rightarrow E$ cofiring relationships in a network; 3) that remapping maybe a misnomer for what is a reregistration between an internally-rigid neural activity representation and external inputs; 4) that place field remapping/reregistration depends on the environment-specific anti-cofiring $E \rightarrow E$ relationships; a

consequence of necessarily altered I→E (and E→I) functional connections, which have been observed as experience- (Mongillo et al., 2018) and learning-dependent changes in network inhibition (Chung et al., 2021). If correct, these ideas will require changes to standard ideas, analyses, and interpretations of the spatial tuning of hippocampal neurons, and perhaps cells in other brain regions too (Ebitz and Hayden, 2021).

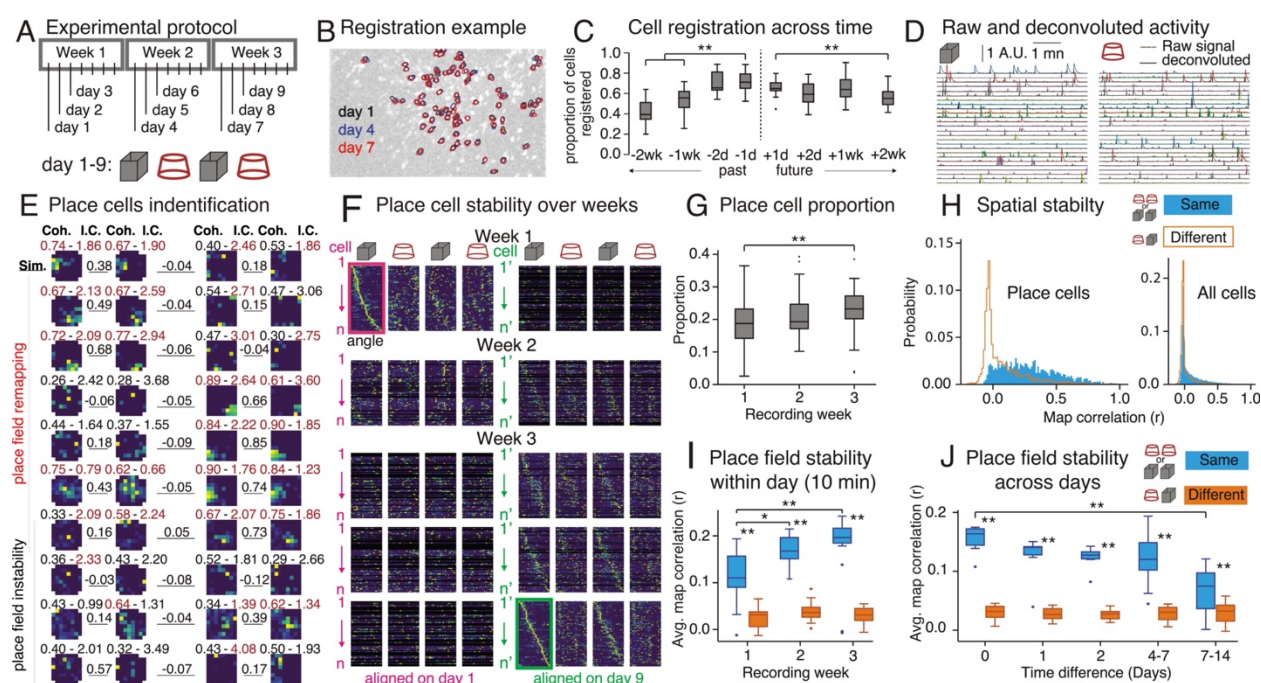


Figure 2. Location-specific CA1 activity imaged across weeks is transient, describes a minority of the CA1 population, and remaps between environments. (A) Experimental protocol. (B) Example and (C) quantification of cell registration across weeks (Past: Time $F_{3,187} = 56.89$, $p = 10^{-26}$, $\eta^2 = 0.48$; Dunnett's tests: -1 day > -2, -1 week(s); Future: Time $F_{3,187} = 7.31$, $p = 10^{-4}$, $\eta^2 = 0.12$; Dunnett's tests: 1 day > 2 weeks). (D) Activity time series in two environments illustrated by a single 25-cell example from a 546-cell ensemble. (E) Activity rate map examples from 10 cells from a 588-ensemble across two environments on day 7. Place field quality metrics, spatial coherence (Coh.; left) and information content (I.C.; right) are indicated above

each map, red indicating significance; average rate map similarity (Sim.; r) is indicated between maps. (F) Linearized ensemble activity rate maps across three weeks (week 1: day 1; week 2: day 4; week 3: days 7,8,9), matched and aligned to activity on day 1 (left, 312 cells) and day 9 (right, 460 cells). (G) Proportions of cells with significant spatial coherence and information content, classified as place cells (Week: $F_{2,64} = 43.70$, $p = 10^{-12}$, $\eta^2 = 0.085$; Dunnett's tests: week 1 < week 3). (H, left) Place map similarity across environments on weeks 2 and 3 illustrates remapping of place cells but (right) not the CA1 population. (I) Place map similarity within day improves across weeks (Week: $F_{2,28} = 20.02$, $p = 10^{-6}$, $\eta^2 = 0.051$; Env. Change: $F_{2,29} = 80.55$, $p = 10^{-10}$, $\eta^2 = 0.65$; Interaction: $F_{2,28} = 4.83$, $p = 0.016$, $\eta^2 = 0.026$; Dunnett's tests: week 1 < weeks 2, 3 for similar environments; post-hoc: Same vs. Different is significant each week, $**p's \leq 10^{-5}$). (J) Place map similarity degrades across days (Time difference: $F_{4,132} = 12.68$, $p = 10^{-9}$, $\eta^2 = 0.11$; Env. Change: $F_{1,132} = 228.37$, $p = 10^{-30}$, $\eta^2 = 0.53$; Interaction: $F_{4,132} = 9.32$, $p = 10^{-6}$, $\eta^2 = 0.12$; Dunnett's tests: day 0 > 7-14 days for similar environments; post-hoc: Same vs. Different is significant at each of the five-time intervals, $**p's \leq 10^{-3}$).

Calcium imaging in CA1

To examine the validity of dedicated place field and ensemble cofiring hypotheses, we infected mouse CA1 with a virus expressing GCaMP6f (Cai et al., 2016; Chen et al., 2013). While the mice were freely exploring, we recorded CA1 ensemble activity with a miniature microscope (<http://miniscope.org>; Cai et al., 2016) placed over a chronically implanted gradient-index (GRIN) lens (Ziv et al., 2013). During the recordings, the animals ($n=6$) explored two physically distinct environments, a cylinder and box. The mice were habituated to the miniscope in their home cage and only experienced the two environments during recordings according to the protocol in Fig. 2A. Raw calcium traces were extracted and segmented from videos using the

CNMF-E algorithm (Pnevmatikakis et al., 2016; Zhou et al., 2018), and, when possible, registered to individual cells across the three-week protocol (Fig. 2B). Over 50% of individual cells could be identified two weeks apart, though not unsurprisingly the registration was greater across shorter intervals (Fig. 2C). Raw calcium traces were converted with the CNMF-E algorithm to spiking activity for analysis (Fig. 2D).

Place field measurements of remapping

In each environment, cells expressed place fields (Fig. 2E) that were stable within a day and unstable across days as previously reported (Lee et al., 2020; Ziv et al., 2013). To better visualize the organization of ensemble of place response, the environment was linearized by plotting activity rates as a function of angle (Fig. 2F). The location-specific pattern expressed in week 3 (aligned on recording day 9) was observed during week 2, indicating the location-specific activity was more stable than it was in week 1 (aligned on recording day 1). These observations indicate that place tuning was learned during the first week. While consistently constituting a minority of the population (Stefanini et al., 2020; Wilson and McNaughton, 1993), the percentage of place cells increased with experience from 20% to 25% (Fig. 2G). Each day, place cells expressed stable firing fields between trials in the same environment, but location-specificity changed between trials of different environments (Fig. 2H, left). Such remapping is not population wide as the measurable difference vanishes when all recorded cells are considered. This is largely because the place tuning of non-place cells is unstable between any pair of trials (Fig. 2H, right). The average within-day place field stability of place cells grows with experience from 0.1 during week 1 to 0.2 during week 3 (Fig. 2I), about 6 times larger than for distinct environments on week 3. Yet, this place field stability degrades across days (Fig. 2J), as previously described (Lee et al., 2020; Ziv et al., 2013). These measurements demonstrate standard mouse place cell phenomena in the data set.

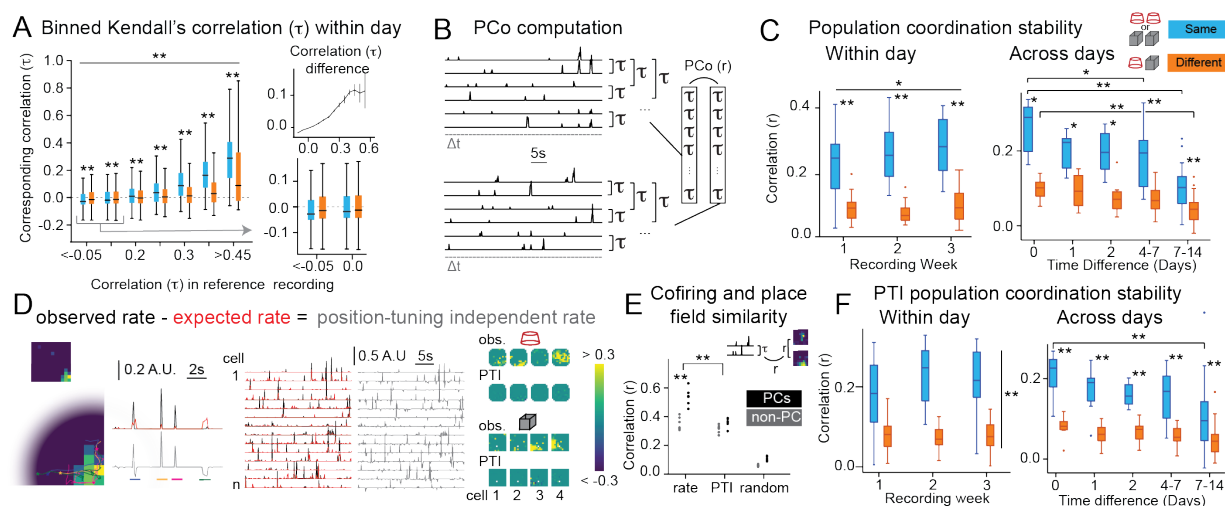


Figure 3. Ensemble cofiring distinguishes two environments independently of place

tuning. (A) Comparison of coactivity (τ) stability of all cell pairs, from weeks 2 and 3, across same-day recordings, for different initial coactivity values. This illustrates that coactive cell pairs tend to remain coactive across the same, but not different environments (Coactivity Level:

$F_{6,3285501} = 9825.46$, $p = 0.00$, $\eta^2 = 0.018$; Env. Change: $F_{1,3285501} = 9.82$, $p = 0.0017$, $\eta^2 = 10^{-6}$;

Interaction: $F_{6,3285501} = 3395.58$, $p = 0.00$, $\eta^2 = 0.0061$; post-hoc: Same vs. Different is significant at each coactivity level, $**p's < 10^{-15}$) and (bottom inset) anti-coactive cell pairs remain anti-

coactive. (top inset) Average signed differences confirm the effect of environment change. (B)

Illustration of Population Coordination (PCo) computation from cofiring distributions. The time intervals ($\Delta\tau = 1$ s) used to discretize activity are illustrated as grey dashes, to scale. (C, left)

Within-day coactivity stability improves across weeks (Week: $F_{2,27} = 4.23$, $p = 0.025$, $\eta^2 = 0.015$;

Env. Change: $F_{1,28} = 92.32$, $p = 10^{-10}$, $\eta^2 = 0.56$; Interaction: $F_{2,27} = 0.99$, $p = 0.38$, $\eta^2 = 0.022$; post-

hoc: Same vs. Different is significant each week, $**p's < 10^{-4}$). (C, right) Across-day cofiring

stability degrades over time (Time Difference: $F_{4,132} = 16.79$, $p = 10^{-11}$, $\eta^2 = 0.21$; Env. Change:

$F_{1,132} = 94.66$, $p = 10^{-17}$, $\eta^2 = 0.32$; Interaction: $F_{4,132} = 4.04$, $p = 0.004$, $\eta^2 = 0.072$; Dunnett's tests

vs day 0, $*p < 0.0058$, $**p \leq 10^{-5}$; post-hoc: Same vs. Different is significant at each time interval, $**p \leq 10^{-3}$). (D) Illustration that subtracting expected position-tuning from observed activity captures activity fluctuations independent of place tuning. (Left to right) Example activity map in the square environment and magnified version with four passes across the place field (star = path start). Observed, expected and PTI activity times series with the color-coded passes marked and additional examples from 14 cells. Observed and PTI activity maps from 4 Fig. 1 example cells. (E) Correlation between place field similarity and coactivity, computed from observed, PTI and randomized PTI activities, for place cells and non-place cells (Time-Series Type: $F_{2,30} = 315.13$, $p = 10^{-21}$, $\eta^2 = 0.84$; Functional Cell Class: $F_{1,30} = 56.74$, $p = 10^{-8}$, $\eta^2 = 0.076$; Interaction: $F_{2,30} = 14.72$, $p = 10^{-5}$, $\eta^2 = 0.039$; post-hoc: PC vs. non-PC for Rate and Rate vs. PTI for Place Cells significant, $**p \leq 10^{-4}$). (F, left) Within-day PTI coactivity stability across weeks (Week: $F_{2,27} = 2.28$, $p = 0.12$, $\eta^2 = 0.009$; Env. Change: $F_{1,28} = 46.67$, $p = 10^{-7}$, $\eta^2 = 0.52$; Interaction: $F_{2,27} = 2.38$, $p = 0.11$, $\eta^2 = 0.028$; post-hoc: Same vs. Different: $t_{67.44} = 10.14$, $p = 10^{-15}$). (F, right) Across-day PTI coactivity stability also degrades over time (Time Difference: $F_{1,132} = 5.23$, $p = 10^{-4}$, $\eta^2 = 0.088$; Env. Change $F_{1,132} = 66.57$, $p = 10^{-13}$, $\eta^2 = 0.33$; Interaction: $F_{4,132} = 1.69$, $p = 0.16$, $\eta^2 = 0.079$; Dunnett's tests: day 0 > days 7-14 for similar environments, $**p = 0.0050$; post-hoc Same vs. Different significant at each time interval, $**p$'s < 0.008).

Ensemble cofiring measurements of remapping between environments

Next, we examined the change in CA1 activity across the two environments without regard to positional information. To evaluate whether distinct environments can be encoded by conjoint discharge, we measured the correlations of cell pairs, as a simple approximation for higher-order correlations (Schneidman et al., 2006). In each trial, for each cell pair we measured the coactivity of their 1-s activity time series by computing their Kendall correlations (Fig. 3A). To

evaluate their stability across environments, we measured the population coordination (PCo) as the similarity of all tau values, computed as the Pearson correlation between two vectors of pairwise tau values (Fig. 3B; Neymotin et al., 2017; Talbot et al., 2018). PCo between trials of the same environments is three times higher than between distinct environments (Fig. 3C), indicating that ensemble coactivity discriminates environments. The individual cell pairs that differ the most between environments are the positively ($\tau > 0.3$) or negatively ($\tau < -0.05$) coactive cell pairs (Fig. 3A). PCo increases with experience (Fig. 3C, left) and deteriorates across days (Fig. 3C, right), like place field measurements of remapping (Fig. 2I, J). These findings indicate that environments are discriminatively encoded in the 1-s coactivity relationships of CA1 neural activity.

