Pre-existing and emergent cortical neuronal assembly sequences during learning

Luke Pemberton[1], Huijeong Jeong[1], Vijay Mohan K Namboodiri[1,2#]

[1] Department of Neurology, University of California, San Francisco, CA, USA

[2] Neuroscience Graduate Program, Weill Institute for Neurosciences, Kavli Institute for Fundamental Neuroscience, Center for Integrative Neuroscience, University of California, San Francisco, CA, USA

[#] Correspondence to VijayMohan.KNamboodiri@ucsf.edu

Neuronal assemblies—groups of co-active neurons—support memory consolidation and retrieval. In the hippocampus, assemblies can pre-exist learning and contribute to memory through sequential activation. Whether similar principles apply to higher cortical areas for flexible memory storage like the orbitofrontal cortex (OFC) remains unclear. Using a novel ground truth-validated clustering approach, we investigated the activity of longitudinally tracked mouse OFC neurons during cue-reward memory acquisition and maintenance. Assemblies active after learning pre-existed the learning and exhibited two distinct sequential dynamics suggestive of memory consolidation or retrieval. Consolidation sequences emerged during learning, while retrieval sequences partly recruited pre-existing reward sequences. These findings demonstrate that OFC learning recruits pre-existing networks flexibly repurposed for new associations, revealing circuit motifs that may enable cortical memory storage.

Neuronal assemblies—groups of neurons with high co-activity reflecting shared connectivity—support perception (*1–3*) and memory (*4–8*). In the visual cortex, distinct assemblies (also called ensembles) represent different sensory percepts (*2*, *9*, *10*), while in the hippocampus, distinct assemblies represent spatial locations (*7*, *11–13*), sensory episodes (*14–17*), or components of a cognitive map (*18–21*). In both sensory cortex and hippocampus, some assemblies pre-exist learning (*9*, *22–25*), though the extent of pre-existence versus emergence during learning remains debated (*26–29*). Assemblies can also form sequences (*9*), particularly in the hippocampus (*30*, *31*). For example, the hippocampus replays sequences of assemblies corresponding to past experiences at compressed (*32*, *33*) or experienced time scales (*34*). Such awake replay is thought to contribute to memory consolidation by repeatedly reactivating prior sensory sequences (*35–37*). During navigation, compressed sequences can sweep ahead of the animal's position (theta sequences), possibly reflecting retrieval of upcoming or planned locations (*38–40*). In sum, assembly sequences may underlie both memory consolidation and retrieval.

Whether assembly sequences that may support memory consolidation and retrieval also exist in higher cortical areas involved in flexible memory storage remains poorly understood. We examined this question in the orbitofrontal cortex (OFC) because similar to the hippocampus, it is important for representing cognitive maps of task states (*41–43*), values (*44–47*), or associations (*48–52*)—functions that may rely on flexible neuronal assembly sequences. To this end, we developed a ground truth–validated clustering approach for identifying pooled assemblies across animals and further validated it using electrophysiological recordings from neurons throughout the visual cortical hierarchy. We then applied this standardized clustering method to analyze assembly sequence dynamics in OFC neurons longitudinally tracked during cue–reward memory acquisition and maintenance.

Results

Assembly detection by clustering of event-triggered activity

Identifying neuronal assemblies by assessing co-activity requires simultaneous recordings from neurons over sufficiently long time periods. As a result, commonly used assembly detection methods operate on the activity of neurons recorded simultaneously during a single session in a single behaving animal (*53–57*). The functional relevance of these assemblies is typically evaluated by analyzing their

activity time-locked to behavioral events. This strategy can then be extended to pool functionally comparable assemblies across animals, though this often results in high redundancy, with many assemblies—sometimes more than 20—relating to a single behavior and overlapping with assemblies associated with other behaviors (*58*).

In contrast, we developed a different approach for assembly detection. We hypothesized that neurons co-active outside of task periods, when shared inputs are minimal, are also likely co-active during task periods when shared inputs are pronounced. Thus, we reasoned that neurons with highly similar peri-event time histograms (PETHs)—reflecting high co-activity around behavioral events—may also be co-active outside those events, thereby forming assemblies. This rationale enables pooling of neurons across animals by evaluating their PETH similarity (**Fig 1A**). After identifying neurons with similar PETHs using dimensionality reduction and clustering, we then assess non-event-aligned co-activity among simultaneously recorded neurons. This allows us to determine whether the clusters represent assemblies, specifically by examining whether within-cluster correlations exceed between-cluster correlations (**Fig 1A**). Using this method, an assembly with high non-event aligned co-activity fluctuations is guaranteed to have non-redundant task activity compared to other assemblies.

To identify the best clustering algorithm for detecting PETH similarity, we simulated realistic ground truth data of clusters with many distinct PETH shapes, and parametrically varied cluster separability (Methods). We tested many common clustering methods, some shown in **Fig 1B** and others in Fig S1. When ground truth clusters are well separated, all methods performed equally well. However, at reduced separability, discrete spectral clustering (*59*) proved to be the most effective (**Fig 1C-E**). This approach was also the best on less realistic simulated ground truth data of different types (Figs S2, S3). Across these ground truth datasets, the benefit of using discrete spectral clustering was higher when the simulated PETHs differed in shape (**Figs 1,** S2) and not just the location of similarly shaped peak activity (Fig S3). Together, our findings indicate that discrete spectral clustering is the best method for identifying neurons with high PETH co-activity.

We next tested whether the clusters identified using this method are assemblies using a large-scale single unit electrophysiology dataset across the visual cortical hierarchy (*60*), similar to what has been reported (*61*) (Fig S4). Instead of aligning activity on an external behavioral event, here we aligned the activity of neurons to hippocampal ripples, an internal brain event. We identified three neuronal clusters based on ripple-aligned PETHs. The activity of these clusters showed higher within-cluster co-activity outside of ripple periods (Fig S4), validating that identified clusters are indeed assemblies. Interestingly, the sequential activity of these assemblies was conserved across both ripple and non-ripple periods, proceeding in a non-time compressed manner. Thus, event-aligned co-activity using discrete spectral clustering can indeed identify true neuronal assemblies and enable investigation of assembly sequences.

OFC neuronal assemblies detected after cue-reward learning pre-exist the learning

We next applied this method to identify neuronal assemblies in the OFC during cue-reward learning (*48*, *62*). Briefly, head-fixed mice learned the association between an auditory tone lasting 2s (CS+) and a reward after a 1s gap ("trace interval"); another auditory tone (CS-) was not paired with reward (**Fig 2A**). Trials were followed by variable intertrial intervals (ITI; constant 3s plus an exponential distribution of mean 30s). Mice learned this association in a week on average. After a few sessions of stable behavior, two forms of contingency degradation were performed in a random order (either reduction of reward probability to 50%, or unpredicted rewards in the inter-trial interval, see Methods). After retraining to full contingency, mice finally underwent extinction of the cue-reward pairing. Conditioned behavior, as measured by anticipatory licking, was low on early days and high after learning as well as the day before extinction, and low after extinction (*48*) (**Fig 2B**). OFC neurons were targeted using an AAV expressing GCaMP6s under the CaMKIIα promoter (**Fig 2C**). Average neuropil corrected CS+ trace interval fluorescence across the entire longitudinally tracked neuronal population (n=1,435 neurons from n=5 mice) was low on early days of learning, high after learning and low after

extinction (**Fig 2D**). Collectively, these results demonstrate that animals did not exhibit evidence of behavioral or neural learning on day 1 of conditioning, allowing a timepoint to evaluate pre-existing network structure.

We next used discrete spectral clustering to cluster neurons based on their PETHs on CS+ and CS- trials after removing non-responsive neurons (Methods). We identified seven neuronal clusters with distinct PETHs (**Fig 2E**). This contrasts with a conventional assembly detection algorithm, which exhibited highly redundant task activity (Fig S5). The activity of these clusters on day 1 did not differ between CS+ and CS- trials (**Fig 2F**). We next assessed whether these identified clusters were assemblies by measuring within-cluster vs between-cluster correlations during the ITI. When measuring ITI correlations, we used deconvolved fluorescence signals that have been shown to approximate spiking using ground truth *ex vivo* OFC recordings (Supplementary Fig 1 in (*48*)) and using a meta-analysis of deconvolution methods (*63*). On the trained day, clusters 1-4 and 6 showed higher within-cluster than between-cluster correlation, indicating that they are assemblies (**Fig 2G**, all p values were obtained by hierarchical bootstrapping (*64*) and corrected for multiple comparisons (see Methods and Fig S6). Clusters 1-4 also showed higher within-cluster than between-cluster correlation on day 1, indicating that the assemblies pre-existed cue-reward learning (**Fig 2G**). We next plotted the average ITI correlation between simultaneously recorded neurons between all cluster pairs and found that this overall network correlation structure looked remarkably similar between day 1 and trained days (**Fig 2H**). Indeed, the correlation between ITI correlations of every cluster pair across the two days was extremely high, with a slope less than 1 indicating that the network correlation structure strengthened the higher within-cluster correlations and weakened the lower between-cluster correlations (**Fig 2H**). In sum, OFC neuronal assemblies identified based on their PETHs on the trained day pre-existed and strengthened with learning.

OFC neuronal assembly sequences consistent with memory consolidation and retrieval

OFC maintains cue-reward memories in a cognitive map. While memory retrieval is triggered by the cue, memory consolidation, by definition, requires cue-reward memory representations to be maintained outside the cue-reward experiences (e.g., like hippocampal replay during rest (*35*)). We next investigated whether there were any sequential dynamics of the OFC neuronal assemblies consistent with memory consolidation and retrieval. Here, we focused on the four main clusters with strong cue and reward evoked responses: 1, 2, 3, and 4 (**Fig 3A**). Clusters 1 and 2 strongly respond to the cue, while clusters 3 and 4 strongly respond to reward. Visualization of the CS+ PETHs suggested a sequence of activity among these clusters from 2 to 1 to 3 to 4. However, this apparent sequence on CS+ trials may just be an artifact of the fact that clusters 1 and 2 mainly respond to cues, and 3 and 4 mainly respond to rewards. Indeed, the ordering between reward responses for clusters 3 and 4 appears to be maintained on day 1 with cluster 3 exhibiting short latency reward responses and cluster 4 exhibiting long latency reward responses (**Fig 2F**).

