

Organizing the coactivity structure of the hippocampus from robust to flexible memory

Giuseppe P. Gava*, Laura Lefèvre*, Tabitha Broadbelt, Stephen B. McHugh, Vítor Lopes-dos-Santos, Demi Brizee, Katja Hartwich, Hanna Sjöberg, Pavel V. Perestenko, Robert Toth, Andrew Sharott, and David Dupret.

*These authors contributed equally.

Medical Research Council Brain Network Dynamics Unit, Nuffield Department of Clinical Neurosciences, University of Oxford, Oxford, United Kingdom.

Science 2024, Vol 385, Issue 6713, pp. 1120-1127, DOI: 10.1126/science.adk9611

Corresponding author. Email: david.dupret@bndu.ox.ac.uk.

This article is licensed under [Creative Commons Attribution 4.0 International \(CC BY\)](https://creativecommons.org/licenses/by/4.0/).

New memories are integrated into prior knowledge of the world. But what if consecutive memories exert opposing demands on the host brain network? We report that acquiring a robust (food-context) memory constrains the hippocampus within a population activity space of highly correlated spike trains that prevents subsequent computation of a flexible (object-location) memory. This densely correlated firing structure developed over repeated mnemonic experience, gradually coupling neurons of the superficial CA1 *pyramidale* sublayer to whole population activity. Applying hippocampal theta-driven closed-loop optogenetic suppression to mitigate this neuronal recruitment during (food-context) memory formation relaxed the topological constraint on hippocampal coactivity and restored subsequent flexible (object-location) memory. These findings uncover an organizational principle for the peer-to-peer coactivity structure of the hippocampal cell population to meet memory demands.

Every day, we use our existing knowledge to guide the actions we make in our environment, integrating new information to gain further knowledge about the world. Therefore, building new memories does not take place in a state of *tabula rasa*, but against a background of prior experiences that have been accumulated across the lifespan and have shaped their host brain networks (1, 2).

The hippocampus network uses the collective activity of the population of its neurons to support everyday memory (3, 4). In principle, the level and structure of the activity coupling between individual neurons could reflect a critical tradeoff between the robustness versus the flexibility of the whole population in processing information. That is, strong peer-to-peer coupling could yield highly correlated spike trains, increasing the consistency of activity patterns within the population for robust memory expression. In contrast, weaker population coupling could release network activity space for new patterns, allowing more diverse mnemonic representations for dynamically adaptable behavior. Owing to convergent innervation on post-synaptic targets (e.g., subiculum, entorhinal cortex), adjusting population coupling to ongoing demand would influence information transmission of hippocampal inputs to downstream reader neurons (3, 5, 6). However, the hippocampus may have to switch between robust versus flexible computations depending on current demands. What are the consequences of placing the hippocampal population into a robust computational mode for subsequent memories that instead require flexible information processing?

To address this question, we first trained six mice (always fed *ad libitum*) to acquire a strong contextual memory. On 16 consecutive days, mice explored two arenas (Fig. 1A, B). During the first 10 days ('Food-context conditioning'), we paired one arena (context X) with two regular chow pellets (fig. S1A). The other arena (context Y) contained one chow pellet and one high-fat-diet (Hfd) pellet (fig. S1A), which mice encountered for the first time and did not eat much on day 1 (Fig. 1C). By repeating these foraging sessions on each subsequent conditioning day, mice showed escalated food intake in context Y (Fig. 1C; from day 1 to day 10, a fold-change of 21.62 ± 6.70 versus 0.63 ± 0.24 in context Y versus X; mean \pm s.e.m.). Food locations were randomized every day in each context to promote mnemonic association of food items not to fixed places but whole contexts (fig. S1B). To probe discriminative food-context association, we then measured the propensity of these mice to express context-biased feeding. By providing both arenas with new food items in post-conditioning days ('Novel food test'; days 11 and 12; Fig. 1B and fig. S1C), we observed higher novel food intake in context Y compared to context X (Fig. 1D; a fold-change of 3.47 ± 1.40 in context Y versus X; mean \pm s.e.m.). Thus, mice in the Hfd-conditioned context readily overcame the rodent natural tendency to express food neophobia (7, 8). Mice also exhibited lower Hfd intake when provided in a third arena (context Z) never paired with any food (fig. S1D). Body weight remained stable across days (fig. S1E).

We next switched task demand to assess novelty detection in these contexts

paired with food. For this, we used a hippocampus-dependent continuous Novel Object Recognition task ('cNOR test'; from day 13 to day 16; Fig. 1A, B and fig. S2A-C). On each cNOR day, we re-exposed mice to either context X or Y ('Re-exposure'; without food) before they encountered four novel objects ('Sampling'; Fig. 1A, B and fig. S2A, B). Across four more exploration sessions that day, we iteratively replaced one of the initially sampled objects with a new one ('Testing'). This procedure yielded a set of three familiar (already-seen) objects and one novel (first time-seen) object in each cNOR test (Fig. 1B and fig. S2A). We measured novelty detection in each cNOR test using the time spent exploring the novel object over the total time spent on all four objects, thereby probing memory for objects explored in session. Mice showed novel object preference in context X but not in context Y (Fig. 1E). Exploratory behavior measured by locomotor speed, distance travelled, and time spent with objects did not differ across contexts (fig. S2D-H). Altogether, the first demand in this 2-memory paradigm was for animals to repeatedly associate food resources to selective contexts across many days, yielding a robust memory able to shape contextual feeding. This later prevented mice from coping with a second demand: to dynamically update another memory for continually detecting novel items across many sessions each day. Thus, the robust (food) memory acquired in context Y interfered with the subsequent flexible (object-location) memory in that context.

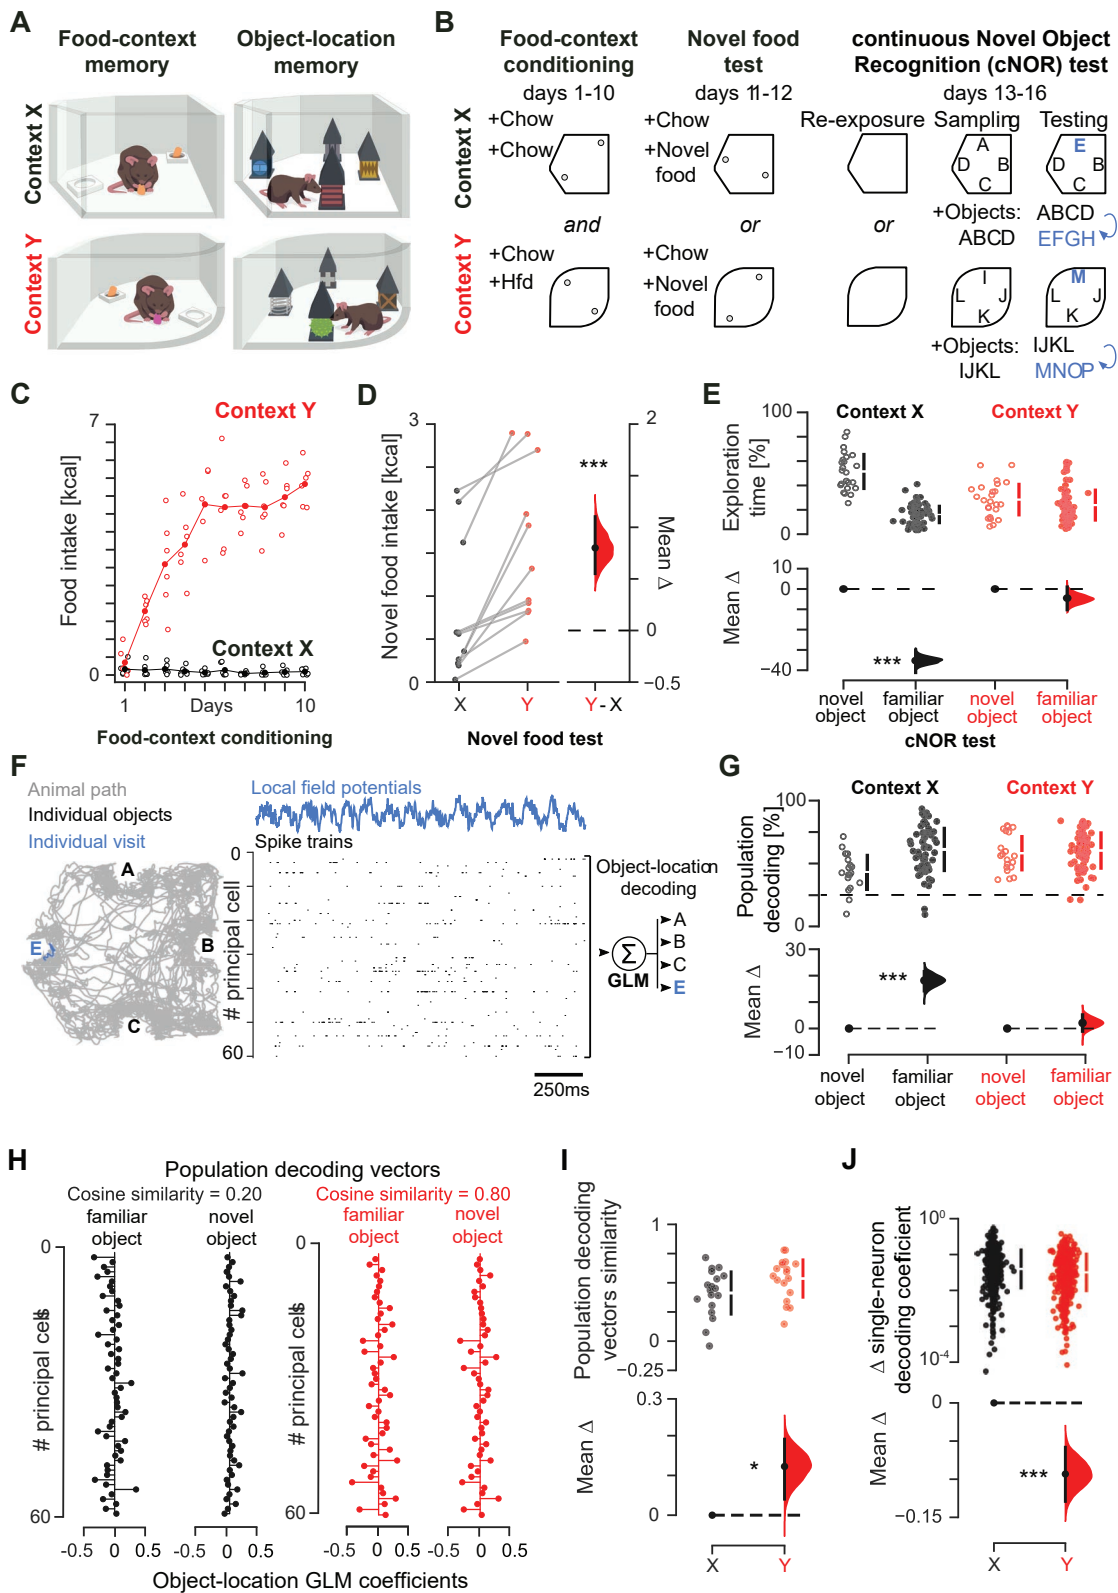


Fig. 1. Robust contextual (food) memory prevents subsequent flexible (object) memory.

(A, B) Behavioral tasks with open-field contexts (A) and multiday layout (B) for a two-memory paradigm. (C) Animals' food intake during contextual conditioning (each data point represents one mouse). (D) Estimation plot showing the effect size for the difference in novel food intake between context X and Y after conditioning. (E) Percentage of time exploring novel versus familiar objects during cNOR tests in context X or Y. (F) Schematic of the GLM predicting the identity of each object-location compound from population vectors of theta-nested principal cell spiking. A sample of CA1 ensemble spike data for one object visit (blue trace) is shown. (G) Estimation plot showing the classification accuracy of object-location compound in test n by GLM trained in session $n-1$ (each data point represents one object). (H) Example pairs of population decoding vectors containing the neuron-wise GLM coefficients for the familiar versus a novel object at the same location in context X versus Y (each data point represents one neuron). (I, J) Estimation plots showing the cosine similarity between familiar and novel object-location GLM vectors (I) and the change (update) in individual neuron contribution to population object-location decoding (J) across two consecutive cNOR tests in either context. For each estimation plot: *Upper (Left for D)*, raw data (points) with mean \pm SD (vertical lines); *Lower (Right for D)*, mean difference (black-dot) with 95% CI (black-ticks) and bootstrapped sampling-error distribution (filled-curve) with respect to the (left-most) group-reference (horizontal dashed line; see Methods). *** $P < 0.001$, * $P < 0.05$, two-sided paired permutation test. $N = 6$ mice, 2506 CA1 principal cells.

Robust memory increases neuronal coactivity and population coupling in the hippocampus.

We aimed to identify the neuronal correlates of this cross-memory interference. During active behavior, groups of principal cells recruited from the population of hippocampal neurons cooperate within the timeframe of theta-band (5–12Hz) oscillations to support codes and computations for memory (3). We recorded cell ensembles and local field potentials in the CA1 *stratum pyramidale* of these mice. Using the action potentials discharged by principal cells in theta cycles during exploration of each object in cNOR days (Fig. 1F and fig. S2I), a generalized linear model (GLM) trained on session and applied in test identified each object-location compound with up to 93.5% accuracy (range 9.4 – 93.5%; mean 55.0% compared to a chance level of 25.0%; mean \pm s.e.m. number of principal cells = 51.1 ± 3.7 per GLM), consistent with work showing population-level object representation in the hippocampus (9, 10). In context X, the mean accuracy of this population-level decoding started at $38.6 \pm 3.2\%$ for the novel object-location compound (Fig. 1G) and improved by the following trial (fig. S2J), indicating rapid mnemonic integration of each object. This across-test gain in object-location representation did not occur in context Y (Fig. 1G and fig. S2J). In fact, the mean decoding accuracy there started at higher levels for the novel object-location ($54.9 \pm 3.4\%$; $P < 0.001$, two-sided paired permutation test, compared to context X), without significant changes in the following tests (Fig. 1G and fig. S2J). In both context X and Y, object-location memory was not affected by provision of Hfd (in neutral context Z) beforehand that day (fig. S3). Compared to context X, the contribution of individual cells to each novel object-location decoding during test in context Y resembled its previous one expressed in session at that location while having another (familiar) object. This was reported by the higher similarity between the population decoding vector that contained the set of neuron-wise GLM coefficients representing the novel object during test in context Y versus that representing the familiar object at the same location during session (Fig. 1H, I). In line with this observation, the across-test modulation in single-neuron contributions to population decoding vectors when encountering a novel object (i.e., the changes in the magnitude of individual GLM coefficients) was weaker in context Y compared to X (Fig. 1J).

We next computed the CA1 place maps expressed during the cNOR task (Fig. 2A and fig. S4A). Context Y did not exhibit an over-representation of the randomized Hfd location (fig. S4B). By quantifying the cross-session similarity of place maps from re-exposure to sampling to individual

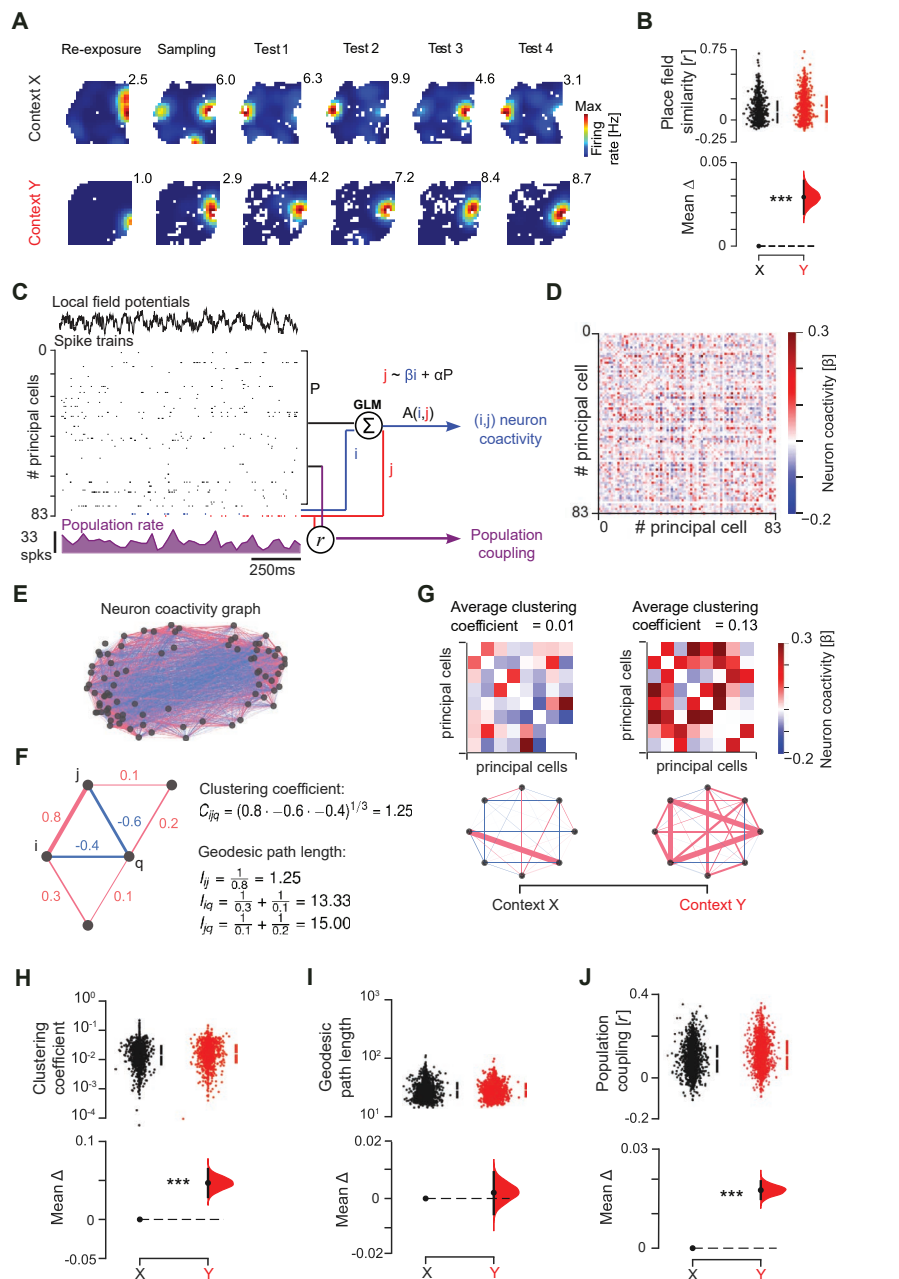


Fig. 2. Robust memory increases neuronal coactivity and population coupling. (A) Example firing maps across the consecutive cNOR sessions for one mouse day in context X (top) versus context Y (bottom). Each row shows one principal cell (numbers indicate peak rate for each map). (B) Estimation plot showing the place field similarity for the pairs of place maps expressed by individual cells across two contiguous cNOR sessions (e.g., Sampling and Test 1) in context X versus context Y (each data point represents one cell). (C) Schematic of the population-level analyses (see methods). Coactivity between any two neurons measured as the regression weight from the GLM assessing their firing relationship while accounting for network-level modulation using the sum of the remaining cells in the population (to estimate neuron pair coactivity beyond the population rate). Population coupling of each cell measured as the Pearson correlation coefficient between its theta-binned spike train and the cumulative activity of the remaining cells. (D-G) Example adjacency matrix of regression weights (D) and corresponding coactivity graph (E) using the procedure depicted in (C) to access the neuron pairwise coactivity structure of the population in each context. Example subset of a coactivity graph (F) with five neurons (nodes) and their pairwise coactivity values (edges with numbers); the clustering coefficient of neuron forming one example triad with neurons i and j is calculated along with the geodesic path lengths. Shown in (G) is a subset of adjacency matrices representing contexts X and Y (top), along with their average clustering coefficients and motifs of coactivity (bottom). (H-J) Estimation plots showing that the population coactivity structure is tighter in context Y than X, as reported by the higher clustering coefficient of neuronal graphs containing stronger triads of coactive neurons (H), without a significant change in geodesic path length (I), along with stronger population coupling (J). For each estimation plot: Upper, raw data (points) with mean \pm SD (vertical lines); Lower, mean difference (black-dot) with 95% CI (black-ticks) and bootstrapped sampling-error distribution (filled-curve) with respect to the (left-most) group-reference (horizontal dashed line; see Methods). *** $P < 0.001$, two-sided paired permutation test. $N = 6$ mice, 2506 CA1 principal cells.

object recognition tests on each day, we observed that CA1 principal cells exhibited higher place map stability (i.e., lower spatial remapping) across contiguous sessions in context Y (Fig. 2B and fig. S4C). With the GLM object-location decoding, this result on the hippocampal place code supported the notion of a more rigid memory in the Hfd-paired context Y, but a flexible (cross-session updated) memory in context X.

We hypothesized that this representational rigidity reflects the organization of the population activity into a non-permissive structure. An operational principle serving many brain functions, including memory, is to leverage the collective activity of neural populations as an emergent property beyond that of individual cells (11–14). We reasoned that Hfd context-conditioning yielded highly correlated firing patterns that created a dense network activity space for strong contextual (food) memory. But this later conflicted with the switch to a different demand where continually processing familiar versus novel stimuli would instead require sparser, weakly correlated patterns for disentangling discrete (object-location)

representations. The correlational structure of the population activity was markedly different in context Y compared to X (Fig. 2C–J). We quantified the coactivity association of each cell pair by predicting the theta-nested spike discharge of neuron from the activity of neuron while regressing out the activity of the remaining population (Fig. 2C). This procedure returned a matrix of regression weights (Fig. 2D) that represented the neurons pairwise coactivity structure of the population in each context. For both context X and Y we constructed weighted neuronal graphs (with no self-connections) where each node is a cell and the edge linking any two nodes represents the coactivity of that cell pair (Fig. 2D–G). Neuronal graphs contained stronger triads of coactive nodes in context Y than X, as reported by higher clustering coefficients (Fig. 2H and fig. S5A; mean increase (95% CI): 11.0% (7.3–14.1%)). The population coactivity strength level, calculated for each node as the average weight of all its edges, was higher in Y (fig. S5B), with no difference in the mean neurons' firing rate across contexts (fig. S5C). The hippocampal population exhibited this denser coactivity structure in context Y without a reduction

in geodesic path length, calculated as the mean shortest path between any two nodes (Fig. 2I and fig. S5A) (15). This suggested that the hippocampus, which usually displays the features of a small-world network (fig. S6A–C) that could allow for flexible information updating through efficient synchronization, has in fact acquired in context Y the rigidity of a more coherent, lattice-like network (fig. S6D–F) (16–18). These neuronal graphs that are composed of both correlated and anti-correlated spike trains (i.e., positive and negative edges) indeed showed more stable population activity patterns in context Y, as suggested by higher structural balance (fig. S6G–I). These topological alterations developed across the 10 conditioning days (fig. S7A) to continue altering the level and structure of population coactivity in post-conditioning days (fig. S7B), affecting the baseline re-exposure to context Y prior to any testing. Hippocampal graphs yet shared some common correlation structure across contexts (fig. S8), suggesting a coactivity backbone for cross-context generalization.