Ensemble cofiring measurements of remapping between environments independently of position tuning

Because CA1 activity is influenced by the animal's position, we evaluated the extent to which position tuning and overlapping place fields explains the 1-s coactivity relationships of CA1 activity. We calculated a position-tuning independent rate (PTI-rate) by subtracting the expected rate from the cell's observed rate, as schematized in Fig. 3D. The expected rate was computed from the position of the animal and the session-averaged firing-rate map (Fenton et al., 2010; Fenton and Muller, 1998). As a direct consequence, place fields computed from PTI-rate disappear (Fig. 3D, right). Nonetheless, correspondence is high ($0.4 < r < 0.6$) between cell-pair correlations computed from the observed rates of the cell pairs and those computed from the PTI-rates of these pairs, indicating activity correlations in cofiring that exist beyond position-tuning (Fig. S3B), as also observed ubiquitously in neocortex with cells tuned to different variables (Churchland et al., 2010; Hennequin et al., 2018; Lee et al., 1998; Schneidman et al., 2003). The cell-pair correlations computed from observed rates covary with spatial firing

similarity for both place cells and non-place cells. The cell-pair correlations computed from PTI also covary with spatial firing similarity for both place cells and non-place cells, but at the same level as calculated with the observed rates for non-place cells (Fig. 3E). Taken together, this indicates that overlapping place fields are responsible for a small fraction of the relationship between cell-pair correlation and spatial firing similarity in place cells (Fig. 3E). PCo computed from PTI (PTI-PCo) decreases by approximately 20%, remains higher for trials between the same environment compared to between different environments, but does not change from week 1 to week 3 (Fig. 3F, left). Like PCo computed from observed rates (Fig. 3C, right), PTI-PCo also degrades over time (Fig. 3F, right). These findings indicate that the information contained in the cell's position tuning is, in theory, not necessary for discriminating two environments using the activity of CA1 principal cells, including place cells, motivating us to examine how coactivity relationships could be used for representing spatial information.

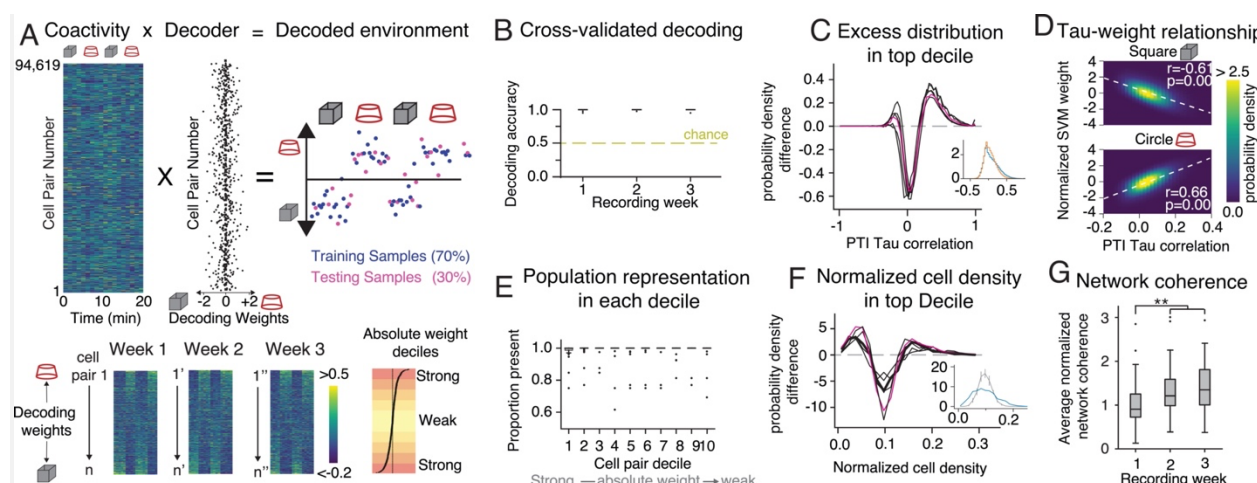


Figure 4. Decoding environments from PTI coactivity. (A) Illustration of SVM decoding from PTI coactivity; (top) PTI coactivity, computed every 60 s, is projected with SVM weights onto a single dimension that separates the two environments. (Bottom) Cell-pair coactivity from an example ensemble ordered by SVM weight each week; weights' absolute value, indicative of separation strength, are classified by decile. (B) Cross-validated decoding accuracy is very high

for all ensembles. (C) Difference between overall and first decile distribution of cell pair correlations (example distributions in inset and shown in purple in main figure) indicates importance of coactive and anti-coactive cell pairs, confirmed by (D) 2-D distribution of SVM weights against PTI coactivity. (E) Proportion of individual cells present in at least one cell pair in each SVM-ordered cell pair decile. (F) Difference between random and first decile, normalized distributions of individual cell participation (example distributions in inset and shown in purple in main figure) indicates overrepresentation of some cell in the first decile. (G) PTI network consistency normalized by the standard deviation of the cell-pair randomized distribution (Week: $F_{2,60} = 15.76$, $p=10^{-6}$, $\eta^2 = 0.094$; Dunnett's tests: week 0 < weeks 2, 3).

Decoding environments from position-tuning independent ensemble coactivity

We then investigated whether the current environment could be reliably decoded from the PTI ensemble coactivity, from which we eliminated all position tuning, in order to discover whether there is a crucial contribution of coactivity to decoding. Coactivity was computed at 1-s resolution over the period of a minute, yielding up to five independent estimates of coactivity per trial. We trained a support-vector machine (SVM) decoder to project the coactivity along a single composite dimension that separates the two environments (Fig. 4A). The decoder almost always correctly identifies the current environment (Fig. 4B), demonstrating that the coactivity itself carries discriminative information, independent of position tuning. Remarkably, we find that the most discriminative cell pairs are either strongly coactive or strongly anti-coactive (Fig. 4C, D). Most individual cells contributed to the cell pairs in each decile (Fig. 4E), but some cells are overrepresented in the subset of discriminative cell pairs that constitutes the first decile (Fig. 4F). Taken together, these two characteristics of cell-pair activity co-fluctuations beyond positional tuning indicate that population correlation dynamics can be described as scale-free, strikingly similar to models that describe flocks of birds (Hemelrijk and Hildenbrandt, 2015).

Indeed, power laws fit the distribution of activity correlations as a function of distance between soma; the best fit exponents are negative and slightly smaller than 1 (average \pm s.d.: -0.90 ± 0.05 ; see Fig. S3D). Accordingly, the correlation length (distance at which the average correlation = 0) is large, indicating that on average, the activity of each cell influences the activity of every other cell. The environment-discriminating activity of the neural population is driven by a subset of the individual cells but impacts the whole population, which is reminiscent of the flocking and schooling behaviors of birds and fish (Hemelrijk and Hildenbrandt, 2011, 2012; Reynolds, 1987). We consequently calculated the network consistency (van Dijk and Fenton, 2018) as an estimate of the overall correspondence between cell pairs' (i and j) PTI coactivity ($\tau_{i,j}$) and the pairs' instantaneous PTI rate (r_i and r_j). Network consistency increases with experience (Fig. 4G) indicating an increased alignment of the 1-s short- and 300-s long-timescales of cell pair cofiring in the network. Repeating the set of coactivity calculations without first eliminating the place fields yielded similar, but less discriminative results (Fig. S4, Week: $F_{2,19} = 34.96$, $p = 10^{-7}$, $\eta^2 = 0.023$; Time-Series Type: $F_{1,20} = 14.66$, $p = 0.001$, $\eta^2 = 0.35$; Interaction: $F_{2,19} = 26.61$, $p = 10^{-6}$, $\eta^2 = 0.011$; post-hoc: PTI > rate on weeks 1 and 3, $p \leq 0.004$; $p = 0.027$ on week 2). This is consistent with the possibility that the environment-specific information signaled by the coactivity of cells carries a separate type of information than the positional tuning measured by firing fields (Huxter et al., 2003).

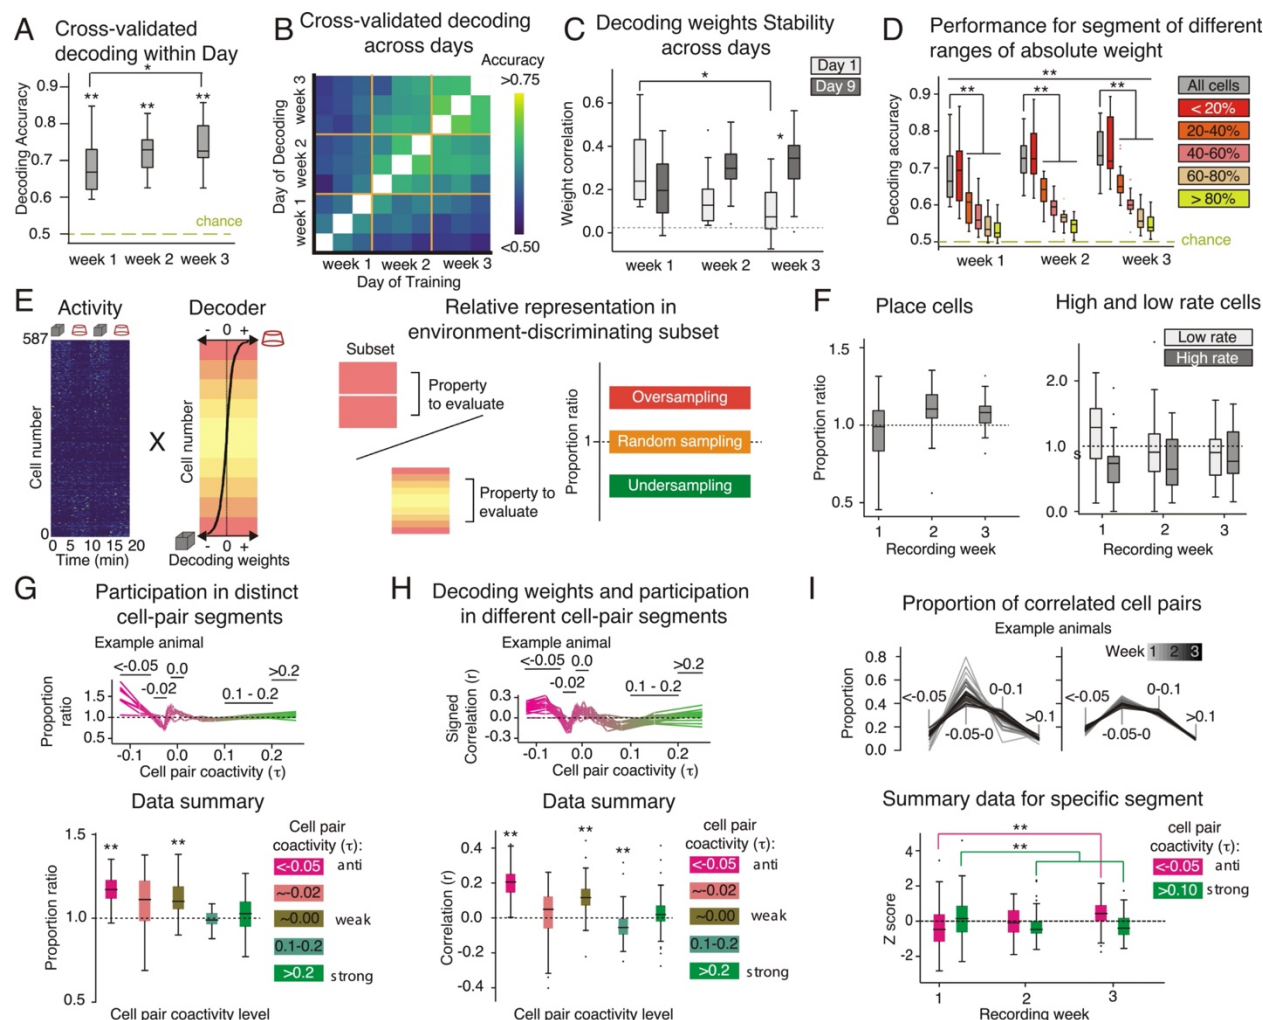


Figure 5. Decoding of 1-s activity vectors relies on a population subset and highlights the role of anti-cofiring cell pairs. (A) Cross-validated decoding performance is above chance and increases each week (Week: $F_{2,12} = 12.63$, $p=0.0011$, $\eta^2 = 0.11$; Dunnett's tests: week 1 < week 3; post-hoc t tests vs. 0, ** $p<10^{-7}$) (B) Heatmap of cross-validated decoding performance across days. (C) Correlation of weights trained on day 1 and 9 with weights trained on other days, with weights from the same day excluded (No effect of Week: $F_{2,9} = 2.11$, $p = 0.18$, $\eta^2 = 0.011$; no effect of Day: $F_{1,10} = 1.71$, $p=0.22$, $\eta^2 = 0.08$; Interaction: $F_{2,9} = 8.03$, $p = 0.01$, $\eta^2 = 0.14$; Dunnett's tests: week 0 < week 2 on day 1; post-hoc t tests: day 1 vs day 9, * $p < 0.017$) (D) Cells are ordered by SVM absolute weights and separated into quintiles (20%) whose cross-validated performance is measured independently (Quintile: $F_{4,65} = 38.58$, $p = 10^{-16}$, $\eta^2 = 0.65$;

Week: $F_{2,64} = 29.86$, $p = 10^{-10}$, $\eta^2 = 0.022$; Interaction: $F_{8,89.135} = 3.69$, $p = 10^{-4}$, $\eta^2 = 0.026$; post-hoc t tests: All Cells vs Quintiles, $**p \leq 0.001$; Dunnett's tests: week 0 < week 3 for quintile 1 and 2). (E) Illustration of activity vectors with cells ordered by weights. Weights' absolute value, indicative of separation strength, are classified by quintile. The properties of the top quintile, the environment-discriminating subset, are compared to the properties of the population to evaluate how cells from this subset differ from the population. (F, left) Ratio of proportion of cells classified as a place cell in either environment in top the quintile versus the overall population (No effect of Week: $F_{2,12} = 3.58$, $p = 0.061$, $\eta^2 = 0.13$; t test vs. unity (1): $t_{48} = 1.38$, $p = 0.18$). (right) Ratio of proportion of cells with high or low rate in top quintile versus overall population (High Rate: No effect of Week: $F_{2,12} = 0.65$, $p = 0.54$, $\eta^2 = 0.015$; t test vs. unity (1): $t_{48} = 3.19$, $p = 0.0025$; Low Rate: No effect of Week: $F_{2,12} = 3.47$, $p = 0.065$, $\eta^2 = 0.064$; t test vs. 1: $t_{48} = -0.07$, $p = 0.95$). (G) Ratio of average participation in cell pairs of different coactivity levels for top quintile versus overall population (top) in an example animal and (bottom) for specific levels of coactivity (Coactivity Level: $F_{4,65} = 5.71$, $p = 10^{-4}$, $\eta^2 = 0.094$; No effect of Week: $F_{2,65} = 0.19$, $p = 0.83$, $\eta^2 < 10^{-3}$; No Interaction: $F_{8,128} = 1.14$, $p = 0.34$, $\eta^2 = 0.031$; post-hoc t tests vs. 1, $**p \leq 10^{-6}$). (H) Average correlation between SVM weight and participation in cell pairs of different coactivity levels (top) in an example animal and (bottom) for specific coactivity levels (Coactivity Level: $F_{4,135} = 55.53$, $p = 10^{-27}$, $\eta^2 = 0.43$; No effect of Week: $F_{2,134} = 1.57$, $p = 0.21$, $\eta^2 = 0.0012$; No Interaction: $F_{8,189.12} = 1.82$, $p = 0.076$, $\eta^2 = 0.018$; post-hoc t tests vs. 0, $**p \leq 10^{-7}$). (I) Proportion of cell pairs with different coactivity levels across days (top) in two example animals and (bottom) for cofiring and anti-cofiring cell pairs (No effect of Coactivity Level: $F_{1,122} = 0.74$, $p = 0.39$, $\eta^2 = 0.0043$; Week: $F_{2,121} = 3.91$, $p = 0.023$, $\eta^2 = 0.0075$; Interaction: $F_{2,121} = 10.12$, $p = 10^{-5}$, $\eta^2 = 0.07$; Dunnett's tests vs week 1, $**p \leq 0.009$).