To rule out trivial sequences observed only during the experimentally induced cue-reward pairing, we first assessed sequences during the ITI. We measured sequences between a pair of assemblies during the CS+ trials and the ITIs following CS+ trials (ITI+) using cross-correlation analysis (Fig S7). We identified the asymmetry in the cross-correlogram by measuring the difference between the negative lag area under the curve and the positive lag area under the curve (ΔAUC) (Fig S3, S7). Zero ΔAUC indicates no sequential relationship between the pair of clusters and high ΔAUC suggests strong sequential relationship. We plotted the correlation between ΔAUC for every pair of clusters (including 5-7) between CS+ and ITI+ periods on the trained day and found a significant, but weak correlation of slope less than 1 (**Fig 3B**). This suggests that some cluster pairs exhibited similar sequential relationships between CS+ and ITI+ (consistent with replay), while others exhibited sequential relationships only during CS+ trials. We identified that the assembly pairs (1,2), (2,3) and (2,4) exhibited similar sequences between CS+, ITI+, CS-, and ITI- periods (ITI- indicates ITIs following CS- trials). We refer to these as replayed sequences (**Fig 3C**). We also measured the centroid lag of the cross-

correlogram (Fig S8) and found that the lag between these sequences was similar, but slightly smaller, between CS+ periods with cue-reward pairing and the other periods, suggesting that the CS+-reward memory may be replayed by these sequences at approximately the experienced cue-reward lag, consistent with a memory consolidation mechanism (**Fig 3C**).

On the other hand, the (1,3), (1,4) and (3,4) sequences were much stronger in the CS+ periods compared to CS-, ITI+, or ITI- periods. We refer to these as CS+ sequences (**Fig 3D**). The CS+ sequences could trivially be present only during the experimentally induced cue-reward pairing. To rule this out, we measured ΔAUC on CS+ unrewarded trials in the 50% reward probability session (**Fig 3E**). All these sequences maintained ΔAUC on CS+ unrewarded trials (**Fig 3F-H**). The maintenance of the same sequences on CS+ unrewarded trials is consistent with cue-reward memory retrieval.

We next assessed sequence dynamics during and after cue-reward learning (**Fig 4**). We discovered that replayed sequences emerged during learning, i.e., were weak on days 1-3, and were thereafter maintained on CS+, ITI+, CS-, and ITI- periods even after extinction (**Fig 4A-D**). Thus, these sequences are consistent with a mechanism that consolidates an acquired cue-reward memory. Among the CS+ sequences, (1,3), and (1,4) sequences emerged during learning on CS+ trials, but the (3,4) sequence pre-existed learning and was stably maintained during and after learning (**Fig 4E-H**). Interestingly, though the (3,4) sequence is visible after reward on day 1 (**Fig 2F**), it is co-opted as a memory sequence after learning, as evident from its presence on CS+ unrewarded trials (**Fig 3E**) and after cue-reward behavioral extinction (**Fig 4E**). Thus, sequences consistent with memory retrieval partly repurpose a pre-existing reward sequence. Interestingly, consolidation vs retrieval sequences were cleanly separated by the presence or absence, respectively, of cluster 2—the cluster that responds similarly to both CS+ and CS- (**Fig 2E**). This suggests that a generic cue encoding assembly forms learned sequences with other reward-related assemblies to consolidate cue-reward memory. Finally, there was also a drifting set of sequences reflecting idle states (Fig S9).

Discussion

Cue-reward memories are thought to be encoded in higher-order cortical regions like the orbitofrontal cortex (OFC) by the strength of neuronal activation (*48*, *49*, *65*, *66*). Here, we show that OFC also repurposes pre-existing neuronal assemblies into sequences consistent with memory retrieval—i.e., induced by the reward predicting cue—and memory consolidation—i.e., induced both by the cue and replayed at moments without the cue. Most of these sequences emerged during learning and were maintained thereafter, although one memory retrieval sequence utilized a pre-existing reward sequence.

Unlike conventional methods, we identified comparable assemblies across animals via task-aligned activities, enabling natural pooling of similar assemblies. Discrete spectral clustering proved optimal for this approach, as task-defined clustering after learning revealed neuronal groups that correspond to pre-existing assemblies during non-task periods before and after learning. A potential limitation is that this method only detects non-overlapping assemblies, while traditional techniques can reveal overlapping membership. However, interpreting task-relevant activity remains a challenge for conventional assembly detection due to highly redundant task-aligned assemblies (Fig S5).

These findings directly relate to two key debates in neuroscience. The first concerns representational drift versus stability (*67*–*72*). While representational drift has been reported in several regions even during identical behavioral tasks, this study finds largely stable assemblies and learned sequences persisting across diverse behavioral contingency changes over days. Even after probability or contingency degradations and extinction, many learned sequences remain stable: (1,2), (2,3), (2,4), and (3,4) (**Fig 4**). This stability aligns with recent findings of persistent assemblies in V1 during passive visual exposure (*73*), and extends them by demonstrating that OFC assembly sequences are stable despite multiple interceding task changes.

The second debate addresses categorical versus non-categorical neural coding. Although OFC can exhibit neuronal clusters (*48*, *74*–*76*), various studies emphasize the prevalence and computational

advantage of non-categorical coding without distinct clusters (*77–79*). Indeed, a recent large-scale survey suggests that categorical coding is uncommon outside sensory cortex (*80*). Notably, however, the presence of clusters in this study is analyzed using time-averaged activity across many task conditions. Here, in contrast, we show that clustering based on PETHs reliably identifies neuronal clusters that even pre-exist learning and form assembly sequences. We therefore propose temporal activity profiles as critical when assessing the presence of activity-based neuronal clusters. This is consistent with reports of neuronal clusters in many regions identified via PETHs aligned to internal or external events (*48*, *61*, *81–85*). In sum, neuronal assemblies with co-active responses may represent a general organizational principle throughout the brain, even if their time-averaged task activity often appears non-categorical.

How are cue-outcome memories stored in the brain? The engram hypothesis, formulated mainly in hippocampus and amygdala, posits that a set of excitable neurons activated during initial cue-outcome experience are captured into a memory representation that can later be retrieved by reactivating the same cells (*86*). Here, we observed more complex dynamics: although assemblies themselves were present on the first day, the assembly sequences associated with memory consolidation and retrieval largely emerged only after learning, being absent on the first day of experience. This may reflect systems consolidation where memories recruit hippocampus during the early phase of learning but recruit cortex only later (*87*). Notably, a pre-existing reward sequence present on day 1 later became incorporated into a memory retrieval sequence, suggesting that some sequences present initially are relevant for future memory retrieval. Additionally, we identified a drifting sequence representing idle states, which may be recruited for encoding newly learned memories, consistent with interpretations from the engram literature (*88*).

Cortical assembly sequences identified here, as in the hippocampus, are consistent with a role in memory consolidation and retrieval. However, as with the hippocampus, current technology does not allow direct causal tests of the function of individual assembly sequences in memory. Demonstrating causality would require tools capable of expressing distinct opsins in activity-defined assemblies with unknown genetic or circuit markers—something that remains technically infeasible. The correspondence between activity-defined assemblies and genetically or circuit-defined cell types is unclear, with some studies suggesting a weak link (*48*, *89*). As a result, existing genetic or viral methods cannot label the subsecond-timescale, activity-defined assemblies described here. Further, even though holographic optogenetic stimulation has been used to trigger individual activity-defined assemblies (*2*, *6*), playing back or selectively removing assembly sequences at scale to assess behavioral effects remains technically challenging. Major technological advances will therefore be necessary to replay or remove arbitrary neuronal assembly sequences to directly assess their function. Despite these technical limitations, the present findings demonstrate that the OFC flexibly repurposes pre-existing neuronal assemblies into sequences that may support cue-reward memory consolidation and retrieval.

Acknowledgments: We thank members of the Namboodiri laboratory for helpful discussions. This project was supported by NIH R01MH129582, R01AA029661, the Scott Alan Myers Endowed Professorship, Alfred P Sloan Fellowship, Pew Biomedical Scholarship, and Klingenstein-Simons Fellowship to V.M.K.N. We thank Garret D Stuber in whose lab V.M.K.N. collected the OFC imaging data.