To explore the development of a dense population activity structure in context Y, we

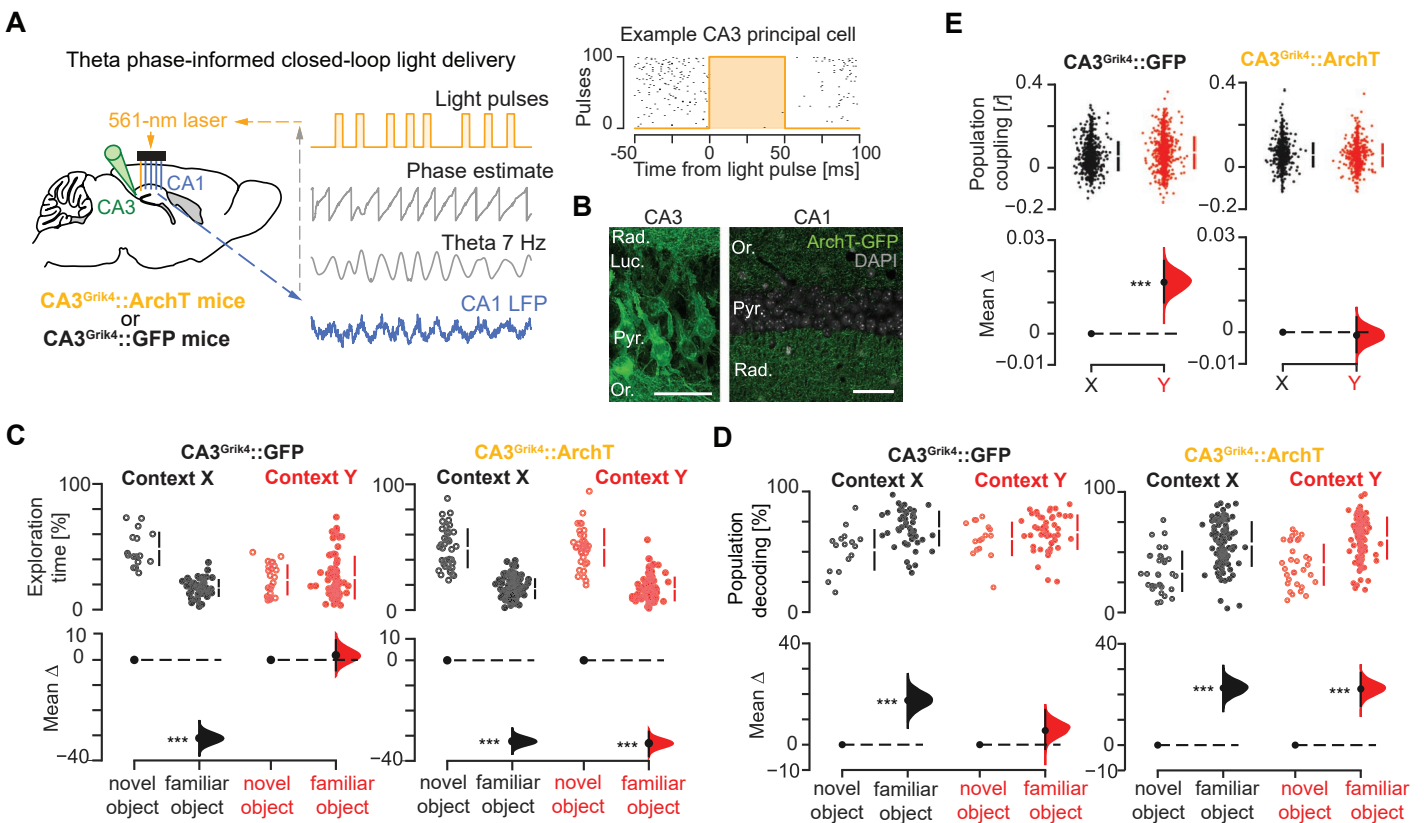


Fig. 3. CA1 population coupling requires CA3 during food-context conditioning.

(A, B) Optogenetic targeting of CA3 with either ArchT-GFP (in CA3^{Grik4}::ArchT mice) or GFP-only (in CA3^{Grik4}::GFP mice) combined with theta phase-informed light delivery for closed-loop suppression of CA3 during Hfd-context Y conditioning (A). CA3 principal cells transduced with ArchT-GFP (B, left) and projecting to CA1 (B, right; Stratum oriens, Or.; pyramidale, Pyr.; radiatum, Rad.; lucidum, Luc.; cell nuclei stained with DAPI; scale bar, 50 μ m). (C–E) Estimation plots showing that in CA3^{Grik4}::ArchT mice, but not in control CA3^{Grik4}::GFP mice, applying this optogenetic intervention throughout the 10-day Hfd conditioning subsequently restored in context Y the behavioral cNOR performance (C) with CA1 population object-location decoding (D) and activity coupling (E) to levels seen in context X. For each estimation plot: Upper, raw data (points) with mean \pm SD (vertical lines); Lower, mean difference (black-dot) with 95% CI (black-ticks) and bootstrapped sampling-error distribution (filled-curve) with respect to the (left-most) group-reference (horizontal dashed line; see Methods). *** $P < 0.001$, two-sided paired permutation test. $N = 1548$ CA1 neurons from 4 CA3^{Grik4}::ArchT mice versus 881 CA1 neurons from 2 CA3^{Grik4}::GFP mice.

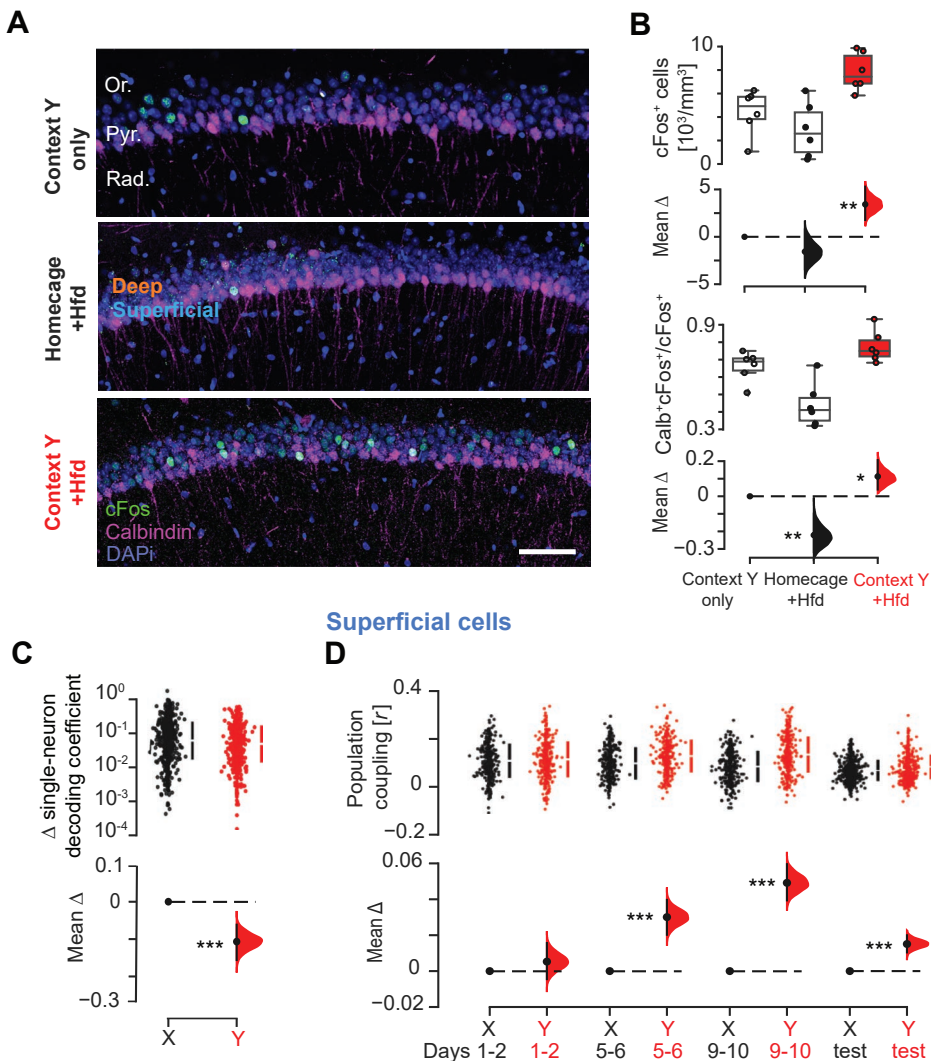


Fig. 4. Contextual food memory recruits CA1 superficial pyramidal sublayer cells.

(A, B) cFos-expressing CA1 neurons with Calbindin-1 delineated superficial pyramidal sublayer for three mice after 10-day exposure to either context Y only, homeage with Hfd, or context Y with Hfd (A; stratum oriens, Or.; pyramidal, Pyr.; radiatum, Rad.; cell nuclei stained with DAPI; scale bar, 50 μm ; see also fig. S12), and quantification of cFos⁺ cell density in CA1 pyramidal (B; top; each datum represents one mouse; N = 18 mice, 6 mice per condition) with corresponding proportion of cFos⁺ Calbindin⁺ cells in the CA1 superficial pyramidal sublayer (B; bottom). (C, D) Estimation plots showing that CA1 superficial pyramidal sublayer cells have reduced cross-test change (update) in their contribution to object-location decoding (C) and increased population coupling (D) in context Y compared to context X (N = 6 mice, 1871 CA1 superficial cells). ***P < 0.001, **P < 0.01, *P < 0.05, two-sided paired permutation test.

investigated the one-to-many relationship between individual neurons and the rest of the population. We measured the coupling of each principal cell instantaneous firing rate in theta cycles to the concomitant summed activity of all other recorded cell members of the population ('population rate,' Fig. 2C). Consistent with the higher topological clustering (Fig. 2H), the average population coupling of individual neurons was stronger in context Y (Fig. 2J and fig. S5A; mean increase (95% CI): 16.6% (12.6–20.8%)). This increased population coupling reflected a stronger cross-neuron spiking relationship: shuffling the spike times across neurons and theta cycles, while preserving each neuron's mean rate and the population

rate distribution, cancelled the increased population coupling seen in context Y (fig. S9A). This heightened coupling developed across the 10 conditioning days to mark the re-exposure to context Y during post-conditioning days even before any test (fig. S9B). Restricting food-context conditioning to two days in five additional mice did not alter subsequent cNOR memory (fig. S10A–D), allowing successful object-location decoding in both contexts with similar CA1 population activity level and structure (fig. S10E–J). Contextual food conditioning thus seemed to increase the recruitment of principal cells as 'choristers of a larger hippocampal orchestra' (19) when daily repeated for well over a week.

Mitigating neuronal recruitment during robust memory formation relaxes hippocampal coactivity and restores flexible memory.

We sought to relate hippocampal population activity to memory expression more directly by first manipulating an underlying neural pathway. The CA3 region features extensive excitatory recurrent connections and Hebbian synaptic plasticity (20–23). It could therefore promote population coupling in the downstream CA1, which has little to no recurrent excitation (6, 24). To test this, we transduced CA3 principal cells with the neural silencer Archaeorhodopsin-T in four Grik4-Cre mice versus the GFP-only control in two Grik4-Cre mice (Fig. 3A, B). Bilateral implantation of tetrodes combined with optic fibers allowed monitoring of CA1 ensembles while actuating a theta phase-informed controller for real-time suppression of CA3 principal cells (Fig. 3A and fig. S11A–C). In CA3^{Grik4::ArchT} mice, but not CA3^{Grik4::GFP} mice, applying this closed-loop intervention during each Hfd conditioning session subsequently restored in post-conditioning test days object-location memory with CA1 place map similarity (cross cNOR-session remapping) and population coupling in context Y to levels seen in context X (Fig. 3C–E and fig. S11D–H).

We then examined the CA1 population after contextual food conditioning. Using the activity-dependent immediate-early-gene cFos, we quantified neuron recruitment in the CA1 pyramidal layer of six mice exposed to context Y with Hfd for 10 days (Fig. 4A and fig. S12). In parallel, six control mice explored context Y without food while six others ate Hfd in their homeage. Mice undergoing Hfd-context Y conditioning showed higher density of cFos⁺ neurons compared to both controls (Fig. 4B). Contrasting cFos expression in the superficial versus the deep pyramidal sublayers using the marker Calbindin 1 (Fig. 4A) (25–27) suggested that Hfd-context conditioning preferentially recruited CA1 superficial cells (Fig. 4B). This is consistent with recent studies highlighting that CA1 pyramidal cells segregate along the anatomical axes of the hippocampus (e.g., with respect to molecular markers, neural connectivity, and electrophysiological properties), indicating a functional specialization based on somatic location (25–35).

We identified neurons recorded in the CA1 superficial pyramidal sublayer, using the electrophysiological profile of each tetrode (fig. S13) (15). The representational rigidity affecting object-location memory update in context Y (Fig. 1J) corresponded to reduced modulation in the individual contribution of superficial cells to the population decoding from cNOR session to test (Fig. 4C and fig. S14A). This suggested that superficial cell population did not distinguish the novel object from the

familiar previously encountered at the same location in context Y, only representing the location itself. In line with this, the spatial map rigidity (i.e., higher place field stability) observed across cNOR sessions in the Hfd-conditioned context Y (Fig. 2B) was explained by superficial cell maps (fig. S14B), with no representational bias to the Hfd location (fig. S14C). Superficial cells formed stronger coactivity triads in the network (fig. S15A-E), increasing their population coupling in context Y (Fig. 4D and fig. S15F-H). Deep *pyramidale* sublayer cells did not show such enhanced activity profile (fig. S15). The CA1 representational

rigidity that prevented flexible object-location memory in context Y following Hfd conditioning was thus primarily explained by enhanced recruitment of superficial *pyramidale* sublayer cells.

We hypothesized that harnessing the rise in population coupling during robust contextual conditioning would restore subsequent flexible memory. We used an intersectional optogenetic strategy to target cells recruited in the CA1 superficial *pyramidale* sublayer during Hfd-context Y conditioning. CA1 superficial *pyramidale* sublayer cells are genetically defined by the molecular marker Calbindin-1 (25, 27,

34). They express cFos during contextual Hfd conditioning (Fig. 4A, B). We thus bred double-transgenic Calb1-Cre;cFos-tTA mice and generated a viral construct for the two-term Boolean logic (36) expression of the yellow light-driven neural silencer ArchT-EYFP (or its EYFP-only control) dependent on the two recombinases Cre and FlpO (Fig. 5A). We transduced the CA1 of these mice with this construct together with a second construct allowing the tTA-dependent expression of FlpO (Fig. 5A) for lasting optogenetic tagging of CA1 superficial *pyramidale* sublayer cells with either ArchT-EYFP (in five ContextY::CA1^{Calb1-cFos}::ArchT

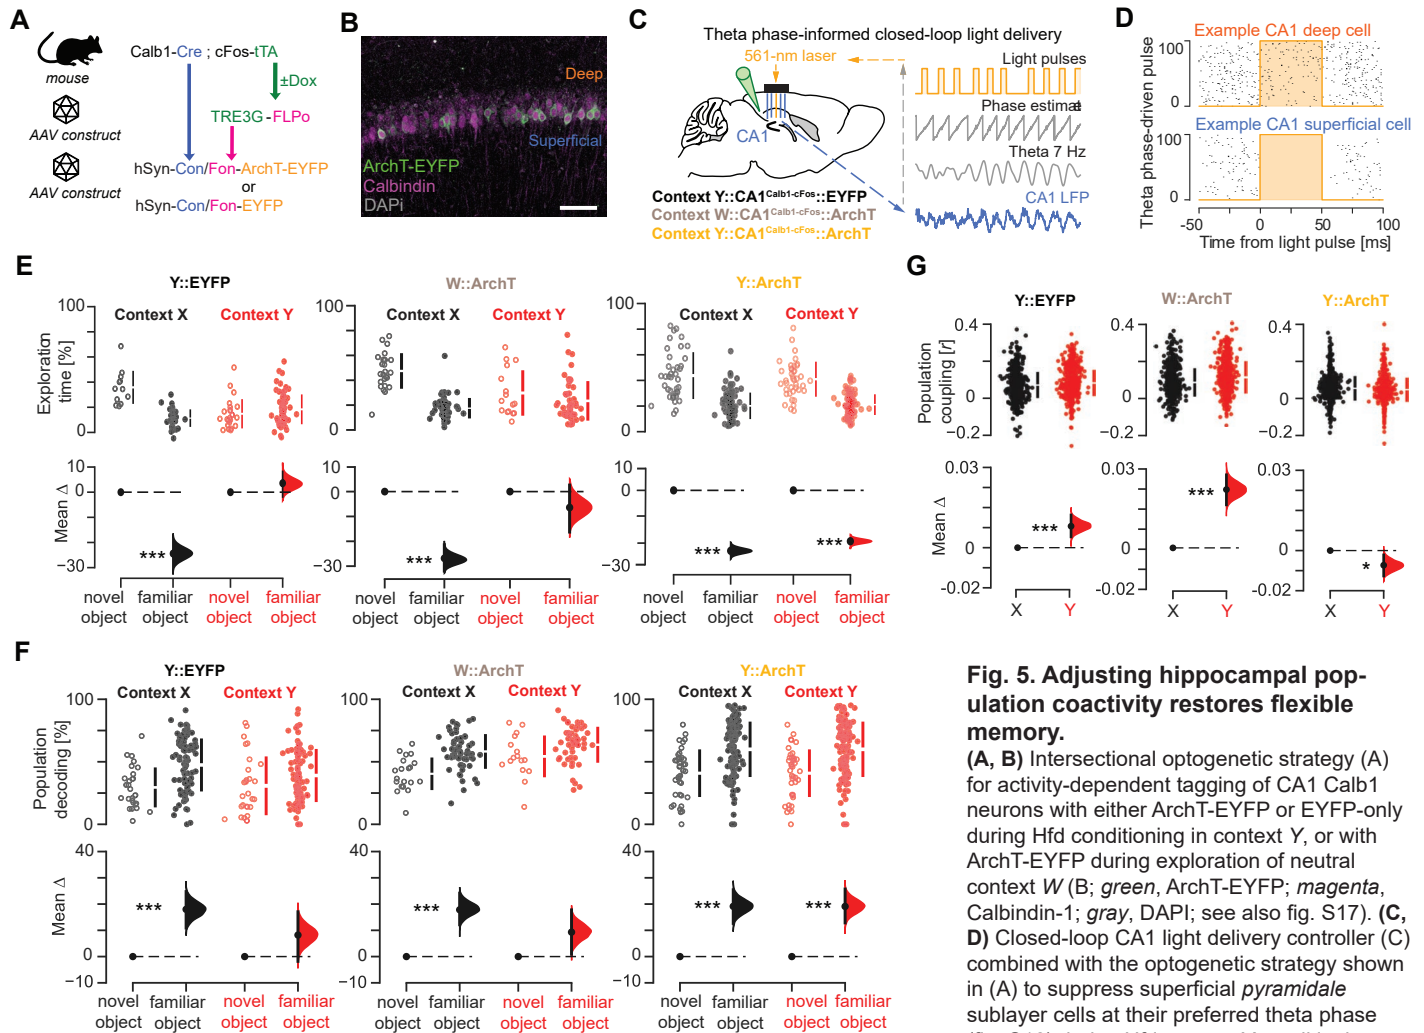


Fig. 5. Adjusting hippocampal population coactivity restores flexible memory.

(A, B) Intersectional optogenetic strategy (A) for activity-dependent tagging of CA1 Calb1 neurons with either ArchT-EYFP or EYFP-only during Hfd conditioning in context Y, or with ArchT-EYFP during exploration of neutral context W (B; green, ArchT-EYFP; magenta, Calbindin-1; gray, DAPI; see also fig. S17). (C, D) Closed-loop CA1 light delivery controller (C) combined with the optogenetic strategy shown in (A) to suppress superficial *pyramidale* sublayer cells at their preferred theta phase (fig. S18) during Hfd-context Y conditioning. (D) shows example spiking activity for a deep (top) versus a superficial (bottom) *pyramidale* sublayer cell with respect to theta phase-driven light delivery. (E-G) Estimation plots. For the context Y-tagged optogenetically adjusted ArchT mice, but not for the context Y-tagged EYFP-only mice nor the context W-tagged ArchT mice, this cell type-selective, network pattern-informed intervention restored cNOR performance (E) along with population object-location decoding (F) and population coupling (G) in context Y. *** $P < 0.001$, * $P < 0.05$, two-sided paired permutation test. N = 1097 CA1 principal cells from 5 ContextY::CA1^{Calb1-cFos}::ArchT mice (Y::ArchT) versus 1007 from 4 ContextY::CA1^{Calb1-cFos}::EYFP mice (Y::EYFP) and 683 from 3 ContextW::CA1^{Calb1-cFos}::ArchT mice (W::ArchT).

mice) or the control EYFP-only (in four ContextY::CA1^{Calb1-cFos}::EYFP mice) from the onset of Hfd conditioning in context Y (Fig. 5B and figs. S16, S17). We also tagged with ArchT-EYFP the set of CA1 superficial cells associated with a task-unrelated arena (context W) in another group of control mice (three ContextW::CA1^{Calb1-cFos}::ArchT mice; fig. S17). CA1 superficial principal cells, which preferentially fire action potentials at the trough of theta cycles (fig. S18A-C) (15, 33, 35), exhibited an increased theta modulation across conditioning days (fig. S18D). We thus combined this intersectional strategy with our theta phase-informed controller for closed-loop silencing of CA1 superficial cells during Hfd conditioning

sessions (Fig. 5C,D and fig. S18E,F). Despite undergoing 10-day Hfd feeding in context Y (fig. S19A), this cell type-defined, network pattern-informed intervention subsequently restored in ContextY::CA1^{Calb1-cFos}::ArchT mice natural behavioral response to novel food (fig. S19B) and object-location memory (Fig. 5E). This was not the case in ContextY::CA1^{Calb1-cFos}::EYFP mice and ContextW::CA1^{Calb1-cFos}::ArchT mice that showed impaired object-location memory in context Y (Fig. 5E). In line with this behavioral outcome, ContextY::CA1^{Calb1-cFos}::ArchT mice recovered functional object-location decoding and population coupling (Fig. 5F,G and figs. S19C-E, S20).