The environment-discriminating subset of CA1 cells

It is commonplace to represent hippocampal population activity patterns as activity vectors.

These are distinct between the two environments, yet this difference is remarkably small (Fig. S5), consistent with the notion that the activity of only a minority of individual cells contribute to discriminating the environments. Because only a quarter of CA1 cells are place cells, and because only place cells change their position tuning and the firing of even fewer turns on or off across environments, this small change in population activity across environments might be expected if place cells drive the discrimination. Consequently, we ask what subset of the CA1 population contributes to discriminating the two environments. We trained an SVM decoder to project the data along a single composite dimension that separates the two environments. Using 1-s activity vectors recorded on the same day, the decoder correctly identifies the current environment well above chance; performance increases with the animal's experience (Fig. 5A). After one week, decoding across days also performs above chance (Fig. 5B). Indeed, after one week, the SVM weights are stable across weeks; the weights obtained from the last day of the third week are strongly correlated to the weights obtained up to 9 days earlier, consistent with the expectation that SVM weight stability reflects learning, much like firing field stability has been traditionally interpreted (Fig. 5C). To evaluate which subset of cells contributes to discriminating the environments, we trained a decoder using select portions of the population to which the initial SVM decoder gave different weights. Decoders using just the 20% of cells with the largest weights perform indistinguishably from decoders that use the entire population. On the other hand, decoders using the 40% of cells with the smallest SVM weights perform close to chance (Fig. 5D). Decoders using cells with the 20% highest weights also display the largest increase in performance with experience, while for the bottom 60%, performance does not increase, consistent with the expectation that the largest SVM weights estimate the strength of spatial coding and learning, much like place field quality has been interpreted.

We call this 20% of cells the “environment-discriminating subset” and evaluated their properties (Fig. 5E). The place cells, strongly active cells, and weakly active cells are not more likely than chance to be in the environment-discriminating subset (Fig. 5F). We then evaluated whether the environment-discriminating subset is comprised of cells that tend to be coactive with other cells, or alternatively, tend to be anti-coactive with other cells. Only the anti-coactive cells are more likely than chance to be part of the environment-discriminating subset (Fig. 5G). Furthermore, the SVM decoding weights are predicted by the number of negatively correlated cell pairs to which a cell belongs. By contrast, participation in positively correlated pairs is not related to either being in the environment-discriminating subset or SVM weights (Fig. 5H). Consistent with anti-coactivity being an important feature for discriminating the two environments, we observe an increase in the number of negatively correlated cell pairs and a decrease in the number of positively correlated cell pairs, as mice learn to discriminate the two environments (Fig. 5I).

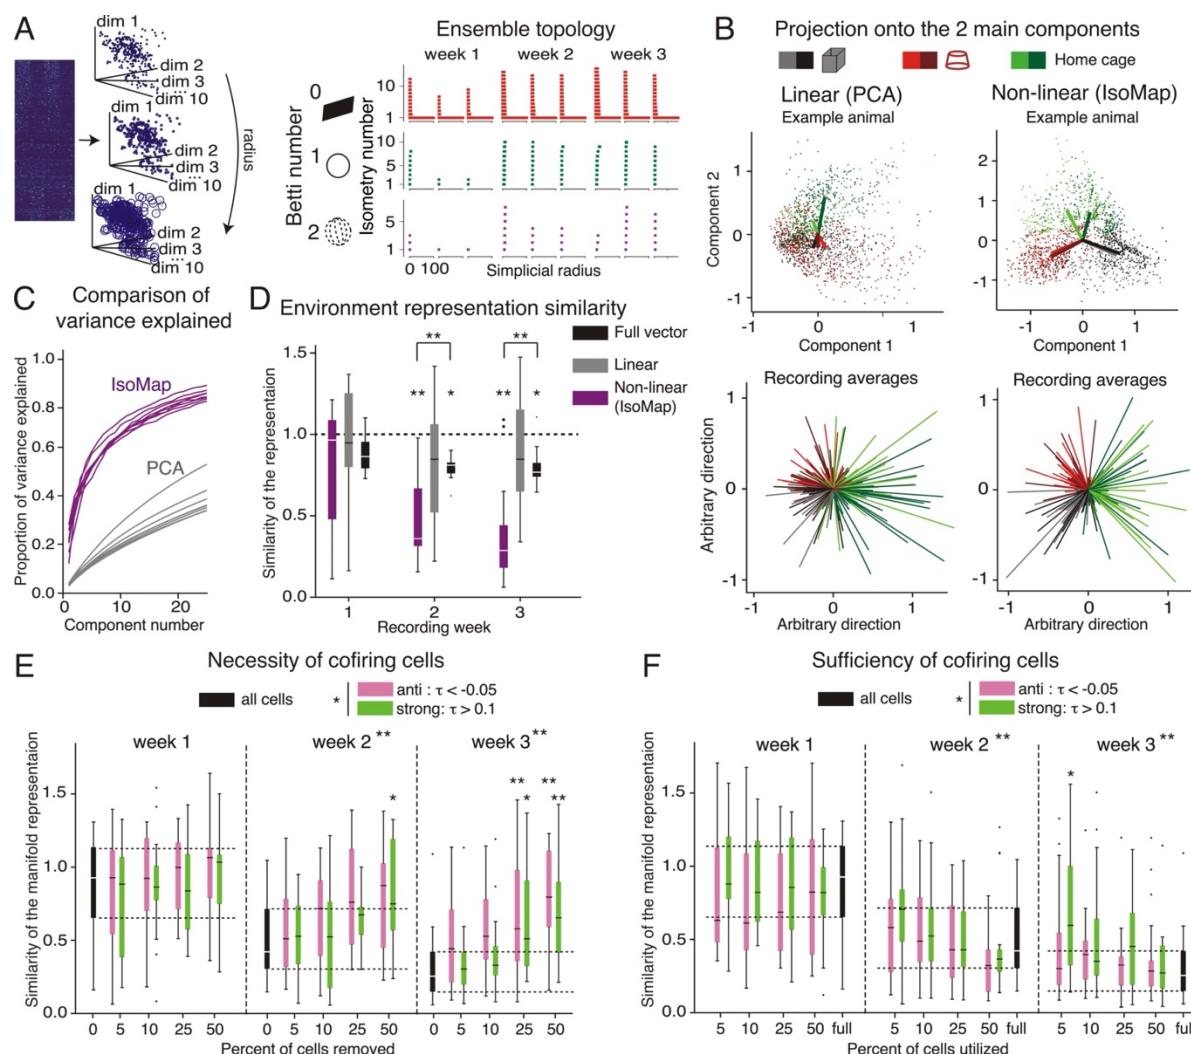


Figure 6. Non-Linear embedding distinguishes between environments and relies on anti-cofiring cell pairs. (A, right) Illustration of the algorithm used to determine the Betti numbers that describe the topology of ensemble activity. Activity vectors are projected onto a 10-dimensional space using IsoMap and points are connected to neighbors with increasing radius of each point. Isometries are determined at each radius. (right) Example of Betti barcode showing that only one isometry of Betti number 0 persists, indicating that the population activity appears constrained to a single connected subspace. (B) Activity vectors projected onto a 2-D subspace using a linear (PCA) or non-linear (IsoMap) dimensionality reduction algorithm. In each recording the average 2-D vector is computed as shown in (top) an example animal with

individual points, projections of individual 1-s activity vectors, and (bottom) for all recordings and animals. Repetitions of the same environments on the same day are given distinct color shades.

(C) Projections onto 3-D and higher dimension subspaces were also performed. The proportion of variance explained as a factor of the number of dimensions for PCA and IsoMap is given, indicating that the non-linear 2-D IsoMap subspace explains substantially more variance than the linear PCA projections. (D) Ratio of the difference between average vector from similar environment and different environments using IsoMap, PCA, or the full, non-dimensionally reduced, activity vectors (Method: $F_{2,39} = 10.38$, $p = 10^{-4}$, $\eta^2 = 0.19$; Week: $F_{2,38} = 6.67$, $p = 0.003$, $\eta^2 = 0.076$; Interaction: $F_{4,44.58} = 2.96$, $p = 0.03$; post-hoc Dunnett's tests: IsoMap < Full Vector, $**p \leq 0.046$; post-hoc Dunnett's tests: week 1 > weeks 2, 3, $*p < 0.03$, $**p < 0.01$). (E) Similar ratio after removing different percentages of the population as a function of their contribution to cell pairs of different coactivity levels (Coactivity Level: $F_{1,99} = 6.06$, $p = 0.016$, $\eta^2 = 0.012$; Percentage: $F_{1,99} = 17.86$, $p = 10^{-5}$, $\eta^2 = 0.078$; Week: $F_{2,98} = 22.08$, $p = 10^{-8}$, $\eta^2 = 0.13$; no Interaction, all p 's > 0.18; post-hoc: Anti > Strong $t_{380.93} = 2.11$, $p = 0.036$; post-hoc Dunnett's tests: week 1 > weeks 2, 3, $**p \leq 10^{-7}$; post-hoc t tests vs. All Cells, $*p$'s < 0.05, $**p$'s < 0.01). (F) Similar ratio using only percentages of the population as a function of their contribution to cell pairs of different coactivity levels (Coactivity Level: $F_{1,99} = 5.27$, $p = 0.24$, $\eta^2 = 0.014$; Percentage: $F_{1,99} = 4.71$, $p = 0.032$, $\eta^2 = 0.025$; Week: $F_{2,98} = 13.00$, $p = 10^{-6}$, $\eta^2 = 0.17$; Week x Percentage x Coactivity Levels: $F_{2,98} = 3.10$, $p = 0.049$, $\eta^2 = 0.0055$; no other significant interaction, all p 's > 0.12; post-hoc: Anti < Strong $t_{380.40} = 2.33$, $p = 0.020$; post-hoc Dunnett's test: week 1 > weeks 2, 3, $**p < 0.01$; post-hoc t tests vs. All Cells, $*p$'s < 0.05).

Planar manifold topology of ensemble firing distinguishes environments

Place fields on a surface are planar, making it straightforward to imagine that place cell ensemble activity generates representations with 2-D structure that can differentially represent distinct 2-D environments. What is the dimensionality of the n-cell dimensions of ensemble activity in which coactivity relationships differentially represent environments? To evaluate the representational topology of the data we used persistent homology and computed Betti number barcodes (Ghrist, 2008) which indicate that the topology of 1-s CA1 ensemble activity patterns has one component and no holes. We assumed this might approximate a 2-D surface, and asked whether an unsupervised, non-linear dimensionality reduction algorithm can readily distinguish environments on a 2-D manifold (Fig. 6A; Video S1). The IsoMap algorithm uses the distances between activity vectors to non-linearly project the data (Balasubramanian and Schwartz, 2002), so we compared it with PCA, a linear transformation of the data. We applied both algorithms to 1-s activity vectors recorded in the cylinder, box, and home cage environments. IsoMap distinguishes the ensemble activity vectors recorded in the three environments when the 1-s ensemble activity vectors are projected onto the two main IsoMap dimensions (Fig. 6B, right). The environments are harder to discriminate when the identical activity vectors are projected onto the two main PCA dimensions (Fig. 6B, left). These impressions are quantified using explained variance. The first 10 IsoMap dimensions explain more than 60% of the variance, compared to about 20% with the first 10 PCA components (Fig. 6C). To quantify environment discrimination, we calculated for each recording the average projection onto the 2 main dimensions of IsoMap or PCA (Fig. 6D). The difference between projections from the same environments was conservatively normalized by the difference between projections from different environments (Fig. S6). This discrimination ratio will decrease from 1, when all environments are similarly represented, to zero, when the same environments are identically encoded, and different environments are well discriminated. Using IsoMap, this normalized vector difference decreases from 1 to about 0.3 with the animal's

experience (Fig. 6D). This contrasts with the modest decrease of the discrimination ratio when the activity vectors are projected onto the PCA dimensions, which stays above 0.8 on average. The discrimination ratio also only decreased modestly when it was computed on the raw activity vectors without any dimensionality reduction (Fig. 6D). Finally, to elucidate whether the contribution of coactivity is necessary and sufficient for the IsoMap discrimination of the environments, we recomputed the activity vectors after systematically removing or including only the top 5, 10, 25, and 50% of cells that participated most in the coactive and anti-coactive cell pairs. The surviving ensemble activity vectors were then projected on the two main IsoMap dimensions. Removal of the most coactive or anti-coactive cells selectively degraded the IsoMap discrimination (Fig. 6E), and inclusion of only the most coactive and anti-coactive cells was uniquely sufficient for the discrimination (Fig. 6F). These effects were stronger for the anti-coactive cells. These findings motivate a novel “reregistration” hypothesis that CA1 ensemble coactivity can discriminatively represent distinct environments through a relatively conserved neuronal ensemble activity pattern that is constrained on a non-linear surface manifold that is at once conserved across environments, and distinctively registered to each environment as a result of the experience-dependent development of anti-coactive neural discharge (Fig. 1B; 7A). Importantly, the hypothesis asserts that environment-specific differences in cofiring relationships amongst the cells as shown in Fig. 3, are largely due to the anti-cofiring cells, that this subset of cells is responsible for the differences between environment-specific neuronal activity manifolds, and that the topology of the environment-specific manifolds is invariant (Fig. 7B). Anti-cofiring power was computed for each cell in an ensemble as the proportion of the ensemble cells that cofire such that $\tau < -0.05$ (Fig. 7C).