Figures

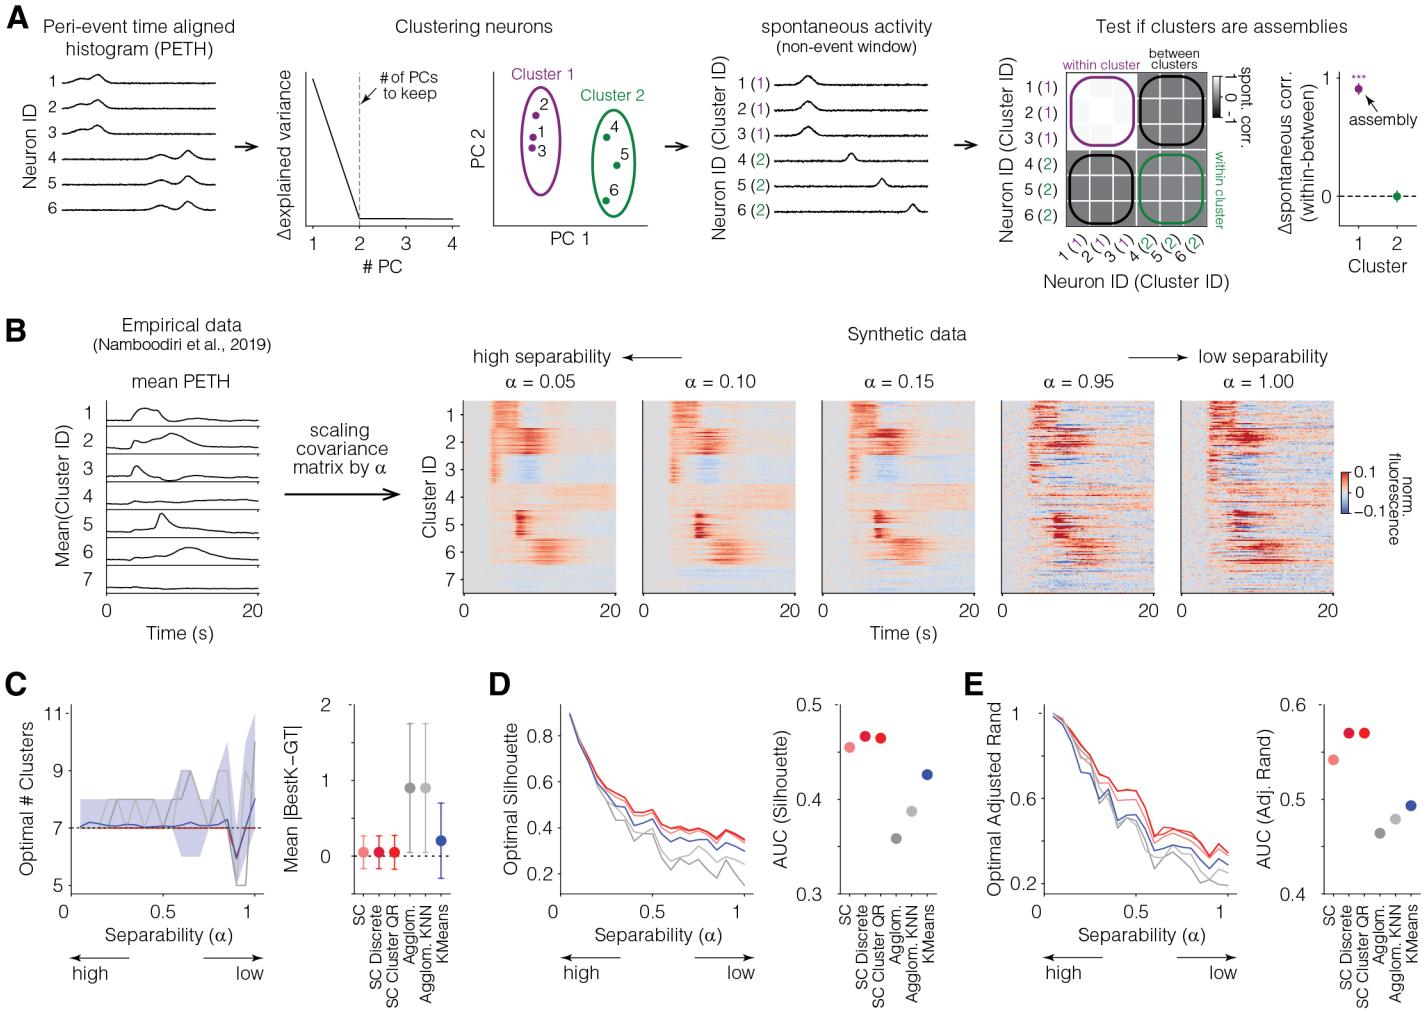

Fig 1: Discrete spectral clustering is the best ground truth validated clustering method for identifying PETH similarity

A. Schematic of clustering pipeline for assembly detection involving clustering of task-related PETHs and subsequent testing of whether these clusters are assemblies during non-task aligned activities. See Fig S1 for all clustering algorithms tested here.
B. First ground truth simulated dataset (other datasets in Fig S2, S3). Here, we used a previously published dataset to simulate seven possible PETHs (*48*). The original covariance matrix across the seven possible clusters was multiplied by the scale factor α. Low values of α increase cluster separability by retaining mean locations in activity space while reducing variance. 100 neurons with each PETH type were simulated with random Gaussian noise added to each neuron's PETH (Methods). α=1 represents separability comparable to the original data (but in a ground truth simulation).
C. Identified number of clusters based on six different clustering algorithms to separate neurons by their PETH types. Seven is the ground truth number of clusters. Different algorithms are represented by different colors, as shown in the plot on the right measuring the deviation between the identified number of clusters ("BestK") and the ground truth number of clusters ("GT").
D. The silhouette score, a measure of cluster separation, for the identified optimal number of clusters for each algorithm.
E. The adjusted rand index, a measure of how correctly each neuron was assigned to its ground truth cluster, for each algorithm.

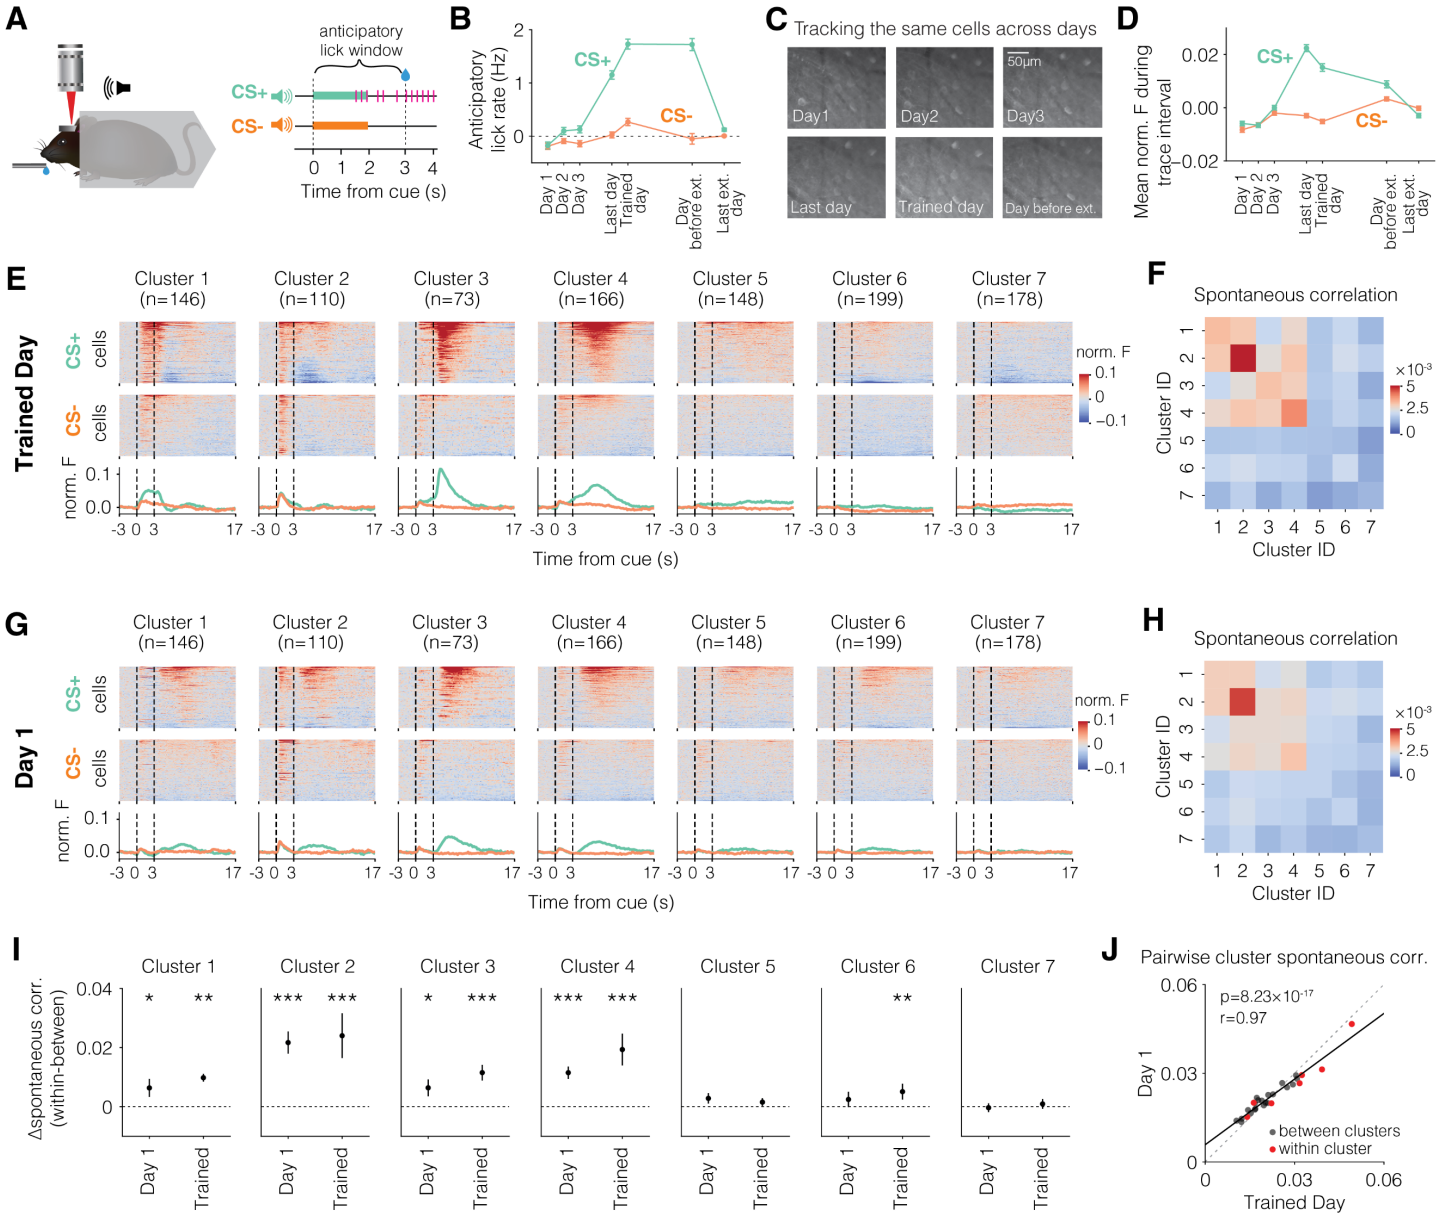

Fig 2: OFC neuronal assemblies identified after learning pre-exist learning.

A. Schematic of head-fixed two-photon imaging and trace conditioning paradigm. Anticipatory lick rate is measured during the cue-outcome period to assess behavioral evidence reflecting cue-reward memory.
B. Baseline subtracted anticipatory lick rate across various task days, demonstrating clear behavioral evidence of acquired cue-reward memory on the last day of learning (i.e., day before trained day), trained day, and day before extinction. Data from n=5 mice. During extinction, reward was not delivered after CS+, which eliminates anticipatory licking. Last day of learning was day 7.2 on average across mice. Trained day was day 8.2 on average across mice. Day before extinction was day 21.8 on average across mice. Last day of extinction was day 24.4 on average across mice (see Methods for detailed timeline of experiments).
C. Example cropped field of view showing longitudinal tracking of the same neurons across days.
D. Mean normalized fluorescence (see Methods) during the trace interval (2-3 s after CS+ onset) across all longitudinally tracked neurons (n=1,435 neurons across 5 mice), showing that there is no evidence of neural learning on day 1.
E. Results of discrete spectral clustering applied to OFC neurons based on their CS+ and CS- evoked PETHs. Only neurons that were longitudinally tracked throughout learning are

included. Top two rows show CS+ and CS- aligned activity of all neurons within each cluster. The neurons are sorted by their total activity during CS+ trials and arranged in the same order for CS-.

F. Heatmap showing spontaneous correlations between each pair of clusters based on their inter-trial interval (ITI) activity outside of CS+ or CS- trials.
G. Same as E but for day 1 of conditioning.
H. Same as F but measured using ITIs on day 1 of conditioning.
I. The difference in within-cluster versus between-clusters spontaneous correlations during the ITI on both day 1 and trained day. Clusters 1-4 are assemblies on both days.
J. Scatter plot of the spontaneous correlations between each pair of clusters measured on day 1 and trained day. Within-cluster and between-clusters pairs are separated.