Discussion

Our results show that a robust (food-context) memory raises the population coupling of CA1 superficial *pyramidale* sublayer neurons, creating a dense network coactivity structure. CA1 pyramidal cells, the primary output of the hippocampus, segregate along the anatomical axes with different activity behaviors, indicating that cells arranged along the radial axis contribute differentially to information processing (25–31, 35). Compared to the deep, the superficial sublayer is enriched in context-modulated cells with slower response dynamics to environmental changes (28, 29), being more active in cue-poor environments and preferentially using a rate code driven by intra-hippocampal inputs; while deep cells are more active in cue-rich environments and use a phase code driven by entorhinal inputs (30). Our findings show that the hippocampus network preferentially engages superficial cells in strong peer-to-peer coactivity for robust contextual memory, at the risk of subsequent mnemonic rigidity. This finding is consistent with recent studies showing that CA1 superficial cells are more recruited into replay events and show stronger synaptic potentiation after novel experiences compared to deep cells (37, 38). The dense coactivity structure emerged over many conditioning days and could prevent the network to switch between alternative coding schemes or cell assemblies, and override representations by the deep sublayer counterparts (28, 30). The development of this coactivity structure required upstream CA3 activity, consistent with the observation that Schaffer collateral excitation is stronger in calbindin-expressing CA1 principal cells (25). With respect to subsequent behavior, this acquired hippocampal topology of heightened coactivity related to strong contextual (feeding) response and impaired novel information processing. Neural networks have been proposed to fall into the broad class of ‘small-world’ networks, a middle ground between regular and random networks where the combination of high clustering of elements (a property of regular networks) with short path lengths between elements (a property of random networks) would allow important properties of complex networks such as increased computational power and effective synchronizability (16, 17, 39–41). Our observation that repeated food-context conditioning affects the coactivity structure of the network by increasing neuronal clustering without shortening node-to-node paths suggests that the hippocampal topology can deviate from a small-world toward a more coherent, regular lattice. This way, the joint activity of an increased number of neurons operating as a cohesive population would permit robust information flow to downstream receiver neurons, possibly at the expense of other (e.g., novel) input channels.

The hippocampus could broadcast this heightened coactivity to several recipient circuits and reader neurons. For instance, recent work points to a contribution of the nucleus accumbens in translating hippocampal dynamics of appetitive memory into a behavioral readout (42) or the hypothalamus in driving non-homeostatic contextual feeding (43). The clustered spiking activity that developed in the hippocampus across food-context conditioning days is also likely to influence neocortical circuits for memory storage via systems consolidation (44, 45). With hippocampal support neuronal ensembles in prefrontal cortex can be rapidly formed to then undergo a process of functional maturation over weeks (46). This maturation could allow prefrontal cortex ensembles

to converge onto a lower-dimensional activity space to extract latent rules and common relational features across multiple experiences, gradually developing a knowledge structure of the world (47, 48). Importantly, the instantiation of this highly clustered topology can be prevented: applying cell type-selective, network pattern-informed neuronal suppression during contextual learning rebalances population activity and restores flexible memory. Together, these findings suggest that the plastic organization of hippocampal coactivity supports a network tradeoff between robust and flexible computations, shaping continual integration of new memories and their adaptability to cognitive demands.

References and Notes

1. D. L. Schacter, D. R. Addis, R. L. Buckner, *Nat. Rev. Neurosci.* **8**, 657–661 (2007).
2. G. Brod, M. Werkle-Bergner, Y. L. Shing, *Frontiers in Behavioral Neuroscience*. **7** (2013).
3. G. Buzsáki, *Neuron*. **68**, 362–385 (2010).
4. W. Mau, M. E. Hasselmo, D. J. Cai, *eLife*. **9**, e63550 (2020).
5. C. A. Barnes, B. L. McNaughton, S. J. Y. Mizumori, B. W. Leonard, L.-H. Lin. *Progress in Brain Research*, **83**, 287–300.
6. N. M. van Strien, N. L. M. Cappaert, M. P. Witter, *Nat. Rev. Neurosci.* **10**, 272–282 (2009).
7. J.-P. Kronenberger, J. Médioni, *Behavioural Processes*. **11**, 53–59 (1985).
8. S. A. Barnett, *British Journal of Psychology*. **49**, 195–201 (1958).
9. A. Nagelhus, S. O. Andersson, S. G. Cogno, E. I. Moser, M.-B. Moser, *Neuron*. **111**, 2091–2104.e14 (2023).
10. B. Vandrey, S. Duncan, J. A. Ainge, *Hippocampus*. **31**, 881–896 (2021).
11. M. El-Gaby *et al.*, *Nat Neurosci.* **24**, 694–704 (2021).
12. G. M. van de Ven, S. Trouche, C. G. McNamara, K. Allen, D. Dupret, *Neuron*. **92**, 968–974 (2016).
13. C. Liu, R. Todorova, W. Tang, A. Oliva, A. Fernandez-Ruiz, *Science*. **382**, eadi8237 (2023).
14. S. Panzeri, M. Moroni, H. Safaai, C. D. Harvey, *Nat Rev Neurosci.* **23**, 551–567 (2022).
15. G. P. Gava *et al.*, *Nature Neuroscience*. **24**, 326–330 (2021).
16. D. J. Watts, S. H. Strogatz, *Nature*. **393**, 440–442 (1998).
17. M. E. J. Newman, *SIAM Rev.* **45**, 167–256 (2003).
18. S. F. Muldoon, E. W. Bridgford, D. S. Bassett, *Scientific Reports*. **6**, 22057 (2016).
19. M. Okun *et al.*, *Nature*. **521**, 511–515 (2015).
20. L. Wittner, D. A. Henze, L. Záborszky, G. Buzsáki, *Brain Struct Funct.* **212**, 75–83 (2007).
21. N. Ishizuka, J. Weber, D. G. Amaral, *J Comp Neurol.* **295**, 580–623 (1990).
22. J. S. Bains, J. M. Longacher, K. J. Staley, *Nat Neurosci.* **2**, 720–726 (1999).
23. R. A. Zalutsky, R. A. Nicoll, *Science*. **248**, 1619–1624 (1990).
24. D. Amaral, P. Lavenex, in *The Hippocampus Book*, P. Andersen, R. Morris, D. Amaral, T. Bliss, J. O’Keefe, Eds. (Oxford University Press, 2006), pp. 37–114.
25. I. Soltesz, A. Losonczy, *Nat Neurosci.* **21**, 484–493 (2018).
26. M. S. Cembrowski, N. Spruston, *Nature Reviews Neuroscience*. **20**, 193–204 (2019).
27. M. Valero *et al.*, *Nature Neuroscience*. **18**, 1281–1290 (2015).
28. N. B. Danielson *et al.*, *Neuron*. **91**, 652–665 (2016).
29. T. Geiller, M. Fattahi, J.-S. Choi, S. Royer, *Nat Commun.* **8**, 14531 (2017).
30. F. Sharif, B. Tayebi, G. Buzsáki, S. Royer, A. Fernandez-Ruiz, *Neuron*. **109**, 363–376.e6 (2021).
31. A. V. Masurkar *et al.*, *Cell Reports*. **18**, 148–160 (2017).
32. A. V. Masurkar *et al.*, *Journal of Neurophysiology*. **123**, 980–992 (2020).
33. A. D. Grombark, G. Buzsáki, *Science*. **351**, 1440–1443 (2016).
34. M. S. Cembrowski *et al.*, *Neuron*. **89**, 351–368 (2016).
35. A. Navas-Olive *et al.*, *Nature Communications*. **11**, 2217 (2020).

36. L. E. Fenno *et al.*, *Nature Methods*. **11**, 763–772 (2014).
37. M. Berndt, M. Trusel, T. F. Roberts, B. E. Pfeiffer, L. J. Volk, *Neuron*. **111**, 2984-2994.e4 (2023).
38. R. E. Harvey, H. L. Robinson, C. Liu, A. Oliva, A. Fernandez-Ruiz, *Neuron*. **111**, 2076-2090.e9 (2023).
39. E. Bullmore, O. Sporns, *Nat Rev Neurosci*. **10**, 186–198 (2009).
40. D. S. Bassett, O. Sporns, *Nature Neuroscience*. **20**, 353–364 (2017).
41. M. Barahona, L. M. Pecora, *Phys. Rev. Lett.* **89**, 054101 (2002).
42. S. Trouche *et al.*, *Cell*. **176**, 1393-1406.e16 (2019).
43. H. Mohammad *et al.*, *Nat Neurosci*, 1–10 (2021).
44. D. R. Euston, A. J. Gruber, B. L. McNaughton, *Neuron*. **76**, 1057–1070 (2012).
45. P. W. Frankland, B. Bontempi, *Nat Rev Neurosci*. **6**, 119–130 (2005).
46. T. Kitamura *et al.*, *Science*. **356**, 73–78 (2017).
47. M. D. Morrissey, N. Insel, K. Takehara-Nishiuchi, *eLife*. **6**, e22177 (2017).
48. J. Zhou *et al.*, *Nature*. **590**, 606–611 (2021).
49. L. Lefèvre, G. P. Gava, Hippocampal ensemble recordings in Contextual Feeding and cNOR tasks from mice (2024), doi:10.60964/BNDU-MZG5-V129.
50. G. P. Gava, L. Lefèvre, Analysis of hippocampal ensembles during Contextual Feeding and cNOR tasks (2024), doi:10.60964/BNDU-R97N-BG26.
51. P. Perestenko, 3D printable files for objects to use in mouse object recognition tasks (2024), doi:10.60964/BNDU-GQ68-JZ91.
52. K. Nakazawa *et al.*, *Science*. **297**, 211–218 (2002).
53. L. Drane, J. A. Ainsley, M. R. Mayford, L. G. Reijmers, *Front Mol Neurosci*. **7**, 82 (2014).
54. S. Trouche *et al.*, *Nat. Neurosci*. **19**, 564–567 (2016).
55. V. Lopes-dos-Santos, D. Brizee, D. Dupret, Spatio-temporal organization of network activity patterns in the hippocampus (2023), p. 2023.10.17.562689, doi:10.1101/2023.10.17.562689.
56. X. Han *et al.*, *Frontiers in Systems Neuroscience*. **5** (2011).
57. C. G. McNamara, M. Rothwell, A. Sharott, *Cell Reports*. **41** (2022).
58. J. Magland *et al.*, *Elife*. **9** (2020).
59. M. Pachitariu, N. A. Steinmetz, S. N. Kadir, M. Carandini, K. D. Harris, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., 2016), vol. 29.
60. V. Lopes-Dos-Santos *et al.*, *Neuron*. **100**, 940-952.e7 (2018).
61. J.-P. Onnela, J. Saramäki, J. Kertész, K. Kaski, *Physical review. E, Statistical, nonlinear, and soft matter physics*. **71**, 065103 (2005).
62. G. Costantini, M. Perugini, *PLOS ONE*. **9**, e88669 (2014).
63. J. Saramäki, M. Kivelä, J.-P. Onnela, K. Kaski, J. Kertész, *Phys. Rev. E*. **75**, 027105 (2007).
64. R. W. Floyd, *Commun. ACM*. **5**, 345- (1962).
65. B. Roy, *C. R. Acad. Sci. Paris*. **249**, 216–218 (1959).
66. S. Warshall, *J. ACM*. **9**, 11–12 (1962).
67. E. Estrada, *Discrete Applied Mathematics*. **268**, 70–90 (2019).
68. J. Ollion, J. Cochenec, F. Loll, C. Escudé, T. Boudier, *Bioinformatics*. **29**, 1840–1841 (2013).
69. J. Ho, T. Tumkaya, S. Aryal, H. Choi, A. Claridge-Chang, *Nature Methods*. **16**, 565–566 (2019).
70. F. Pedregosa *et al.*, *Journal of Machine Learning Research*. **12**, 2825–2830 (2011).
71. A. A. Hagberg, D. A. Schult, P. J. Swart, *Proceedings of the 7th Python in Science Conference*, 11–15 (2008).
72. R. A. A. Ince, R. S. Petersen, D. C. Swan, S. Panzeri, *Front. Neuroinform*. **3** (2009).
73. K. Mizuseki, K. Diba, E. Pastalkova, G. Buzsáki, *Nat. Neurosci*. **14**, 1174–1181 (2011).
74. A. Oliva, A. Fernández-Ruiz, G. Buzsáki, A. Berényi, *Hippocampus*. **26**, 1593–1607 (2016).

Acknowledgements. We thank R. Lambiotte and S.R. Schultz for commenting on a previous version of the manuscript; B. Micklem and A. Desaulniers for technical assistance; all members of the Dupret lab for feedback during the project.

Funding: This work was supported by the Biotechnology and Biological Sciences Research Council UK (awards BB/S007741/1 and BB/N002547/1 to DD) and the Medical Research Council UK (programme MC_UU_00003/4 and award MR/W004860/1 to DD; and MC_UU_00003/6 to AS).

Author contributions: Conceptualization, GPG, LL, and DD; Investigation, GPG, LL, TB, SBMc, DB, KH, HH; Analysis, GPG, LL, and DD; Methodology, GPG, LL, SBMc, VLdS, and DD; Resources, PVP, RT, AS, and DD; Visualization, GPG and DD; Funding acquisition: DD; Writing – Original Draft, GPG and DD; Writing – Reviewing & Editing, GPG, LL, TB, SBMc, VLdS, DB, KH, HH, PVP, RT, AS and DD; Supervision, DD.

Competing interests: A.S. is an inventor on a pending patent application related to the phase-tracking algorithm used in this paper. All other authors declare no competing interests.

Data and materials availability: All data are available in the manuscript or the supplementary materials. The c-fos-tTA mouse line was generated at The Scripps Research Institute and obtained from Dr. L.G. Reijmers at Tufts University under a material transfer agreement with The Scripps Research Institute. Hippocampal recordings (49) and Jupyter notebooks for ensemble analysis (50), and 3D printable objects for cNOR task (51) are available via the Medical Research Council (MRC) Brain Network Dynamics Unit (BNDU) [data sharing platform](#).

The journal has requested that the following text be included:

This is the author's version of the work. It is posted here by permission of the AAAS for personal use, not for redistribution. The definitive version was published in *Science* on 6 September 2024, DOI: 10.1126/science.adk9611.

Hyperlink to the Final Published Version on the *Science* website: <https://www.science.org/doi/10.1126/science.adk9611>

Materials and Methods

Animals. These experiments used adult (4–6 months old) C57BL/6J wild-type mice (Charles River Laboratories, UK), transgenic hemizygous Grik4-Cre mice (The Jackson Laboratories; C57BL/6-Tg(Grik4-cre)G32-4Stl/J, stock number 006474, RRID: IMSR_JAX:006474; (52)) and double-transgenic mice obtained by crossing transgenic mice heterozygous for the transgene expressing the Cre recombinase under the control of the Calbindin 1 (Calb1) promoter (Jackson Laboratories; obtained from C57BL/6J crossed with Calb1-IRES2-Cre-D B6.129S-Calb1^{tm2.1(cre)Hze}/J, stock number 028532, RRID: IMSR_JAX:028532) with c-fos-tTA transgenic male mice heterozygous for the transgene carrying the c-fos promoter-driven tetracycline transactivator (tTA) (53, 54). The c-fos-tTA (a.k.a. TetTag) mouse line was generated at The Scripps Research Institute and obtained from Dr. L.G. Reijmers at Tufts University under a material agreement with The Scripps Research Institute for shipment to the MRC BNDU at the University of Oxford where they were bred from c-fos-tTA mice crossed with C57Bl6/J mice. Animals were housed with their littermates up until the start of the experiment. All mice held in IVCs, with wooden chew stick and nestlets in a dedicated housing facility with a 12/12 h light/dark cycle (lights on at 07:00), 19–23°C ambient temperature and 40–70% humidity. They had free access to water and food *ad libitum* throughout the experiment. Experimental procedures performed on mice in accordance with the Animals (Scientific Procedures) Act, 1986 (United Kingdom), with final ethical review by the Animals in Science Regulation Unit of the UK Home Office.

Surgical procedure. All surgical procedures were performed under deep anesthesia using isoflurane (0.5–2%) and oxygen (2 l/min), with analgesia provided before (0.1 mg/kg vetergesic) and after (5 mg/kg metacam) surgery.

For CA3 optogenetic silencing, we bilaterally transduced the dorsal CA3 of Grik4-Cre mice with a Cre-dependent ArchT-GFP viral construct (200 nl per site; 2 injection sites per side; see section “Viral constructs”) using stereotaxic coordinates (site 1: –1.7 mm anteroposterior and ±1.8 mm lateral from bregma, and –1.8 mm ventral from dura; site 2: –2.3 mm anteroposterior and ±2.6 mm lateral from bregma, and –2.1 mm ventral from dura). The viral vector was delivered at a rate of 100 nl/min using a glass micropipette lowered to the target site and held in place for 5 min after virus delivery before being withdrawn. Likewise, we injected a Cre-dependent GFP-only construct in the dorsal CA3 of other Grik4-Cre mice to generate the corresponding controls.

To generate expression of either ArchT-EYFP or EYFP-only in principal cells of the CA1 superficial *pyramidale* sublayer, we combined Cre-dependent, FlpO-dependent, and tTA-dependent approaches (see section “Viral constructs”). We mixed in a 1:5 ratio a tTA-inducible AAV carrying a TRE3G-FlpO construct with a Cre-dependent FlpO-dependent AAV carrying either a hSyn-Con/Fon-ArchT-EYFP or the corresponding control hSyn-Con/Fon-EYFP. Viral injections were targeted bilaterally to dorsal CA1 hippocampus of double-transgenic Calb1-Cre;cFos-tTA mice using stereotaxic coordinates (–1.7 and –2.3 mm anteroposterior from bregma, ±1.25 and ±1.7 mm lateral from bregma, and –1.1 mm ventral from dura; 150 nl per site; 2 injection sites per side) at a rate of 100 nl/min using a glass micropipette lowered to the target site and held in place for 5 min after virus delivery before being withdrawn.

For electrophysiological tetrode recordings, mice were implanted with a single microdrive containing 14 independently movable tetrodes, targeting the *stratum pyramidale* of the dorsal CA1 hippocampus (15). Tetrodes were constructed by twisting together four insulated tungsten wires (12 µm diameter, California Fine Wire) which were briefly heated to bind them together into a single bundle. Each tetrode was loaded in one cannula attached to a 6 mm long M1.0 screw to enable its independent manipulation of depth. The drive was implanted under stereotaxic control in reference to bregma using central coordinates –2.0 mm anteroposterior from bregma, +1.7 mm lateral from bregma as a reference to position each individual tetrode contained in the microdrive, initially implanting tetrodes above the pyramidal layer (–1.0 mm ventral from dura). The distance between neighboring tetrodes was 350 µm. Following the implantation, the exposed parts of the tetrodes were covered with paraffin wax, after which the drive was secured to the skull using dental cement and stainless-steel anchor screws inserted into the skull. Two of the anchor screws, both above the cerebellum, were attached to a 50 µm tungsten wire (California Fine Wire) and served as ground. For the recordings, each tetrode was lowered along the vertical axis to reach the *pyramidale* layer using the rotations applied to its tetrode cannula-holding screw and the electrophysiological profile of the local field potentials in the hippocampal ripple frequency band, with final depth position subsequently confirmed by histology of anatomical tracks.

For optogenetic manipulations, two optic fibers (230 µm diameter, Doric Lenses, Canada) were incorporated into the 14-tetrode microdrive designed to bilaterally deliver light to dorsal CA1 (or CA3 in Grik4-Cre mice) and implanted 10 days after viral injections.

To identify the location of CA1 principal cells in the deep versus the superficial sublayers of the CA1 *stratum pyramidale*, we leveraged silicon probe recordings with known spacing between the recording sites along a linear shank (55). Silicon probes were implanted following the same surgical procedure to span the somato-dendritic axis of CA1 principal cells and establish the laminar profile of the sharp-wave ripples detected in the local field potentials. These silicon probe recordings allowed estimating the position (depth) of individual tetrode-recorded principal cell soma with respect to the deep versus the superficial sublayers of the dorsal CA1 *stratum pyramidale* (fig. S13).

For the hippocampal lesion surgery (fig. S2C), mice in the lesion group underwent the same anesthetic induction as above, then scalp incision and craniotomy, followed by N-methy-D-aspartate (NMDA; 10 mg / mL) injections directly into the hippocampus at four sites per hemisphere using a modified Hamilton 36G syringe needle (anterior-posterior: –1.7, –2.3, –2.8, –3.1; mediolateral: ±1.2, ±1.7, ±2.2, ± 2.8; dorsoventral: –1.9, –1.9, –2.0, –4.0, respectively; all from bregma; 100 to 200 nL per site at the infusion / diffusion rates described above), and were then sutured. Midazolam (5 mg / kg, sub-cutaneous) was used to prevent seizures in hippocampal-lesioned mice. Mice receiving sham surgery were incised and then sutured. All mice had at least 2 weeks recovery before behavioral testing.

Viral constructs. To optogenetically target CA3 neurons genetically defined to express Grik4, we either injected the Cre-dependent AAV9-CAG-Flex-ArchT-GFP viral vector (UNC Vector Core, #AV6222b) to generate CA3^{Grik4}::ArchT mice, or the corresponding AAV9-CAG-Flex-GFP control vector (UNC Vector Core, #AV5220C) to generate CA3^{Grik4}::GFP mice.

To optogenetically target cFos-expressing CA1 superficial *pyramidale* sublayer neurons genetically defined to express Calbindin 1, we first produced a TRE3G-FlpO AAV carrying the optimised FlpO recombinase under the control of the third generation of tetracycline responsive element containing promoter (TRE3G, Clontech Laboratories). For this, we exchanged ArchT-GFP open

reading frame (ORF) (cut with NcoI and EcoRV) in pAAV-Tre3G-ArchT-GFP vector (54) with the Myc-FLPo ORF (cut with NcoI and Klenow blunted Ascl) from pAAV-EF1a-DIO-FLPo-Myc vector (42) (Addgene plasmid # 124641; <http://n2t.net/addgene:124641>; RRID:Addgene_124641).