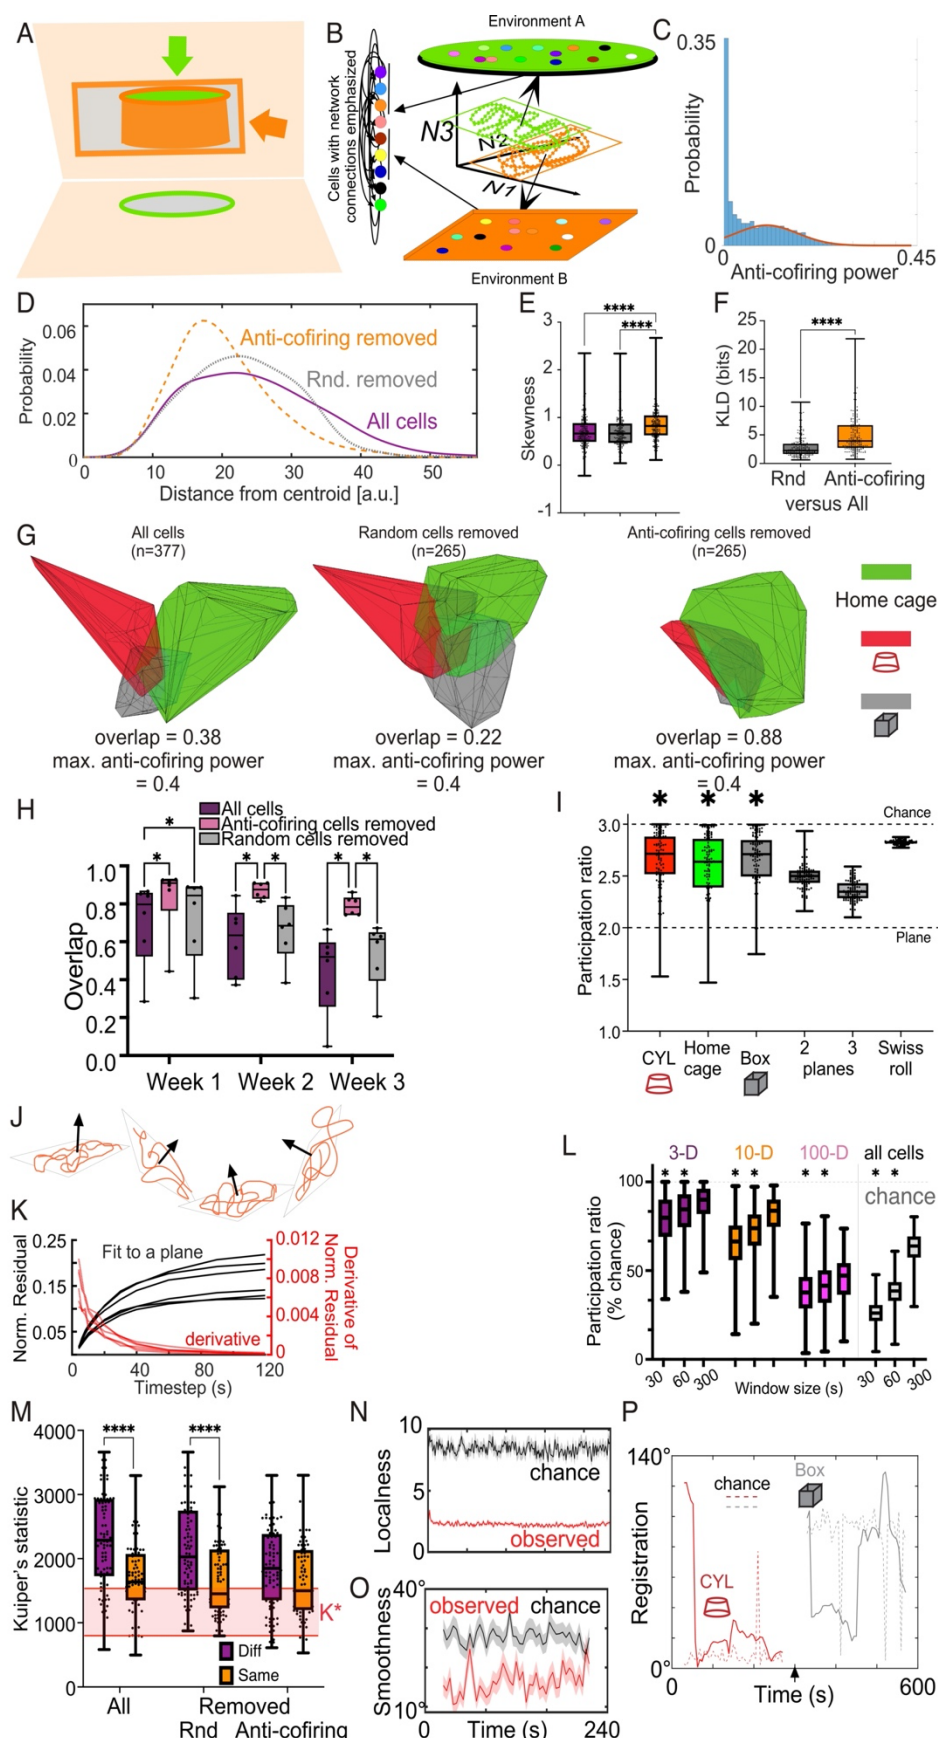


Fig. 7. Reregistration, reconceptualizing how hippocampus represents multiple environments.

(A) Schematic illustrating that projections of an invariant object (a solid cylinder) at distinct angles can generate distinct representations (circle or rectangle). (B) The reregistration hypothesis illustrated. Intrinsically-organized neuronal network activity organizes as smooth trajectories on a low-dimensional surface. A low-dimensional (2-D) manifold (green and orange rectangles) is depicted in the 3-D subspace. Environment-specific representations result from projecting the environment onto the manifold, which does not reorganize cofiring relationships, but reorganizes place cell firing fields (color ellipses). Environment-specific registrations of the manifold with external cues (different angles in the 3-D subspace) result from taking perceptual fixes that orient manifold activity with respect to external cues. (C) Distribution (blue) with Gaussian fit of the anti-cofiring tail (orange) of coactivity power computed for cells from each recording in the dataset. (D) A day-7 recording example histogram of 300 1-s activity vector distances from the centroid after projecting a 377-cell (All) ensemble into a 60-D IsoMap subspace. The 112 (30%) most anti-cofiring and an equivalent random set of cells were removed, and the remaining 265-cell ensembles were projected into a 60-D IsoMap subspace. The histogram comparisons suggest that the anti-cofiring cells are disproportionately far from the centroid. The example is from the second recording in the rectangular environment during the third week, for mouse M39. Summary quantification of (E) Skewness and (F) Kullback-Liebler divergence of the 3 histograms from each of all recordings ($n = 200$) confirm anti-cofiring cells are distinctive and significantly distant from the centroid. (G) Visualization as a volume, after projecting the 1-s ensemble activity vectors from panel D into a 3-D subspace defined by the first 3 IsoMap components. The component coordinate system was determined from two 5-min recordings of the ensemble population in the home cage, box, and cylinder. The volumes occupied by activity vectors were compared for all cells (left), a subset of randomly selected cells (middle) or after removing the most anti-cofiring cells (right). Removing the 30% most anti-cofiring cells collapsed the distinction between the cylinder and box volumes. (H) Summary data

from all ensembles recorded across the 3 weeks illustrates that the distinction between the environment-specific manifold volumes increases with experience, and that the anti-cofiring subset of cells are crucial for the distinction (two-way Ensemble x Week ANOVA, Ensemble: $F_{2,45} = 14.04$, $p = 10^{-5}$; Week: $F_{2,45} = 7.33$, $p = 0.002$; Interaction $F_{4,45} = 0.32$, $p = 0.9$; Friedman statistic = 30.79, Anti-cofiring removed > All cells = Randomly removed). (I) Participation ratios (PR) estimate the dimensionality of the ensemble activity projected into the 3-D IsoMap subspace for all the recordings. The dimensionality is significantly different than chance (PR=3) and indistinguishable from planes with multiple shifts as well as a non-linear Swiss Roll (Kruskal-Wallis statistic = 138.2, $p = 0$). (J) Schematic illustrating the hypothesis that ensemble activity in the 3-D IsoMap subspace is organized along planar trajectories that episodically shift orientation in the subspace. (K) Residuals and derivative from best fit planes for each mouse as a function of the data timestep used to estimate the duration ensemble activity is constrained to a plane in the 3-D IsoMap subspace. (L) Participation ratio data dimensionality estimates by projecting different windows of data into IsoMap subspaces of different dimensionality. (M) Kuiper's statistics compared the set of average angles of the best fit planes in the 3-D IsoMap subspace for each of the daily recordings in the same or different box and cylinder environments. (N) Summary measures of the localness and (O) smoothness timeseries of the 1-s ensemble activity trajectories in the 3-D subspace compared to chance and (P) an example timeseries (4-s timestep) of the angular registrations of the best fit plane to 60-s of data along the ensemble trajectory through the 3-D IsoMap subspace.

We then investigated whether the anti-cofiring cells caused the ensemble activity to be distinctive by projecting the data into a 60-D IsoMap subspace. The 60-D choice facilitated statistical comparisons by allowing all ensembles to be assessed in a subspace of the same dimensionality that is sufficient to account for most of the variance in activity vectors (Fig. 6C).

The 60-D centroid of the 1-s activity vectors from each recording was computed and the distribution of the Euclidean distance from the centroid for each activity vector was examined. Fig. 7D, the example of a 377-cell ensemble recorded on day 7, illustrates the distribution has a long tail of ‘outlier’ vectors that are far from the centroid. We tested if the anti-correlated cells disproportionately contribute to the outlier by removing the 30% (112) most anti-cofiring cells from the ensemble and as a control, removed an equivalent random sample of cells then recomputed the histogram. Removing the anti-cofiring cells reduced the number of ‘outliers’ and strongly shifted the distribution leftward more than removing an equivalent random number of cells. This impression was quantified by computing the skewness of the three distributions for each recording across the mice (Fig. 7E; Friedman statistic = 69.12, $p < 0.0001$; 3 distributions, each with 200 samples; anti-cofiring cells removed > all cells = random cells removed, p ’s < 0.0001). Using the Kullback-Liebler divergence (KLD) we then compared how different the histograms of the distances from the centroid became after removing the anti-cofiring and random cells. Consistent with the example in panel D, removing anti-cofiring cells increased the difference from the histogram of the full ensemble far more than removing random cells (Fig. 7F; Wilcoxon $W=18260$, $p < 0.0001$). These findings from the 60-D IsoMap subspace suggest that anti-cofiring cells cause distinctive activity vectors.

Next, to visualize and test the possibility that the anti-cofiring cells are responsible for the distinction between the environmental representations, we projected into a 3-D IsoMap subspace, the activity vectors from the pair of cylinder and box recordings on a particular day. The three principal IsoMap components were determined by concatenation of the home cage, box, and cylinder recordings of 1-s activity vectors, and projected into the same 3-D sub-space. Figure 7G shows the volumes occupied by the 600 (2 x 5 min) 1-s ensemble activity vectors from the 377-cell ensemble recorded on day 5 in the home cage, cylinder, and box for mouse

M39. The overlap of the cylinder and box activity is modest (0.38), consistent with ensemble activity being sufficiently distinctive to discriminate the two environments (Fig. 6). Removing a random 30% of cells did not increase the overlap (0.22) whereas removing 30% of the most anti-cofiring cells reduced the outliers (by inspection), did not remove enough cells to change the maximum anti-cofiring power, and importantly, collapsed the distinction between the cylinder and box activity patterns (0.88). We repeated this analysis for the second or third recording during each week to determine whether or not the key role of the anti-cofiring cells in distinguishing the cylinder and box ensembles developed with experience. Figure 7H illustrates that during the first week, the cylinder and box activity vectors occupied a largely overlapping subspace but with experience, from week 2, became more distinct. Importantly, removing the anti-cofiring cells, which reduced the maximum anti-cofiring power from 0.20 ± 0.013 to 0.15 ± 0.013 , reduced the distinction between the two environments, but not when an equivalent subset of cells was randomly removed (maximum anti-cofiring power 0.19 ± 0.013). This pattern is predicted by the reregistration hypothesis.

Multistable manifold representations of environments

Projecting the activity vectors into the volume of the first three IsoMap components, hinted they may organize on 2-D surfaces (see Videos S1,S2), which is corroborated by computing participation ratios to estimate the dimensionality ($PR = 2.65 \pm 0.30$, $n = 258$) (Gao et al., 2017; Recanatesi et al., 2019). These ratios were significantly lower than chance ($PR = 3$), but indistinguishable from a non-linear Swiss roll surface, or 2-D planes that change orientation within the 3-D space (Fig. 7I).

We've previously observed that internally-organized entorhinal and hippocampal spatial representations are multistable, they spontaneously change registration to the environment on the time scale of ~10 seconds during free navigation (Chung et al., 2021; Dvorak et al., 2021; Dvorak et al., 2018; Fenton et al., 2010; Kelemen and Fenton, 2016; Park et al., 2019; Talbot et al., 2018). Accordingly, we considered the possibility that CA1 activity is multistable, in that it intermittently registers to distinct environmental features, which would change the orientation of a fundamentally lower-dimensional (2-D) surface within the higher-dimensional (3-D) IsoMap space, as cartooned in Fig. 7J.

Consistent with the multistable conjecture, 30-60 s of activity vectors are best fit to a plane in the 3-D IsoMap subspace (Fig. 7K). We animated the activity vector timeseries projected into the IsoMap subspace (Video S2), which upon inspection shows the high-dimensional ensemble activity takes non-random, smooth trajectories through a 2-D subspace that intermittently changes orientation. Analysis of participation ratios computed on 30, 60, and 300-s time windows indicate that the dimensionality is significantly lower for the shorter time windows (Fig. 7L) suggesting that the hypothesized multistability occurs on a time scale of ~60 s.

We proceeded by computing Kuiper's statistic to measure whether the set of 60-s planar angles differed for daily recordings in the same and different box and cylinder environments. Fig. 7M shows that the angles were more similar in two recordings in the same environment than they were in different environments, consistent with Figure 6, and indicating that manifold orientation in the subspace can distinguish the two environments. Importantly, this distinction between recordings in the two environments is lost when anti-cofiring, but not random cells are removed (Two-way Environment x Ensemble ANOVA, Environment: $F_{1,534} = 56.87$, $p = 10^{-13}$, Ensemble: $F_{2,534} = 6.77$, $p = 0.001$; Interaction $F_{2,534} = 5.06$, $p = 0.007$; post-hoc Tukey: Different

environments: Anti-cofiring removed < All cells = Randomly removed p 's < 0.02; Same environments: Anti-cofiring removed = All cells = Randomly removed, p 's > 0.98). We quantified the localness (Fig. 7N; 2.26 ± 0.01 vs. chance: 8.44 ± 0.04 , Wilcoxon $W = 8.77 \times 10^8$, $p < 10^{-4}$), and smoothness (Fig. O; $16.5^\circ \pm 0.4^\circ$ vs. chance: $28.0^\circ \pm 0.39^\circ$, Wilcoxon $W = 1.68 \times 10^7$, $p < 10^{-4}$) of the ensemble trajectory in the IsoMap subspace, which are distinct from chance and together indicate that ensemble activity projected into the 3 best components of the IsoMap subspace is largely constrained to a 2-D plane for several 10's of seconds, and then within 1-2 seconds, the plane reorients in the IsoMap subspace (Fig. 7P).

Analogous to firing rate maps, the spatial distribution of each of the three principal IsoMap components of the ensemble activity vector projections were computed to investigate the possibility of time-averaged spatial tuning (Fig. S7A). By inspection, no obvious patterns were recognized, but they were also not random because within week 2 recordings (days 4,5, or 6), the patterns across the same environments were significantly correlated for 4/6 mice (10/36 correlations were significant), and also across the different environments for 5/6 mice (7/72 correlations were significant). Finally, we examined the activity vectors around the times of the reorientations but did not recognize any relationships (r 's < 0.14, p 's > 0.2; Fig. S7B).