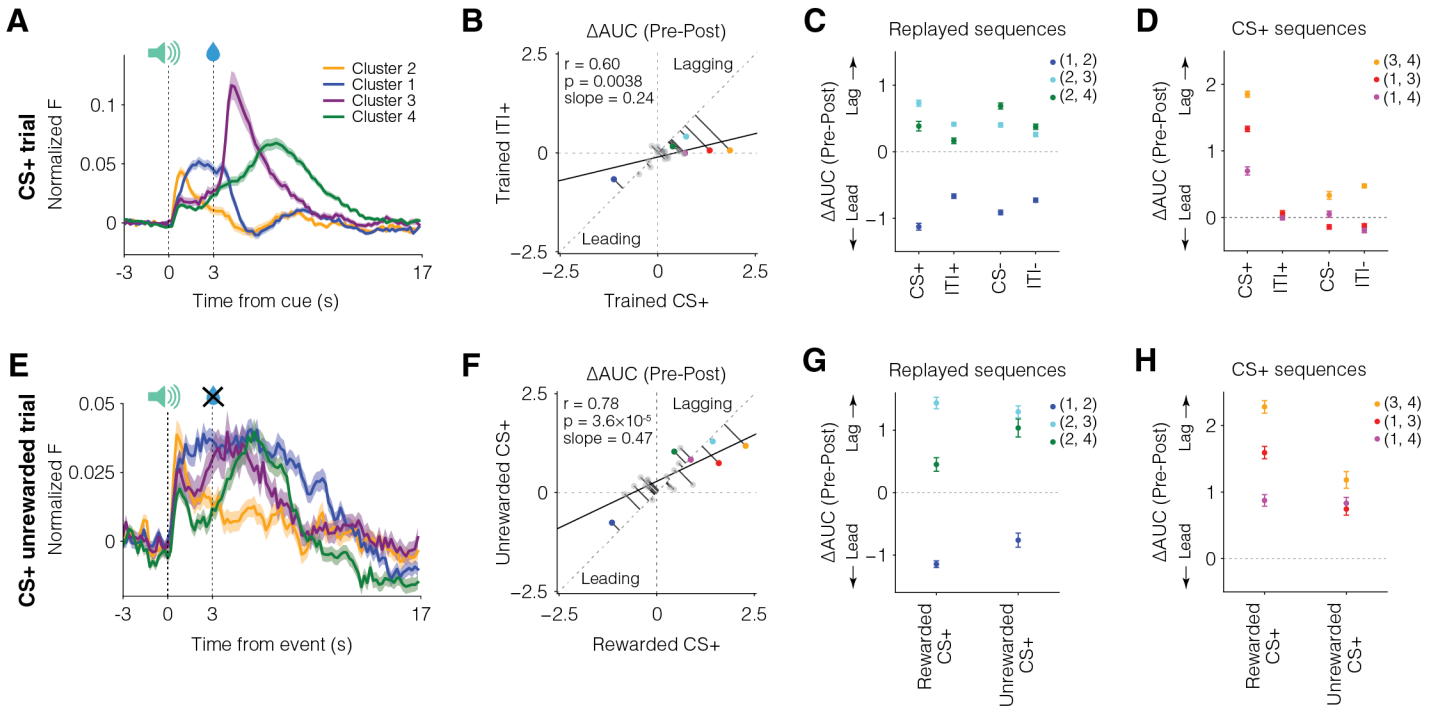

Fig 3: OFC assembly sequence dynamics after cue-reward learning are consistent with memory consolidation and retrieval.

A. CS+ aligned PETHs for clusters 1-4 on the trained day. The sequence 2→1→3→4 is visible.
B. The sequential relationship between assemblies was measured using the area under the curve of the cross correlogram (CCG) over negative lags minus the corresponding area over positive lags (referred to as ΔAUC) (Fig S7). Here, ΔAUC is plotted for all cluster pairs during CS+ periods (0-17 s following CS+ onset) and ITI+ periods (ITIs following CS+ trials). Weak but significant correlation in these ΔAUCs imply that some sequences are replayed during ITI while others are not.
C. ΔAUC of "replayed sequences" during CS+, ITI+, CS- and ITI- periods, showing maintenance of ΔAUCs across all time periods. Note that all replayed sequences include cluster 2.
D. ΔAUC of "CS+ sequences" during CS+, ITI+, CS- and ITI- periods, showing much larger ΔAUCs during CS+ compared to other time periods. Note that no CS+ sequence includes cluster 2.
E. CS+ aligned PETHs on unrewarded trials during a 50% reward probability session after the trained day (see Methods). The sequence 2→1→3→4 is still visible.
F. ΔAUC is plotted for all cluster pairs during CS+ rewarded vs unrewarded periods on the 50% reward probability session.
G. ΔAUC of "replayed sequences" during CS+ rewarded and unrewarded periods.
H. ΔAUC of "CS+ sequences" during CS+ rewarded and unrewarded periods.

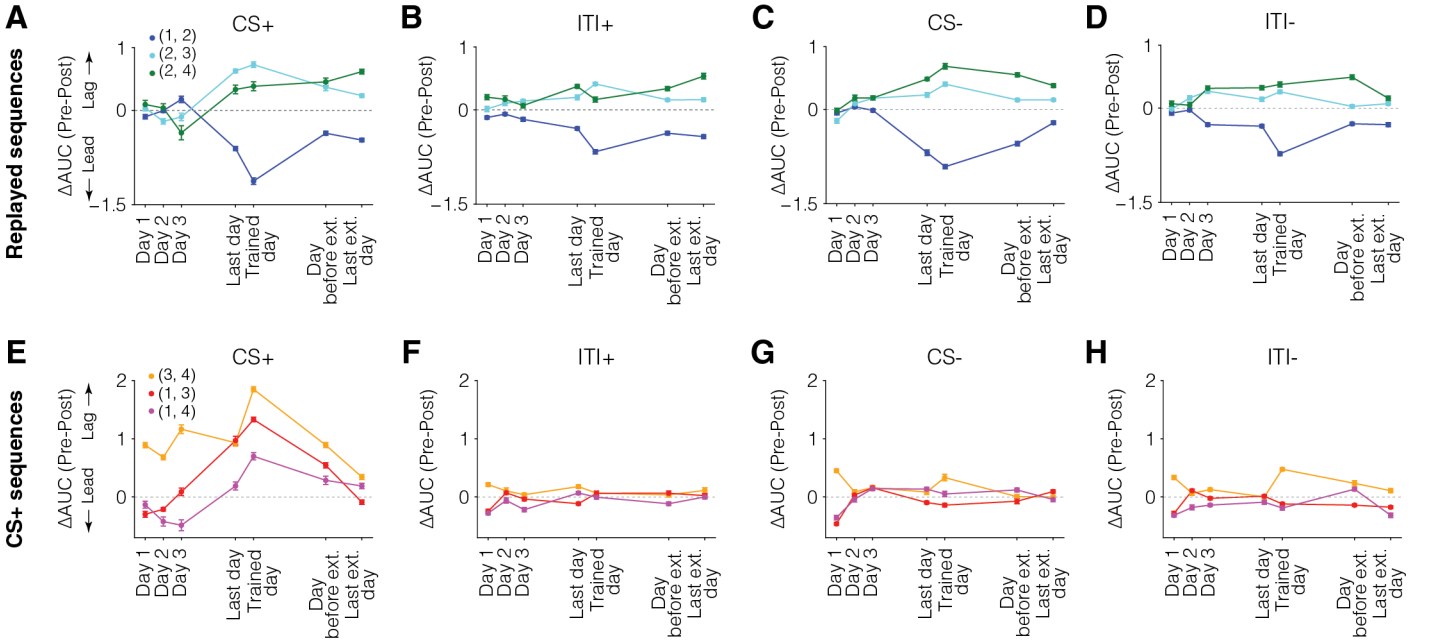

Fig 4: Pre-existing and emergent OFC assembly sequences during cue-reward learning

A-D. ΔAUC of replayed sequences across days of interest for all time periods showing emergence during learning and subsequent maintenance even despite behavioral extinction.

E-H. ΔAUC of CS+ sequences across days of interest for all time periods showing selective enhancement during CS+ but not other time periods. (3,4) sequence pre-exists learning. While (1,3) and (1,4) also pre-exist learning (ΔAUC<0), their sequential order flips after learning (ΔAUC>0).

Methods

Ground truth clustering simulations (Fig 1)

To evaluate which clustering algorithm to use for PETH separation in the assembly discovery pipeline, we generated three simulated ground truth datasets: one based on real data, and two hypothesized types of co-activity patterns. We refer to the first as the "Fake OFC" datasets (**Fig 1**) and the latter two as the "Shape" datasets (Fig S2) and "Peak" datasets (Fig S3).

To generate the "Fake OFC" dataset, we merged OFC neural clusters 7, 8, and 9 from previous data (cite). We kept clusters 1-6 as is, for 7 total clusters. We calculated the sample mean and covariance matrices of the PSTHs calculated from cue aligned activity spanning [-3,17] s from neurons in each cluster. We then generated several, entirely new, synthetic datasets using the python library NumPy random.multivariate_normal function. To generate a dataset, we scaled each covariance matrix per cluster by a coefficient alpha, parameterizing the separability of neurons in data space. Then we generated 100 new simulated "neurons" per cluster based on that cluster's true sample mean in the data and the scaled covariance matrix. We sampled values of alpha between [0.05, 1] in increments of 0.05 to generate 20 datasets.

We generated the "Shape" simulation datasets to model neurons that share a common temporal shape of activity within each cluster, while different clusters differ in the number of distinct peaks in that shape. We simulated 5 clusters of 25 neurons each with 125 time bins for each neuron's timeseries with a bin size of 0.5. For cluster *k*, each neuron's time series was constructed from *k* Gaussian components (i.e., cluster 1 has one Gaussian, cluster 2 has two Gaussians, … cluster 5 has five Gaussians). The peak times were arranged as such: 1.5, 12.5, 25, 37.5, and 50. All Gaussian components shared the same standard deviation, and their peak times were arranged so that the first component's peak time was identical across clusters (so cluster 2's first peak aligns with cluster 1's peak), then subsequent peaks were placed with adjacent offsets. To control separability, we scaled the amplitudes of all Gaussian components *except* the first by a common multiplier α, and generated datasets for α ranging [0, 0.66] in increments of 0.06. We added stationary Gaussian noise using NumPy's random.normal function with default parameters and uniformly sampled noise NumPy's random.uniform with range [0.01*M, 0.1*M] (M is the maximum value of a given time series) to each time series and normalized time series to sum to 1.