We next produced a viral vector allowing the Cre-dependent and Flp-dependent expression of ArchT-eYFP. The corresponding pAAV-hSyn-Con/Fon-ArchT-eYFP construct was cloned in two stages. First, pAAV-hSyn-Coff/Fon-ArchT-eYFP has been cloned. Plasmid vectors pAAV-CamKII-ArchT-GFP (a gift from Edward Boyden, Addgene plasmid #37807; <http://n2t.net/addgene:37807>; RRID: Addgene_37807)(56), pAAV-hSyn-Coff/Fon-hChr2(H134R)-eYFP (a gift from Karl Deisseroth, Addgene plasmid #55648; <http://n2t.net/addgene:55648>; RRID: Addgene_55648) (36) and the combination of the PCR products was used to assemble two inserts that were then subcloned into the pAAV-hSyn Coff/Fon hChr2(H134R)-eYFP vector to substitute the corresponding hChr2-eYFP coding exons with the ArchT-eYFP ones. Primers and the template plasmid DNA for the first insert (exon 1): GTTCTGCTAGCAACCCCGACACTTACCTTAGCCAGCAGGGCCAG, GTTCTGAGCTCGCCACCATGGACCCCATC, plasmid#37807. The PCR product was then cloned into the plasmid #55648 using NheI and SacI recognition sites thus forming the intermediate vector. For the following subcloning of the exon 2 the three PCR products were generated with the primers and the corresponding template DNAs: GTTCTACTAGTCCTCCTGTACTCACC, GTGAGCAAGGGCGAGGAG, plasmid #55648; CTCCTCGCCCTTGCTCACTGCTACTACCGGTCGGGG, GACTCTATTTCTCATGTGTTAGGTGGACAGGGTGAGCATCG, plasmid #37807; CCTAAACACATGAGAAATAGAGTC, CGAAGTTATGGTACCTGTGCCCCCCC, plasmid #55648. These three products were then combined in the overlapping PCR and inserted into the intermediate vector using SpeI and KpnI cloning sites forming pAAV-hSyn-Coff/Fon-ArchT-eYFP. pAAV-hSyn-Coff/Fon-ArchT-eYFP vector then was used to produce pAAV-hSyn-Con/Fon-ArchT-eYFP by inverting the sequence containing part of ArchT-eYFP exon and flanked with the SpeI and KpnI restriction enzymes. The corresponding insert was produced by the PCR with the primers GTTCTACTAGTTGTGCCCCCTTTTTTTTAT and GTTCTGGTACCCCTCCTGTACTCACCTTGCC using pAAV-hSyn-Coff/Fon-ArchT-eYFP vector as a template.

The control vector expressing eYFP in the Cre-dependent and Flp-dependent manner was a gift from Karl Deisseroth (Addgene plasmid # 55650; <http://n2t.net/addgene:55650>; RRID:Addgene_55650) (36).

We mixed in a 1:5 ratio the AAV carrying TRE3G-FlpO and one of the two Cre-dependent Flp-dependent AAVs (for ArchT-eYFP versus eYFP-only expression) in the CA1 of double-transgenic Calb1-Cre;cFos-tTA mice or their mono-transgenic (i.e., Calb1-Cre and c-fos-tTA) control mice (Fig. 5 and fig. S16).

Contextual feeding task. Following full recovery from the surgery, each mouse was first handled in a dedicated handling cloth and connected to the recording system to be familiarized with the recording procedure over a period of one week prior to the start of the experiment itself. During this period, tetrodes were gradually lowered to the CA1 *stratum pyramidale* using their estimated depth location, local field potentials, and neuronal spike waveforms. All mice were fed *ad libitum* throughout the entire experiment. Our contextual feeding task involved three open-field arenas (referred to as context X, Y, and Z) that were all novel environments for the mouse on conditioning day 1. These enclosures (46 cm width; 38 cm height) differed in their shape (Fig. 1A and fig. S1B, D) and wall-attached cue cards. The first stage of this task ("Food-context conditioning") spanned either 10 (Fig. 1B) or 2 (fig. S10) consecutive days. On each conditioning day, mice explored the arenas in a random order (30-min exploration each). Context X contained two food containers, each provided with one regular diet pellet (Chow; fig. S1A), identical to those present in the mouse homecage. Context Y contained two similar food containers that allowed a choice between one chow pellet versus one high-fat diet pellet (fig. S1A; 45% Hfd; Research Diets Inc.; catalog number #D12451). Each mouse provided one data point showing its food intake per conditioning day in context X and context Y. Context Z did not contain food during conditioning. On each conditioning day, mice were allowed to rest for 30 minutes between two arena explorations. For the histological assessment of cFos expression, we either exposed mice to context Y without food, context Y with Hfd, or their homecage with Hfd (all 30-min experience every day for 10 days; Fig. 4A and fig. S12).

The second stage of this task ("Novel food test") allowed probing discriminative food-context association while assessing changes in baseline population-level activity following the 10-day (versus 2-day) conditioning stage (Fig. 1 and fig. S10). For this, we first re-exposed mice to either context X or Y ("Re-exposure"; 30-min) before we then measured mouse propensity to express context-biased feeding ("Novel food test"; 30-min). In each novel food test, the arena contained two containers that allowed a choice between two food items: one had a regular chow pellet, the other had a never-seen-before resource (e.g., apricot or blackcurrant fruit jams; chocolate spread or Biscoff food pastes; fig. S1C). The energy value and palatability of novel food resources were matched across test days and contexts. Novel food intake was measured as the novel food kCal intake in context Y minus that in context X. Each mouse provided one pair of data points showing its food intake in context Y versus X per novel food test day. We further allowed mice to eat Hfd in context Z that then contained two food containers (one with a chow pellet, the other with a Hfd pellet) on some cNOR days ("Hfd test"). We first used this to assess the contextual Hfd feeding bias of each mouse by measuring its Hfd intake over 30 minutes in context Z minus that measured in context Y in the previous day, over their sum (fig. S1D). We also used this Hfd test in context Z before engaging mice in the cNOR task to evaluate whether they would express a defective object-location memory in context Y, versus a successful object-location memory in context X, regardless of whether they were fed or not with Hfd just before that day (fig. S3).

We adapted our contextual feeding paradigm in Calb1-Cre;cFos-tTA mice to optogenetically tag the set of CA1 superficial *pyramidale* sublayer cells that are genetically defined by Calbindin 1 and are expressing cFos with their recruitment in a given context. This activity-dependent tagging of CA1 superficial cells was performed in either the Hfd-paired context Y or in a fourth, neutral arena (referred to as context W) unrelated to the contextual feeding task (Fig. 5 and fig. S17). For this, Calb1-Cre;cFos-tTA mice were fed for 10 days with doxycycline-containing pellets (Envigo Ltd., Catalog number #TD120240) prior to viral construct injections and microdrive implantation (fig. S17A). In this set of experiments, we replaced the Chow pellets with Dox pellets in the mouse homecage and both context X and Y, except when conducting the neuronal tagging during the off-Dox days in context Y versus W, as appropriate (fig. S17A). Before these tagging days, we conducted the first 2 days of conditioning in context X (fig. S17A). For the optogenetic tagging, we then ceased the dox-containing pellet homecage feeding for two days when we exposed mice to either context Y or W (fig. S17A), after which they received high-doxycycline containing pellets (Envigo Ltd., Catalog number #TD120658) as homecage feeding for 24 hours before returning to normal doxycycline containing pellets for the remaining of the experiment.

Continuous Novel Object Recognition task. On each cNOR task day (Fig. 1B and figs. S2A, S10A, S11A, S17A), mice first re-explored context X or Y ("Re-exposure"; without any food; 15-min exploration). We used these re-exposure sessions to evaluate

enduring changes in baseline population-level activity following food-context conditioning. In this context, we next inserted four novel objects (“Sampling”; 15-min exploration). Mice continued to explore this context for four more sessions (“Tests”; 10-min exploration each) that day (figs. S2A, S10A, S11A, S17A). Each object (outer dimension: 50 x 50 mm width; 55 mm height; example objects shown in fig. S2B) was positioned midway along a given wall. From the sampling session to the first cNOR test and then for each other cNOR test, we replaced one of the objects initially sampled with a novel object. This procedure thus allowed the mouse to explore three familiar (previously seen) and one completely novel object on each cNOR test that day. On each test, we measured the time spent exploring each object and we calculated the percentage time spent with the novel object versus the (mean) percentage time spent with the familiar objects. Each mouse provided one data point showing its time spent exploring a given object per cNOR test. The interval between two cNOR test sessions was 5-min, during which the mouse was in a sleep/rest box temporarily placed in the recording arena. Mice could eat regular-diet (Chow) food before each cNOR day experiment, as any other day of our paradigm to ensure they would not perform tests while hungry. We finished each cNOR day by re-conditioning mice, exposing them again to context Y with Hfd (30-min session).

Multichannel data acquisition, position tracking and light delivery. The extracellular signals from each tetrode channel were amplified, multiplexed, and digitized using a single integrated circuit located on the head of the animal (RHD2164, Intan Technologies; http://intantech.com/products_RHD2000.html; pass band 0.09 Hz to 7.60 kHz). The amplified and filtered electrophysiological signals were digitized at 20 kHz and saved to disk along with the synchronization signals (transistor-transistor logic digital pulses) reporting the animal’s position tracking and laser activations. The location of the animal was tracked using three differently colored LED clusters attached to the electrode casing and captured at 39 frames per second by an overhead color camera (<https://github.com/kevin-allen/positrack/wiki>). The LFPs were down-sampled to 1,250 Hz for all subsequent analyses. For optogenetic intervention, a 561-nm diode-pumped solid-state laser (Crystal Laser, model CL561-100; distributor: Laser 2000, Ringstead, UK) was used to feed implanted optic fibers with light (~15–20mW input power; ~9–11mW output power) bilaterally to the dorsal CA3 (Fig. 3) or CA1 (Fig. 5) hippocampus via a 2-channel rotary joint (Doric Lenses Inc.). This was performed using a closed-loop system to deliver in real time 561-nm light 50-ms pulses using dynamic tracking of ongoing theta phase in CA3^{Grik4::ArchT} mice and CA3^{Grik4::GFP} mice (Fig. 3 and fig. S11), CA1^{Calb1-cFos::ArchT} mice and CA1^{Calb1-cFos::EYFP} mice (Fig. 5 and fig. S18). The real-time phase estimate was obtained using the OscillTrack algorithm (57) implemented in the field-programmable gate array of the Intan Technologies interface board. Phase detection was obtained by continuously operating on the data stream coming from an input channel containing the CA1 pyramidal layer LFPs. This input channel used as the phase reference was high-pass filtered using a 1st order digital infinite impulse response filter with a corner frequency of 0.4 Hz to remove amplifier offset and electrode drift, then down-sampled 125-fold to a rate of 160 Hz for processing. The phase estimation was operated with a loop-gain of 0.0625 at a center frequency of 7 Hz. Stimulation was triggered with a phase lead to align the target phase with the middle of the 50-ms light pulse.

Spike detection and unit isolation. Spike sorting and unit isolation were performed with an automated clustering pipeline using Kilosort (<https://github.com/cortex-lab/KiloSort>) via the SpikeForest framework (<https://github.com/flaitionstitute/spikeforest>) (58, 59). To apply KiloSort to data acquired using tetrodes, the algorithm restricted templates to channels within a given tetrode bundle, while masking all other recording channels. All sessions recorded on a given day were concatenated and cluster cut together to monitor cells throughout the experiment day. The resulting clusters were verified by the operator using cross-channel spike waveforms, auto-correlation histograms, and cross-correlation histograms. Each unit used for analyses showed throughout the entire recording day stable spike waveforms, clear refractory period in their auto-correlation histogram, and absence of refractory period in its cross-correlation histograms with the other units. The trough-to-peak width of the spike waveform was used to distinguish hippocampal principal cells from interneurons, as described previously (55). In total, this study includes 9,111 CA1 principal cells: 2,506 using 6 wild-type mice for the original 10-day conditioning dataset (Figs. 1, 2, and 4); 1,389 using 5 wild-type mice in the additional 2-day conditioning dataset (fig. S10); 1,548 using 4 CA3^{Grik4::ArchT} mice and 881 using 2 CA3^{Grik4::GFP} control mice for the CA3 silencing dataset (Fig. 3); 1,097 using 5 CA1^{Calb1-cFos::ArchT} mice tagged with ArchT in context Y, 1,007 using 4 CA1^{Calb1-cFos::EYFP} mice tagged with EYFP in context Y, and 683 using 3 CA1^{Calb1-cFos::ArchT} mice tagged with ArchT in context W for the optogenetic tagging dataset (Fig. 5).

Theta oscillations, sharp-wave ripples detection, and tetrode depth estimation.

To detect theta oscillations from the LFPs, we applied Ensemble Empirical Mode Decomposition and selected bouts of at least 5 cycles during active exploratory behavior (animal speed > 2 cm/s) (60). To detect sharp-wave ripples, the LFP signal of each CA1 pyramidal layer channel was subtracted by the mean across all channels/sites (common average reference). These re-referenced signals were then filtered for the ripple band (110 to 250 Hz; 4th order Butterworth filter) and their envelopes (instantaneous amplitudes) were computed by means of the Hilbert transform. The peaks (local maxima) of the ripple band envelope signals above a threshold (5 times the median of the envelope values of that channel) were regarded as candidate events. Further, the onset and offset of each event were determined as the time points at which the ripple envelope decayed below half of the detection threshold. Candidate events passing the following criteria were determined as SWR events: (i) ripple band power in the event channel was at least 2 times the ripple band power in the common average reference (to eliminate common high frequency noise); (ii) an event had at least four ripple cycles (to eliminate events that were too brief); (iii) ripple band power was at least 2 times higher than the supra-ripple band defined as 200-500 Hz (to eliminate high frequency noise, not spectrally compact at the ripple band, such as spike leakage artefacts). We estimated the tetrode depth within the pyramidal layer using SWR and theta waveforms, as previously described (55). Raw local field potentials of each tetrode were averaged around the SWR power peaks for sharp-wave waveforms and around the descending zero crossings of theta oscillations for theta waveforms. The channel with the highest ripple power was used as theta phase reference. These waveforms were then projected onto a feature embedding template created using silicon probe recordings (55). Deep sublayer tetrodes (closer to oriens) showed positive sharp-waves, while superficial ones (closer to radiatum) displayed negative sharp-waves (fig. S13). Additionally, superficial channels exhibited stronger theta-nested slow-gamma components.

Object-location population decoding. We trained a General Linear Model (GLM) to classify each object-location compound explored during cNOR tests, from the recorded neuronal spiking activity nested in theta cycles associated with these visits (Fig. 1F). That is,

we used the population vectors of action potentials discharged by CA1 principal cells in the theta cycles associated with the mouse active exploration of each object-location compound. In doing so, the animal's body was within a radius of 5 cm from the object being explored. By training a model in session and testing it in test we obtained one novel and three familiar object scores per cNOR session (four session pairs per day; e.g., Fig. 1G).

A given object-location compound pair was then associated to a population decoding vector composed of each neuron's GLM coefficient (β), reporting the individual weight contribution of neuron to the population representation of (Fig. 1H). By taking the cosine distance metric between pairs of decoding vectors across subsequent sessions representing the same location but different objects (familiar object versus novel object), we quantified the change in population representation (decoding vectors similarity (Fig. 1I)). The representational change of individual neurons induced by the insertion of novel objects, was instead computed as (referred to as "single-neuron decoding coefficient", Fig. 1J).

To build the GLM model, we employed the function *LogisticRegression* from the *sklearn* package (v .1.2.2) with L2 regularization (inverse strength $C=1e-2$, selected by maximizing the classification scores across sessions).

Spatial rate maps. The horizontal plane of the recording arena was divided into bins of 2×2 cm to generate spike count maps (number of spikes fired in each bin) for each unit and an occupancy map (time spent by the animal in each bin), considering only active exploratory behavior (speed > 2 cm.s⁻¹), discarding periods of immobility, and further excluding sharp-wave/ripples. Then, spatial rate maps were generated by normalizing (dividing) the spike count maps by the corresponding occupancy map. The spatial rate maps were smoothed by convolution with a Gaussian kernel having s.d. equal to one bin width.

Single-neuron place field similarity. In this study, the place field similarity (PFS) is a pairwise measure that compares the place maps expressed by a given cell across two contiguous cNOR sessions (e.g., sampling versus test 1) in a given context. It represents the degree to which a cell that fired in a given region of space during one session (e.g., sampling) still fired in that region of space during the subsequent session (e.g., test 1). For each cell, the PFS value for a given pair of cNOR sessions is computed as the Pearson correlation coefficient from the direct bin wise comparisons between its two spatial rate maps (non-smoothed place maps; one map per session), limited to valid bins (occupancy greater than zero).

Neuronal coactivity graphs. We constructed hippocampal population graphs that represent the coactivity relationships between all pairs of principal cell spike trains recorded during a given task session. These coactivity graphs were computed using theta cycles as time windows spanning active exploratory behavior (speed > 2 cm.s⁻¹), discarding periods of immobility and further excluding sharp-wave/ripples. To further control for the shared influence of the general network activity on peer-to-peer coactivity, we used for any two neurons (i, j) the regression coefficient β obtained by fitting the GLM (Fig. 2C):

$$x_j \sim \beta_{ij}x_i + \alpha_{ij}P$$

where x_j, x_i are the z-scored theta-nested spike trains of individual neurons j (the target) and i (the predictor), and P is the summed activity of the other $N - 2$ neurons,

$$P = \sum_{n=0}^{N-\{i, j\}} x_n$$

with β_{ij} weighting the influence of the population contribution to the activity of target neuron j .

The recorded neurons (and their coactivity associations) are therefore the nodes (and their edges) in the coactivity graph of each task session. We described each graph by its adjacency matrix, A , as the $N \times N$ square matrix containing the pairwise coactivity relations within the network, yielding a weighted graph with no self-connections:

$$A = \begin{pmatrix} \beta_{0,0} & \cdots & \beta_{0,N} \\ \vdots & \ddots & \vdots \\ \beta_{N,0} & \cdots & \beta_{N,N} \end{pmatrix}$$

with $\beta_{i,i} = 0 \forall i \text{ in } N$ and the symmetry in the weights of the network being ensured by setting $A = \frac{A+A^T}{2}$ to form an undirected graph.

Clustering coefficient. We computed the clustering coefficient to characterize the network's local coactivity structure by scoring the triadic firing relationships established by each neuron i with the other neurons in the population, using the formula proposed by Onnela et al. (61–63):

$$C_i = \frac{\sum_{jq} (\widehat{\beta}_{ij}\widehat{\beta}_{iq}\widehat{\beta}_{jq})^{1/3}}{k_i(k_i - 1)}$$

where j and q are neighbors of neuron i , all edge weights are normalized by the maximum edge weight in the network $\widehat{\beta} = \beta / \max(\beta)$ and k_i is the degree of neuron i , which in these weighted graphs with no self-connection is equal to the number of neurons minus one. This formula accounts for negative edges, yielding a negative value when there is an odd number due to the negative edges in the triad; it is positive otherwise. For this reason, this quantity can also be interpreted as a measure of the structural balance around a node, in the sense that its neighbourhood presents coherent patterns of firing relationships (see below).

Geodesic path length. We measured the geodesic (i.e., shortest) path length to estimate the ability of the network to efficiently switch from the activity of one node to another (Fig. 2F). In a binary graph, this functional measure would represent the smallest number of

edges connecting two nodes. Here, we define the length between two nodes i and j as the inverse of their coactivity: $l = 1 / \beta_{ij}$ discarding all negative edges (18). We then computed the geodesic path length between any two nodes in each neuronal graph using the Floyd-Warshall algorithm (64–66). We report the average path length for each principal cell in the coactivity graph to evaluate the easiness of transitioning across nodes in the hippocampus network.

Single-neuron coactivity strength. We defined the single-neuron coactivity strength as the average pairwise coactivity relation of a given node in a weighted graph (fig. S5B). As a reference, the strength in a weighted graph can be compared to the degree in a binary graph, which accounts for the number of the node's neighbors. Here, the strength S_i of a node i is the average across all the weights β_{ij} of the edges projected from that node:

$$S_i = \frac{\sum_{j=0}^N \beta_{ij}}{N}$$

where N is the number of neurons j that node i projects to.

Null models of coactivity graphs. To characterize further the structure of coactivity graphs, we compared them with equivalent null models of known topology, namely their random and lattice surrogates (fig. S6A). For each real-data neuronal graph, we aimed to construct a random equivalent graph where the average geodesic path length would be shorter while the original distribution of the observed coactivity strength value of each node would be preserved. Interestingly, we noted that the control graphs obtained by randomly shuffling the edge weights of the real-data coactivity graphs however displayed longer geodesic path lengths than those of the real-data graphs (18). Thus, we randomly reassigned the edges of each neuronal node to preserve its coactivity strength by first extracting the upper-triangular matrix U of the coactivity graph A , then row-wise shuffle U to obtain U_r and construct the random equivalent graph $R = U_r + U_r^T$ where R is the random equivalent graph of A that preserved the original single-neuron coactivity values.

For each real-data neuronal graph, we also constructed a lattice equivalent graph by re-arranging its original coactivity edge weights onto a one-dimensional (circular) lattice, such that the edges that corresponded to the smallest Euclidean distance between nodes were assigned the highest weights. We reasoned that in a 1-D lattice with N nodes and unit spacing between nodes, there are N edges with a Euclidean distance of $d = 1$ between nodes. We therefore ranked the observed edge weights by decreased strength and randomly distributed the top N connections among the edges representing $d = 1$, the next $N-1$ weights to the $d = 2$ edges, etc. We continued in this manner until all the edges in the original coactivity graph were placed in the lattice (18).

Structural balance. The combination of positive and negative edges in a network can give rise to both stable and unstable patterns of relationships. We applied structural balance analysis to neuron coactivity relationships to assess their coherence, and therefore the stability of hippocampal graphs. In this network analysis, triads of neurons are classified as balanced or unbalanced (fig. S6G). An intuition about such triadic relationships arises from social networks where three nodes form a balanced motif either by having three positive edges (e.g., “the friend of my friend is my friend”), or by having two negative ones (e.g., “the enemy of my friend is my enemy”) (67). Any triad with an odd number of negative relations would make the motif unbalanced. We inspected the triadic coactivity relationships of hippocampal graphs and assessed their stability by considering for each graph the top 10% of edges with the highest absolute magnitude (both positives and negatives), finding all triads, and classifying them into balanced or unbalanced (as per the sign of their three edges). We defined the structural balance of the graph as the proportion of its balanced triads. With this, we also characterized the structure of the equivalent null models (fig. S6H). We found that random graphs are less stable, displaying a structural balance at chance level, while lattice graphs contain only balanced motifs and are therefore more stable.

Population coupling. We defined the coupling of neuron i to the rest of the population of N units (population coupling, PC) as the Pearson correlation between its theta-binned activity x_i and that of all remaining neurons summed together $P^{N-\{i\}}$:

$$PC_i = \text{Pearson}(x_i, P^{N-\{i\}})$$

To assess whether the increase in population coupling in context Y was due to the fine temporal structure of the neural population activity, we independently circularly shuffled the neurons' theta-binned activity by random delays between 10 to 100 theta cycles, therefore preserving each neurons' rate and autocorrelation. The final value was taken as the average over 100 iterations.

