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

Eliott R.J. Levy[1], Eun Hye Park[1], William T. Redman[2], André A. Fenton[1,3*]

[1]Center for Neural Science, New York University, New York, NY 10003, USA
[2] Interdepartmental Graduate Program in Dynamical Neuroscience, University of California, Santa Barbara, Santa Barbara, CA 93106, USA.
[3]Neuroscience 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 CA1 place cells express a spatial neural code by discharging action potentials in cell-specific locations ('place fields'), but their discharge timing is also coordinated by multiple mechanisms, suggesting an alternative 'ensemble cofiring' neural code, potentially distinct from place fields. We compare the importance of these distinct information representation schemes for encoding environments. Using miniature microscopes, we recorded the ensemble activity of mouse CA1 principal neurons expressing GCaMP6f across a multi-week experience of two distinct environments. We find that both place fields and ensemble coactivity relationships are similarly reliable within environments and distinctive between environments. Decoding the environment from cell-pair coactivity relationships is effective and improves after removing cell-specific place tuning. Ensemble decoding relies most crucially on anti-coactive cell pairs distributed across CA1 and is independent of place cell firing fields. We conclude that ensemble cofiring relationships constitute an advantageous neural code for environmental space, independent of place fields.

Keywords:
Hippocampus, place cell, coactivity, synapsemble, environment, context

Introduction

The hippocampus is crucial for spatial navigation as well as memory, and although hippocampal place cells are tuned to discharge at specific locations within an environment called place fields, 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, which suggests the ensemble pattern of activity is unique and environment-specific, constituting a place 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). Because both memory and environmental context representations rely on the hippocampus, it is standardly hypothesized 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 memories that include information about environments (Kentros, 2006; Mizumori, 2006). 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) because almost all place cells have multiple place fields in environments bigger than 1m^2 (Fenton et al., 2008; Harland et al., 2021; Muller and Kubie, 1987), because discharge in firing fields is extremely variable during the 1-5 seconds it takes to cross the place field (Fenton et al., 2010; Fenton and Muller, 1998; Jackson and Redish, 2007), because the rate in a place field varies systematically across behavioral episodes (Leutgeb et al., 2005b), and because only a minority of place cells have stable place fields across days in familiar environments (Ziv et al., 2013). Indeed, having multiple place fields can degrade the ability to decode environments from ensemble discharge (Fig. S1). Studies to support the hypothesis that remapping is related to memory have been undecisive (Duvelle et al., 2019; Jeffery et al., 2003; Leutgeb et al., 2005a; van Dijk and Fenton, 2018).

One possible reason for the inability to relate remapping to memory may be how the field has assumed that 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 typically average the activity of individual cells over minutes, discarding its discharge in time and cofiring relationships to other cells to extract the cell's relationship to a certain variable, such as its location-specific tuning. This approach, a dedicated-rate "place field" hypothesis assumes the cell's momentary firing rate independently carries information that can be adequately extracted from

analysis of each cell's tuning, and that the firing relations to other neurons are 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 cofounded. 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. S2; Fenton, 2015a; Harris et al., 2003; Kelemen and Fenton, 2013). This motivates an independent 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 (each dimension being a neuron's activity) neural population activity typically manifests as if constrained to a manifold, a low-dimensional sub-space defined by recurring across-cell patterns of coactivity (Chaudhuri et al., 2019; Jazayeri and Afraz, 2017; McNaughton et al., 2006; Park et al., 2019; Peyrache et al., 2015; Rubin et al., 2019; Umakantha et al., 2021). Here we repeated a standard place cell experimental procedure to evaluate the dedicated place field and ensemble cofiring hypotheses for representing distinct environments in the hippocampus.

Results

Figure 1. Two neural coding hypotheses for representing environments (A) Illustration of hypothesized (left) dedicated place field and (right) ensemble cofiring codes for environment. (B) Schematic of neural network encoding of circular environments with STDP learning rules. (C, top) Learned network representation of the circular environment after (black) enabling or (grey) disabling STDP learning at cell-type specific connections, (middle) as measured on time-averaged rate maps by (bottom) activity similarity. 'Nearest neighbor' measures the difference between the peak activity and the activity at the closest positions to the location of the peak; 'halfway' measures the difference between the peak activity and the activity at the position halfway across track. (D) Evaluation of (top) dedicated place field and (bottom) ensemble cofiring codes in two different 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 but is most sensitive to inhibitory plasticity for the ensemble cofiring code.