We constructed the "Peak" datasets to simulate neurons that have a Gaussian shaped curve with 6 clusters of simulated neurons with 20 neurons each. We used 200 time bins for each neuron's time series with a bin size of 1. Neurons in each cluster had a shared peak time of activity. Each cluster had a unique peak time. We selected peak times of [50, 60, 95, 105, 140, 150], and parameterized data generation by mean noise, controlling the within cluster variability of single neuron peak times, and standard deviation, controlling the spread of individual neurons' Gaussian shaped activities. As mean noise increases, cluster identity of neurons becomes harder to correctly recover, and when mean noise is greater than 5, peak times between adjacent clusters can overlap. We generated datasets with mean noise values between [0, 6] in increments of 1, and for each mean noise parameter value, we sampled standard deviation values of [0.6, 3.6] in increments of 0.6, then applied the exponential function to the sampled standard deviation value. We added stationary Gaussian and uniformly sampled noise as described above to each time series and normalized time series to sum to 1.

In Fig S1, we tested several algorithms (80%/90%/95% explained variance, scree plot elbow, Marchenko-Pastur cutoff, Kaiser rule, Broken Stick rule, Minka's MLE rule, Optimal Hard Threshold, Horn's parallel analysis) that automatically selected the number of principal components to retain after PCA in combination with several clustering algorithms and variants available in scikit learn (KMeans, Agglomerative across all linkage types, KMeans/Discrete/ClusterQR Spectral Clustering, DBSCAN, HDBSCAN, OPTICS) on these three simulated ground truth datasets. The entire pipeline was tested

to identify the appropriate number of retained principal components and clustering algorithms and associated parameters that maximized performance metrics described later.

The approaches to identify number of principal components are described in detail below.

80%/90%/95% explained variance: Choose the smallest number of principal components so that the cumulative proportion of variance explained exceeds a threshold (e.g., 80%, 90%, 95%).

Scree-plot "elbow": Plot the eigenvalues (or variance explained) in descending order and look for the "elbow" point where the drop flattens out. Then retain all components before the elbow. This is a heuristic reflecting diminishing returns in variance explained beyond a point where additional components are likely due to noise in the data (*90*).

Marchenko–Pastur cutoff: Compute the theoretical upper bound of the bulk eigenvalue distribution under a null (noise) model and retains only components whose eigenvalues exceed that bound. Above the bound, components are treated as noise (*91*).

Kaiser rule (Kaiser-Guttman criterion): Retain those components whose eigenvalue is equal to or greater than 1 (for correlation-matrix PCA) by the logic that each component should explain at least as much as one original variable (*92*).

Broken Stick rule: Compare the observed variances explained by each component to the expected variances from a "broken-stick" distribution, generated by randomly distributing total variance amongst components. Retain components whose observed variance exceeds the broken-stick expectation (*93*).

Thomas Minka's Maximum Likelihood Estimation rule: Treats PCA as probabilistic/density estimation. First compute the Bayesian model evidence for different numbers of components, then select the k that maximizes evidence; effectively an automatic "maximum-likelihood" choice of dimension (*94*).

Optimal Hard Threshold for singular values: Derived in high-dimensional settings, this method sets a fixed threshold on singular-values above which components are retained. Below that threshold, components are considered noise (*95*). We used the optht python library optht.optht.

Horn's Parallel Analysis: Create many random datasets (same size & variable count) of uncorrelated variables, compute their eigenvalues, and compare to observed eigenvalues. Then retain components whose eigenvalues exceed the corresponding percentile of the random-data eigenvalues (*96*). We used the horns python library function horns.parallel_analysis.

For clustering followed by the above PCA dimensionality reduction, we used the scikit-learn sklearn.cluster library to run and test all clustering algorithms listed below. Unless otherwise specified below, default parameters were used. We varied all n_cluster parameters from 1 to 10.

KMeans: Partitions the data into *k* clusters by repeatedly assigning points to the nearest cluster centroid and updating centroids to minimize the within-cluster sum of squared distances. The n_clusters parameter selects the number of clusters. The n_init selects the number of times to run with different initial seeds and was set to 1 (*97*).

Agglomerative/Hierarchical Clustering: Starts with each data point as its own cluster and successively merges the two closest clusters according to a chosen linkage criterion until a stopping condition. We tested each linkage type ('ward', 'complete', 'average', 'single'), treating it like a parameter and selecting the best performing linkage type for each dataset. The n_clusters parameter selects the number of clusters. There is also the option to restrict the data points that get merged via the nearest neighbor's connectivity graph. We did so for each linkage type, varying the number of nearest neighbors by a factor of 20 (*98*, *99*).

Spectral Clustering: Constructs a similarity (affinity) graph of the data, computes an eigenvector embedding of the graph Laplacian, then clusters points in that embedding via KMeans, discretization, or cluster qr methods to capture non-convex structure (*59*, *100*, *101*). The n_clusters parameter selects the number of clusters. We varied the number of nearest neighbors used by a factor of 20 when computing the nearest neighbor's graph for use as the affinity graph. We tested all three methods of assigning cluster points. We used the eigen_solver 'arpack'.

DBSCAN: Finds clusters as regions of high point-density by grouping points that have at least a minimum number of neighbors within a radius ε and marks low-density points as noise. We systematically varied core parameters eps (maximum radius between neighbors) and min_samples (minimum number of points within that radius to be a "core" point) which control how dense a region must be to form a cluster and how many clusters and outliers result by factors of 10 (*102*, *103*).

HDBSCAN: Builds on DBSCAN by constructing a hierarchical tree of density-based clusters over varying ε values and then selecting flat clustering based on cluster stability (*104*, *105*). We varied the min_cluster_size (minimum size for any cluster) and min_samples (minimum number of neighbours for a core point, defaults to min_cluster_size if unspecified) parameters by factors of 10.

OPTICS: Orders points by reachability distance to reveal the clustering structure at multiple density levels and can extract clusters at varying densities (thus generalising DBSCAN's fixed ε approach). We varied the parameters min_samples (minimum number of points in a neighbourhood for a core point), xi (steepness threshold controlling cluster boundary extraction) and min_cluster_size (minimum size for extracted clusters) which govern core-point selection, maximum radius, and extraction sensitivity by factors of 10 (*106*).

The reason we tested several criteria for number of retained principal components is that we observed that retaining too many PCs can result in over-clustering. We observed that the number of clusters is directly proportional to the number of PCs retained for a wide range. Therefore, retaining too many PCs causes the number of recovered clusters to go well beyond the ground truth number of clusters. Hence, selecting the correct number of PCs is crucial.

We also selected the number of nearest neighbors in a range without values close to 0 or number of samples. This is because these end ranges result in significantly lower performance metrics.

We measured three metrics for dimensionality reduction and clustering algorithm performance: BestK, Silhouette Score, and Adjusted Rand Index. BestK is the number of identified clusters (number of unique cluster labels). Silhouette Score is a commonly accepted intrinsic measure of cluster separation. We utilized the cosine distance metric for measuring silhouette score for all analyses. Adjusted Rand Index measures the similarity of clustering labels to the ground truth labels. We used the scikit-learn sklearn.metrics.silhouette_score and sklearn.metrics.adjusted_rand_score functions to compute them.

We observed that several dimensionality reduction methods (excluding the 80%/90%/95% empirical rules) similarly selected the correct number of PCs for downstream clustering, so we used the widely employed and computationally simple scree plot elbow method. For a given algorithm on a dataset, the optimal parameters were determined by the highest silhouette score over grid search. Then, the BestK and Adjusted Rand index values were calculated for those parameters. We measured overall performance for each clustering algorithm across the three datasets by calculating the area under the curve of each metric across dataset parameters.

Electrophysiology data

We used an Allen Institute Neuropixels dataset (Allen Brain Observatory – Neuropixels Visual Coding) freely available to the public (https://allensdk.readthedocs.io/en/latest/visual_coding_neuropixels.html). Specifically, we used the 'Functional Connectivity' dataset of neural activity recorded from up to six Neuropixels probes while head-fixed mice were exposed to a variety of visual stimuli, including a 30-

minute block of mean luminance gray screen. Analyses were conducted on this 30-minute block. We followed the same data preprocessing and exclusion criteria as previously published (cite). In total, 20 mice (16 males) were used for the clustering, including 13 wild-type (WT), two Pvalb-IRES-Cre/wt;Ai32(RCL-ChR2(H134R) EYFP)/wt, three Sst-IRES-Cre/wt;Ai32(RCL-ChR2(H134R) EYFP)/wt, and two Vip-IRES-Cre/wt;Ai32(RCL-ChR2(H134R) EYFP)/wt mice. We also used a previously published pipeline for hippocampal sharp wave ripple detection (cite). Three types of ripples were identified: dorsal hippocampal ripple, intermediate hippocampal ripple, and global hippocampal ripple. Clustering was performed on PETHs of visual cortical neurons aligned to dorsal and intermediate ripples, but cross-correlation analysis during ripples was performed on global ripples. Details of clustering and cross-correlation analysis are described later.

Behavior and imaging data (Fig 2-4)

All experimental procedures were approved by the Institutional Animal Care and Use Committee of the University of North Carolina and accorded with the Guide for the Care and Use of Laboratory Animals (National Institutes of Health). These data were collected as described previously (*48*). Briefly, adult male wild type C57BL/6J mice (Jackson Laboratories, 6-8 weeks, 20-30 g) were group housed with littermates and acclimatized to the animal housing facility until surgery. Survival surgeries were stereotaxically performed while maintaining sterility. Following recovery, animals used for behavioral studies were water deprived to reach 85-90% of their pre-deprivation weight and maintained in a state of water deprivation for the duration of the behavioral experiments.

Water deprived mice were first trained to lick for random unpredictable sucrose (10-12.5%, ~2.5 μL) deliveries in a conditioning chamber. On subsequent days, mice received one of two possible auditory tones (3 kHz pulsing tone or 12 kHz constant tone, 75-80 dB) that lasted for 2 seconds. A second after the cues turned off, the mice received a sucrose reward following one of the tones (CS+), while the other tone resulted in no reward (CS-). The identity of the tones was counterbalanced across mice. A total of 50 CS+ and CS- each were presented in a pseudorandom order. The intertrial interval between two consecutive presentations of the cues was drawn from a truncated exponential distribution with mean of 30 s and a maximum of 90 s, with an additional 6 s constant delay. Animals were considered trained after the area under a receiver operating curve between anticipatory licks on CS+ vs CS- trials became larger than 0.7 on at least 2 consecutive sessions or larger than 0.85. Two contingency degradation experiments were performed as described previously (*48*), with reward probability reduced 50% in one (**Fig 3E**) and background unpredictable rewards introduced in the intertrial interval in the other while CS+ fully predicted reward. After retraining with full contingency, the CS+-reward pairing was extinguished for at least two days.