Anatomy. Mice were deeply anesthetized with isoflurane/pentobarbital and transcardially perfused with PBS followed by cold 4 % PFA dissolved in PBS. The brains were extracted, kept in 4% PFA for 24 h, and sliced into 50 μm thick coronal sections. For immunostaining, free-floating sections were rinsed extensively in PBS with 0.25 % Triton X-100 (PBS-T) and blocked for 1 h at room temperature in PBS-T with 10 % normal donkey serum (NDS). Sections were then incubated at 4°C for 48 h with primary antibodies (rabbit anti-cFos, 1:500, Synaptic Systems, cat# 226 003; goat anti-calbindin 1, 1:1,000, Nittobo Medical, cat# Af1040) diluted in PBS-T with 3 % NDS blocking solution. Sections were then rinsed three times for 10 min in PBS and incubated for 24 h at 4°C in secondary antibodies (donkey anti-rabbit Alexa Fluor 488, 1:250, Thermo Fisher Scientific, cat# A-21206, RRID:AB_2535792; donkey anti-goat Cy3, 1:1000, Jackson ImmunoResearch, cat# 705-165-147, RRID: AB_2307351) diluted in PBS-T with 1 % NDS. This step was followed by three rinses for 15 min in PBS. Sections were then incubated for 1 min with 4',6-diamidino-2-phenylindole (DAPI; 0.5 $\mu\text{g}/\text{ml}$, Sigma-Aldrich, cat# D8417) diluted in PBS to label cell nuclei, before undergoing three additional rinse steps of 10 min each in PBS. Sections were finally mounted on slides, cover-slipped with Vectashield mounting medium (Vector Laboratories) and stored at 4 °C. Sections were also used for anatomical verification of the tetrode tracks. Mice with incorrect tetrode or viral targeting were excluded. The CA1 pyramidal layer was imaged in all 3 channels (DAPI-A405, cFos-A488, and Calb1-CY3) as a mosaic stack across the full z-range using a Zeiss confocal microscope (LSM 880 Indimo, Axio Imager 2) with a Plan-Apochromat 20x/0.8 M27 Objective and the ZEN software (Zeiss Black 2.3, step size: 10 μm) using 2 sections at the same rostro-caudal position (-2.0mm and -2.5 mm)

for each mouse. Automated counting was performed using Fiji and the 3D ImageJ Suite (68). A CA1 stratum pyramidale mask was generated using the DAPI channel, subdividing it into the deep versus the superficial pyramidale sublayers using the Calbindin 1 channel. To automatically detect single cFos cells, a 3D iterative thresholding (minimum threshold: RenyEntropy, pixel-volume: 2000-3000) was used on the cFos stack. A cFos positive cell was also positive for Calb1 when the average intensity within the cFos segment was above the automatic threshold (3D intensity measure, threshold: Huang). Cell densities are number of cells divided by the volume of the pyramidale layer. For the hippocampal lesion experiment, mice were perfused transcardially with physiological saline (0.9% NaCl) followed by 10% formol saline (10% formalin in physiological saline). The brains were then removed and placed in 10% formol saline and 72 hours later transferred to 30% sucrose-formalin. Coronal sections (50 μ m) were cut on a freezing microtome and stained with cresyl violet to enable visualization of lesion extent.

Data and statistical analyses. Data were analysed in Python 3.6 and using the packages DABEST v0.3.0 (69), scikit-learn v0.23.2 (70), NetworkX v2.4 (71), Numpy v1.18.1, Scipy v1.4.1, Matplotlib v3.1.2, Pandas v0.25.3 and Seaborn v0.11.0. All statistical tests related to a symmetric distribution were performed two-sided using permutation tests and visualized with Gardner-Altman plots (to compare 2 groups) and Cumming plots (for more groups) (69). These estimation plots allow visualizing the effect size by plotting the data as the mean or median difference between one of the groups (the left-most group of each plot, used as group-reference) and the other groups (to the right, along the x-axis of each plot). For each estimation plot: (i) the upper panel shows the distribution of raw data points for the entire dataset, superimposed on bar-plots reporting group mean \pm SD, unless stated otherwise; and (ii) the lower panel displays the difference between a given group and the (left-most) group-reference, computed from 5,000 bootstrapped resamples and with difference-axis origin aligned to the mean or the median of the group-reference distribution. For each estimation plot: *black-dot*, mean (for normal distributions) or median (for skewed distributions) as indicated; *black-ticks*, 95% confidence interval; *filled-curve*: bootstrapped sampling-error distribution. Data distributions were assumed to be normal, but this was not formally tested. We also used the t-test to compare two conditions; the Wald test for assessing the significance of regression lines; and the Kolmogorov–Smirnov test for comparing probability distributions. Neural and behavioral data analyses were conducted in an identical way regardless of the identity of the experimental condition from which the data were collected, with the investigator blind to group allocation during data collection and/or analysis.

Supplementary Figures and Legends

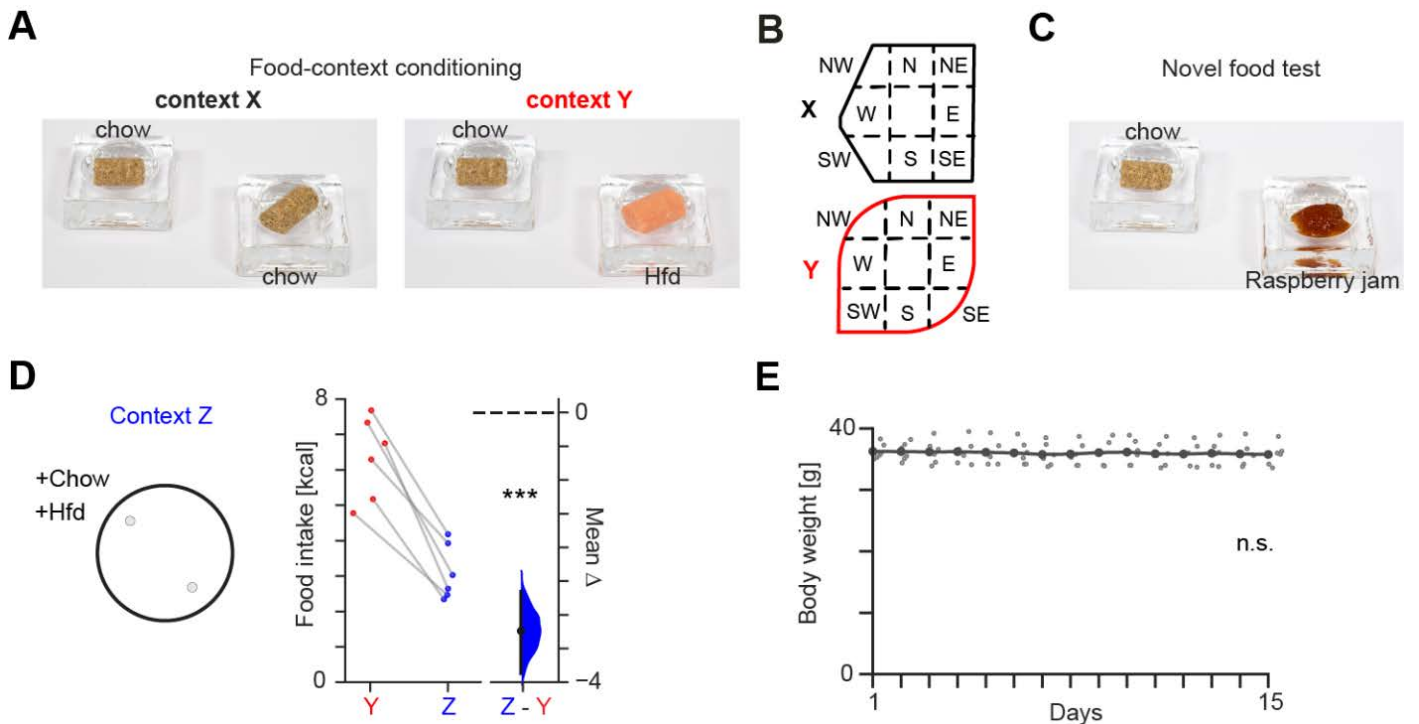


Fig. S1. Contextual feeding task.

(A) Pictures of the two food containers used to provide mice with a choice of two chow pellets in context X (left) and a choice between one chow pellet versus one high-fat diet (Hfd) pellet in context Y (right). These were used on each day of the food-context conditioning.

(B) Schematic of each context arena (top view) showing the eight possible food locations (cardinal and intercardinal directions: *N*, north; *NE*, northeast; *E*, east; *SE*, southeast; *S*, south; *SW*, southwest; *W*, west; and *NW*, northwest) used throughout the whole paradigm. For each behavioral session and context, we randomly selected two of these locations to position the food containers where we provided the regular-diet Chow pellets in context X, and the Chow versus Hfd pellets in context Y. We used this randomization of food locations across all behavioral sessions to promote mnemonic association of food items to whole context as opposed to any discrete place. This randomization of food locations was also applied to the novel food tests and the re-conditioning session performed at the end of each testing day (to prevent behavioral extinction).

(C) Shown are the two food containers used to provide mice with a choice between one chow pellet versus a new food resource (e.g., raspberry jam) during the post-conditioning novel food tests (e.g., days 11 and 12 in the 10-day conditioning paradigm).

(D) Layout and intake for the post-conditioning Hfd food test in context Z. Shown on the left is a top-view schematic of context Z as a circular-shaped arena that mice also explored on each conditioning day. For the post-conditioning Hfd food test, mice re-explored this neutral context Z now with two food containers providing a choice between one chow pellet versus one Hfd pellet (as in context Y during conditioning; see A). Shown on the right is the corresponding estimation plot (see methods) showing the effect size for the difference in Hfd intake between context Y versus context Z. Note the significantly stronger Hfd food intake in context Y compared to Z ($p < 0.001$, two-sided paired permutation test).

(E) Mice body weight across experiment days remained stable (each data point represents one mouse; $t = -0.68$, $p = 0.5$; multiple regression).

*** $P < 0.001$, two-sided paired permutation test, 6 mice.

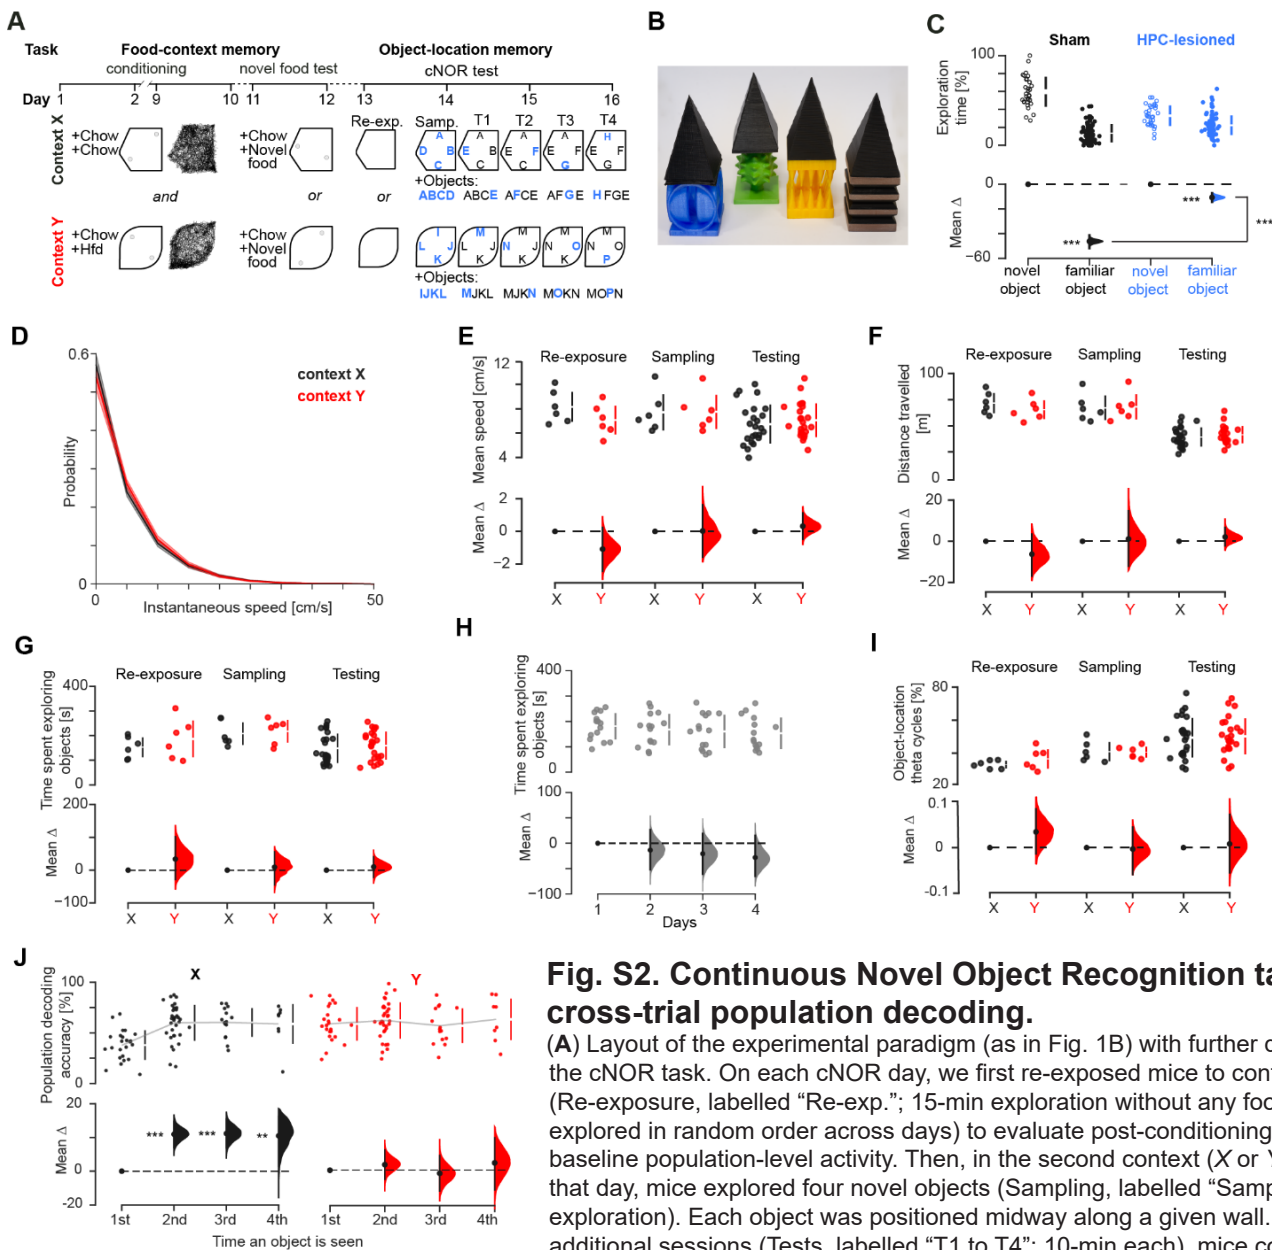


Fig. S2. Continuous Novel Object Recognition task and cross-trial population decoding.

(A) Layout of the experimental paradigm (as in Fig. 1B) with further details about the cNOR task. On each cNOR day, we first re-exposed mice to context X and Y (Re-exposure, labelled “Re-exp.”; 15-min exploration without any food; X and Y explored in random order across days) to evaluate post-conditioning changes in baseline population-level activity. Then, in the second context (X or Y) explored that day, mice explored four novel objects (Sampling, labelled “Samp.”; 15-min exploration). Each object was positioned midway along a given wall. Over four additional sessions (Tests, labelled “T1 to T4”; 10-min each), mice continued to

explore this context where one of the initial objects from the sampling session was replaced by a novel one. That is, by the fourth cNOR test, all initial novel objects (e.g., A, B, C, and D) were replaced by another one (e.g., E, F, G, and H). Using this procedure, mice were exposed to three familiar (previously seen) and one completely novel object on each cNOR test. On each test, we measured the time spent on each object and we calculated the percentage time spent investigating the novel object versus the (mean) percentage time spent investigating the familiar objects.

(B) Example set of four objects. The “pyramid hat” on each object was used to prevent mice from climbing and staying onto the objects.

(C) Hippocampal lesions impair cNOR performance. Estimation plot showing the effect size in the difference in time spent exploring novel versus familiar objects. Mice with sham lesions showed a strong preference for the novel object over the familiar ($p < 0.001$, two-sided paired permutation test; $n = 8$ mice; 2 cNOR test days per mouse). Mice with hippocampal lesions showed significantly weaker cNOR performance ($p < 0.001$, two-sided paired permutation test; $n = 8$ mice; 2 cNOR days per mouse). Direct comparison showed significantly stronger novelty preference in sham compared to hippocampal-lesioned mice ($p = 0.005$, two-sided permutation test).

(D-F) Locomotor speed and distance travelled did not differ between context X and Y. This is reported in the distribution of instantaneous speed values (D); the estimation plot showing the mean animal speed during re-exposure, sampling, and test sessions (E); and the corresponding distance travelled (F).

(G-I) Estimation plots showing that the time spent exploring the four object locations was not different across contexts during the re-exposure, sampling and test sessions (G) and across the four cNOR days (H). Likewise, the percentage of theta cycles during exploration of each object-location was not different across contexts (I).

(J) Decoding accuracy of each object-location across the four cNOR tests. The 4 objects initially sampled are replaced by a new object, one at a time across tests. Thus, only one of the initially sampled objects is seen 4 times; and the last one being replaced is only seen once. Decoding increased with the number of times an object was encountered in context X (suggesting a gain in familiarity), but this was not the case in context Y.

** $P < 0.01$, *** $P < 0.001$, two-sided paired permutation test, 6 mice.

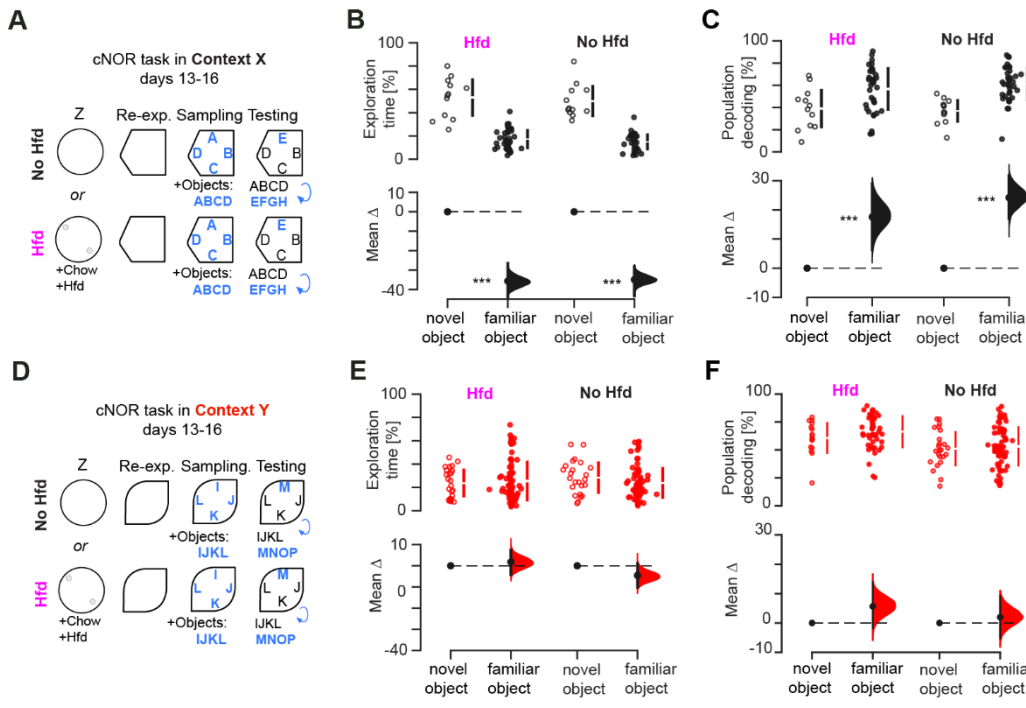


Fig. S3. Eating Hfd prior to cNOR does not change contextual object-location memory.

Following 10 days conditioning, mice were provided on some cNOR days with Chow and Hfd pellets in the neutral context Z (fig. S1D) prior to performing the cNOR task in either context X (A-C) or context Y (D-F) that day. Shown are the layouts of the corresponding sessions (A, D), the percentage of exploration time with novel and familiar objects during cNOR tests (B, E; as in Fig. 1E), and the classification accuracy of object-location compound in test by the GLM trained in session (C, F; as in Fig. 1G). Note that prior Hfd feeding neither changed the successful object-location memory in context X nor the defective one in context Y. *** $P < 0.001$, two-sided paired permutation test. $N = 3$ mice with Hfd and 3 mice without Hfd prior to cNOR in X (A-C); 2 mice with Hfd and 3 mice without Hfd prior to cNOR in Y (D-F).

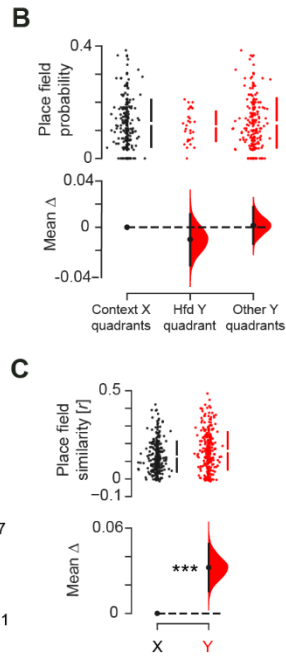
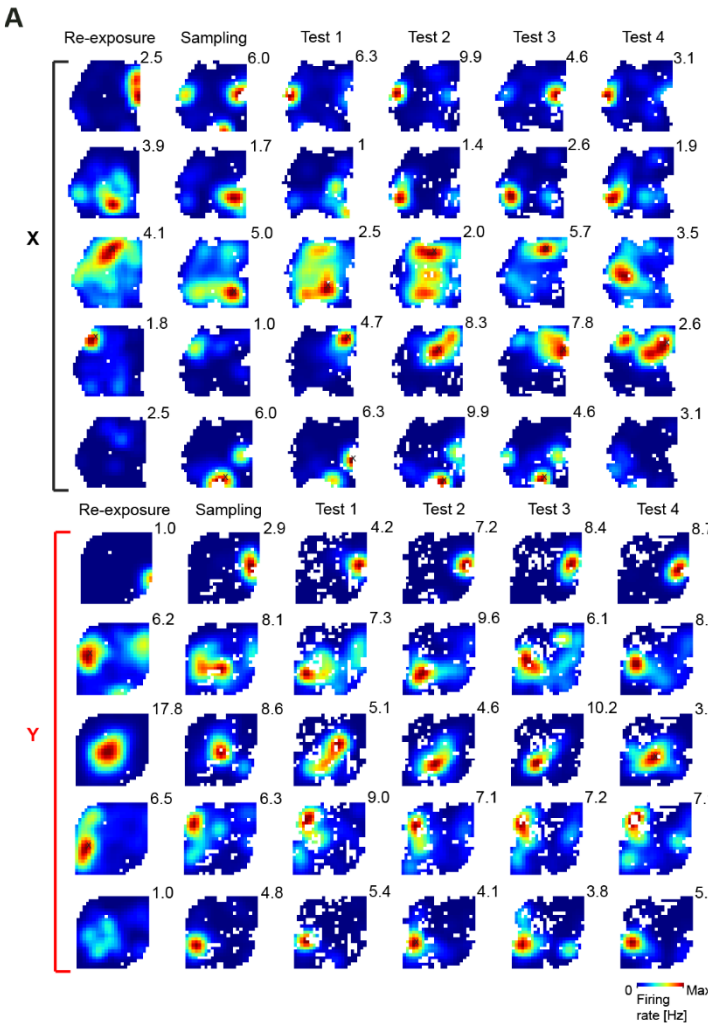


Fig. S4. CA1 place maps are more stable across cNOR sessions in context Y.