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

We formalized the place field and cofiring hypotheses by examining a neural network model designed to evaluate the contributions of place tuning (Fig. 1A, left) and cofiring (Fig. 1A, right) to representing environments. The recurrently connected excitatory-inhibitory (E-I) network receives excitatory input from units tuned to randomly sampled positions of a circular track (Fig. 1B, left). The recurrent network weights change according to standard spike-timing dependent plasticity (STDP) rules (Fig. 1B, right) that are sufficient to generate place-selective activity (place fields, Fig. 1C,

middle) in the recurrent network after ~30 laps of experience on the track (Fig. 1C, bottom). This could be assessed by the similarity of the place fields (Fig. 1C), as well as the ability to decode current position from momentary network activity (data not shown). Remarkably, E→I and I→E plasticity are necessary for developing the place tuning, but E→E plasticity is not (Fig. 1C), which only illustrates that tunable excitation-inhibition coordination can be important for network and memory function (Caroni, 2015; Dvorak et al., in press; Fenton, 2015b; Mongillo et al., 2018; van Dijk and Fenton, 2018). To simulate 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 an experience of 30 laps. As expected, the place fields of the network units changed, they 'remapped,' on the new track (Fig. 1D, top). In contrast, the cofiring relationships amongst network cell pairs changes with time but only when I→E plasticity is enabled, irrespective of whether the position-to-input mapping is the same or different from the initial mapping, indicating that I→E plasticity continually alters network cofiring (Fig. 1D, bottom). This apparent 'remapping' in the cofiring code is thus independent of the spatial input but sensitive to changes in the I→E weights. This network simulation demonstrates that a place field code and a cofiring code could be distinct and operate in parallel while highlighting the potential importance of experience-dependent changes in network inhibition (Mongillo et al., 2018).

Figure 2. Place cell 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}, $\eta^2 = 0.48$; post-

hoc Dunnett's tests: -1 day > -2, -1 week(s); Future: Time $F_{3,187}$ = 7.31, p = 10^{-4}, η^2 = 0.12; post-hoc Dunnett's tests: 1 day > -2 weeks). (D) Activity time series in two environments illustrated by a 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 (left) and information content (right) are indicated above each map, red indicating significance; average rate map stability (r) is indicated between maps. (F) Linearized ensemble activity rate maps across three weeks, 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}, η^2 = 0.085; post-hoc 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 population. (I) Place map similarity within day improves across weeks (Week: $F_{2,28}$ = 20.02, p = 10^{-6}, η^2 = 0.051; Env. Change: $F_{2,29}$ = 80.55, p = 10^{-10}, η^2=0.65; Interaction: $F_{2,28}$ = 4.83, p = 0.016, η^2=0.026; post-hoc Dunnett's tests: week 1 < weeks 2, 3 for similar environments; post-hoc: Same vs. Different is significant each week, **p's =< 10^{-5}). (J) Place map similarity degrades across days (Time: $F_{4,132}$ = 12.68, p = 10^{-9}, η^2=0.11; Env. Change: $F_{1,132}$ = 228.37, p = 10^{-30}, η^2=0.53; Interaction: $F_{4,132}$ = 9.32, p = 10^{-6}, η^2=0.12; post-hoc Dunnett's tests: day 0 > 7-14 days for similar environments; post-hoc: Same vs. Different is significant at each time interval, **p's =< 10^{-3}).

Calcium imaging in CA1

To examine the validity of dedicated place field and ensemble cofiring hypotheses, we infected mouse CA1 principal cells with GCaMP6f (Chen et al., 2013) and, 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 cells could be identified two weeks apart (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. 1E) that were stable within day and, as previously reported (Ziv et al., 2013), unstable across days (Fig. 2F). 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 they remapped between trials of different environments (Fig. 2H, left). This remapping is not population wide as the measurable difference vanishes when all recorded cells are considered 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 (Ziv et al., 2013). These measurements demonstrate the standard mouse place cell phenomena in the data set.