Discussion

Place field and coactivity-based neural codes for space

We find that the coactivity relationships amongst CA1 principal cells discriminate environments and stabilize with experience (Fig. 3), identifying a learned neural code that can be independent of the place field-based code conventionally used to understand the hippocampal role in spatial cognition. Instead of single cell properties, our analyses focused on the set of pairwise activity correlations that approximate the higher-order activity correlations that may define informative

neural population activity patterns (Ebitz and Hayden, 2021; Jazayeri and Afraz, 2017; Schneidman et al., 2006). We took advantage of the miniature microscope technology to monitor the activity and positions of large numbers of the same cells over weeks (Fig. 2), an ability which was not available with the electrophysiological techniques that established the importance of place fields and remapping for the representation of space. There are of course limitations to this technology that impact the present study. The timescale of analysis was restricted to 1-s as processes faster than ~300s are precluded by the slower time scale of GCaMP6f transients. Nonetheless, 1-s is the behavioral timescale of CA1 synaptic plasticity mechanisms (Bittner et al., 2017; Milstein et al., 2021), and sufficient time resolution for computing the cofiring relationships because the statistics of cofiring relationships measured at 25 and 40 ms are equally well captured by the statistics of 1-5 s cofiring relationships (Olypher et al., 2006; Park et al., 2019). We restricted analyses to anatomically separated cell pairs to avoid spuriously high activity correlation, but this prevents analysis of the fine-scale topographical organization of the neural circuit (Hampson et al., 1999; Pavlides et al., 2019). Place cell stability was variable with some cells unstable (similarity ~ 0) and others highly stable (similarity > 0.5). Stability increased after the first 3 days (30-min total) of exposure (Fig. 2F,G) reaching a magnitude that is typical of mouse place cells recorded with tetrodes during purposeful behavior (Kentros et al., 1998; Muzzio et al., 2009). Network consistency, which estimates the alignment of higher-order activity correlations, also increases with experience and with coactivity discrimination of environments, further highlighting the importance of coactivity (Fig. 4), consistent with the modeling results (Figs. 1, S2), and reports that CA1 place field plasticity relies more population cofiring statistics than on the activity of individual cell pairs (Milstein et al., 2021). Coactivity relationships were scale-free (Fig. S3) and dominated by a minority of anti-cofiring cells for discriminating environments (Figs. 4,5,7), features that describe the behavior of flocks and their interactions with external features (Hemelrijk and Hildenbrandt, 2011, 2012;

Reynolds, 1987). Support vector machine decoding objectively measured the globally optimal contribution of each cell or cell pair to distinguishing the two environments (Figs. 4,5). While this was effective on a 1-s timescale, it only demonstrates the availability of discriminative information, rather than how information is encoded and extracted by neural circuits, not unlike how the finding of place fields has not determined how the hippocampus uses them. The findings identify the importance of coactivity for representing environments, however these data in no way preclude place field-based codes for representing locations. Rather, both coding schemes may operate in parallel. We note that the possibility of having multiple firing fields in larger environments (Fenton et al., 2008; Harland et al., 2021; Muller and Kubie, 1987; Park et al., 2011) can degrade decoding of environments (Fig. S1) and the literature reports diverse additional phenomena that may also complicate such dedicated place codes. These phenomena include fundamental features of these cells such as firing-rate overdispersion (Fenton et al., 2010; Fenton and Muller, 1998; Jackson and Redish, 2007; Nagele et al., 2020; Poucet et al., 2012), mixed tuning to extra-positional variables (Fusi et al., 2016; Hardcastle et al., 2017; Stefanini et al., 2020), theta-phase temporal coding (Buzsaki and Chrobak, 1995; Harris et al., 2003; Hirase et al., 1999; Huxter et al., 2003), and firing-field instability (Fig. 2; Muzzio, 2018; Ziv et al., 2013), which should be properly evaluated by simulation studies (Kang et al., 2021). We also relied on the IsoMap algorithm to discover that the somewhat variable ensemble activity was constrained to low-dimensional subspaces where distinct environment-specific patterns relied on coactivity (Fig. 6,7). We analyzed ensemble activity within abstract 60-D IsoMap subspaces that account for almost all the variance (Fig. 6C), as well as assuming a temporally-local 2-D organization within a 3-D IsoMap subspace to evaluate the multistability conjecture without assuming knowledge of the true geometry, or that the true dimensionality is 2-D. Nonetheless, others have argued the neural topology reflects the spatial geometry of the environments from which place cells are recorded (Chen et al., 2014). Taken together, these

findings constrain the possible neural mechanisms for reading out the information represented in coactivity and point to synapse properties that have been proposed (Buzsaki, 2010; Buzsaki and Tingley, 2018) and are demonstrated, in artificial neural networks, to be effective mechanisms for transforming variable input activities to a reliable readout (Heeger and Mackey, 2019).

The roles of feature tuning, coactivity, recurrence, and stability in the neural space code

Our experimental design (Fig. 2A) is the classic remapping experiment repeated across weeks, with recordings during all exposures to the environments. We replicated the fundamental findings that single cells change their place fields between environments and reinstate the place fields upon imminent return to an environment (Fig. 2F,H). While recurrence manifests as this place field stability and is standardly interpreted as memory persistence (Leutgeb et al., 2005c; Wills et al., 2005), it is hard to reconcile with the replicated finding that place fields are also unstable across weeks even in familiar environments (Fig. 2F,J; Lever et al., 2002; Ziv et al., 2013). This would require numerous ($\sim n$) cell-specific rules to map the place fields from n cells in order to maintain, from one experience to the next, the coherent representation of the environment that a memory system requires (Lever et al., 2002; O'Keefe and Burgess, 1996). Furthermore, for such a scheme to be useful for discriminating environments, place cells should discharge predictably in their firing fields, but they do not (Fenton et al., 2010; Fenton and Muller, 1998; Jackson and Redish, 2007; Leutgeb et al., 2005b). In contrast, we identified recurrence of moment-to-moment coactivity relationships amongst CA1 cells, that the recurrence is modest (Fig. 3C), but sufficient for differentially representing environments (Fig. 4B).

Importantly, recent findings demonstrate that variable, distributed activity can organize within the stability of a manifold representation that explicitly does not rely on the reliability of either single cell or population activity (Figs. 6,7; Chaudhuri et al., 2019; Nieh et al., 2021; Rubin et al., 2019), the way that conventional place field-based and activity vector-based spatial codes have been conceived. Our analyses focus on the internal organization of neuronal activity, based on the firing of other neurons, rather than external variables like places, as demonstrated effective in MEC, a major input to hippocampus (Gardner et al., 2022; Park et al., 2019; Yoon et al., 2013). The present findings offer an alternative concept for how hippocampal activity can be persistently informative without firing field stability (Fig. 2), or even population activity stability (Fig. 3). A connectivity weight matrix, in other words the hypothesized synapse ensemble (Buzsaki, 2010; Buzsaki and Tingley, 2018), can transform the non-linear low-dimensional projections of neuronal activity (Fig. 6,7) to a steady readout (e.g., environment identity), so long as the connectivity matrix is invariant to shifts in the manifold axes, in which case the manifold would subserve memory persistence despite instability of neural activity, as demonstrated using artificial neural networks (Heeger and Mackey, 2019).

Reregistration instead of remapping to represent spaces and memory

The remapping concept is founded on the notion that spatial and thus temporal firing relationships amongst hippocampal neurons change across environments and memories. This interpretation is based on the observation that firing fields change across, but not within environments, and the inference that ensemble temporal firing relationships change in correspondence (Colgin et al., 2008; Kubie et al., 2020; Muller and Kubie, 1987). As reviewed above, these assumptions and related inferences are violated (Fig. 2). Indeed, CA1 ensemble activity is only modestly different across environments (Fig. 3, Fig. 7) and within environments, can exhibit multistable dynamics that resemble remapping expectations (Dvorak et al., 2018;

Fenton, 2015a; Jackson and Redish, 2007; Kao et al., 2017). Taken together, the data are consistent with the alternative concept of “reregistration.” Reregistration assumes CA1 activity is internally-organized in a relatively invariant manner such as on a low-dimensional manifold (Fig. 7A,B). Instead of external stimuli causing the rearrangement of neuronal activity, the internally-organized activity maintains but variably registers to external stimuli (Rule et al., 2019). This process may resemble fitting a model to observations (Fig. 7B). Cells that are most strongly influenced by currently salient stimuli are most likely to change their firing and because of the scale-free network correlations (Fig. S3) they influence and are influenced by the activity of all other cells. In flocks this allows the activity of most individuals to maintain locally while the entire flock changes globally (Hemelrijk and Hildenbrandt, 2011, 2012; Reynolds, 1987). The result of reregistration is that only a minority of cells change their activity distinctly from the activity of the other cells in the network, and according to present findings (Figs. 4,5,6,7), these we speculate are the anti-cofiring cells. The global network activity retains its low-dimensional organization but is registered (i.e. attached) uniquely to the environment in a manner that persists, especially we speculate, if synaptic weights are preferentially adjusted between the anti-cofiring cells and their environment-associated inputs. Because place cell firing is fundamentally multimodal, the activity only appears unimodal in small environments (Fenton et al., 2008; Harland et al., 2021), the ~25% of CA1 cells that express firing fields change firing field locations with reregistration, although the internal organization of activity changes much less, as confirmed by simulations (Fig. 1, S2). Reconceptualizing remapping as reregistering identifies the importance of the anti-cofiring cell subset, not only as crucial for discriminating environments (Fig. 4,5,6,7), but also for orienting the relatively invariant manifold activity patterns within the neuronal activity subspace (Fig. 7).

CA1 activity has been recorded during explicit memory tasks to observe how place fields change with memory (Frank et al., 2004; Jeffery et al., 2003; Lenck-Santini et al., 2001). Although the field has struggled to establish a firm relationship to firing fields, memory correlates have been identified in cofiring relationships between excitatory as well as inhibitory cells in sleep and active behavior (Dvorak et al., 2021; O'Neill et al., 2008; Skaggs and McNaughton, 1996; van Dijk and Fenton, 2018). This is in line with the original conceptualization of remapping as a reorganization of the temporal discharge properties within an ensemble (Kubie et al., 2020; Kubie and Muller, 1991). Such changes in coactivity are consistent with synaptic plasticity studies that demonstrate balance (Okun and Lampl, 2008), bidirectionality (Milstein et al., 2021), and involvement of both excitatory and inhibitory cells (Basu et al., 2016; Caroni, 2015; Chung et al., 2021; Mongillo et al., 2018; Ruediger et al., 2011). We find that the strongly coactive and anti-coactive cell pairs are particularly discriminative at the 1-s timescale and, remarkably, anti-coactive cells are both more discriminative and more important for organizing the activity on a manifold (Fig. 6,7; Okun et al., 2015). This is consistent with the increase of interneuron-principal cell cofiring that has been observed at moments of memory discrimination and recollection (Dvorak et al., 2021; van Dijk and Fenton, 2018). Using PTI, we analyzed activity fluctuations around place tuning (Agarwal et al., 2014), in effect removing the place signal from the time series. We did this to compute cell-pair coactivity and found that it improved the ability to discriminate environments (Fig. 4 and Suppl. Fig. 1), hinting at the possibility of relative independence of the place and environment codes. Although we did not study the role of place cells and coactivity in location-specific place coding, there are precedents of parallel codes for distinct spatial variables (Harris et al., 2003; Huxter et al., 2003; Meshulam et al., 2017; Poucet et al., 2012; Sarel et al., 2017; Stefanini et al., 2020; Tanaka et al., 2018). The coexistence of place field and coactivity codes for different features of space predicts that the explicit memory-related information in CA1 activity that has evaded the field's efforts, would manifest in the

position-tuning independent cell-pair coactivity introduced here, although in our view the manifold organization of observed ensemble coactivity relationships, replete with spatial modulation, is more likely the proper description of the hippocampal spatial code. We recognize that while we have identified a correspondence between patterns of neural coactivity and our ability to discriminate environments, we, like the rest of the field, have not determined what is actually represented in neural activity nor whether this correspondence causes the subject to understand its environments as distinct, or anything else for that matter. Nonetheless, our findings demonstrate the value of a conceptual shift towards a serious consideration of such vexing problems from the vantage of the collective and inherently temporally-structured behavior of neural activity (Brette, 2019; Buzsaki, 2019).

References

- Agarwal, G., Stevenson, I.H., Berenyi, A., Mizuseki, K., Buzsaki, G., and Sommer, F.T. (2014). Spatially distributed local fields in the hippocampus encode rat position. *Science* 344, 626-630.
- Aharoni, D., and Hoogland, T.M. (2019). Circuit Investigations With Open-Source Miniaturized Microscopes: Past, Present and Future. *Front Cell Neurosci* 13, 141.
- Alme, C.B., Miao, C., Jezek, K., Treves, A., Moser, E.I., and Moser, M.B. (2014). Place cells in the hippocampus: eleven maps for eleven rooms. *Proc Natl Acad Sci U S A* 111, 18428-18435.
- Balasubramanian, M., and Schwartz, E.L. (2002). The isomap algorithm and topological stability. *Science* 295, 7.
- Basu, J., Zaremba, J.D., Cheung, S.K., Hitti, F.L., Zemelman, B.V., Losonczy, A., and Siegelbaum, S.A. (2016). Gating of hippocampal activity, plasticity, and memory by entorhinal cortex long-range inhibition. *Science* 351, aaa5694.
- Bittner, K.C., Milstein, A.D., Grienberger, C., Romani, S., and Magee, J.C. (2017). Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science* 357, 1033-1036.
- Bostock, E., Muller, R.U., and Kubie, J.L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus* 1, 193-205.
- Brette, R. (2019). Is coding a relevant metaphor for the brain? *Behavioral and Brain Sciences* 42, e215.
- Buzsaki, G. (2010). Neural syntax: cell assemblies, synapsembles, and readers. *Neuron* 68, 362-385.
- Buzsaki, G. (2019). *The Brain From Inside Out* (New York: Oxford University Press).

- Buzsaki, G., and Chrobak, J.J. (1995). Temporal structure in spatially organized neuronal ensembles: a role for interneuronal networks. *Curr Opin Neurobiol* 5, 504-510.
- Buzsaki, G., and Tingley, D. (2018). Space and Time: The Hippocampus as a Sequence Generator. *Trends Cogn Sci* 22, 853-869.
- Cai, D.J., Aharoni, D., Shuman, T., Shobe, J., Biane, J., Song, W., Wei, B., Veshkini, M., La-Vu, M., Lou, J., *et al.* (2016). A shared neural ensemble links distinct contextual memories encoded close in time. *Nature* 534, 115-118.
- Caroni, P. (2015). Inhibitory microcircuit modules in hippocampal learning. *Curr Opin Neurobiol* 35, 66-73.
- Chaudhuri, R., Gercek, B., Pandey, B., Peyrache, A., and Fiete, I. (2019). The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nat Neurosci* 22, 1512-1520.
- Chen, T.-W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., *et al.* (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* 499, 295-300.
- Chen, Z., Gomperts, S.N., Yamamoto, J., and Wilson, M.A. (2014). Neural representation of spatial topology in the rodent hippocampus. *Neural Comput* 26, 1-39.
- Chung, A., Jou, C., Grau-Perales, A., Levy, E.R.J., Dvorak, D., Hussain, N., and Fenton, A.A. (2021). Cognitive control persistently enhances hippocampal information processing. *Nature* 600, 484-488.
- Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, M.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., *et al.* (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience* 13, 369-378.
- Colgin, L.L., Moser, E.I., and Moser, M.B. (2008). Understanding memory through hippocampal remapping. *Trends Neurosci* 31, 469-477.
- Duvelle, E., Grieves, R.M., Hok, V., Poucet, B., Arleo, A., Jeffery, K.J., and Save, E. (2019). Insensitivity of Place Cells to the Value of Spatial Goals in a Two-Choice Flexible Navigation Task. *J Neurosci* 39, 2522-2541.
- Dvorak, D., Chung, A., Park, E.H., and Fenton, A.A. (2021). Dentate spikes and external control of hippocampal function. *Cell Rep* 36, 109497.
- Dvorak, D., Radwan, B., Sparks, F.T., Talbot, Z.N., and Fenton, A.A. (2018). Control of recollection by slow gamma dominating mid-frequency gamma in hippocampus CA1. *PLoS Biol* 16, e2003354.
- Ebitz, R.B., and Hayden, B.Y. (2021). The population doctrine in cognitive neuroscience. *Neuron*.
- Farnebäck, G. (2003). Two-Frame Motion Estimation Based on Polynomial Expansion. (Berlin, Heidelberg, Springer Berlin Heidelberg), pp. 363-370.
- Fenton, A.A. (2015a). Coordinating with the "Inner GPS". *Hippocampus* 25, 763-769.
- Fenton, A.A. (2015b). Excitation-inhibition discoordination in rodent models of mental disorders. *Biol Psychiatry* 77, 1079-1088.
- Fenton, A.A., Kao, H.-Y., Neymotin, S.A., Olypher, A.V., Vayntrub, Y., Lytton, W.W., and Nandor, L. (2008). Unmasking the CA1 ensemble place code by exposures to small and large environments: more place cells and multiple, irregularly-arranged, and expanded place fields in the larger space. *J Neurosci* 28, 11250-11262.