(A) Example CA1 firing maps across the consecutive cNOR sessions for one mouse day in context X (top) and one day in context Y (bottom). One principal cell per row (with numbers indicating peak firing rate for each map).

(B) The defective object-location population decoding expressed in context Y (Fig. 1F-J) was not explained by a CA1 place field over-representation of the last Hfd location. In our paradigm, food locations (fig. S1B) were randomized across days to promote mnemonic association of food to whole context as opposed to a fixed, discrete location. This randomization was applied on each conditioning session and on the re-conditioning session conducted at the end of each test day (to prevent extinction). Consistent with this, the probability for CA1 place fields to represent the previously experienced Hfd location while performing cNOR tests in context Y ("Hfd Y quadrant") was not significantly different from the probability to represent any other

locations in context Y ("Other Y quadrants"). This is also compared to the probability to observe place field peaks in any quadrants of context X where Hfd has never been provided ("Context X quadrants"). We found this by comparing the amount of place field peaks in the last encountered Hfd quadrant with that in the other quadrants. This indicates that no over-representation of the past Hfd location could explain the defective object-location memory (Fig. 1E) and population decoding (Fig. 1G) in context Y.

(C) Mean place field similarity (PFS) for the pairs of place maps expressed by a given CA1 principal cell across two contiguous cNOR sessions in context X versus those in context Y. Here, the results of this analysis are using the mean PFS across all pairs of contiguous cNOR sessions for each cell. Each data point represents one CA1 place cell. This PFS analysis shows that the hippocampal CA1 place code exhibited greater stability across cNOR sessions when performed in context Y, but stronger spatial remapping when in context X. See Fig. 2B for this analysis applied to the individual PFS for each pair of contiguous cNOR sessions for each place cell. *** $P < 0.001$, two-sided paired permutation test, 6 mice, 2506 CA1 principal cells.

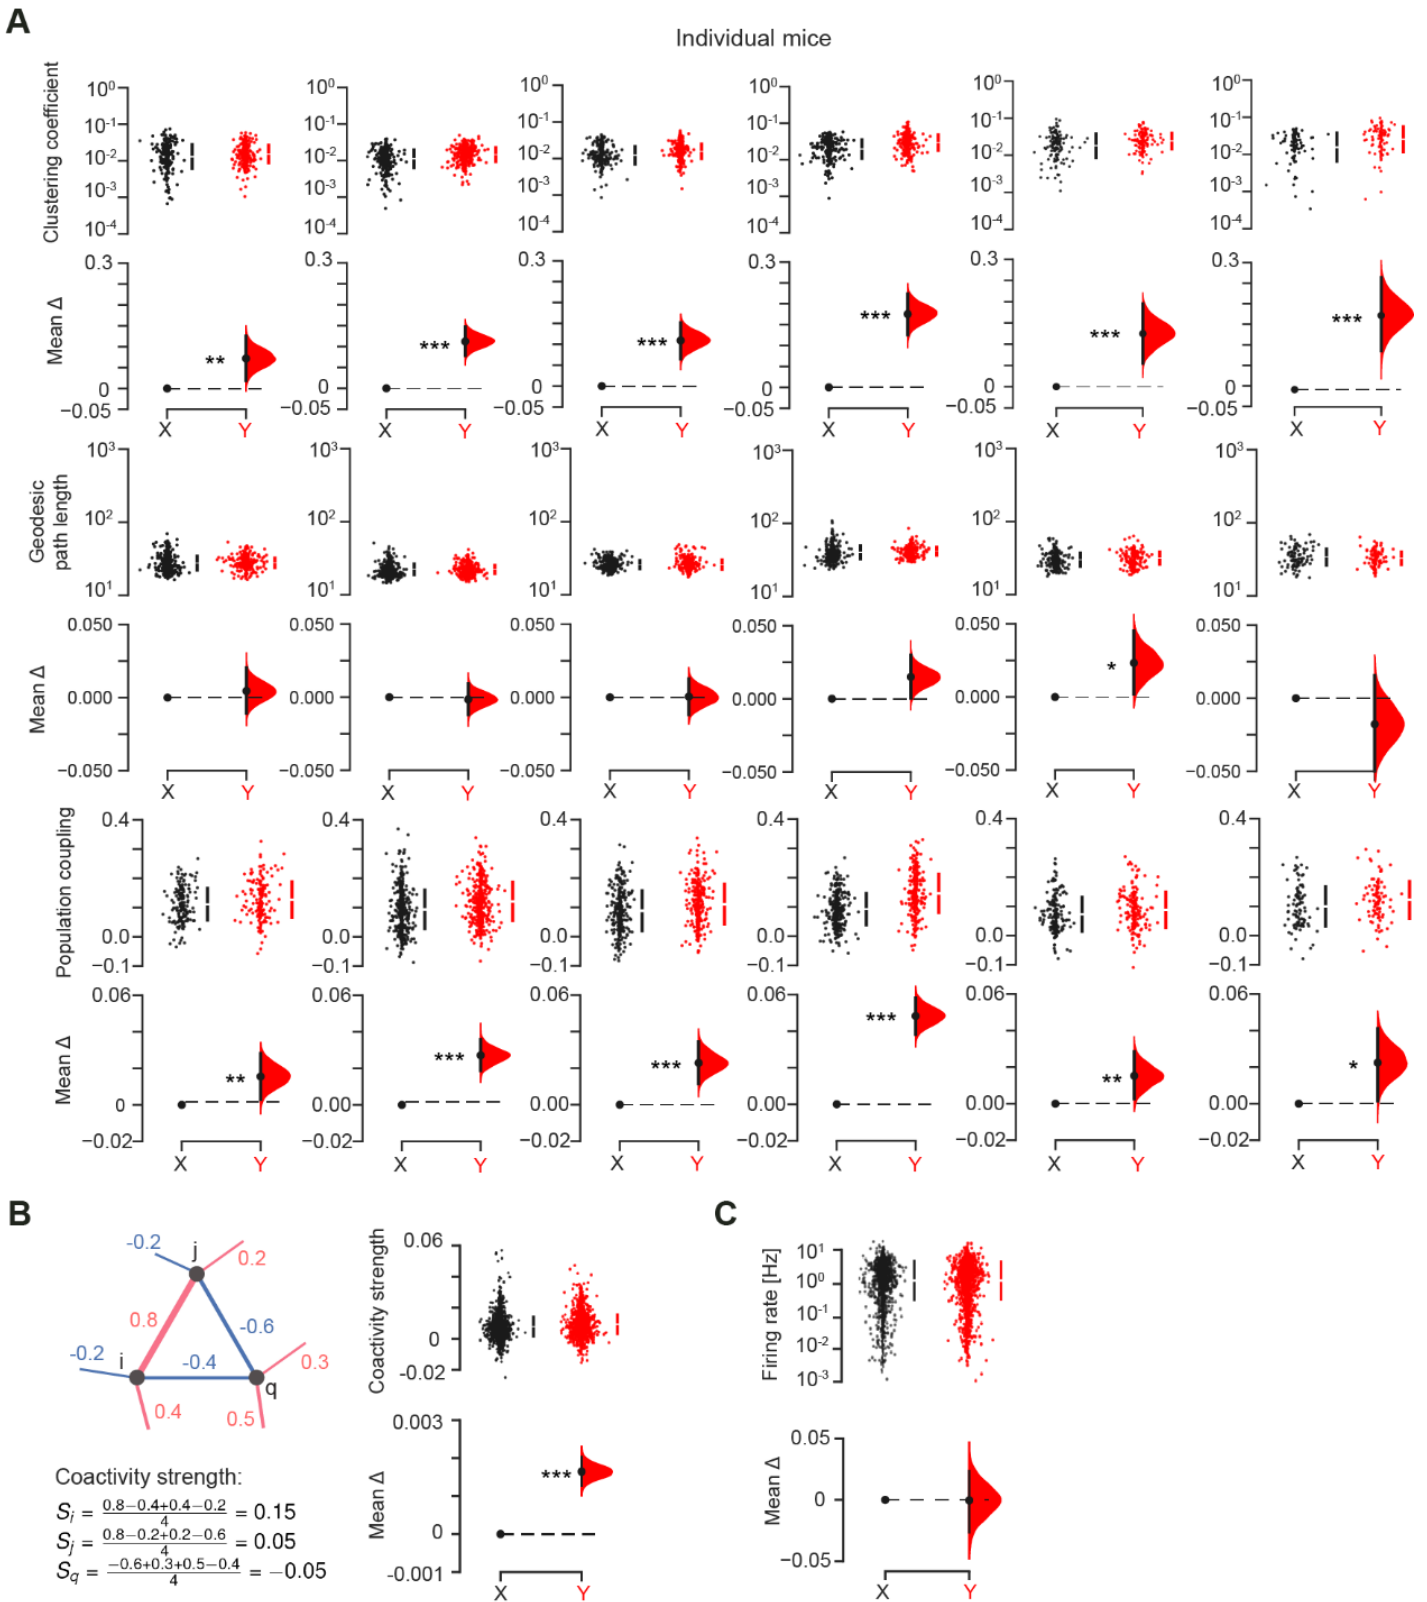


Fig. S5. Animal-wise neuronal coactivity and population coupling were higher in context Y.

(A) Estimation plots showing the effect size for the differences in mean clustering coefficient (first row), geodesic path length (second row), and population coupling (third row) between the neuronal coactivity graphs representing context X and those representing context Y across individual animals (one mouse per column).

(B) Coactivity strength analysis. *Left:* for clarity, an example subset of a coactivity graph is shown (top) with three neurons (nodes) along with their edges (lines) and signed pairwise coactivity values (numbers); for each node, the average weight of its edges is calculated (bottom). *Right:* estimation plot showing the effect size for the difference in coactivity strength levels in context Y compared to context X.

(C) The higher mean pairwise coactivity strength (B) occurred while the mean firing rate of CA1 principal cells was not different across contexts.

*** $P < 0.001$, two-sided paired permutation test, 6 mice, 2506 CA1 principal cells.

Null model characterization of coactivity graphs

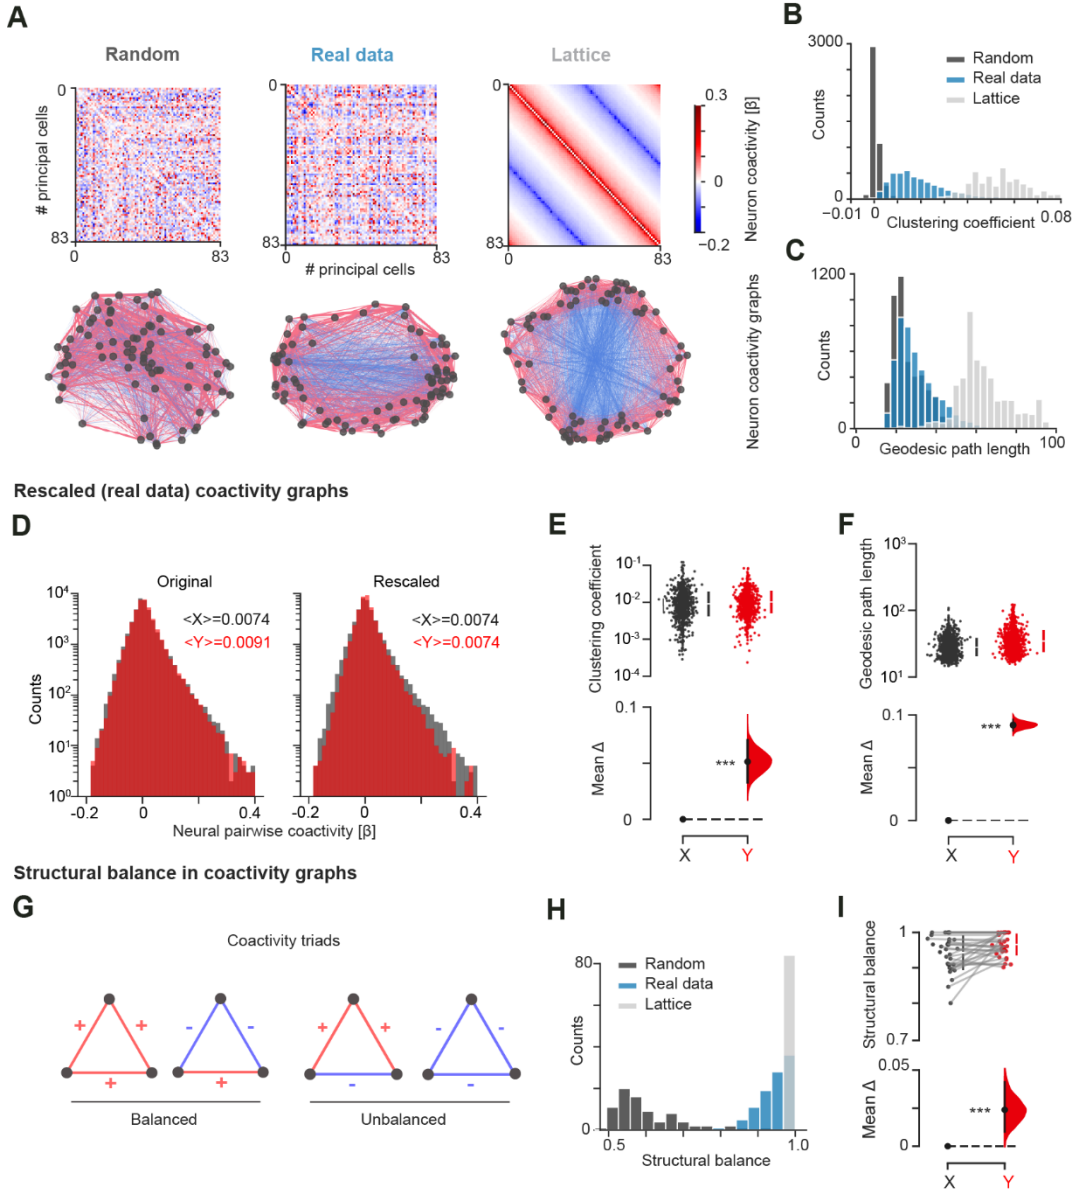


Fig. S6. The topology of hippocampal coactivity in context Y converged towards that of a regular lattice.

We conducted three additional analyses to further test the notion that the topology of hippocampal coactivity graphs developed in context Y with Hfd conditioning progressed towards the structure of a regular lattice.

(A-C) Null models (random and lattice equivalent graphs) characterization. We compared the real data networks obtained from our recordings to both their lattice and their random equivalents (see Methods). For one example session in context Y (**A**) are shown the principal cell coactivity matrices (top) and the corresponding graphs (bottom) for the random equivalent (left column), the real data (middle column; another example day of real data graph is shown in Fig. 2E), and the lattice equivalent (right column) of the CA1 network recorded that day. The distribution of clustering coefficients (**B**) and geodesic path lengths (**C**) are shown for the real data graphs obtained for all contexts, along with their random shuffled and lattice equivalents. Note that real data networks are overall featured by clustering coefficients closer to those of lattice networks while their geodesic path lengths resembled those of random networks, thus indicating a small-world topology (16, 18).

(D-F) Rescaled (real data) coactivity graphs. We investigated whether the lattice-like topology of the neuronal coactivity structure observed in context Y (Fig. 2H, I) is merely explained by an increase in the coactivity values (Fig. S5B). To assess this, we analytically rescaled the edges of Y graphs to match the coactivity values of X graphs (**D**). With this, neuronal graphs in Y continued to display a higher clustering coefficient (**E**) but also a longer geodesic path length (**F**). Notably, this result suggested that the topological changes developed in context Y with Hfd conditioning (i.e., higher clustering coefficient and no significant change in geodesic path length; Fig. 2H, I) are not mere consequences of increased pairwise coactivity in the hippocampal network.

(G-I) Structural balance in coactivity graphs. Hippocampal coactivity graphs are composed of both correlated and anti-correlated spike trains (i.e., positive and negative edges). We thus further evaluated the structural balance (67) of these (signed) graphs, thresholding them to keep the top 10% of the highest magnitude links (both positive and negative) to then quantify the proportion of balanced triads (with an even number of negative edges) (**G**). As expected, the lattice equivalent graphs were fully balanced while the random ones were the most unbalanced; we observed that the real data graphs lied in between, closer to the lattice ones (**H**). Furthermore, real data coactivity graphs of context Y displayed a higher structural balance than those of context X (**I**), supporting the idea that the hippocampus network shifted towards a more rigid, lattice-like topology in context Y.

*** $P < 0.001$, two-sided paired permutation test, 6 mice, 2506 CA1 principal cells.

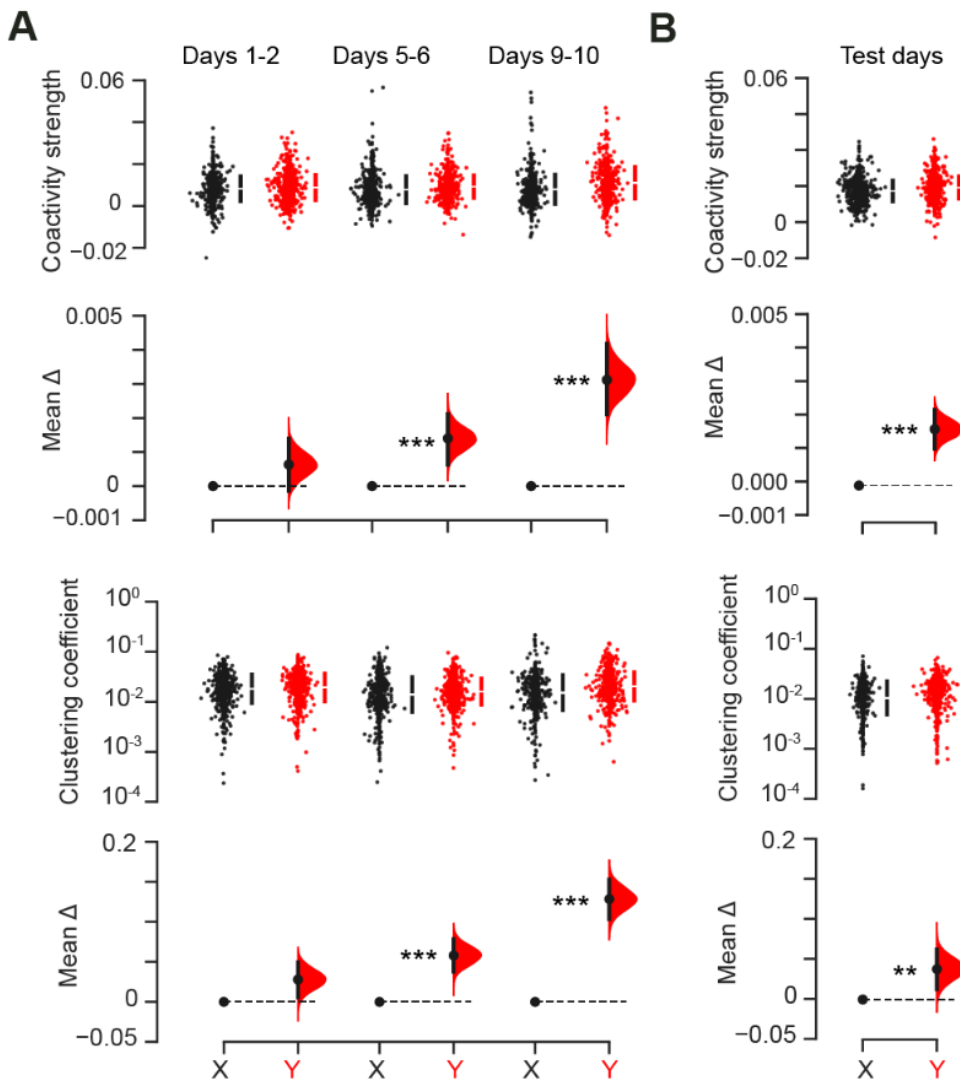


Fig. S7. The topological changes affecting coactivity graphs develop during conditioning.

We observed that the mean clustering coefficient (top) and the individual neuron coactivity strength (bottom) increased in context Y compared to X over the conditioning days (A) to continue affecting neuronal graphs during the re-exposure session in post-conditioning test days (B).

*** $P < 0.001$, ** $P < 0.01$, two-sided paired permutation test, 2506 CA1 principal cells from 6 mice.

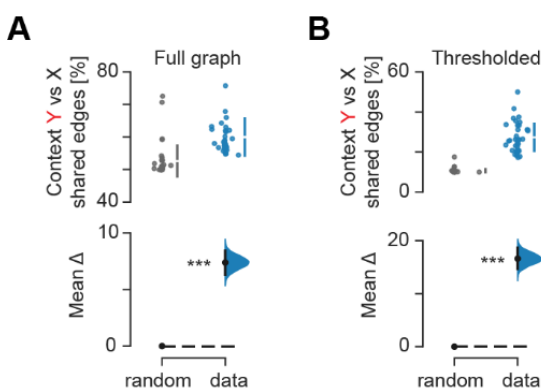


Fig. S8. Neuronal graphs share some common coactivity structure across contexts.

We further explored the relation between the population correlation structure underlying the neuronal graphs of context X versus those of context Y. Shown are the proportion of shared coactivity edges (i.e., common node-to-node same-sign edges) between the original (full dense) graphs of both context X and Y (A); and the proportion of these shared edges that remain in the corresponding binarized graphs (B; after thresholding the graphs to maintain only the 10% most positive links). For both analyses, we compared the results from the real data (population recorded in both context X and Y) graphs with their control graphs obtained by randomizing node-to-node edges in the original graphs. With both approaches, real data graphs for context X and Y (“data”) contained a higher proportion of shared structure than their random counterparts do (“random”). This shared structure might represent a backbone of correlated activity expressed across contexts. *** $P < 0.001$, two-sided paired permutation test, 6 mice.