Figure 3. Ensemble cofiring distinguishes the 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, η^2 = 0.018; Env. Change: $F_{1,3285501}$ = 9.82, p = 0.0017, η^2 = 10^{-6}; Interaction: $F_{6,3285501}$ = 3395.58, p = 0.00, η^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, which is confirmed by (top inset) the average signed difference for each animal. (B) Illustration of Population Coordination (PCo) computation from cofiring distributions. (C, left) Within-day coactivity stability improves across weeks (Week: $F_{2,27}$ = 4.23, p = 0.025, η^2=0.015; Env. Change: $F_{1,28}$ = 92.32, p = 10^{-10}, η^2=0.56; Interaction: $F_{2,27}$ = 0.99, p = 0.38, η^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: $F_{4,132}$ = 16.79, p = 10^{-11}, η^2=0.21; Env. Change: $F_{1,132}$ = 94.66, p = 10^{-17}, η^2=0.32; Interaction: $F_{4,132}$ = 4.04, p = 0.004, η^2=0.072; post-hoc Dunnett's tests vs day 0, *p<0.0058, **p=<10^{-5}; post-hoc: Same vs. Different is significant at each time interval, **p =< 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 figure 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}, η^2 = 0.84; Functional Cell Class: $F_{1,30}$ = 56.74, p = 10^{-8}, η^2 = 0.076; Interaction: $F_{2,30}$ = 14.72, p = 10^{-5}, η^2 = 0.039; post-hoc: PC vs. non-PC for Rate and Rate vs. PTI for Place Cells significant, **p =< 10^{-4}). (F, left) Within-day PTI coactivity stability across weeks (Week: $F_{2,27}$ = 2.28, p = 0.12, η^2 = 0.009; Env. Change: $F_{1,28}$ = 46.67, p = 10^{-7}, η^2 = 0.52; Interaction: $F_{2,27}$ = 2.38, p = 0.11, η^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: $F_{1,132}$ = 5.23, p = 10^{-4}, η^2 = 0.088; Env.

Change $F_{1,132} = 66.57$, $p = 10^{-13}$, $\eta^2 = 0.33$; Interaction: $F_{4,132} = 1.69$, $p = 0.16$, $\eta^2 = 0.079$; post-hoc Dunnett's tests: day 0 > days 7-14 for similar environments, **$p = 0.0050$; Same vs. Different significant at each time interval, **p's < 0.008).

Ensemble cofiring measurements of remapping between environments

Next, we examined remapping without tracking location, under the assumption that, in a fundamental way, the neural elements also are ignorant of location; neurons only have knowledge of the inputs they receive from other neurons. To evaluate how distinct environments can be encoded in the network dynamics, we measured the correlations of cell pairs, as the simplest technically feasible proxy for measuring 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 correlation (τ), some of which are distinct in the two environments (Fig. 3A). To evaluate the stability of these network dynamics across conditions, we measured the population coordination (PCo) between trials. PCo measures the similarity of all tau values and is computed as the Pearson correlation between two vectors of pairwise tau values (Fig. 2A; 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 also discriminates environments. The individual cell pairs that differ the most between environments are the positively ($\tau > 0.3$) or negatively ($\tau < -0.1$) coactive cell pairs (Fig. 3A). PCo increases with experience (Fig. 3C, left) and deteriorates across days (Fig. 3C, right), similar to 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. 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 from the cell's observed rate the expected rate, computed from the position of the animal and the session-averaged positional tuning measured by the firing-rate map, as schematized in Fig. 3D (Fenton et al., 2010; Fenton and Muller, 1998). As a direct consequence, place fields computed from PTI-rate disappear (Fig. 3D, right). Yet, 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 (Fig. S2). Cell-pair activity rate correlations covary with spatial firing similarity for both place cells and non-place cells. For PTI-rate, cell-pair correlations covary with spatial firing similarity for both place cells and non-place cells but at the same level as calculated with activity rate 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 3 (Fig. 3F, left). Similar to rate-computed PCo (Fig. 3C, right), PTI-PCo also degrades over time (Fig. 3F, right). These findings indicate that information contained in the cell's position tuning is not necessary for discriminating two environments from the activity of CA1 principal cells, including place cells.