GCaMP6s was virally expressed under a CaMKIIα promoter to image single-cell calcium changes using 2-photon microscopy through a gradient refractive index (GRIN) lens implanted in the OFC (*48*). Data were acquired with a 6 frame online averaging at an effective frame rate of 5 Hz. Motion correction and cell extraction were performed as explained previously (*48*). Imaging acquisition was triggered by a custom Arduino code and a TTL output of every frame was sent as an input to the Arduino. The imaging acquisition was triggered off at the end of the behavioral session. In every mouse, one z-plane was imaged throughout acquisition so that the same cells could be tracked through learning. Only such longitudinally tracked neurons are included here.

Data Analysis

All statistical tests involving multiple comparisons utilized the Benjamini-Hochberg correction.

Identifying task responsive OFC neurons

We used a Gaussian family GLM with an identity link function to filter out neurons that were inactive during the task. Continuous fluorescence was modeled under a normality assumption for simplicity. Each neuron's median normalized fluorescence [-3,17] s from cue onset was selected and

concatenated. Median normalization was done as (F(t)-median(F(t)))/(max(F(t))-min(F(t)) such that normalized fluorescence was close to 0 except during calcium events. The design matrix included an indicator variable and an intercept term. The indicator represented CS+ cues with a 1 and CS- trials with a 0 at each timepoint throughout the entire window. Parameters were estimated by maximum likelihood via iteratively reweighted least squares using the statsmodels python library GLM function, and parameter inference used t-tests with two-sided *p*-value. We identified active OFC neurons by keeping neurons whose indicator variable parameter had an associated p value less than 0.05. The GLM equation for a single neuron is as shown as below:

$$F(t) = F_0 + \beta_{cue}\mathrm{I}_{cue}(t) + \varepsilon$$

Where F(t) represents median normalized fluorescence on frame t, ß represents the coefficients, I represents the indicator variable, and ε represents the error term. A total of 925 task responsive longitudinally tracked OFC neurons were used for analysis.

Clustering of OFC neurons

We clustered all OFC neurons pooled across animals that were recorded on the trained day (i.e., not just the longitudinally tracked neurons analyzed hereafter) by their median normalized fluorescence after filtering task inactive neurons using the above GLM. After filtering, the averaged fluorescence during the [-3,17] second window relative to cue onset was calculated separately for CS+ (rewarded) and CS- (unrewarded) trial types. The average fluorescence timeseries for each trial type was concatenated and z-normalized, comprising a 200-column vector of datapoints for each neuron.

As determined by our analysis in **Fig 1**, we applied principal component analysis for dimensionality reduction and kept the number of principal components determined by the scree plot elbow. We subsequently clustered neurons using the discrete spectral clustering algorithm variant using the scikit-learn function sklearn.cluster.SpectralClustering with the affinity matrix calculated using a k-nearest neighbor connectivity matrix. We tested the following numbers of nearest neighbors based on earlier ground truth simulations: 100, 200, 300, 500, 750, 1000, 1500, 2000, 2500. The number of clusters was also systematically varied from 2 to 20. The best combination of parameters was chosen by maximizing the silhouette score over a grid search of parameters. Each combination of parameters was run 100 times with a different random seed due to the stochasticity of the discrete spectral clustering algorithm. All subsequent analyses were repeated across the 100 iterations of cluster labels that were generated for the optimal parameters.

Clustering of Visual cortex neurons

The averaged firing rates of visual cortical neurons during the [-1.5, 1.5] second window relative to ripple onset were calculated separately for dCA1 and iCA1 ripple types. The averaged firing rate timeseries for each ripple type were concatenated and z-normalized for each neuron, comprising a single dataset. We followed the same dimensionality reduction and cluster procedure as above. Labels were shown to be stable across random iterations.

Correlation and Cross Correlation Analysis

To estimate the lag 0 correlation during task periods and inter trial intervals within and between assemblies, we calculated all pairwise correlations between neurons using the deconvolved fluorescence-inferred spiking activity during all time windows for each animal. Time windows for task periods spanned [0,17] s from cue onset, and time windows for inter trial intervals spanned 17 seconds after cue onset to the time of next cue onset. Inter trial intervals that were shorter than three frames were excluded. We then employed the hierarchical bootstrap sampling method directly on these pairwise correlations to get an estimate for the true lag 0 correlations for a given assembly. For each assembly, we hierarchically sampled animals and then neuron pairs for 10,000 bootstrapped samples. To estimate the within-assembly correlation, we sampled all pairs of neurons where both neurons

belonged to a given assembly. To estimate between-assembly correlations, we sampled neuron pairs where one neuron belonged to an assembly, and the other neuron did not. For computing between-assembly lag 0 correlations for a specific assembly pair, we sampled neuron pairs where one neuron belonged to the first assembly, and the other neuron belonged to the second.

To estimate cross correlations, the area under the curve (AUC), and centroid time derived from the cross-correlation curves between pairs of assemblies, we averaged the deconvolved activity within a single time window of all neurons grouped in an assembly and then computed the cross correlation between pairs of assemblies' average activity with up to 4 seconds of lag on individual time windows. We used the same time windows as above for task periods and ITI periods. ITI periods shorter than 4 seconds were excluded from the analysis. Then we computed the cross-correlation AUC and centroid time for each time window. We z-normalized each cluster time series then computed the cross correlogram with the NumPy function correlate in default mode 'valid'. Next, we hierarchically sampled animals and then time windows to come up with a bootstrapped estimate of AUC or centroid time for each assembly pair. We generated 10,000 bootstrapped samples to estimate means and 95% confidence intervals. 95% confidence intervals were calculated as the 2.5 and 97.5 percentiles of the bootstrapped distribution.

Lag 0 and cross correlation analyses were repeated using the same hierarchical bootstrap procedure for visual cortex neurons from the Functional Connectivity dataset. We calculated Lag 0 correlations within and between pairs of visual cortex cluster single neurons using activity during all time windows outside of any [-1.5,1.5] s window surrounding ripples. We term these windows outside of ripple windows as spontaneous windows. Adjacent ripples with overlapping windows were joined together. Cross correlations were computed using 4 seconds of lag during ripple windows and spontaneous windows.

References

1. R. Yuste, R. Cossart, E. Yaksi, Neuronal ensembles: Building blocks of neural circuits. *Neuron* **112**, 875–892 (2024).
2. L. Carrillo-Reid, S. Han, W. Yang, A. Akrouh, R. Yuste, Controlling Visually Guided Behavior by Holographic Recalling of Cortical Ensembles. *Cell* **178**, 447-457.e5 (2019).
3. L. Carrillo-Reid, R. Yuste, Playing the piano with the cortex: role of neuronal ensembles and pattern completion in perception and behavior. *Current Opinion in Neurobiology* **64**, 89–95 (2020).
4. D. O. Hebb, *The Organization of Behavior; a Neuropsychological Theory* (Wiley, Oxford, England, 1949)*The organization of behavior; a neuropsychological theory*.
5. C. Liu, R. Todorova, W. Tang, A. Oliva, A. Fernandez-Ruiz, Associative and predictive hippocampal codes support memory-guided behaviors. *Science* **382**, eadi8237 (2023).
6. N. T. M. Robinson, L. A. L. Descamps, L. E. Russell, M. O. Buchholz, B. A. Bicknell, G. K. Antonov, J. Y. N. Lau, R. Nutbrown, C. Schmidt-Hieber, M. Häusser, Targeted Activation of Hippocampal Place Cells Drives Memory-Guided Spatial Behavior. *Cell* **183**, 1586-1599.e10 (2020).
7. K. D. Harris, J. Csicsvari, H. Hirase, G. Dragoi, G. Buzsáki, Organization of cell assemblies in the hippocampus. *Nature* **424**, 552–556 (2003).
8. E. A. Aery Jones, L. M. Giocomo, Neural ensembles in navigation: From single cells to population codes. *Current Opinion in Neurobiology* **78**, 102665 (2023).
9. L. Carrillo-Reid, J.-E. K. Miller, J. P. Hamm, J. Jackson, R. Yuste, Endogenous sequential cortical activity evoked by visual stimuli. *J Neurosci* **35**, 8813–8828 (2015).
10. R. N. Ramesh, C. R. Burgess, A. U. Sugden, M. Gyetvan, M. L. Andermann, Intermingled Ensembles in Visual Association Cortex Encode Stimulus Identity or Predicted Outcome. *Neuron* **100**, 900-915.e9 (2018).
11. A. A. Fenton, H.-Y. Kao, S. A. Neymotin, A. Olypher, Y. Vayntrub, W. W. Lytton, N. Ludvig, 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 (2008).
12. E. Park, D. Dvorak, A. A. Fenton, Ensemble Place Codes in Hippocampus: CA1, CA3, and Dentate Gyrus Place Cells Have Multiple Place Fields in Large Environments. *PLOS ONE* **6**, e22349 (2011).
13. J. O'keefe, L. Nadel, *The Hippocampus as a Cognitive Map* (Oxford university press, 1978; https://discovery.ucl.ac.uk/id/eprint/10103569/1/HCMComplete.pdf).
14. P. R. Shirvalkar, Hippocampal Neural Assemblies and Conscious Remembering. *J Neurophysiol* **101**, 2197–2200 (2009).
15. H. Gelbard-Sagiv, R. Mukamel, M. Harel, R. Malach, I. Fried, Internally generated reactivation of single neurons in human hippocampus during free recall. *Science* **322**, 96–101 (2008).
16. E. Pastalkova, V. Itskov, A. Amarasingham, G. Buzsáki, Internally Generated Cell Assembly Sequences in the Rat Hippocampus. *Science* **321**, 1322–1327 (2008).
17. C. J. MacDonald, K. Q. Lepage, U. T. Eden, H. Eichenbaum, Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron* **71**, 737–749 (2011).
18. E. Kelemen, A. A. Fenton, Dynamic Grouping of Hippocampal Neural Activity During Cognitive Control of Two Spatial Frames. *PLOS Biology* **8**, e1000403 (2010).
19. J. Jackson, A. D. Redish, Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks. *Hippocampus* **17**, 1209–1229 (2007).
20. A. M. Wikenheiser, A. D. Redish, Decoding the cognitive map: ensemble hippocampal sequences and decision making. *Curr Opin Neurobiol* **32**, 8–15 (2015).
21. E. B. Knudsen, J. D. Wallis, Hippocampal neurons construct a map of an abstract value space. *Cell* **184**, 4640-4650.e10 (2021).