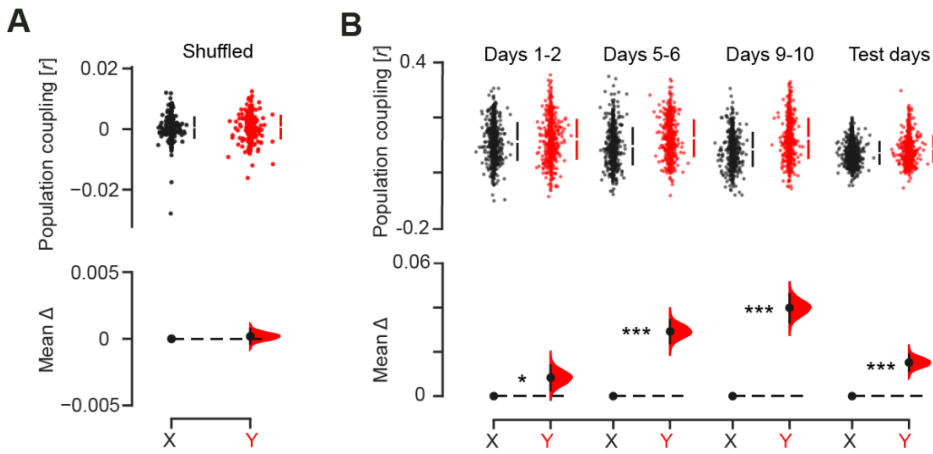


Fig. S9. Population coupling of individual neurons increases during the 10-day conditioning and is maintained thereafter. (A) Shuffling the spike times across neurons and theta cycles, while preserving each neuron's mean rate and the population rate distribution, cancelled the increased average population coupling of individual neurons seen in context Y (Fig. 2J). For this control analysis, the activity of each neuron was circularly shuffled independently with a delay drawn from a uniform distribution between 10 and 100 theta cycles. This result indicates that the increased population coupling seen in context Y compared to X (Fig. 2J) reflected stronger cross-neuron spiking relationship. (B) Heightened population coupling developed across conditioning days to continue marking the re-exposure to context Y before any test during post-conditioning days. * $P < 0.05$, *** $P < 0.001$, two-sided paired permutation test, 2506 CA1 principal cells from 6 mice.

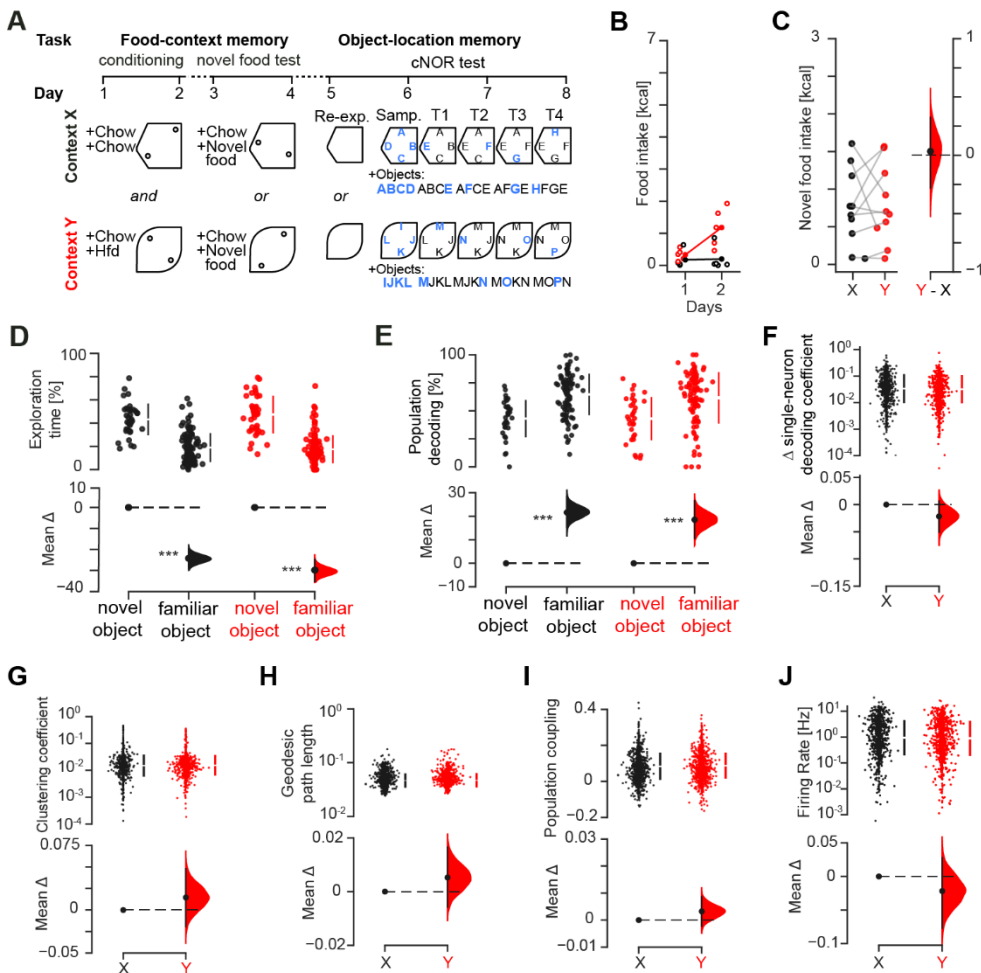


Fig. S10. A 2-day food-context conditioning did not affect novel object preference and CA1 population activity. (A) Layout of the experimental paradigm (as in Fig. 1B and fig. S2A) but using only 2 days of food-context conditioning in five additional mice. (B) Corresponding food intake during conditioning (each data point represents one mouse). (C) Estimation plot showing the effect size for the difference in novel food intake across context X and Y after 2-day conditioning. Mice did not show feeding bias in context Y over X. (D) Percentage of exploration time with novel and familiar objects during the subsequent cNOR tests in context X (black) versus context Y (red). Mice exhibited successful novel object preference in both contexts. (E) Population classification accuracy of object-location compound in test by GLM trained in session (as in Fig. 1G; each data point represents one object). The object-location decoding using the spiking activity of CA1 principal cells was similar across the two contexts. (F) The magnitude of single-neuron contribution to the population object-location classification was similar in both contexts. This contrasted with the difference obtained after 10 days of conditioning (Fig. 1J).

(G, H) The population coactivity structure remained similar in context Y versus X following the 2-day conditioning, as reported by the similar clustering coefficient of the corresponding neuronal graphs (G). The geodesic path length was also similar across contexts (H). (I, J) Single-neuron coupling to population (I) and firing rate (J) also remained similar across the two contexts. *** $P < 0.001$, two-sided paired permutation test, 5 mice, 1389 CA1 principal cells.

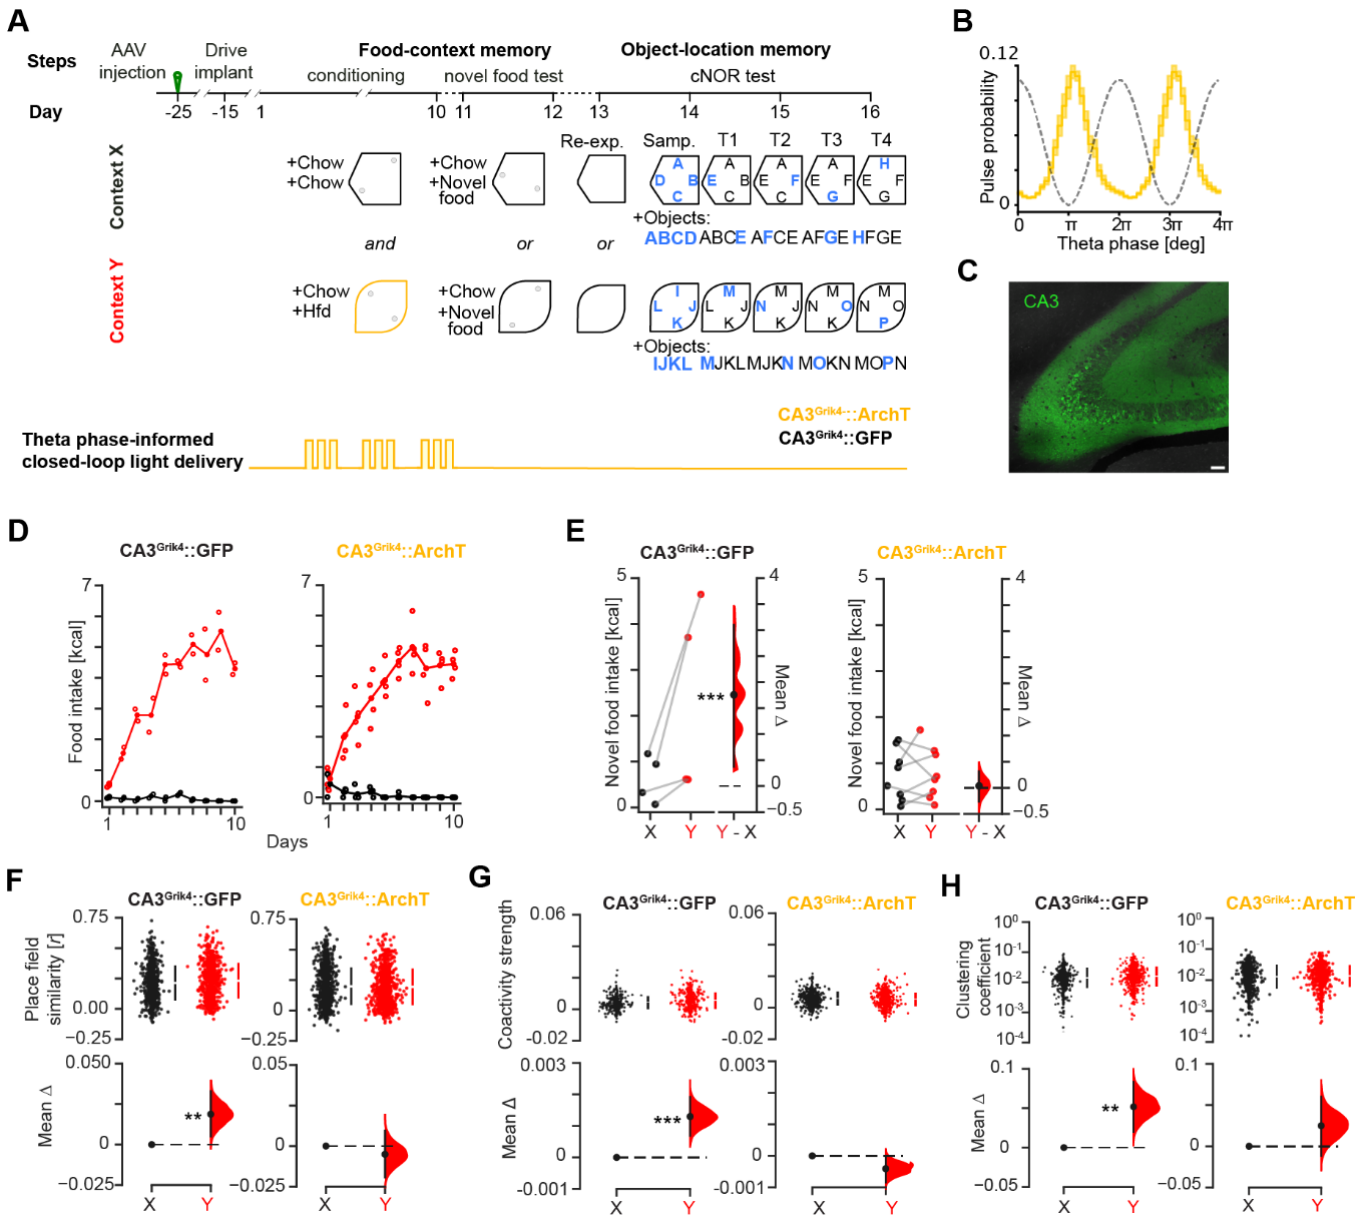


Fig. S11. Contribution of CA3 during conditioning to subsequent CA1 coactivity.

(A) Experimental layout. We used Grik4-Cre transgenic mice to target CA3 with either the ArchT-GFP construct or its GFP-only control; subsequent implantation of a 14-tetrode microdrive including two optic fibers allowed monitoring CA1 neuronal ensembles with theta phase-informed bilateral light delivery to CA3.

(B) Light pulse midpoint probability as a function of ongoing theta phase for closed-loop CA3 light delivery during each Hfd conditioning session in context Y.

(C) Low magnification confocal image showing CA3 principal cells transduced with ArchT-GFP (green). One optical section; scale bar, 100 μ m.

(D) Food intake over the 10-day conditioning in context X (black) versus context Y (red) for the CA3^{Grik4::}GFP control mice (left) and the CA3^{Grik4::}ArchT mice (right). Each data point represents one mouse.

(E) Estimation plots showing the effect size for the difference in novel food intake across context X and Y in the post-conditioning novel food test for the CA3^{Grik4::}GFP control mice (left) and the CA3^{Grik4::}ArchT mice (right). Each data point represents one mouse.

(F) Place field similarity between the place maps expressed by individual CA1 principal cells in each pair of contiguous cNOR sessions in context X versus context Y (each data point represents one cell; as in Fig. 2B). This PFS analysis indicates that applying theta phase-informed closed-loop suppression of CA3 principal cells during each Hfd conditioning session subsequently restored in CA3^{Grik4::}ArchT mice (right), but not CA3^{Grik4::}GFP mice (left), similar levels of spatial remapping across cNOR sessions in context Y compared to X.

(G) In CA3^{Grik4::}ArchT mice, we observed that following CA3 suppression throughout conditioning days, the CA1 coactivity strength in context Y was similar to that in context X during post-conditioning test days. This was not the case in CA3^{Grik4::}GFP mice, which showed higher CA1 coactivity strength in context Y compared to X.

(H) Likewise, we observed that this intervention leveled the CA1 mean clustering coefficient across the two contexts in CA3^{Grik4::}ArchT mice (right), but not in CA3^{Grik4::}GFP mice (left), during subsequent post-conditioning test days.

***P<0.001, two-sided paired permutation test, 1548 CA1 principal cells from 4 CA3^{Grik4::}ArchT mice versus 881 CA1 principal cells from 2 CA3^{Grik4::}GFP mice.




Fig. S12. Endogenous cFos expression in the CA1 *pyramidale* layer with respect to context Y and Hfd.

Confocal pictures showing cFos-expressing neurons in the CA1 *pyramidale* layer for fifteen additional mice (one picture per mouse) after either 10-day exploration of context Y without Hfd (left), 10-day consumption of Hfd in the homecage (middle), or 10-day exploration of context Y with Hfd (right). Each mouse session day was 30 minutes. Cell nuclei stained with DAPI (blue). Scale bar: 50 μm . CA1 layers: *Oriens*, Or.; *Pyramidale*, Pyr.; *Radiatum*, Rad.

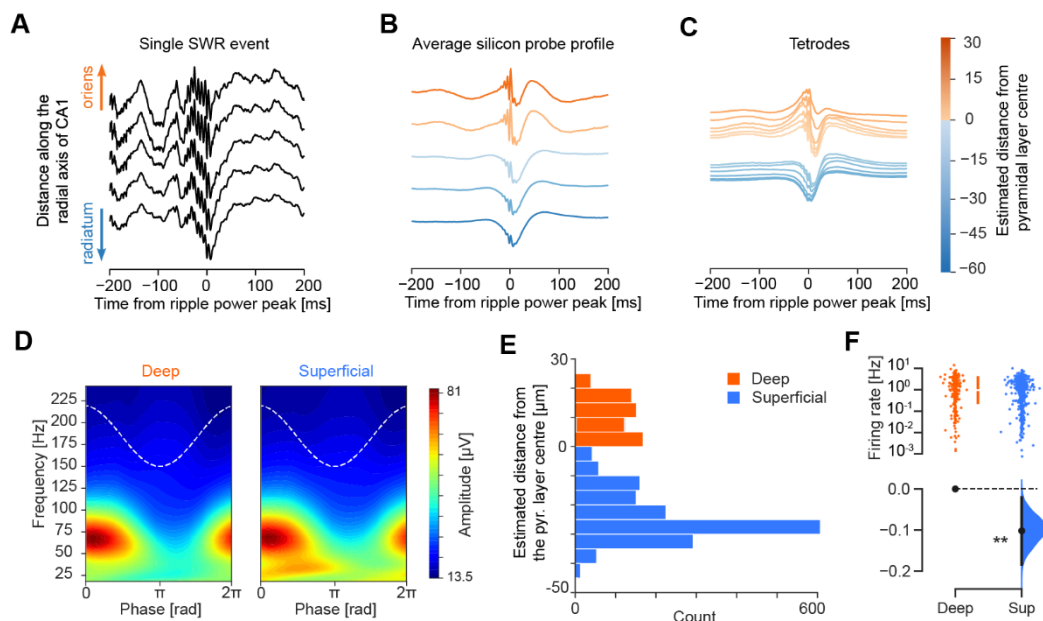


Fig. S13. Identification of deep versus superficial CA1 *pyramidale* sublayer cells.

(A-C) We estimated the position (depth) of individual tetrode-recorded principal cell soma by leveraging silicon probe recordings with known spacing between the recording sites along a linear shank (55). From these silicon probe recordings ($n=6$ mice), we first acquired the laminar profile of both sharp-wave ripples (SWRs) and theta oscillations detected in the LFPs. Using these average LFP waveforms, we then established a feature trajectory along the radial axis of the CA1 hippocampus (55). Within the CA1 pyramidal layer, the information about depth primarily relies on the laminar profile of the SWR waveforms (A, B). Single tetrodes (and thus their recorded neurons) are then projected onto this trajectory to estimate their depth (C).

(D) We further computed the average spectral profile for deep and superficial *pyramidale* tetrodes. This independent analysis showed that the spectral content of the LFP signals recorded from deep and superficial tetrodes differed with respect to the presence of slow gamma oscillations.

(E) Corresponding distribution of CA1 principal cells with respect to the estimated depth of their recording tetrodes.

(F) CA1 principal cells recorded from the superficial *pyramidale* sublayer (closer to *radiatum*) showed lower firing rates (measured during theta oscillations of active exploration) compared to their deep sublayer counterparts (closer to *oriens*), in line with previous work (15, 28, 73, 74). $**P<0.01$, two-sided paired permutation test. $N = 635$ deep sublayer cells and 1,871 superficial sublayer cells from 6 mice.

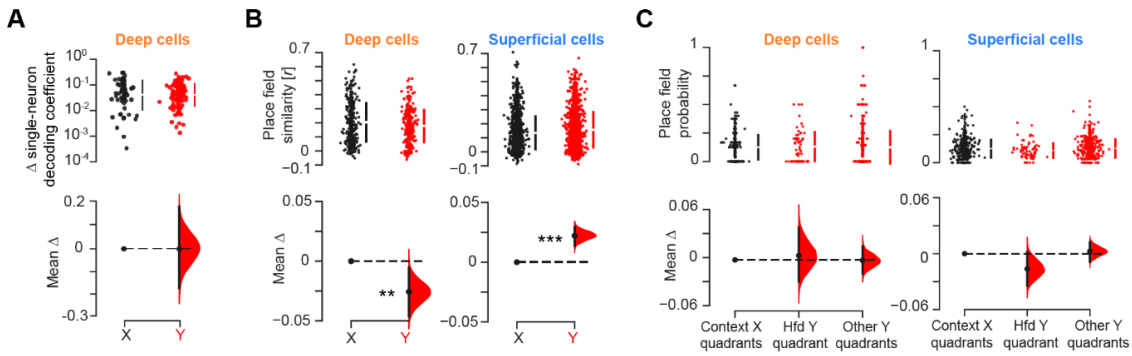


Fig. S14. Object-location decoding contribution of deep *pyramidale* sublayer cells and place map stability across cNOR sessions for the two CA1 principal cell subpopulations.

(A) Estimation plot showing the change in the magnitude of single-neuron classification contribution to novel object-location population decoding for CA1 deep *pyramidale* sublayer cells, quantified as the difference in their GLM coefficients in either context X or Y (see Fig. 4C for the superficial cells). (B) Likewise, shown for both context X and Y is the place field similarity of the spatial maps expressed by principal cells in the deep versus the superficial sublayers was assessed across cNOR sessions in either context X or context Y (as in Fig. 2B) following the 10-day conditioning. This shows that the greater stability of the CA1 place maps across cNOR sessions in context Y (Fig. 2B) is strongly explained by the higher place field similarity (i.e., lower cross-session remapping) of superficial cells. (C) The probability for deep and superficial cell to represent the previously experienced Hfd location while performing cNOR tests in context Y (“Hfd Y quadrant”) was not significantly different from the probability to represent any other locations in context Y (“Other Y quadrants”) or in context X (“Context X quadrants”). Same place field analysis as in fig. S4B. *** $P < 0.001$, ** $P < 0.01$, two-sided paired permutation test. $N = 635$ deep sublayer cells and 1,871 superficial sublayer cells from 6 mice.