Figure 4. Decoding environments from PTI coactivity. (A) Illustration of SVM decoding from PTI coactivity; (top) PTI coactivity, computed every 60s, 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, top) Cross-validated decoding accuracy is very high for all ensembles and (bottom) for the first 3 cell-pair deciles. (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 cell present in at least one cell pair in each SVM-ordered cell pair decile. (F) Difference between random and first decile normalized distribution of individual cell participation in decile subset (example distributions in inset and shown in purple in main figure) indicates overrepresentation of some cell in the first decile. (G) PTI network coherence normalized by the standard deviation of the cell-pair randomized distribution (Week: $F_{2,60}$ = 15.76, p=10^{-6}, η^2 = 0.094; Dunnett's tests: week 0 < weeks 2, 3).

Decoding environments from position-tuning independent ensemble coactivity

After showing that the metric can discriminate environments, we investigated whether the current environment could be reliably decoded from the PTI ensemble coactivity. 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), and to investigate which cell pairs are discriminative we sorted the cell pairs into deciles according to the SVM-assigned weights (Fig. 4A, bottom). The decoder correctly determines the current environment almost always (Fig. 4B,

top) demonstrating that the coactivity itself carries discriminative information, independent of position tuning, and find that the first three decile are responsible for the decoder's success (Fig. 4B, bottom). The most discriminative 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 co-fluctuations beyond positional tuning indicate that population correlation dynamics can be described as scale-free. Reminiscent of the flocking and schooling behavior of birds and fish, the discriminative activity of the neural population is driven by a subset of the individual cells but impacts the whole population (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 a cell pair's PTI-PCo and the pair's momentary activity co-fluctuations. Network consistency increases with experience (Fig. 4G) indicating increased alignment of short- and long-timescale 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. S3, 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 week 1 and 3, p=<0.004; p = 0.027 on week 2), consistent with the possibility that positional tuning measured by firing fields carries a separate type of information than the environment-specific information that is signaled by the coactivity of cells (Huxter et al., 2003).

Figure 5. Decoding of 1-s activity vectors relies on population subset and highlight 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, η^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) and 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, η^2 = 0.011; no effect of Day: $F_{1,10}$ = 1.71, p=0.22, η^2 = 0.08; Interaction: $F_{2,9}$ = 8.03, p=0.01, η^2 = 0.14; post-hoc Dunnett's tests: week 0 < week 2 on day 1; post-hoc t tests: day 1 vs day 9, *p<0.0167) (D) Cells are ordered by SVM absolute weights and separated in quintiles (20%) whose cross-validated performance is measured independently. (Quintile: $F_{4,65}$ = 38.58, p = 10^-16, η^2 = 0.65; Week: $F_{2,64}$ = 29.86, p = 10^-10, η^2 = 0.022; Interaction: $F_{8,89.135}$ = 3.69, p = 10^-4, η^2 = 0.026; post-hoc t tests: All Cells vs Quintiles, **p=<0.001; post-hoc 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 are compared to the properties of the population to evaluate how cells from the top quintile differ from the population. (F, left) Ratio of proportion of cells classified as place cell in either environment in top quintile versus overall population (No effect of Week: $F_{2,12}$ = 3.58, p = 0.061, η^2 = 0.13; t test vs. 1: t_{48} = 1.38, p = 0.18) and (right) ratio of proportion of cells with high or low rate of activation in top quintile versus overall population (High Rate: No effect of Week: $F_{2,12}$ = 0.65, p = 0.54, η^2 = 0.015; t test vs. 1: t_{48} = -3.19, p = 0.0025; Low Rate: No effect of Week: $F_{2,12}$ = 3.47, p = 0.065, η^2 = 0.064; t test vs. 1: t_{48} = -0.07, p = 0.95). (G) Ratio of average participation in cell pairs of different correlation values for top quintile versus overall

population (top) in example animal and (bottom) for specific levels of coactivity for all animals (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=<10^{-6}). (H) Average correlation between SVM weight and participation in cell pairs of different correlation values (top) in example animal and (bottom) for specific correlation values for all animals (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=<10^{-7}). (I) Proportion of cell pairs with different coactivity levels across days (top) in two example animals and (bottom) z-score of proportion of 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: $F_{2,121} = 10.12$, $p = 10^{-5}$, $\eta^2 = 0.07$; post-hoc Dunnett's tests vs week 1, **p=<0.0087).