- Fenton, A.A., Lytton, W.W., Barry, J.M., Lenck-Santini, P.P., Zinyuk, L.E., Kubik, S., Bures, J., Poucet, B., Muller, R.U., and Olypher, A.V. (2010). Attention-like modulation of hippocampus place cell discharge. *J Neurosci* 30, 4613-4625.
- Fenton, A.A., and Muller, R.U. (1998). Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proc Natl Acad Sci U S A* 95, 3182-3187.
- Frank, L.M., Stanley, G.B., and Brown, E.N. (2004). Hippocampal plasticity across multiple days of exposure to novel environments. *J Neurosci* 24, 7681-7689.
- Fusi, S., Miller, E.K., and Rigotti, M. (2016). Why neurons mix: high dimensionality for higher cognition. *Curr Opin Neurobiol* 37, 66-74.
- Gao, P., Trautmann, E., Yu, B., Santhanam, G., Ryu, S., Shenoy, K., and Ganguli, S. (2017). A theory of multineuronal dimensionality, dynamics and measurement. *bioRxiv*, 214262.
- Gardner, R.J., Hermansen, E., Pachitariu, M., Burak, Y., Baas, N.A., Dunn, B.A., Moser, M.-B., and Moser, E.I. (2022). Toroidal topology of population activity in grid cells. *Nature*.
- Ghrist, R. (2008). Barcodes: The persistent topology of data. *B Am Math Soc* 45, 61-75.
- Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Brown, B.L., Koay, S.A., Taxidis, J., Najafi, F., Gauthier, J.L., Zhou, P., *et al.* (2019). CalmAn an open source tool for scalable calcium imaging data analysis. *Elife* 8.
- Guzowski, J.F., Knierim, J.J., and Moser, E.I. (2004). Ensemble dynamics of hippocampal regions CA3 and CA1. *Neuron* 44, 581-584.
- Hampson, R.E., Simeral, J.D., and Deadwyler, S.A. (1999). Distribution of spatial and nonspatial information in dorsal hippocampus. *Nature* 402, 610-614.
- Hardcastle, K., Maheswaranathan, N., Ganguli, S., and Giocomo, L.M. (2017). A Multiplexed, Heterogeneous, and Adaptive Code for Navigation in Medial Entorhinal Cortex. *Neuron* 94, 375-387 e377.
- Harland, B., Contreras, M., Souder, M., and Fellous, J.M. (2021). Dorsal CA1 hippocampal place cells form a multi-scale representation of megaspace. *Curr Biol*.
- Harris, K.D., Csicsvari, J., Hirase, H., Dragoi, G., and Buzsaki, G. (2003). Organization of cell assemblies in the hippocampus. *Nature* 424, 552-556.
- Heeger, D.J., and Mackey, W.E. (2019). Oscillatory recurrent gated neural integrator circuits (ORGaNICs), a unifying theoretical framework for neural dynamics. *Proc Natl Acad Sci U S A* 116, 22783-22794.
- Hemelrijk, C.K., and Hildenbrandt, H. (2011). Some causes of the variable shape of flocks of birds. *PLoS One* 6, e22479.
- Hemelrijk, C.K., and Hildenbrandt, H. (2012). Schools of fish and flocks of birds: their shape and internal structure by self-organization. *Interface Focus* 2, 726-737.
- Hemelrijk, C.K., and Hildenbrandt, H. (2015). Scale-Free Correlations, Influential Neighbours and Speed Control in Flocks of Birds. *Journal of Statistical Physics* 158, 563-578.
- Hennequin, G., Ahmadian, Y., Rubin, D.B., Lengyel, M., and Miller, K.D. (2018). The Dynamical Regime of Sensory Cortex: Stable Dynamics around a Single Stimulus-Tuned Attractor Account for Patterns of Noise Variability. *Neuron* 98, 846-860.e845.
- Hirase, H., Czurko, A., Csicsvari, J., and Buzsaki, G. (1999). Firing rate and theta-phase coding by hippocampal pyramidal neurons during 'space clamping'. *Eur J Neurosci* 11, 4373-4380.
- Huxter, J., Burgess, N., and O'Keefe, J. (2003). Independent rate and temporal coding in hippocampal pyramidal cells. *Nature* 425, 828-832.

- Jackson, J., and Redish, A.D. (2007). Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks. *Hippocampus* 17, 1209-1229.
- Jazayeri, M., and Afraz, A. (2017). Navigating the Neural Space in Search of the Neural Code. *Neuron* 93, 1003-1014.
- Jeffery, K.J., Gilbert, A., Burton, S., and Strudwick, A. (2003). Preserved performance in a hippocampal-dependent spatial task despite complete place cell remapping. *Hippocampus* 13, 175-189.
- Kang, L., Xu, B., and Morozov, D. (2021). Evaluating State Space Discovery by Persistent Cohomology in the Spatial Representation System. *Frontiers in Computational Neuroscience* 15.
- Kao, H.Y., Dvorak, D., Park, E., Kenney, J., Kelemen, E., and Fenton, A.A. (2017). Phencyclidine discoordinates hippocampal network activity but not place fields. *J Neurosci* 37, 12031–12049.
- Kelemen, E., and Fenton, A.A. (2013). The organization of neuronal discharge on timescales of milliseconds and seconds is related to the spatial response properties of hippocampal neurons, 1 edn (New York: Springer-Verlag).
- Kelemen, E., and Fenton, A.A. (2016). Coordinating different representations in the hippocampus. *Neurobiol Learn Mem* 129, 50-59.
- Kentros, C. (2006). Hippocampal place cells: The “where” of episodic memory? *Hippocampus* 16, 743-754.
- Kentros, C., Hargreaves, E., Hawkins, R.D., Kandel, E.R., Shapiro, M., and Muller, R.V. (1998). Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science* 280, 2121-2126.
- Kubie, J.L., Levy, E.R.J., and Fenton, A.A. (2020). Is hippocampal remapping the physiological basis for context? *Hippocampus* 30, 851-864.
- Kubie, J.L., and Muller, R.U. (1991). Multiple representations in the hippocampus. *Hippocampus* 1, 240-242.
- Kuhn, H.W. (1955). The Hungarian Method for the assignment problem. *Naval Research Logistics Quarterly* 2, 83-97.
- Kullmann, D.M., and Lamsa, K.P. (2011). Interneurons go plastic. *Neuropharmacology* 60, 711.
- Lamsa, K., Heeroma, J.H., and Kullmann, D.M. (2005). Hebbian LTP in feed-forward inhibitory interneurons and the temporal fidelity of input discrimination. *Nat Neurosci* 8, 916-924.
- Lamsa, K.P., Kullmann, D.M., and Woodin, M.A. (2010). Spike-timing dependent plasticity in inhibitory circuits. *Front Synaptic Neurosci* 2, 8.
- Ledochowitsch, P., Huang, L., Knoblich, U., Oliver, M., Lecoq, J., Reid, C., Li, L., Zeng, H., Koch, C., Waters, J., *et al.* (2019). On the correspondence of electrical and optical physiology in *in vivo* population-scale two-photon calcium imaging. *bioRxiv*, 800102.
- Lee, D., Port, N.L., Kruse, W., and Georgopoulos, A.P. (1998). Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J Neurosci* 18, 1161-1170.
- Lee, J.S., Briguglio, J.J., Cohen, J.D., Romani, S., and Lee, A.K. (2020). The Statistical Structure of the Hippocampal Code for Space as a Function of Time, Context, and Value. *Cell* 183, 620-635 e622.

- Lenck-Santini, P.P., Save, E., and Poucet, B. (2001). Evidence for a relationship between place-cell spatial firing and spatial memory performance. *Hippocampus* 11, 377-390.
- Leutgeb, J.K., Leutgeb, S., Treves, A., Meyer, R., Barnes, C.A., McNaughton, B.L., Moser, M.B., and Moser, E.I. (2005a). Progressive transformation of hippocampal neuronal representations in "morphed" environments. *Neuron* 48, 345-358.
- Leutgeb, S., Leutgeb, J.K., Barnes, C.A., Moser, E.I., McNaughton, B.L., and Moser, M.B. (2005b). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* 309, 619-623.
- Leutgeb, S., Leutgeb, J.K., Moser, M.B., and Moser, E.I. (2005c). Place cells, spatial maps and the population code for memory. *Curr Opin Neurobiol* 15, 738-746.
- Lever, C., Wills, T., Cacucci, F., Burgess, N., and O'Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature* 416, 90-94.
- McHugh, T.J., and Tonegawa, S. (2007). Spatial exploration is required for the formation of contextual fear memory. *Behav Neurosci* 121, 335-339.
- McNaughton, B.L., Battaglia, F.P., Jensen, O., Moser, E.I., and Moser, M.B. (2006). Path integration and the neural basis of the 'cognitive map'. *Nat Rev Neurosci* 7, 663-678.
- Meshulam, L., Gauthier, J.L., Brody, C.D., Tank, D.W., and Bialek, W. (2017). Collective Behavior of Place and Non-place Neurons in the Hippocampal Network. *Neuron*.
- Milstein, A.D., Li, Y., Bittner, K.C., Grienberger, C., Soltesz, I., Magee, J.C., and Romani, S. (2021). Bidirectional synaptic plasticity rapidly modifies hippocampal representations. *Elife* 10.
- Mizumori, S.J.Y. (2006). Hippocampal place fields: A neural code for episodic memory? *Hippocampus* 16, 685-690.
- Moita, M.A., Rosis, S., Zhou, Y., LeDoux, J.E., and Blair, H.T. (2004). Putting fear in its place: remapping of hippocampal place cells during fear conditioning. *J Neurosci* 24, 7015-7023.
- Mongillo, G., Rumpel, S., and Loewenstein, Y. (2018). Inhibitory connectivity defines the realm of excitatory plasticity. *Nat Neurosci* 21, 1463-1470.
- Muller, R.U., and Kubie, J.L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J Neurosci* 7, 1951-1968.
- Muller, R.U., and Kubie, J.L. (1989). The firing of hippocampal place cells predicts the future position of freely moving rats. *J Neurosci* 9, 4101-4110.
- Muller, R.U., Kubie, J.L., and Ranck, J.B., Jr. (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *J Neurosci* 7, 1935-1950.
- Muzzio, I.A. (2018). Spatial Instability: The Paradox of Place Cell Remapping. *Curr Biol* 28, R1306-R1307.
- Muzzio, I.A., Levita, L., Kulkarni, J., Monaco, J., Kentros, C., Stead, M., Abbott, L.F., and Kandel, E.R. (2009). Attention enhances the retrieval and stability of visuospatial and olfactory representations in the dorsal hippocampus. *PLoS Biol* 7, e1000140.
- Nagele, J., Herz, A.V.M., and Stemmler, M.B. (2020). Untethered firing fields and intermittent silences: Why grid-cell discharge is so variable. *Hippocampus* n/a.
- Neymotin, S.A., Talbot, Z.N., Jung, J.Q., Fenton, A.A., and Lytton, W.W. (2017). Tracking recurrence of correlation structure in neuronal recordings. *J Neurosci Methods* 275, 1-9.
- Nieh, E.H., Schottorf, M., Freeman, N.W., Low, R.J., Lewallen, S., Koay, S.A., Pinto, L., Gauthier, J.L., Brody, C.D., and Tank, D.W. (2021). Geometry of abstract learned knowledge in the hippocampus. *Nature* 595, 80-84.

- O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Exp Neurol* 51, 78-109.
- O'Keefe, J. (1979). A review of the hippocampal place cells. *Prog Neurobiol* 13, 419-439.
- O'Keefe, J., and Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature* 381, 425-428.
- O'Neill, J., Senior, T.J., Allen, K., Huxter, J.R., and Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nat Neurosci*.
- Okun, M., and Lampl, I. (2008). Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat Neurosci* 11, 535-537.
- Okun, M., Steinmetz, N.A., Cossell, L., Iacarus, M.F., Ko, H., Bartho, P., Moore, T., Hofer, S.B., Mrsic-Flogel, T.D., Carandini, M., and Harris, K.D. (2015). Diverse coupling of neurons to populations in sensory cortex. *Nature* 521, 511-515.
- Olyphar, A.V., Klement, D., and Fenton, A.A. (2006). Cognitive disorganization in hippocampus: a physiological model of the disorganization in psychosis. *J Neurosci* 26, 158-168.
- Park, E., Dvorak, D., and Fenton, A.A. (2011). Ensemble Place Codes in Hippocampus: CA1, CA3, and Dentate Gyrus Place Cells Have Multiple Place Fields in Large Environments. *PLoS ONE* 6, e22349.
- Park, E.H., Keeley, S., Savin, C., Ranck, J.B., Jr., and Fenton, A.A. (2019). How the Internally Organized Direction Sense Is Used to Navigate. *Neuron* 101, 1-9.
- Pavlidis, C., Donishi, T., Ribeiro, S., Mello, C.V., Blanco, W., and Ogawa, S. (2019). Hippocampal functional organization: A microstructure of the place cell network encoding space. *Neurobiol Learn Mem* 161, 122-134.
- Peyrache, A., Lacroix, M.M., Petersen, P.C., and Buzsaki, G. (2015). Internally organized mechanisms of the head direction sense. *Nat Neurosci* 18, 569-575.
- Pnevmatikakis, E.A., and Giovannucci, A. (2017). NoRMCorre: An online algorithm for piecewise rigid motion correction of calcium imaging data. *J Neurosci Methods* 291, 83-94.
- Pnevmatikakis, E.A., Soudry, D., Gao, Y., Machado, T.A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., *et al.* (2016). Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. *Neuron* 89, 285-299.
- Poucet, B., Hok, V., Sargolini, F., and Save, E. (2012). Stability and variability of place cell activity during behavior: functional implications for dynamic coding of spatial information. *J Physiol Paris* 106, 62-71.
- Recanatesi, S., Ocker, G.K., Buice, M.A., and Shea-Brown, E. (2019). Dimensionality in recurrent spiking networks: Global trends in activity and local origins in connectivity. *PLOS Computational Biology* 15, e1006446.
- Reynolds, C.W. (1987). Flocks, herds and schools: A distributed behavioral model. In *SIGGRAPH, Computer Graphics (Anaheim)*, pp. 25-34.
- Rubin, A., Sheintuch, L., Brande-Eilat, N., Pinchasof, O., Rechavi, Y., Geva, N., and Ziv, Y. (2019). Revealing neural correlates of behavior without behavioral measurements. *Nat Commun* 10, 4745.
- Ruediger, S., Vittori, C., Bednarek, E., Genoud, C., Strata, P., Sacchetti, B., and Caroni, P. (2011). Learning-related feedforward inhibitory connectivity growth required for memory precision. *Nature* 473, 514-518.