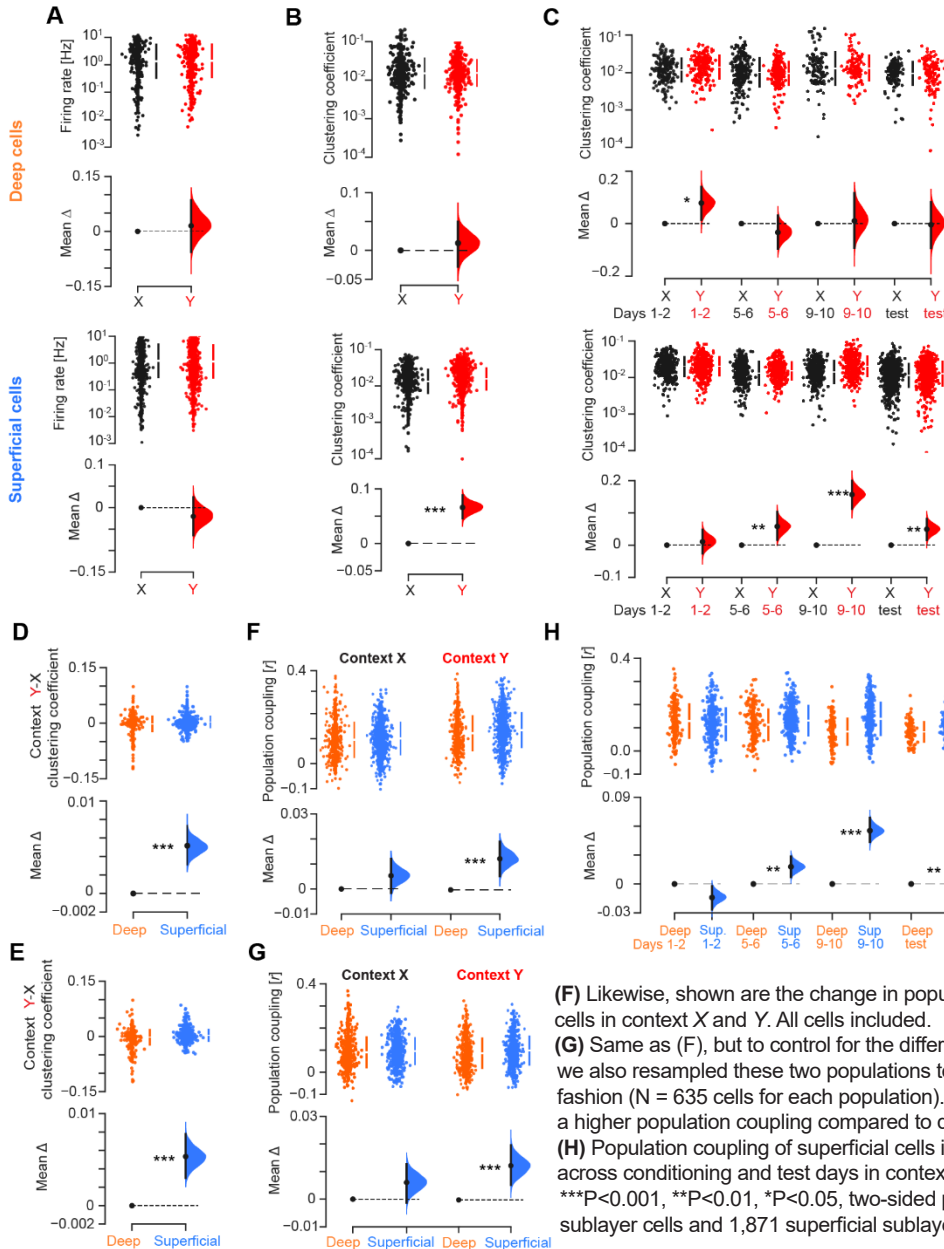


Fig. S15. Comparison of deep versus superficial CA1 cell activity.

(A) Estimation plots showing no significant change in the firing rate of principal cells in the deep (top) and the superficial (bottom) CA1 *pyramidale* sublayers between the Hfd-paired context Y compared to the chow-paired context X. (B) Likewise, shown is the coactivity clustering coefficient for deep (top) and superficial (bottom) *pyramidale* sublayer cells between context X and Y. (C) The clustering coefficient for deep cells (top) tended to be similar between context X and Y from conditioning to testing, albeit higher in Y during the first two days of conditioning. The coactivity clustering coefficient of superficial cells (bottom) increased in context Y compared to X across task days. (D) Change in clustering coefficient for deep versus superficial cells in context Y minus that in context X. All cells included. (E) Same as (D), but to control for the different numbers of deep versus superficial cells (fig. S13F), we resampled these two populations to match their sample size in a day-by-day fashion ($N = 635$ cells for each population). This confirmed that superficial cells display a higher clustering coefficient compared to deep *pyramidale* sublayer cells. (F) Likewise, shown are the change in population coupling for deep versus superficial cells in context X and Y. All cells included. (G) Same as (F), but to control for the different numbers of deep versus superficial cells, we also resampled these two populations to match their sample size in a day-by-day fashion ($N = 635$ cells for each population). This confirmed that superficial cells display a higher population coupling compared to deep *pyramidale* sublayer cells. (H) Population coupling of superficial cells increased compared to that of deep cells across conditioning and test days in context Y. *** $P < 0.001$, ** $P < 0.01$, * $P < 0.05$, two-sided paired permutation test. $N = 635$ deep sublayer cells and 1,871 superficial sublayer cells from 6 mice.

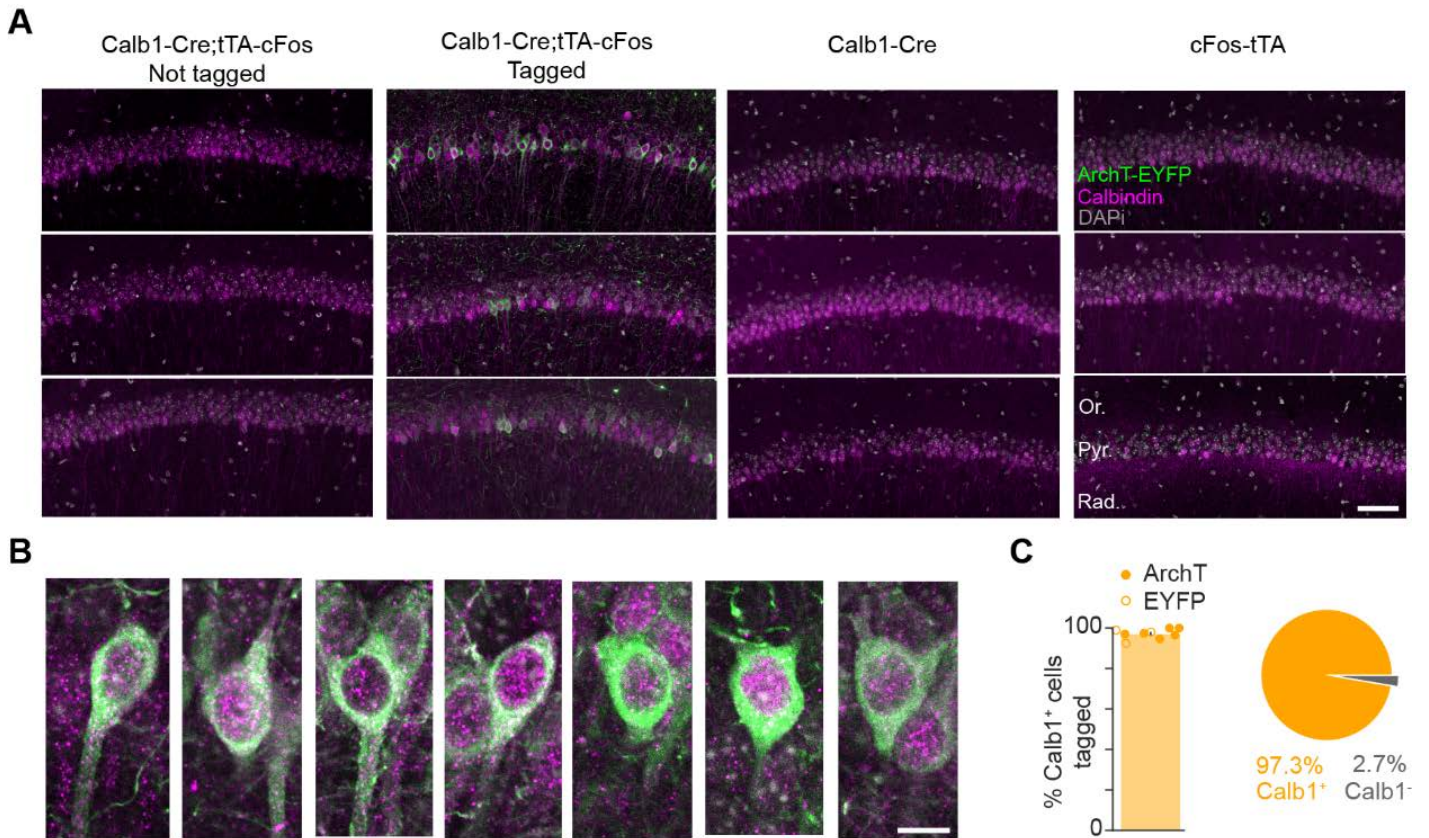


Fig. S16. Histological validation of the intersectional optogenetic strategy used to tag CA1 superficial pyramidale sublayer cells recruited during Hfd-context Y conditioning.

(A, B) To validate the intersectional optogenetic targeting of CA1 superficial pyramidale sublayer cells during Hfd-context Y conditioning (Fig. 5), we injected the hSyn-Con/Fon-ArchT-EYFP and associated TRE3G-FlpO constructs in four groups of mice (from left to right): (1) the double-transgenic Calb1-Cre;cFos-tTA mice exposed to Context Y with Hfd while always on homecage Dox diet (i.e., “not tagged”); (2) the double-transgenic Calb1-Cre;cFos-tTA mice exposed to Context Y with Hfd while transiently not on homecage Dox diet (i.e., “tagged”); (3) the mono-transgenic Calb1-Cre mice and (4) the cFos-tTA mice, also both exposed to Context Y with Hfd. Shown in (A) are low magnification confocal pictures for ArchT-EYFP expression (green) and Calbindin immunostaining (magenta) in the CA1 *pyramidale* layer for twelve mice (one picture per mouse; each column represents one group) one week after 30-min exposure to context Y with Hfd. ArchT-EYFP was selectively expressed in the pyramidale layer of the Context Y “tagged” Calb1-Cre;cFos-tTA mice but not the other three groups. Shown in (B) are high magnification confocal pictures for representative examples of CA1 cells expressing ArchT-EYFP (green) and immuno-positive for Calbindin (magenta) from the double-transgenic Calb1-Cre;cFos-tTA mice tagged in Context Y with Hfd (second column from the left in A). Scale bar: 50 μm (A) and 10 μm (B). Cell nuclei stained with DAPI (gray). CA1 layers: *Oriens*, Or.; *Pyramidale*, Pyr.; *Radiatum*, Rad.

(C) In Context Y-tagged Calb1-Cre;cFos-tTA mice, $97.26 \pm 0.77\%$ of the tagged CA1 *pyramidale* layer cells (with ArchT-EYFP in 6 mice, or EYFP-only in 3 mice) were immuno-stained with the Calb1 antibody (with $90.98 \pm 0.93\%$ of these cells located in the superficial sublayer and $9.02 \pm 0.93\%$ in the deep sublayer).

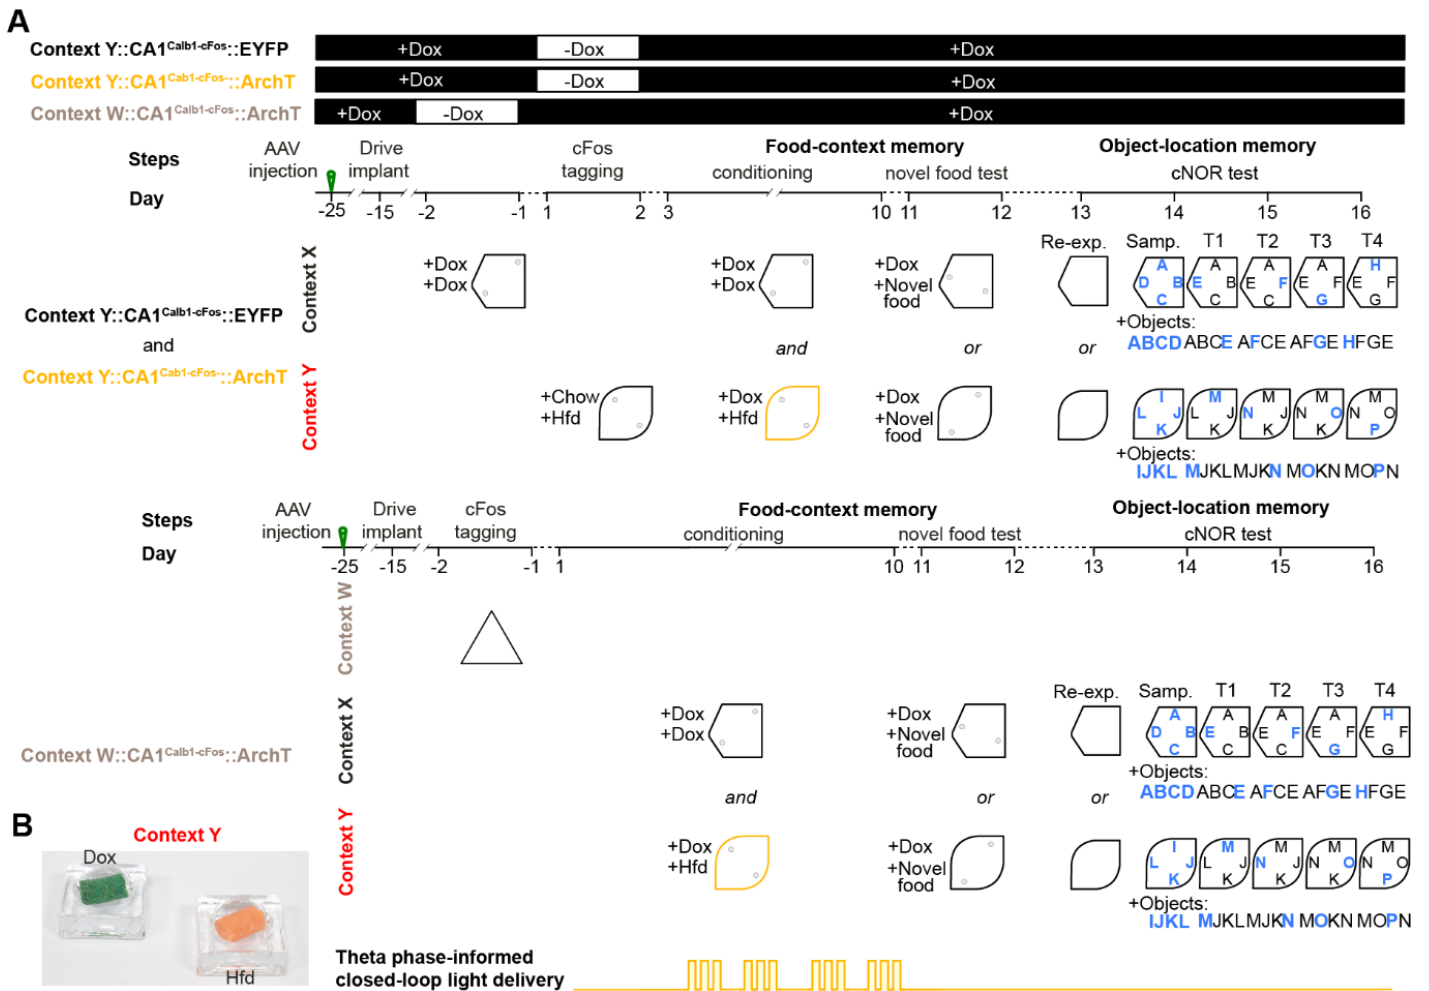


Fig. S17. Experimental framework for theta phase-informed optogenetic feedback in CA1^{Calb1-cFos} mice.
 (A) Experimental layout. We used double-transgenic Calb1-Cre;cFos-tTA mice fed with doxycycline (Dox) containing food pellets in their homecage. Using an intersectional targeting strategy (Fig. 5A), we targeted the CA1 of these mice with either the ArchT-EYFP construct or its EYFP-only control; subsequent implantation of a 14-tetrode microdrive including two optic fibers allowed monitoring CA1 neuronal ensemble with theta phase-informed bilateral light delivery. To tag CA1 neurons in the Hfd-paired context Y with ArchT-EYFP (in ContextY::CA1^{Calb1-cFos}::ArchT mice) or EYFP-only (in ContextY::CA1^{Calb1-cFos}::EYFP mice), the homecage feeding with Dox pellets was replaced with regular Chow pellets during the 48-hour period corresponding to days 1 and 2 of Hfd-context Y conditioning. To restrict this tagging to context Y, mice were not exposed to the other contexts during this off-Dox period. For this reason, the context Y tagged mice started their day 1 and 2 conditioning of Chow in context X in the two days before, while still on Dox diet (i.e., they were not exposed to context X for the two days corresponding to Hfd-context Y tagging). CA1 light delivery was then actuated by real time tracking of theta phase while mice continued to explore context Y with Hfd from conditioning day 3 onward. In a third group of mice, we tag CA1 neurons recruited in a neutral (i.e., unrelated to the contextual feeding task) context W with ArchT-EYFP (ContextW::CA1^{Calb1-cFos}::ArchT mice). For this, the homecage feeding with Dox pellets was transiently replaced with regular Chow pellets for the 48-hour period corresponding to days -1 and -2 with respect to the Hfd-context Y conditioning. To restrict neuronal tagging to context W, these mice were not exposed to the other contexts during this off-Dox period. Context W tagged mice then started their conditioning in context X versus context Y in the following days while back on Dox diet, with CA1 light delivery actuated by real time tracking of theta phase while mice explored context Y with Hfd in each conditioning session.
 (B) The two food containers used to provide mice with a choice between one Dox-diet pellet versus one Hfd pellet in context Y.

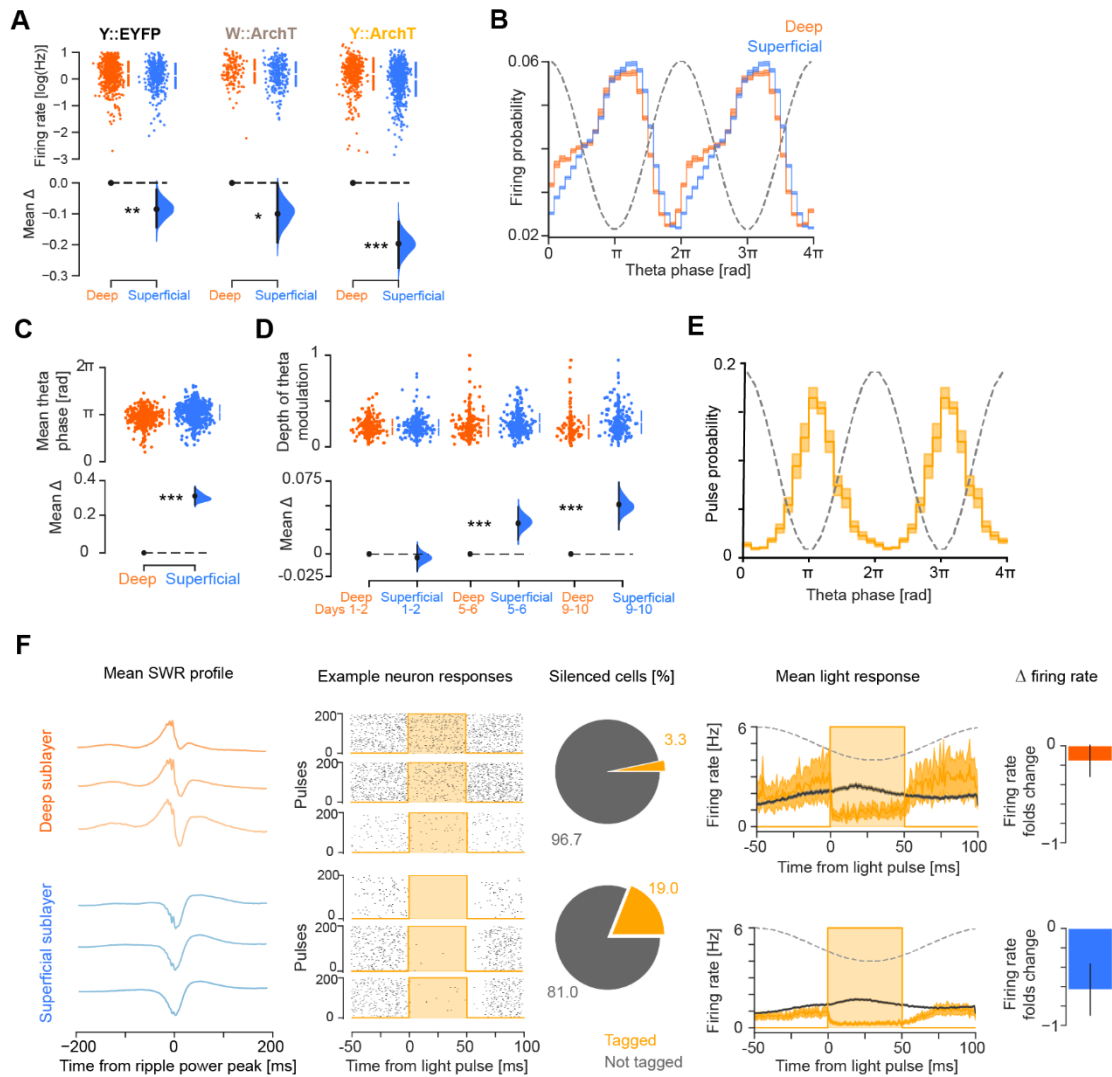


Fig. S18. Firing activity of deep and superficial CA1 principal cells related to the theta phase-informed manipulation.

(A) Estimation plot showing that the baseline firing rate of CA1 superficial cells (blue) is lower than CA1 deep cells (orange) in the three groups of CA1^{Calb1-cFos} mice engaged in the theta phase-informed optogenetic feedback framework (fig. S17). Baseline firing rate calculated in the absence of light delivery. This result replicates the difference observed in the initial cohort of mice (fig. S13F). (B-D) Basic theta-paced firing activity of CA1 principal cells recorded in the original, non-optogenetically modulated mice. Firing probability of deep versus superficial CA1 cells as a function of theta phase (B) and mean theta phase coupling difference between these two neuronal populations (C). Note that superficial cells tend to fire slightly later in the theta cycle. The theta modulation of CA1 superficial cell spiking increases compared to CA1 deep cells as the Hfd conditioning progresses in context Y (D). (E) Probability of real time light delivery (50-ms pulse midpoint) as a function of ongoing theta phase during Hfd conditioning in context Y.

(F) Effect of the closed-loop intervention on the activity of CA1 principal cells recorded in the deep versus superficial CA1 pyramidal sublayer (top and bottom rows, respectively) of the CA1^{Calb1-cFos} mice tagged with ArchT. From *Left to Right*: (i) Example SWR profiles used to locate individual tetrodes along the CA1 radial axis (15). (ii) Examples raster plots for principal cells recorded from deep versus superficial tetrodes and showing single-neuron firing responses to light delivery. (iii) Pie charts showing the percentage of ArchT-tagged principal cells (i.e., cells silenced by light delivery) per CA1 sublayer; note that the few silenced cells recorded in the deep sublayer are likely to be Calbindin-1 cells (fig. S16C). (iv) With respect to the onset of light pulse delivery, shown is the time course of the average firing activity of tagged (yellow) and non-tagged (gray) principal cells; for clarity, cosines (dashed) indicate theta phase reference using one cycle (note the theta modulation of the firing rate). (v) Corresponding folds change in firing rate for the deep versus superficial cell subpopulations.

*** $P < 0.001$, ** $P < 0.01$, * $P < 0.05$, two-sided paired permutation test; $n = 1097$ CA1 principal cells (413 deep with 684 superficial) from 5 ContextY::CA1^{Calb1-cFos}::ArchT mice (Y::ArchT), 1007 (416 deep with 591 superficial) from 4 ContextY::CA1^{Calb1-cFos}::EYFP mice (Y::EYFP), 683 (231 deep with 452 superficial) from 3 ContextW::CA1^{Calb1-cFos}::ArchT mice (W::ArchT) (A, F); and 635 deep sublayer cells with 1,871 superficial sublayer cells from 6 mice (B-D).