The environment-discriminating subset of CA1 cells

It is commonplace to represent hippocampal population activity patterns as activity vectors. As reported, these are distinct between the two environments, yet this difference is remarkably small (Fig. S4), suggesting that only a minority of cells contribute to discriminating the environments. Because only a quarter of CA1 cells are place cells, and only place cells change their position tuning across environments, this small change in population activity across environments might be expected if place cells drive the discrimination. We thus asked what subset of the CA1 population contributes to discriminating the 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 determines the current environment well above chance and 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 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 portions of the population to which the initial SVM decoder gave different weights. A decoder using just the 20% of cells with the largest weights performs indistinguishably from the decoder that uses the entire population. On the other hand, a decoder using the 40% of cells with the smallest SVM weights performs close to chance (Fig. 5D). Decoders using cells with the 20% highest weights also display the largest increase in performance with experience, while the bottom 60% performance do 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 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 cells belong. 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 naturally distinguishes between environment and relies on anti-cofiring cell pairs. (A, right) Illustration of the algorithm used to determined Betti numbers. Activity vectors are projected onto a 10-dimensional space using IsoMap and points are connected to neighbors from increasing distance. At each distance isometries are determined. (right) Example of Betti barcode showing that only one isometry of Betti number 1 is not transient indicating that points live in a 2-D subspace. (B) Activity vectors are projected onto a 2-dimensional space 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 shown distinct color shades. (C) Proportion of variance explained as a factor of the number of dimensions for PCA and IsoMap. (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=<0.046; post-hoc Dunnett's tests: week 1 > weeks 2, 3, *p<0.03, **p<0.01). (E) Similar ratio removing different percent 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>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=<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 interaction, all p > 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 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 ensemble activity in which coactivity relationships differentially represent environments? Computing Betti number barcodes (Ghrist, 2008) indicates that the representational topology of 1-s CA1 ensemble activity patterns is organized on a 2-D manifold (Fig. 6a; Chaudhuri et al., 2019). We thus asked whether an unsupervised, non-linear dimensionality reduction algorithm can readily distinguish environments on a 2-D manifold. 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. IsoMap distinguishes the ensemble activity vectors recorded in the cylinder, box, and home cage 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 normalized by the difference between projections from different environments. This discrimination ratio will decrease from 1, when all environments are similarly represented, to 0,

when the same environments are identically encoded and 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% most coactive and anti-coactive cell pairs. The surviving ensemble activity vectors were then projected on the 2 main IsoMap dimensions. Removal of coactivity or anti-coactivity reduced the IsoMap discrimination (Fig. 6E), and inclusion of only coactive and anti-coactive cells is sufficient for the discrimination (Fig. 6F). These effects were stronger for the anti-coactive cells. These findings provide compelling evidence that CA1 ensemble coactivity can discriminatively represent distinct environments in a non-linear spatial information code that becomes more distinctive with experience.

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 is independent of the place field-based code conventionally used to understand 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 activity patterns (Schneidman et al., 2006). We took advantage of the ability 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. 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 (Pavlides et al., 2019). Network coherence, which captures the alignment of higher order activity correlations, also increases with experience and with coactivity discrimination of environments, further highlighting the importance of coactivity (Fig. 4G). Within the limits of our detection, coactivity relationships were scale-free and not topographically organized yet dominated by a minority of cells for discriminating environments (Fig. 4), features that describe the behavior of flocks (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, similar to 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 do not reject 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; 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 findings that fundamental features of these cells include 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 (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 variable ensemble activity was constrained to a 2-D plane where distinct environment-specific patterns relied on coactivity (Fig. 6). Although how this activity is transformed by the neural networks remains to be studied, both IsoMap and the very definition of place field identifies a 2-D topology. Taken together, these findings constrain the possible neural mechanisms for reading out the information represented in coactivity and point to synapsemble 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, and 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. 2H). 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. 2J; Lever et al., 2002; Ziv et al., 2013). This would require n cell-specific rules to map the place fields from n cells from one