- Rule, M.E., O'Leary, T., and Harvey, C.D. (2019). Causes and consequences of representational drift. *Curr Opin Neurobiol* 58, 141-147.
- Sarel, A., Finkelstein, A., Las, L., and Ulanovsky, N. (2017). Vectorial representation of spatial goals in the hippocampus of bats. *Science* 355, 176-180.
- Schneidman, E., Berry, M.J., 2nd, Segev, R., and Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440, 1007-1012.
- Schneidman, E., Bialek, W., and Berry, M.J., 2nd (2003). Synergy, redundancy, and independence in population codes. *J Neurosci* 23, 11539-11553.
- Skaggs, W.E., and McNaughton, B.L. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science* 271, 1870-1873.
- Skaggs, W.E., McNaughton, B.L., Gothard, K.M., and Markus, E.J. (1993). An information theoretic approach to deciphering the hippocampal code. In *Advances in neural information processing*, S.J. Hanson, J.D. Cowan, and C.L. Giles, eds. (San Mateo, CA: Morgan Kaufmann Publishers), pp. 1030-1037.
- Stefanini, F., Kushnir, L., Jimenez, J.C., Jennings, J.H., Woods, N.I., Stuber, G.D., Kheirbek, M.A., Hen, R., and Fusi, S. (2020). A Distributed Neural Code in the Dentate Gyrus and in CA1. *Neuron* 107, 703-716 e704.
- Talbot, Z.N., Sparks, F.T., Dvorak, D., Curran, B.M., Alarcon, J.M., and Fenton, A.A. (2018). Normal CA1 Place Fields but Discoordinated Network Discharge in a Fmr1-Null Mouse Model of Fragile X Syndrome. *Neuron* 97, 684-697.
- Tanaka, K.Z., He, H., Tomar, A., Niisato, K., Huang, A.J.Y., and McHugh, T.J. (2018). The hippocampal engram maps experience but not place. *Science* 361, 392-397.
- Tenenbaum, J.B., de Silva, V., and Langford, J.C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science* 290, 2319-2323.
- Tralie, C., Saul, N., and Bar-On, R. (2018). Ripser.py: A Lean Persistent Homology Library for Python. *Journal of Open Source Software* 3.
- Umakantha, A., Morina, R., Cowley, B.R., Snyder, A.C., Smith, M.A., and Yu, B.M. (2021). Bridging neuronal correlations and dimensionality reduction. *Neuron*.
- van Dijk, M.T., and Fenton, A.A. (2018). On How the Dentate Gyrus Contributes to Memory Discrimination. *Neuron* 98, 832-845.
- Wills, T.J., Lever, C., Cacucci, F., Burgess, N., and O'Keefe, J. (2005). Attractor dynamics in the hippocampal representation of the local environment. *Science* 308, 873-876.
- Wilson, M.A., and McNaughton, B.L. (1993). Dynamics of the hippocampal ensemble code for space. *Science* 261, 1055-1058.
- Yoon, K., Buice, M.A., Barry, C., Hayman, R., Burgess, N., and Fiete, I.R. (2013). Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nat Neurosci* 16, 1077-1084.
- Zhou, P., Resendez, S.L., Rodriguez-Romaguera, J., Jimenez, J.C., Neufeld, S.Q., Giovannucci, A., Friedrich, J., Pnevmatikakis, E.A., Stuber, G.D., Hen, R., *et al.* (2018). Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data. *Elife* 7.
- Ziv, Y., Burns, L.D., Cocker, E.D., Hamel, E.O., Ghosh, K.K., Kitch, L.J., El Gamal, A., and Schnitzer, M.J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nat Neurosci* 16, 264-266.

Acknowledgements

Supported by NIH grants R01MH115304. We are grateful to Dr. John Kubie for valuable comments on an earlier draft of the manuscript, and Kathryn McClain for the initial network simulations that validated the network consistency measure.

Author Contributions

A.A.F. and E.R.J.L. designed research; E.R.J.L., E.P. performed physiological research; W.R., and J.H. performed modeling research, E.R.J.L., S. C.-S., and A.A.F. analyzed data; A.A.F. and E.R.J.L. wrote the paper; A.A.F. supervised research.

Declaration of interest

The authors declare no competing interests.

Methods

Ethics Approval

All work with mice was approved by the New York University Animal Welfare Committee (UAWC) under Protocol ID: 17-1486.

Virus injection and lens implant

To virally infect principal cells to express the fluorescent calcium indicator GCaMP6f, adult C57BL/6J mice (n=41) were anesthetized with Nembutal (i.p. 50mg/kg), one hour after receiving dexamethasone (s.c. 0.1mg/kg). They were mounted in a stereotaxic frame (Kopf, Tujunga, CA) and through a small craniotomy, they were injected into the right CA1 subfield (AP: 2.1mm, DV: 1.65 mm, ML: 2 mm) at a rate of 4 nl/s with 0.5 μ l AAV1.Syn.GCaMP6f.WPRE.SV40 (titre: 4.65×10^{13} GC/ml; Penn Vector Core). The injection pipette (Nanoject III, Drummond) was left in place for 5 min before it was slowly withdrawn. Thirty minutes later, a larger craniotomy was performed with a 1.8-mm diameter trephine drill and the overlying cortex was removed by suction. A gradient-index 'GRIN' lens (Edmund Optics, 1.8 mm diameter, 0.25 pitch, 670 nm, uncoated) was implanted, fixed with cyanoacrylate, and protected by Kwik-Sil (WPI, Sarasota, FL). The skull was covered with dental cement (Grip Cement, Dentsply, Long Island City, NY). The mice received one slow-release buprenorphine injection (s.c. 0.5 mg/kg), dexamethasone (s.c. 0.1 mg/kg) for 6 days, and amoxicillin in water gel for a week. Three to eight weeks later, once a good field of view was visible through the GRIN lens, the baseplate for a UCLA miniscope (www.miniscope.org; Aharoni and Hoogland, 2019) was implanted and fixed to the skull with dental cement.

Behavior

After recovery from surgery, animals were handled for a couple minutes 1-3 times a week until the baseplate was implanted. The animals were then habituated to wearing the miniscope in their home cage, first in the animal holding room and then in the experimental room. Once the animal was comfortable wearing the miniscope, we started the behavior experiment and recording.

The mice were exposed to two environments, a 32 cm-diameter circle with transparent walls and a 28.5 cm-side square with opaque black walls and distinctive orienting pattern on 3 of the walls. The surface area of the two enclosures were similar, within 2% of 800 cm², but the floor of each environment was also distinctive; the circle's floor was made of red plastic while the rectangle's was of black metal. Orienting cues were also present in the room, and directly visible from the circular enclosure; the door and the animal transport cart were visible, and two salient cues were present on opposite walls of the room.

The mice explore each environment twice, for 5 min each time, in an interleaved fashion. The animal was placed in its home cage between trials for a couple minutes to allow us to change environments. Windows-based software (Tracker 2.36, Bio-Signal Group Corp., Acton, MA) determined the mouse's location in every 30-Hz video frame from an overhead digital video camera.

This protocol was repeated for three consecutive days, every week of three consecutive weeks. For two animals, this was repeated in periods of 5 days (instead of 7). For all animals, we started calcium recording on the day the animal first explored the two environments. We aimed to record the entirety of the 5-min visits to the circle and square environments and 5-min

sessions in the animal's home cage were also recorded each day before and after the 4 sessions in the two experimental environments.

Calcium recording and signal extraction

Neural activity data from 6 of the 41 mice met the quality analysis requirements described below and were analyzed to address the central question. When recording, the UCLA miniscope was attached to the baseplate, and fluorescent images from CA1 were recorded through the GRIN lens using the UCLA miniscope data acquisition hardware and software (www.miniscope.org).

Thirty images were collected each second and analyzed offline. Before analysis, we screened the average calcium level of each recording and removed any day or week where large and abrupt changes to the mean fluorescence were observed. The images from each recording were aligned using the NoRMCorre algorithm (Pnevmatikakis and Giovannucci, 2017). The alignment was done separately for each recording but all recording alignments from a single mouse used the template generated by the alignment of the first video recorded that day.

Aligned videos were then subsampled to 10 Hz by averaging and the recordings from each day were concatenated into a single file. Action potential activity was then estimated separately for each day from the $\Delta F/F$ GCaMP6f signal using the CNMF-E algorithm (Pnevmatikakis et al., 2016; Zhou et al., 2018), which simultaneously separates the cells' fluorescence and deconvolves the calcium transients to infer spiking activity. Accordingly, we refer to 'activity' rather than 'firing', especially because calcium transients are likely to reflect bursts of action potentials rather than single action potentials (Ledochowitsch et al., 2019). The spatial footprint of the cells were seeded using a peak-to-noise ratio determined manually for each recording session. Units with high spatial overlap (>0.65) and high temporal correlation (>0.4) were merged and the CNMF-E algorithm was updated until no pair of units meet the criteria.

Ensembles were evaluated manually for quality control to reduce the likelihood that artifacts

were identified as cells. Furthermore, recordings were examined for evidence of photobleaching and minutes scale downward or upward trends in activity by regression analysis of the activity timeseries of individual cells for every one of the ≤ 54 (cylinder, box, homecage) recordings (6/day X 3 days x 3 weeks). Only recordings with no significant linear trends and without abrupt decreases in activity levels were studied.

To evaluate the integrity of the neuron ensembles we recorded, we used two metrics: temporal independence and spatial dependence. Temporal independence evaluated for each cell the average cell-pair correlation with all cells whose center is within 20- μm of the soma of the cell. Spatial dependence measured the proximity and overlap with nearby cells. For each cell, the CNMF-E algorithm computes a spatial footprint. The spatial dependence is the proportion of the summed footprints within a 20- μm radius area around the soma that is not attributable to the cell itself. These metrics confirmed to the best of our abilities, that no corrupted unit or ensembles were included in the dataset. Ensembles ranged in size from 39 to 588 isolated units, with half larger than 260 cells and less than 10% smaller than 100 cells.

Alignment and cell matching

To match cells across days, every pair of recording days was first aligned to each other using the template used for the within-day alignment. This alignment was done using a non-rigid optical flow algorithm (Farneback, 2003), `calcOpticalFlowFarneback` function in the `cv2` python library) whose parameters (window size, flow levels, number of iterations) were optimized for each animal. Overlapping spatial footprints were then matched conservatively (matches were performed using the Hungarian algorithm (Kuhn, 1955) with a maximum distance considered or cost of 0.7). Alignment and cell-matching algorithms were applied as implemented in the `register_ROIs` function of the CalmAn project (Giovannucci et al., 2019). Day-to-day alignments

were verified by eye and for each alignment an f1-score was computed as the ratio of the number of cells aligned to the total number of cells. Alignments with an f1-score below 0.3 were discarded.

Place Cell Identification

We computed cell-specific spatial activity maps from each cell's activity, extracted at 10 Hz, and the animal's location, subsampled to 10 Hz, by computing the average rate of the cell in each ~2.5 x 2.5 cm bin. Linearized spatial activity maps were computed by changing the Cartesian coordinates to circular coordinates, with the origin at the center of the environment. We then computed the average rate of the cell in each 12° bin with the distance from the origin ignored. To identify place cells, we used two metrics: Spatial coherence and information content (Muller and Kubie, 1989; Skaggs et al., 1993). Spatial coherence was defined for each cell as the correlation between the average rate at each location and the average of the average rate at the locations neighboring this location (up to 8). Information content was defined for each cell as:

$$InfoContent = \sum_{positions\ x} p(x) * \frac{r_x}{R} * \log_2 \frac{r_x}{R}$$

where $p(x)$ is the probability of the animal being at position x , r_x the average rate at position x and R the average rate of the cell.

For each cell we computed both metrics as well as a distribution of randomized values with shuffled activity. The metric *val* was considered significant if

$$\frac{val - \text{mean } val_{random}}{\text{std } val_{random}} > 1.96$$

with val_{random} being the distribution of randomized values.

We classified a cell as a place cell only when both metrics were significant.

Place field similarity

The stability of a cell's spatial tuning was assessed by comparing the spatial activity map in a pair of recordings. The Pearson correlation was computed between the pairs of activity rates at corresponding pixels in the activity rate maps. The average correlation (r) is reported for summary statistics but for statistical comparisons $z = \frac{1}{2} \ln \frac{(1+r)}{(1-r)}$, Fisher's transform was used.

Cell pair correlations

Kendall tau (t) correlations were computed using 1-s time bins because this measure of association is robust to time series with a low range of values and many ties (e.g., 0 and 1). We evaluated cell pair correlations as a function of the distance between neurons and found that neurons within 25 μm of each other were, on average, more strongly correlated. Since our recording method does not allow us to determine whether this is physiological or an artefact, we decided to exclude all such pairs. Therefore, for all analyses using cell pair correlations, we excluded cells pairs that were not spatially separated (distance $< \sim 30\mu\text{m}$). We also excluded cells with very low activity (active less than 2% of 1-s time bins) as Kendall correlations computed on such time series have little meaning.

Position-tuning independent (PTI) rate

From the rate maps (average rate at each position), we computed the expected rate as the average rate at the animal's location / the sampling rate (here 100 ms). The expected rate is then binned into 1-s bins and subtracted from the 1-s binned observed rate.

Support vector machine (SVM) decoding

SVM analysis was performed on concatenated 1-s binned activity time series from both environments using the sklearn python library. Training and decoding are performed 100 times on each dataset using different randomly separated training set (2/3 samples) and a testing set (1/3 samples). Decoding accuracy and features (cells or cell pairs) weights are then averaged across the 100 repetitions. This ensured that the results were not sensitive to random variations in the dataset split. Decoding accuracy is computed as the number of right decisions in the testing set. When decoding and comparing weights across day, no cross validation was used, and the entire reference day dataset was used to train the decoder.

Network consistency

The Network consistency (van Dijk and Fenton, 2018) describes how well the momentary covariance in cell-pair activity fluctuations align to the overall activity correlation of the cell pair. Network consistency of an ensemble of time series is calculated as:

$$NetCo(t) = \sum_{i,j} \tau_{i,j} * r_i(t) * r_j(t)$$

Where $\tau_{i,j}$ is the pairwise Kendall correlation between cell i and cell j and $r_i(t)$ the activity rate of cell i at time t . To compare network consistency values between animals we normalized it in each recording to the standard deviation of the distribution of values obtained when the identities of the cell pairs were shuffled 100 times. In the figure, we report the average network consistency for each recording.

Correlation participation

To quantify the role of cell pair activity correlation in the SVM and IsoMap analyses, we evaluated for each cell their participation in correlations of different ranges of correlation t

values. For each cell, we computed its participation in a range of t values as the number of times a cell participates in a pair with a t value within that range, divided by the total number of pairs in which the cell participates. In effect, it is the proportion of t values falling in a certain range amongst all cell pairs the cell participate in. In practice, this normalization has limited effect but allows us to remove any impact from the exclusion of some cell pairs.

Betti Numbers and dimensionality reduction

Principal Component Analysis (PCA) and IsoMap dimensionality reduction were performed each day separately on concatenated 1-s binned activity time series from both environments as well as the home cage, filtered with a Gaussian filter of 1-s standard deviation. Both transformations used the sklearn python library. IsoMap was performed using Euclidian distance and 5 neighbors (Balasubramanian and Schwartz, 2002; Tenenbaum et al., 2000). Betti barcodes were computed each day on a 10-dimension IsoMap projection using the Ripser algorithm (Tralie et al., 2018) in the scikit python library. To compute isometries of Betti number 3, the data were randomly subsampled to 700 data points (~40%) to allow for computational tractability.

Neural manifold analysis

We start by considering a N -dimensional neural manifold. The dimensionality of the neural manifold is evaluated using the participation ratio (PR), defined as:

$$PR = \frac{[Tr C_{neuron}]^2}{Tr [C_{neuron}]^2}; C_{neuron} = \frac{1}{N_T} XX^T$$

where each column of X represents a pattern of firing rates across M recorded neurons at some specified time, and N_T is equal to the total number of sampled time points.

To quantify the anti-cofiring properties of each cell, we introduce the anti-cofiring power, defined as the number of cell pair correlations with $\tau < -0.05$, normalized with the total number of cell pairs.