22. J. K. Miller, I. Ayzenshtat, L. Carrillo-Reid, R. Yuste, Visual stimuli recruit intrinsically generated cortical ensembles. *Proceedings of the National Academy of Sciences* **111**, E4053–E4061 (2014).
23. G. Dragoi, S. Tonegawa, Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature* **469**, 397–401 (2011).
24. U. Farooq, G. Dragoi, Emergence of preconfigured and plastic time-compressed sequences in early postnatal development. *Science* **363**, 168–173 (2019).
25. G. Dragoi, S. Tonegawa, Distinct preplay of multiple novel spatial experiences in the rat. *Proceedings of the National Academy of Sciences* **110**, 9100–9105 (2013).
26. A. Berners-Lee, T. Feng, D. Silva, X. Wu, E. R. Ambrose, B. E. Pfeiffer, D. J. Foster, Hippocampal replays appear after a single experience and incorporate greater detail with more experience. *Neuron* **110**, 1829-1842.e5 (2022).
27. D. Silva, T. Feng, D. J. Foster, Trajectory events across hippocampal place cells require previous experience. *Nat Neurosci* **18**, 1772–1779 (2015).
28. S. Cheng, L. M. Frank, New Experiences Enhance Coordinated Neural Activity in the Hippocampus. *Neuron* **57**, 303–313 (2008).
29. A. U. Sugden, J. D. Zaremba, L. A. Sugden, K. L. McGuire, A. Lutas, R. N. Ramesh, O. Alturkistani, K. K. Lensjø, C. R. Burgess, M. L. Andermann, Cortical reactivations of recent sensory experiences predict bidirectional network changes during learning. *Nat Neurosci* **23**, 981–991 (2020).
30. K. D. Harris, Neural signatures of cell assembly organization. *Nat Rev Neurosci* **6**, 399–407 (2005).
31. G. Dragoi, Cell assemblies, sequences and temporal coding in the hippocampus. *Current Opinion in Neurobiology* **64**, 111–118 (2020).
32. D. J. Foster, M. A. Wilson, Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**, 680–683 (2006).
33. K. Diba, G. Buzsáki, Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience* **10**, 1241–1242 (2007).
34. E. L. Denovellis, A. K. Gillespie, M. E. Coulter, M. Sosa, J. E. Chung, U. T. Eden, L. M. Frank, Hippocampal replay of experience at real-world speeds. *eLife* **10**, e64505 (2021).
35. M. F. Carr, S. P. Jadhav, L. M. Frank, Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nat Neurosci* **14**, 147–153 (2011).
36. S. P. Jadhav, C. Kemere, P. W. German, L. M. Frank, Awake Hippocampal Sharp-Wave Ripples Support Spatial Memory. *Science* **336**, 1454–1458 (2012).
37. M. P. Karlsson, L. M. Frank, Awake replay of remote experiences in the hippocampus. *Nat Neurosci* **12**, 913–918 (2009).
38. K. Kay, J. E. Chung, M. Sosa, J. S. Schor, M. P. Karlsson, M. C. Larkin, D. F. Liu, L. M. Frank, Constant Sub-second Cycling between Representations of Possible Futures in the Hippocampus. *Cell* **180**, 552-567.e25 (2020).
39. D. J. Foster, M. A. Wilson, Hippocampal theta sequences. *Hippocampus* **17**, 1093–1099 (2007).
40. W. E. Skaggs, B. L. McNaughton, M. A. Wilson, C. A. Barnes, Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**, 149–172 (1996).
41. R. C. Wilson, Y. K. Takahashi, G. Schoenbaum, Y. Niv, Orbitofrontal cortex as a cognitive map of task space. *Neuron* **81**, 267–279 (2014).
42. M. P. H. Gardner, G. Schoenbaum, The orbitofrontal cartographer. *Behav Neurosci* **135**, 267–276 (2021).
43. V. M. K Namboodiri, G. D. Stuber, The learning of prospective and retrospective cognitive maps within neural circuits. *Neuron* **109**, 3552–3575 (2021).
44. E. B. Knudsen, J. D. Wallis, Taking stock of value in the orbitofrontal cortex. *Nat Rev Neurosci* **23**, 428–438 (2022).

45. J. D. Wallis, Cross-species studies of orbitofrontal cortex and value-based decision-making. *Nat Neurosci* **15**, 13–19 (2011).
46. S. Ballesta, W. Shi, K. E. Conen, C. Padoa-Schioppa, Values encoded in orbitofrontal cortex are causally related to economic choices. *Nature* **588**, 450–453 (2020).
47. C. Padoa-Schioppa, J. A. Assad, Neurons in the orbitofrontal cortex encode economic value. *Nature* **441**, 223–226 (2006).
48. V. M. K. Namboodiri, J. M. Otis, K. van Heeswijk, E. S. Voets, R. A. Alghorazi, J. Rodriguez-Romaguera, S. Mihalas, G. D. Stuber, Single-cell activity tracking reveals that orbitofrontal neurons acquire and maintain a long-term memory to guide behavioral adaptation. *Nat. Neurosci.* **22**, 1110 (2019).
49. P. Y. Wang, C. Boboila, M. Chin, A. Higashi-Howard, P. Shamash, Z. Wu, N. P. Stein, L. F. Abbott, R. Axel, Transient and Persistent Representations of Odor Value in Prefrontal Cortex. *Neuron* **108**, 209-224.e6 (2020).
50. B. F. Sadacca, H. M. Wied, N. Lopatina, G. K. Saini, D. Nemirovsky, G. Schoenbaum, Orbitofrontal neurons signal sensory associations underlying model-based inference in a sensory preconditioning task. *Elife* **7** (2018).
51. M. Malvaez, C. Shieh, M. D. Murphy, V. Y. Greenfield, K. M. Wassum, Distinct cortical–amygdala projections drive reward value encoding and retrieval. *Nat Neurosci* **22**, 762–769 (2019).
52. A. C. Sias, A. K. Morse, S. Wang, V. Y. Greenfield, C. M. Goodpaster, T. M. Wrenn, A. M. Wikenheiser, S. M. Holley, C. Cepeda, M. S. Levine, K. M. Wassum, A bidirectional corticoamygdala circuit for the encoding and retrieval of detailed reward memories. *eLife* **10**, e68617.
53. S. A. Romano, G. Sumbre, "Detection of Overlapping Neuronal Assemblies from Activity Recordings of Large Neuronal Populations by Means of Non-orthogonal Low-Dimensional State Spaces" in *Identification, Characterization, and Manipulation of Neuronal Ensembles*, L. Carrillo-Reid, Ed. (Springer US, New York, NY, 2025; https://doi.org/10.1007/978-1-0716-4208-5_6), pp. 139–166.
54. V. Lopes-dos-Santos, S. Ribeiro, A. B. L. Tort, Detecting cell assemblies in large neuronal populations. *Journal of Neuroscience Methods* **220**, 149–166 (2013).
55. J. Mölter, L. Avitan, G. J. Goodhill, Detecting neural assemblies in calcium imaging data. *BMC Biology* **16**, 143 (2018).
56. A. Peyrache, K. Benchenane, M. Khamassi, S. I. Wiener, F. P. Battaglia, Principal component analysis of ensemble recordings reveals cell assemblies at high temporal resolution. *J Comput Neurosci* **29**, 309–325 (2010).
57. J. Wagner-Carena, S. Kate, T. Riordan, R. Abbasi-Asl, J. Aman, A. Amster, A. L. Bodor, D. Brittain, J. A. Buchanan, M. A. Buice, D. J. Bumbarger, F. Collman, N. da Costa, D. J. Denman, S. de Vries, E. Joyce, D. Kapner, C. W. King, J. D. Larkin, J. Lecoq, G. Mahalingam, D. Millman, J. Mölter, C. Morrison, R. C. Reid, C. M. Schneider-Mizell, S. Daniel, S. Suckow, K. T. Takasaki, M. Takeno, R. Torres, D. Vumbaco, J. Waters, D. G. Wyrick, W. Yin, J. Zhuang, S. Mihalas, S. Berteau, Hebb's Vision: The Structural Underpinnings of Hebbian Assemblies. bioRxiv [Preprint] (2025). https://doi.org/10.1101/2025.04.24.649900.
58. K. Watanabe, T. Haga, M. Tatsuno, D. R. Euston, T. Fukai, Unsupervised Detection of Cell-Assembly Sequences by Similarity-Based Clustering. *Front. Neuroinform.* **13** (2019).
59. Yu, Shi, "Multiclass spectral clustering" in *Proceedings Ninth IEEE International Conference on Computer Vision* (2003; https://ieeexplore.ieee.org/document/1238361), pp. 313–319 vol.1.
60. J. H. Siegle, X. Jia, S. Durand, S. Gale, C. Bennett, N. Graddis, G. Heller, T. K. Ramirez, H. Choi, J. A. Luviano, P. A. Groblewski, R. Ahmed, A. Arkhipov, A. Bernard, Y. N. Billeh, D. Brown, M. A. Buice, N. Cain, S. Caldejon, L. Casal, A. Cho, M. Chvilicek, T. C. Cox, K. Dai, D. J. Denman, S. E. J. de Vries, R. Dietzman, L. Esposito, C. Farrell, D. Feng, J. Galbraith, M. Garrett, E. C. Gelfand, N. Hancock, J. A. Harris, R. Howard, B. Hu, R. Hytnen, R. Iyer, E. Jessett, K. Johnson, I. Kato, J. Kiggins, S. Lambert, J. Lecoq, P. Ledochowitsch, J. H. Lee, A. Leon, Y. Li, E.

Liang, F. Long, K. Mace, J. Melchior, D. Millman, T. Mollenkopf, C. Nayan, L. Ng, K. Ngo, T. Nguyen, P. R. Nicovich, K. North, G. K. Ocker, D. Ollerenshaw, M. Oliver, M. Pachitariu, J. Perkins, M. Reding, D. Reid, M. Robertson, K. Ronellenfitch, S. Seid, C. Slaughterbeck, M. Stoecklin, D. Sullivan, B. Sutton, J. Swapp, C. Thompson, K. Turner, W. Wakeman, J. D. Whitesell, D. Williams, A. Williford, R. Young, H. Zeng, S. Naylor, J. W. Phillips, R. C. Reid, S. Mihalas, S. R. Olsen, C. Koch, Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature* **592**, 86–92 (2021).