experience to the next, in order to maintain 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 have to 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 in the stability of moment-to-moment coactivity relationships amongst CA1 cells, and this recurrent coactivity was sufficient for differentially representing environments (Fig. 3). The present findings demonstrate how variable distributed activity can be recurrently organized in the stability of a manifold (Fig. 6; Chaudhuri et al., 2019; Nieh et al., 2021; Rubin et al., 2019). This offers an alternative for how neural activity can be persistently informative without single cell stability of firing fields. More explicitly, a connectivity weight matrix, in other words the synapsemble, can transform the 2-D projection we identified (Fig. 6) to a steady readout (e.g., environment identity), so long as the connectivity matrix is invariant to the single-axis shift, which would subserve memory persistence despite instability of neural activity.

CA1 activity has been recorded during explicit memory tasks to correlate how place cells' 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, we have identified memory correlates in momentary cofiring relationships between excitatory as well as inhibitory cells in sleep and active behavior (Dvorak et al., in press; O'Neill et al., 2008; Skaggs and McNaughton, 1996; van Dijk and Fenton, 2018). Incidentally, 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) and involvement of both excitatory and inhibitory cells (Basu et al., 2016; Caroni, 2015; 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 (Okun et al., 2015). This is consistent with the increase interneuron-principal cell cofiring that has been observed at moments of memory discrimination and recollection (Dvorak et al., in press). We analyzed activity fluctuations around place tuning (Agarwal et al., 2014), in effect removing the place signal from the time series, to compute cell-pair coactivity and found that it improved the ability to discriminate environments (Fig. 4 and Suppl. Fig. 1). This improvement suggests a potential orthogonality between the place and environment codes, although we did not study the role of place cells and coactivity in place coding, there are a number of 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 ensemble activity that has evaded the field's efforts, would manifest in the position-tuning independent cell-pair coactivity that we have introduced here as a proxy for the informative higher order correlations in the network activity (Schneidman et al., 2006). 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).

Acknowledgements:

Supported by NIH grants R01MH115304. We are grateful to Dr. John Kubie for valuable comments on an earlier draft of the manuscript.

Author Contributions:

A.A.F. and E.R.J.L. designed research; E.R.J.L., E.P. and W.R. performed research; E.R.J.L. analyzed data; A.A.F. and E.R.J.L. wrote the paper.

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 μl 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 cm^2, 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.

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 dependance

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 (Farnebäck, 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 CaImAn 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 postion 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.

Cell pair correlations

Kendall tau (τ) 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$30um). 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. 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 coherence

The Network coherence (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 coherence 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 coherence 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 coherence 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 τ values. For each cell, we computed its participation in a range of τ values as the number of times a cell participates in a pair with a τ value within that range, divided by the total number of pairs in which the cell participates. In effect, it is the proportion of τ 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 number were obtained each day from 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.

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. 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.

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). 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 24x24 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).

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_{input} = 1000 neurons that were tuned to locations on the track. This population fed into a layer of n_{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_{inhib} = 50 neurons (W^{EI}). This inhibitory population is connected back to the excitatory layer (W^{IE}). The E→E, E→I, and I→E weights were all plastic.

At initialization, all values of W^{InputE} were uniformly sampled from the interval (0, W_{max}^{InputE}). For W^{EE}, each weight either took a uniformly sampled value from the interval (0, W_{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 $W_{max}^{EE} = W_{max}^{InputE}/p_{EE}$. The same initialization procedure was repeated for W^{EI} and W^{IE} with $W_{max}^{EI} = W_{max}^{InputE}/p_{EI}$ and $W_{max}^{IE} = W_{max}^{InputE}/p_{IE}$.

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_{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 |·| 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

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

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

$$\Delta W_{ij}^{IE}(t) = \eta \left(f_i^E(t) \cdot v_j^I(t) - f_j^I(t) \cdot v_i^E(t) - 0.1 \right).$$

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 $\bar{a}^{control}$ 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 $\Delta t_{Kendall}$ = 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).

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
$\Delta t_{Kendall}$	0.25	size of bin used to average the activity before Kendall correlation
W_{max}^{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.