Further, to describe the distribution of distances to the centroid in a N-dimensional manifold, we employ two different measures. First, we use the skewness as a measure of the asymmetry of the data around the sample mean, i.e., a measure of the tail-ness of the distribution, which contains the ‘outliers’, earlier shown to be most anti-cofiring cells, defined as:

$$skewness = \frac{\frac{1}{n} \sum (x_i - \bar{x})^3}{\left(\sqrt{\frac{1}{n} \sum (x_i - \bar{x})^2} \right)^3}$$

Second, to measure the similarity between distributions for different ensembles, we use Kullback-Liebler divergence (KLD). KLD is a measure of how one probability distribution f_1 is different from a second, reference probability distribution f_2 . This measure can be interpreted as the expected excess surprise from using f_1 as a model when the actual distribution is f_2 .

Then, we construct 2D neural manifolds using IsoMap projected neural activity into a three-dimensional subspace. The 2D manifolds are constructed using the in-built function *alphaShape* in MATLAB software. The algorithm creates a 3D bounding volume using a triangulation algorithm. To identify the overlapping coefficient between two environments, we first identify the number of points from one environment manifold that belongs to the other environment manifold, using the in-built command *inShape*. We then compute the overlapping coefficient (overlap) by normalizing the number of occurrences from Environment 1 inside the 2D manifold of Environment 2, with the total number of points from Environment 1.

To explore the multi-stable planar manifold hypothesis, we compute the best fit plane to the 3D IsoMap projected point cloud data for each environment using an in-house MATLAB function.

The algorithm fits a least square solution for the normal vector that defines the best fit plane.

The algorithm minimizes the sum of the squared dot product of a trial normal vector with a vector passing by the centroid of the data and each data point. The output is a normal vector and the centroid of the 3D point cloud data.

To further quantify the distinction between environments, we compare the list of normal vectors defining subsequent planar fits for a given environment and compared it against arrays of normal vectors from same and different environments using Kuiper's statistic, a rotation-invariant Kolmogorov-type test statistic.

Finally, to describe the trajectories inside the neural manifold, we introduce two quantities: localness, defined as the distance between two temporally adjacent points, and smoothness, defined as the angle difference between the normal vectors of two temporally adjacent planar fits.

Statistics and visualization

All data were plotted using boxplots unless the distribution contains less than 10 data points, in which case individual data points are plotted. For all boxplots, the height is determined by the interquartile range and the median is indicated with a line. The whiskers extend to the furthest data point up to 1.5 times the interquartile range past the boxplot limit. Data points past the whisker limits are either plotted individually or omitted when too large for better visual representation.

Statistically significant differences are indicated on the plots. One asterisk indicates $p < 0.05$ (or the corresponding corrected value) and two asterisks indicates $p < 0.01$ (or corresponding corrected value). When indicated, a statistically significant effect of the factor week is indicated as a line above the graph with asterisk(s).

Within subject comparisons were performed, when possible, using paired statistics such as Student's paired t test, and the corresponding non-parametric test, Wilcoxon test. Comparisons amongst multiple factors or amongst more than two groups were performed using ANOVA, with repeated measures (RM) when the effect of week was being evaluated. The Hotelling-Lawley correction was used when the data being compared by RM ANOVA violated the sphericity assumption as assessed by Mauchly's Test of Sphericity. For analyses in which a particular percentage subset of the data was included or excluded, the factor of percentage was treated as a continuous variable with $df = 1$. Dunnett's tests were used to evaluate post-hoc pairwise comparisons against a control, such as an initial week or control subset, when appropriate. Bonferroni-corrected multiple comparisons were used when only a subset of pairwise comparisons were of interest and when comparing multiple groups against a value or control group (if not present in the initial RM ANOVA). For comparison amongst multiple factors or amongst more than two groups where the data was not normally-distributed we performed Friedman's statistical test and Kruskal Wallis test.

All correlation data were Fisher-z transformed before statistical analysis using parametric methods. The test statistics and degrees of freedom are presented along with the p values and effect sizes for all comparisons, with only power of ten indicated when a value is below 0.001. A value of $p < 0.05$ was considered significant. The statistical analyses of data in a figure are presented in the corresponding figure legend.

Large Environment model

Larger place field maps were created from the concatenations of the place field map of four different cells. Only cells that were categorized as place cells in at least one repetition of the two environments were considered. For each larger map, we obtained four firing field maps, one for

each visit of the two environments. Cells from a single day (day 6) and from all six animals were used to create the 151- 'composite' cell ensemble. Results did not significantly change across days.

A random walk was then simulated across the larger environment but limited to different boundaries depending on the environment size evaluated. To simulate the random walk, we started at a random position within the allowed boundaries and randomly sampled a move every $dt=0.1s$. The size of the large place field space was considered to be 24×24 space units.

Position change was sampled randomly from $[-1, 0, 1]$ and multiplied by a speed variable. The speed variable was initiated at 0.1 and, similarly to the position, was changed at each time step by a random sample from $[-0.1, 0, 0.1]$, with 0.1 and 1 being the speed upper and lower bounds respectively. This corresponds to a maximum speed of 2.5 to 25 cm/s in real space.

Four different random walk were created for the two visits to the two environments and expected activity was created for each place maps as described in the *Position-tuning independent (PTI) rate* method section. Environments were then decoded as described in the *Support vector machine (SVM) decoding* method section.

Spike-time dependent plasticity network model and analyses

The code implementing the spiking neural network model was custom written and has been made freely available (<https://github.com/william-redman/EI-STDP-place-code-network>).

Architecture:

As illustrated in Fig. 1B, the spiking neural network was organized as follows: The input to the network came from a layer of $n_{\text{input}} = 1000$ neurons that were tuned to locations on the track.

This population fed into a layer of $n_{\text{excit}} = 500$ excitatory neurons with weights W^{InputE} . The

connection between the input and excitatory populations were non-plastic. The excitatory population was both recurrently connected to itself (W^{EE}), and to an inhibitory population of $n_{\text{inhib}} = 50$ neurons (W^{EI}). This inhibitory population is connected back to the excitatory layer (W^{IE}). The $E \rightarrow E$, $E \rightarrow I$, and $I \rightarrow E$ weights were all plastic.

At initialization, all values of $W^{\text{Input}E}$ were uniformly sampled from the interval $(0, W_{\text{max}}^{\text{Input}E})$. For W^{EE} , each weight either took a uniformly sampled value from the interval $(0, W_{\text{max}}^{EE})$, with probability $1 - p_{EE}$, or was set to 0, with probability p_{EE} . The diagonal entries of W^{EE} were all set and kept at 0. To ensure that the total amount of initial synaptic weight was conserved across the different weight populations, we set $W_{\text{max}}^{EE} = W_{\text{max}}^{\text{Input}E} / p_{EE}$. The same initialization procedure was repeated for W^{EI} and W^{IE} with $W_{\text{max}}^{EI} = W_{\text{max}}^{\text{Input}E} / p_{EI}$ and $W_{\text{max}}^{IE} = W_{\text{max}}^{\text{Input}E} / p_{IE}$.

Input layer:

The input layer to the spiking neural network was made up of n_{input} neurons, each tuned to at least one position on the linear track. The distribution of number of locations each input neuron was tuned to was binomial, with $p_{\text{input}} = 0.20$. The centers of these tunings took a continuous value from $[0, 10)$. The tuning to each center was taken to be Gaussian. That is, for neuron i with center, c_i , the tuning takes the form

$$T(|x - c_i|) = \exp\left(\frac{-|x - c_i|^2}{\sigma^2}\right)$$

where $| \cdot |$ is circular distance and x is a given location on the track. Neurons with $n > 1$ center take the sum of their tuning, which then takes the form

$$T_i(x) = \sum_{j=1}^n T(|x - c_i^{(j)}|) = \sum_{j=1}^n \exp\left(\frac{-|x - c_i^{(j)}|^2}{\sigma^2}\right) .$$

The probability of input neuron i firing when at position x is proportional to the input with added noise such that, at any given time point, it is given by

$$P_i(x) = [T_i(x) + \alpha \cdot \epsilon_{pink}] \cdot \Delta t \cdot r_{PF}$$

where Δt is the time step length of the simulation, r_{PF} is the “in-field” firing rate, and ϵ_{pink} is a noise term that is drawn from a pink noise distribution (<https://github.com/cortex-lab/MATLAB-tools/blob/master/pinknoise.m>) and is weighted by α . In the case of a negative noise term that is larger than $T_i(x)$, $P_i(x)$ was set to 0.

Neuron model:

The neurons in the excitatory and inhibitory populations were modeled as simplified integrate-and-fire neurons. That is, the voltage of excitatory neuron i at time t is given by the weighted sum

$$V_i^E(t) = V_i^E(t-1) \cdot \exp(-\Delta t) + f^{input} \cdot W_i^{InputE} + f^E \cdot W_i^{EE} - f^I \cdot W_i^{IE} + \alpha \cdot \epsilon_{pink}$$

where f^{input} is the spikes of the input population and W_i^{InputE} are the weights from the input layer to neuron i (and similarly for E and I being the excitatory and inhibitory populations, respectively). As in the case of the input layer, ϵ_{pink} is a noise term that is drawn from a pink noise distribution, which is weighted by α . When the voltage is greater than the threshold T_{IF} , the voltage is set to zero and f_i^E is set to 1 for the next time step. We include an absolute refractory period, τ_{IF} , where V_i^E is kept at 0.

For inhibitory neuron i , the voltage is given by

$$V_i^I(t) = V_i^I(t-1) \cdot \exp(-\Delta t) + f^E \cdot W_i^{EI} + \alpha \cdot \epsilon_{pink}.$$

STDP Plasticity:

All E→E, E→I, and I→E weights could be modified via spike-timing dependent plasticity (STDP)

rules. In the experiments where one (or all) of the plasticity types were removed, then those weights were not subject to STDP, but were instead frozen.

The recurrent excitatory weights were modified at each time step via excitatory STDP that has the following form

$$\Delta W_{ij}^{EE}(t) = \eta \left(f_i^E(t) \cdot v_j^E(t) - f_j^E(t) \cdot v_i^E(t) \right),$$

$$v_i^E(t) = v_i^E(t - \Delta t) \exp(-\Delta t) + \delta \left(1 - f_i^E(t) \right),$$

where the convention that i is the post-synaptic neuron and j is the pre-synaptic neuron is followed. $f_i^E(t)$ is the spiking activity of the excitatory neuron i at time t , η is the learning rate (and maximum amount of change allowed at each time step), and δ is the Kronecker delta function. The second equation keeps track of the spiking history of each neuron.

The weights from the excitatory population to the inhibitory population were similarly modified

$$\Delta W_{ij}^{EI}(t) = \eta \left(f_i^I(t) \cdot v_j^E(t) - f_j^E(t) \cdot v_i^I(t) \right).$$

Finally, the weights that projected from the inhibitory population to the excitatory population were modified at each time step via inhibitory STDP that has the following form

$$\Delta W_{ij}^{IE}(t) = \eta \left(f_i^E(t) \cdot v_j^I(t) - f_j^I(t) \cdot v_i^E(t) - 0.01 \right).$$

As negative weights do not make physical sense in our implementation of the spiking neural network, at each time step we enforced the condition $(W < 0) = 0$ for all weights. Similarly, we enforced that $(W > W_{max}) = W_{max}$.

Single track experiment metrics:

To evaluate the ability of a network with a given set of plasticity conditions, we ran 10 networks, each with a different set of random initial weights and different tuning centers for the input neurons, for 100 laps. To get a qualitative understanding, we split the track into 10 discrete locations and computed the spatial activity rate maps from the experiments using the last 25 laps (examples shown in Fig. 1C top). To make the observations we derived from the rate maps quantitative, we also computed two different similarity metrics for each excitatory neuron every 10 laps (using *only* the activity from those 10 laps). Both have the form

$$S_i = 1 - \frac{\bar{a}_i^{max} - \bar{a}_i^{control}}{\bar{a}_i^{max} + \bar{a}_i^{control}}$$

where \bar{a}_i^{max} is the average rate at the location with maximal activity and $\bar{a}_i^{control}$ is the average rate for a control location. For nearest neighbor similarity, the control locations are the positions on either side of the maximal location (circularly for boundary locations). For halfway similarity, we take the position on the track that is (circularly) the farthest from the location of maximal activity. The mean and standard error of the mean (SEM) for each of these metrics are plotted in Fig. 1C bottom.

Two track experiments:

To investigate how the spiking neural network was affected by exposure to a novel environment, and to see what role plasticity played in the encoding of the novel environment, we initialized

and ran 10 networks under different plasticity conditions. For each experiment, we created two tracks, track A and track B, defined by their random input neuron tunings (i.e., T_i), with the weights that connected the input to excitatory neurons (i.e., W^{InputE}) kept constant for both tracks. We trained a naïve network on track A for 30 laps (the number of laps needed for the similarity metrics to reach a plateau), then transferred that network (that is, used the “learned” W^{EE} , W^{EI} , and W^{IE}) to either track A again, called A’, or to track B to compare the activity obtained from the re-exposure to the “same” environment, track A, to the exposure to a “different” environment, track B. This investigated how much of the difference between the activity on track A and track B was due to noise and the stochastic nature of the spiking network.

To examine the effects of plasticity and novelty quantitatively, we computed two metrics. The first was the Pearson correlation of the rate maps computed on the last 50% (15 laps) of the activity (Fig. 1A left). We call this place field similarity, and the mean and SEM of the 10 simulations are plotted in Fig. 1D, top, for each plasticity condition. The second metric is the Pearson correlation of the Kendall correlations computed on each excitatory neuron cell pairs, again computed on the last 50% of the activity (15 laps) (Fig. 1A right). We call this cofiring similarity, and the mean and SEM of the 10 simulations are plotted in Fig. 1D, bottom, for each plasticity condition. To compute Kendall correlations, we first split the activity into bins of length $\Delta t_{\text{Kendall}} = 0.25$ and then averaged within each bin.

Parameters:

The parameters used for the simulations are listed below. For more details on the implementation of the spiking network, see the freely available commented code (<https://github.com/william-redman/EI-STDP-place-code-network>).

Parameter	Value	Description
n_{input}	1000	# input neurons
n_{excit}	500	# excitatory neurons
n_{inhib}	50	# inhibitory neurons
p_{input}	0.20	parameter of binomial distribution for number of tuning centers of input neuron
p_{InputE}	1.0	parameter of binomial distribution for number of excitatory neurons each input neuron is connected to
p_{EE}	0.25	parameter of binomial distribution for number of excitatory neurons each excitatory neuron is connected to
p_{EI}	0.175	parameter of binomial distribution for number of inhibitory neurons each excitatory neuron is connected to
p_{IE}	0.30	parameter of binomial distribution for number of excitatory neurons each inhibitory neuron is connected to
η	0.0005	synaptic learning rate
α	0.05	strength of pink noise
Δt	0.007	time scale of simulation
r_{PF}	8	firing rate at tuning center
σ^2	0.1	width of tuning center
τ_{IF}	0.02	absolute refractory period of excitatory and inhibitory neurons
T_{IF}	1	spiking voltage threshold
$\Delta t_{\text{Kendall}}$	0.25	size of bin used to average the activity before Kendall correlation
$W_{\text{max}}^{\text{InputE}}$	0.05	maximum initial weight between input neurons and excitatory neurons
W_{max}	0.50	maximum value any weight can take

Histology

Mice were deeply anesthetized with Nembutal (100 mg/kg, i.p.) and perfused through the heart with cold saline followed by cold 4% paraformaldehyde (PFA) in PBS. The brains were

removed, postfixed overnight in 4% PFA, and then cryoprotected in 30% sucrose for a minimum of 3 days. The brains were sectioned at 30–50 μm , stained with DAPI, and examined with a fluorescence microscope.