61. H. Jeong, V. M. K. Namboodiri, M. W. Jung, M. L. Andermann, Sensory cortical ensembles exhibit differential coupling to ripples in distinct hippocampal subregions. *Current Biology* **33**, 5185-5198.e4 (2023).
62. V. M. K Namboodiri, T. Hobbs, I. Trujillo-Pisanty, R. C. Simon, M. M. Gray, G. D. Stuber, Relative salience signaling within a thalamo-orbitofrontal circuit governs learning rate. *Current Biology* **31**, 5176-5191.e5 (2021).
63. P. Berens, J. Freeman, T. Deneux, N. Chenkov, T. McColgan, A. Speiser, J. H. Macke, S. C. Turaga, P. Mineault, P. Rupprecht, S. Gerhard, R. W. Friedrich, J. Friedrich, L. Paninski, M. Pachitariu, K. D. Harris, B. Bolte, T. A. Machado, D. Ringach, J. Stone, L. E. Rogerson, N. J. Sofroniew, J. Reimer, E. Froudarakis, T. Euler, M. R. Rosón, L. Theis, A. S. Tolias, M. Bethge, Community-based benchmarking improves spike rate inference from two-photon calcium imaging data. *PLOS Computational Biology* **14**, e1006157 (2018).
64. V. Saravanan, G. J. Berman, S. J. Sober, Application of the hierarchical bootstrap to multi-level data in neuroscience. *Neuron Behav Data Anal Theory* **3**, https://nbdt.scholasticahq.com/article/13927-application-of-the-hierarchical-bootstrap-to-multi-level-data-in-neuroscience (2020).
65. J. M. Otis, V. M. K. Namboodiri, A. M. Matan, E. S. Voets, E. P. Mohorn, O. Kosyk, J. A. McHenry, J. E. Robinson, S. L. Resendez, M. A. Rossi, G. D. Stuber, Prefrontal cortex output circuits guide reward seeking through divergent cue encoding. *Nature* **543**, 103–107 (2017).
66. Q. Zhang, J. Zhou, Adaptive reward representations integrate expected uncertainty signals in orbitofrontal cortex. *Science Advances* **11**, eadv9590 (2025).
67. L. N. Driscoll, N. L. Pettit, M. Minderer, S. N. Chettih, C. D. Harvey, Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex. *Cell* **170**, 986-999.e16 (2017).
68. M. E. Rule, T. O'Leary, C. D. Harvey, Causes and consequences of representational drift. *Current Opinion in Neurobiology* **58**, 141–147 (2019).
69. Y. Ziv, L. D. Burns, E. D. Cocker, E. O. Hamel, K. K. Ghosh, L. J. Kitch, A. El Gamal, M. J. Schnitzer, Long-term dynamics of CA1 hippocampal place codes. *Nat Neurosci* **16**, 264–266 (2013).
70. K. T. Jensen, N. Kadmon Harpaz, A. K. Dhawale, S. B. E. Wolff, B. P. Ölveczky, Long-term stability of single neuron activity in the motor system. *Nat Neurosci* **25**, 1664–1674 (2022).
71. W. A. Liberti, T. A. Schmid, A. Forli, M. Snyder, M. M. Yartsev, A stable hippocampal code in freely flying bats. *Nature* **604**, 98–103 (2022).
72. N. D. Nguyen, A. Lutas, O. Amsalem, J. Fernando, A. Y.-E. Ahn, R. Hakim, J. Vergara, J. McMahon, J. Dimidschstein, B. L. Sabatini, M. L. Andermann, Cortical reactivations predict future sensory responses. *Nature* **625**, 110–118 (2024).
73. J. Pérez-Ortega, T. Alejandre-García, R. Yuste, Long-term stability of cortical ensembles. *eLife* **10**, e64449 (2021).
74. J. Hirokawa, A. Vaughan, P. Masset, T. Ott, A. Kepecs, Frontal cortex neuron types categorically encode single decision variables. *Nature* **576**, 446–451 (2019).
75. D. Hocker, C. D. Brody, C. Savin, C. M. Constantinople, Subpopulations of neurons in lOFC encode previous and current rewards at time of choice. *bioRxiv*, 2021.05.06.442972 (2021).
76. A. Onken, J. Xie, S. Panzeri, C. Padoa-Schioppa, Categorical encoding of decision variables in orbitofrontal cortex. *PLOS Computational Biology* **15**, e1006667 (2019).
77. T. C. Blanchard, S. T. Piantadosi, B. Y. Hayden, Robust mixture modeling reveals category-free selectivity in reward region neuronal ensembles. *J Neurophysiol* **119**, 1305–1318 (2018).

78. D. Raposo, M. T. Kaufman, A. K. Churchland, A category-free neural population supports evolving demands during decision-making. *Nat Neurosci* **17**, 1784–1792 (2014).
79. M. T. Kaufman, M. K. Benna, M. Rigotti, F. Stefanini, S. Fusi, A. K. Churchland, The implications of categorical and category-free mixed selectivity on representational geometries. *Curr Opin Neurobiol* **77**, 102644 (2022).
80. L. Posani, S. Wang, S. P. Muscinelli, L. Paninski, S. Fusi, Rarely categorical, always high-dimensional: how the neural code changes along the cortical hierarchy. *bioRxiv*, 2024.11.15.623878 (2025).
81. Y. S. Jo, V. M. K. Namboodiri, G. D. Stuber, L. S. Zweifel, Persistent activation of central amygdala CRF neurons helps drive the immediate fear extinction deficit. *Nat Commun* **11**, 422 (2020).
82. S. Lecca, V. M. K. Namboodiri, L. Restivo, N. Gervasi, G. Pillolla, G. D. Stuber, M. Mameli, Heterogeneous Habenular Neuronal Ensembles during Selection of Defensive Behaviors. *Cell Rep* **31**, 107752 (2020).
83. S. C. Piantadosi, E. E. Manning, B. L. Chamberlain, J. Hyde, Z. LaPalombara, N. M. Bannon, J. L. Pierson, V. M. K Namboodiri, S. E. Ahmari, Hyperactivity of indirect pathway-projecting spiny projection neurons promotes compulsive behavior. *Nat Commun* **15**, 4434 (2024).
84. R. I. Grant, E. M. Doncheck, K. M. Vollmer, K. T. Winston, E. V. Romanova, P. N. Siegler, H. Holman, C. W. Bowen, J. M. Otis, Specialized coding patterns among dorsomedial prefrontal neuronal ensembles predict conditioned reward seeking. *eLife* **10**, e65764 (2021).
85. J. Y. Cohen, S. Haesler, L. Vong, B. B. Lowell, N. Uchida, Neuron-type-specific signals for reward and punishment in the ventral tegmental area. *Nature* **482**, 85–88 (2012).
86. S. A. Josselyn, S. Tonegawa, Memory engrams: Recalling the past and imagining the future. *Science* **367** (2020).
87. L. R. Squire, L. Genzel, J. T. Wixted, R. G. Morris, Memory Consolidation. *Cold Spring Harb Perspect Biol* **7**, a021766 (2015).
88. K. Ghandour, T. Haga, N. Ohkawa, C. C. A. Fung, M. Nomoto, M. R. Fayed, H. Asai, M. Sato, T. Fukai, K. Inokuchi, Parallel processing of past and future memories through reactivation and synaptic plasticity mechanisms during sleep. *Nat Commun* **16**, 3618 (2025).
89. I. Shainer, J. M. Kappel, E. Laurell, J. C. Donovan, M. W. Schneider, E. Kuehn, I. Arnold-Ammer, M. Stemmer, J. Larsch, H. Baier, Transcriptomic neuron types vary topographically in function and morphology. *Nature* **638**, 1023–1033 (2025).
90. R. B. Cattell, The Scree Test For The Number Of Factors. *Multivariate Behavioral Research* **1**, 245–276 (1966).
91. V. A. Marčenko, L. A. Pastur, DISTRIBUTION OF EIGENVALUES FOR SOME SETS OF RANDOM MATRICES. *Math. USSR Sb.* **1**, 457 (1967).
92. H. F. Kaiser, The application of electronic computers to factor analysis. *Educational and Psychological Measurement* **20**, 141–151 (1960).
93. D. A. Jackson, Stopping Rules in Principal Components Analysis: A Comparison of Heuristical and Statistical Approaches. *Ecology* **74**, 2204–2214 (1993).
94. T. Minka, “Automatic Choice of Dimensionality for PCA” in *Advances in Neural Information Processing Systems* (MIT Press, 2000; https://papers.nips.cc/paper_files/paper/2000/hash/7503cfacd12053d309b6bed5c89de212-Abstract.html)vol. 13.
95. Inference on the proportion of variance explained in principal component analysis. https://arxiv.org/html/2402.16725v1.
96. J. L. Horn, A rationale and test for the number of factors in factor analysis. *Psychometrika* **30**, 179–185 (1965).
97. J. MacQueen, “Some methods for classification and analysis of multivariate observations” in *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics* (University of California Press, 1967; https://projecteuclid.org/ebooks/berkeley-symposium-on-mathematical-statistics-and-probability/Proceedings-of-the-Fifth-Berkeley-

Symposium-on-Mathematical-Statistics-and/chapter/Some-methods-for-classification-and-analysis-of-multivariate-observations/bsmsp/1200512992)vol. 5.1, pp. 281–298.
98. J. H. Ward Jr., Hierarchical Grouping to Optimize an Objective Function. *Journal of the American Statistical Association* **58**, 236–244 (1963).
99. J. C. Gower, G. J. S. Ross, Minimum Spanning Trees and Single Linkage Cluster Analysis. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **18**, 54–64 (1969).
100. J. Shi, J. Malik, Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**, 888–905 (2000).
101. A. Damle, V. Minden, L. Ying, Simple, direct and efficient multi-way spectral clustering. *Inf Inference* **8**, 181–203 (2019).
102. H.-P. K. Martin Ester, A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise, *AAAI*. https://aaai.org/papers/kdd96-037-a-density-based-algorithm-for-discovering-clusters-in-large-spatial-databases-with-noise/.
103. E. Schubert, J. Sander, M. Ester, H. P. Kriegel, X. Xu, DBSCAN Revisited, Revisited: Why and How You Should (Still) Use DBSCAN. *ACM Trans. Database Syst.* **42**, 19:1-19:21 (2017).
104. R. J. G. B. Campello, D. Moulavi, J. Sander, “Density-Based Clustering Based on Hierarchical Density Estimates” in *Advances in Knowledge Discovery and Data Mining*, J. Pei, V. S. Tseng, L. Cao, H. Motoda, G. Xu, Eds. (Springer, Berlin, Heidelberg, 2013), pp. 160–172.
105. R. J. G. B. Campello, D. Moulavi, A. Zimek, J. Sander, Hierarchical Density Estimates for Data Clustering, Visualization, and Outlier Detection. *ACM Trans. Knowl. Discov. Data* **10**, 5:1-5:51 (2015).
106. M. Ankerst, M. M. Breunig, H.-P. Kriegel, J. Sander, OPTICS: ordering points to identify the clustering structure. *SIGMOD Rec.* **28**, 49–60 (1999).