Supplementary figures

Figure S1. Evaluation of the impact of multiple fields on decoding environments from place cells. (A) Example of corresponding larger maps in two environments. Each map is created by concatenating the place maps of four place cells in the circle and square environment respectively. (B, left) Example of the random walk in the large environment during a simulation corresponding to 1 minute and (B, right) example 16 'composite' cells from a 152-'composite' cell ensemble. From each map, an expected trace is created from the random walk position and the average rate at the corresponding position. (C) Performance for decoding two repetitions of each of the two environments depending on the environment size in a 240-minute-long simulation. Decoding performance was stable for simulations of 60 to 240 minutes.

Figure S2. Related to Figure 3. Additional Rate and PTI Coactivity information
(A, top) Scatter plot of coactivity (τ) stability in (left) the same environments and (right) different environments. Dashed lines are regression lines for individual animals with (bottom) Pearson correlation (r) values plotted in grey over the bar graph indicating the average relationship. Solid lines indicate overall regression, with r values in black in bar graph. (B) Relationship between rate coactivity (τ) and (top) PTI coactivity or (middle) randomized PTI coactivity in scatter plots. Dashed mauve lines indicate regression for individual animals (bottom) with r values plotted in black for place cells and grey for non-place cells. Solid red lines indicate overall regression. (C) Relationship between place field similarity (r) and (top) rate coactivity, (middle) PTI coactivity, or (bottom) randomized PTI coactivity in scatter plots. Dashed mauve lines indicate regression for individual animals while solid red lines indicate overall regression.

Figure S3. Related to Figure 4. Decoding environments from rate coactivity. (A) Illustration of SVM decoding from rate 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; The absolute values of the weights, indicative of separation strength, are classified by decile. (B, top) Cross-validated decoding accuracy is very high for all ensemble and (bottom) for the first two cell-pair deciles. (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 cell pairs, confirmed by (D) 2-D distribution of SVM weights against PTI cofiring. (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 distribution of individual cell participation in decile subset (example distributions in inset and shown in purple in main figure). (G) Rate network coherence normalized by the standard deviation of the cell-pair randomized distribution (Week: $F_{2,60} = 28.49$, $p=10^{-9}$, $\eta^2 = 0.13$; Dunnett's tests: week 0 < weeks 2, 3).

Figure S4. Related to Figure 5. Activity vector stability and environment discrimination
(A) Illustration of activity vector correlation matrices, comparing similar or distinct environments, where each pixel is color-coded by the correlation between the 2 activity vectors observed at the time indicated on the x- and y- axes. (B) Histogram indicating distribution of (left) average correlation between activity vectors of separate recordings and (right) 'correlation similarity', average normalized by average correlation between activity vectors of the same recording. (C) Within-day correlation similarity improves across weeks (Week: $F_{2,28} = 4.47$, p = 0.021, $\eta^2 = 0.081$; Env. Change: $F_{1,29} = 24.10$, p = 10^{-5}, $\eta^2 = 0.20$; No interaction: $F_{2,28} = 0.18$, p = 0.83, $\eta^2 = 10^{-4}$; Dunnett's tests: week 1 < weeks 3 for similar environments; Same vs. Different is significant each week, *p's<0.015, **p = 0.0017). (D) Across-day correlation similarity degrades over time (Time: $F_{4,132} = 26.28$, p = 10^{-16}, $\eta^2 = 0.37$; Env. Change: $F_{1,132} = 32.99$, p = 10^{-8}, $\eta^2 = 0.15$; Interaction: $F_{4,132} = 0.40$, p = 0.81, $\eta^2 = 0.043$; Dunnett's tests vs day 0 **p<0.007; Same vs. Different **p =< 0.0013).

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.

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.

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. (in press). Dentate spikes and external control of hippocampal function. Cell Reports.

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.

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). CaImAn 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.

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.

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.

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).

Kentros, C. (2006). Hippocampal place cells: The "where" of episodic memory? Hippocampus 16, 743-754.

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.

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.

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.

Mizumori, S.J.Y. (2006). Hippocampal place fields: A neural code for episodic memory? Hippocampus 16, 685-690.

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.

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., Schottdorf, 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., Iacaruso, 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.

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.

Pavlides, 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.

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.

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.

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.

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.