Offline hippocampal reactivation during dentate spikes supports flexible memory

Highlights

- Neurons across hippocampal regions synchronize during offline dentate spike events

- Hippocampal population activity differs in dentate spikes versus sharp-wave ripples

- Dentate spikes nest higher-dimensional patterns of hippocampal population activity

- Waking coactivity patterns reactivate in offline DS events to support flexible memory

Authors

Stephen B. McHugh,
Vítor Lopes-dos-Santos,
Manfredi Castelli, ..., Timothy Denison,
Andrew Sharott, David Dupret

Correspondence

stephen.mchugh@bndu.ox.ac.uk
(S.B.M.),
david.dupret@bndu.ox.ac.uk (D.D.)

In brief

The offline reactivation of hippocampal neuronal activity during sleep is thought to be crucial for memory. We investigated the contribution of dentate spikes to offline reactivation of hippocampal activity patterns and flexible recognition memory by combining triple-ensemble (DG-CA3-CA1) recordings and closed-loop optogenetic feedback in mice undergoing sleep/rest following active exploratory behavior. This work establishes dentate spikes as a second network event that, together with sharp-wave ripples, supports offline hippocampal dynamics for memory-guided behavior.

CellPress

Neuron

Article

Offline hippocampal reactivation during dentate spikes supports flexible memory

Stephen B. McHugh,[1,*] Vítor Lopes-dos-Santos,[1] Manfredi Castelli,[1] Giuseppe P. Gava,[1] Sophie E. Thompson,[1] Shu K.E. Tam,[1] Katja Hartwich,[1] Brook Perry,[1] Robert Toth,[1] Timothy Denison,[1] Andrew Sharott,[1] and David Dupret[1,2,*]
[1]Medical Research Council Brain Network Dynamics Unit, Nuffield Department of Clinical Neurosciences, University of Oxford, Oxford OX1 3TH, UK
[2]Lead contact
*Correspondence: stephen.mchugh@bndu.ox.ac.uk (S.B.M.), david.dupret@bndu.ox.ac.uk (D.D.)
https://doi.org/10.1016/j.neuron.2024.08.022

SUMMARY

Stabilizing new memories requires coordinated neuronal spiking activity during sleep. Hippocampal sharp-wave ripples (SWRs) in the cornu ammonis (CA) region and dentate spikes (DSs) in the dentate gyrus (DG) are prime candidate network events for supporting this offline process. SWRs have been studied extensively, but the contribution of DSs remains unclear. By combining triple-ensemble (DG-CA3-CA1) recordings and closed-loop optogenetics in mice, we show that, like SWRs, DSs synchronize spiking across DG and CA principal cells to reactivate population-level patterns of neuronal coactivity expressed during prior waking experience. Notably, the population coactivity structure in DSs is more diverse and higher dimensional than that seen during SWRs. Importantly, suppressing DG granule cell spiking selectively during DSs impairs subsequent flexible memory performance during multi-object recognition tasks and associated hippocampal patterns of neuronal coactivity. We conclude that DSs constitute a second offline network event central to hippocampal population dynamics serving memory-guided behavior.

INTRODUCTION

Memories are stabilized during periods of sleep and rest.[1–4] Decades of work have provided important insights into the underlying brain network mechanisms and have identified offline hippocampal activity as essential for this process.[5,6] Central to our current understanding are hippocampal sharp-wave ripples (SWRs) that feature an intermittent, high-frequency (100–250 Hz) network event detected in the local field potentials (LFPs) of the cornu ammonis (CA)1 region.[7–10] During SWRs, the firing activity of CA1 principal cells is transiently modulated[11,12] and reactivates the population-level firing patterns expressed in previous waking experience.[13] These offline spiking correlates have behavioral significance: suppressing CA1 neurons during SWRs impairs memory recall for recently acquired information.[14–16] Conversely, prolonging SWRs or reinforcing the coordination between SWRs and neocortical activity promotes memory consolidation and subsequent behavioral performance.[17,18] Hippocampal SWRs therefore constitute an offline network event important for memory-guided behavior. However, during sleep/rest periods, the hippocampus exhibits another prominent network event: dentate spikes (DSs), which are seen in the LFPs of the dentate gyrus (DG). To date, DSs have received little attention compared with SWRs. Accordingly, here we characterize the neuronal spiking dynamics nested in DSs with respect to SWRs and evaluate whether DSs constitute a second network

event central to offline reactivation of waking firing patterns and subsequent memory-guided behavior.

The DG gates sensory information to the hippocampus, notably decorrelating these inputs into dissimilar neural patterns.[19–21] This function may be crucial for the hippocampus to integrate multiple items in memory and to flexibly distinguish between stimuli with overlapping features. Within the DG, DSs represent intermittent, large-amplitude network events recorded in the LFPs of the DG granule cell layer and are associated with increased spiking activity in dentate cells.[22–24] However, across the literature, both increased and suppressed spiking activity of CA principal cells have been reported,[22,24–28] although, notably, some of these studies were in anesthetized[25,27] or head-fixed animals.[24,28] Thus, here we further performed a systematic comparative assessment of DG and CA principal cell spiking activity during DSs versus SWRs in non-anesthetized, freely behaving mice.

To investigate the influence of DSs on hippocampal population activity and memory, we combined triple-site (DG-CA3-CA1) extracellular multichannel recordings and closed-loop optogenetic interventions in mice during active exploratory behavior and offline sleep/rest. We observed that during offline DSs, principal cell spiking transiently increased across the DG and CA regions of the hippocampus, nesting offline population-level activity patterns that are distinct from those in SWRs. Further, we report that the cell-to-cell coactivity seen during prior waking experience is reactivated during DSs (as well as SWRs). DS-nested neuronal

activity is relevant to whole-hippocampus population dynamics and memory-guided behavior: closed-loop suppression of DG granule cell spiking selectively during offline DSs, but not SWRs, impairs subsequent flexible memory performance in hippocampal-dependent, multi-object recognition tasks. We propose that DSs constitute a second hippocampus network event that plays a complementary role to that of SWRs by supporting the offline reactivation of diverse population patterns of neuronal coactivity in support of memory-guided behavior.
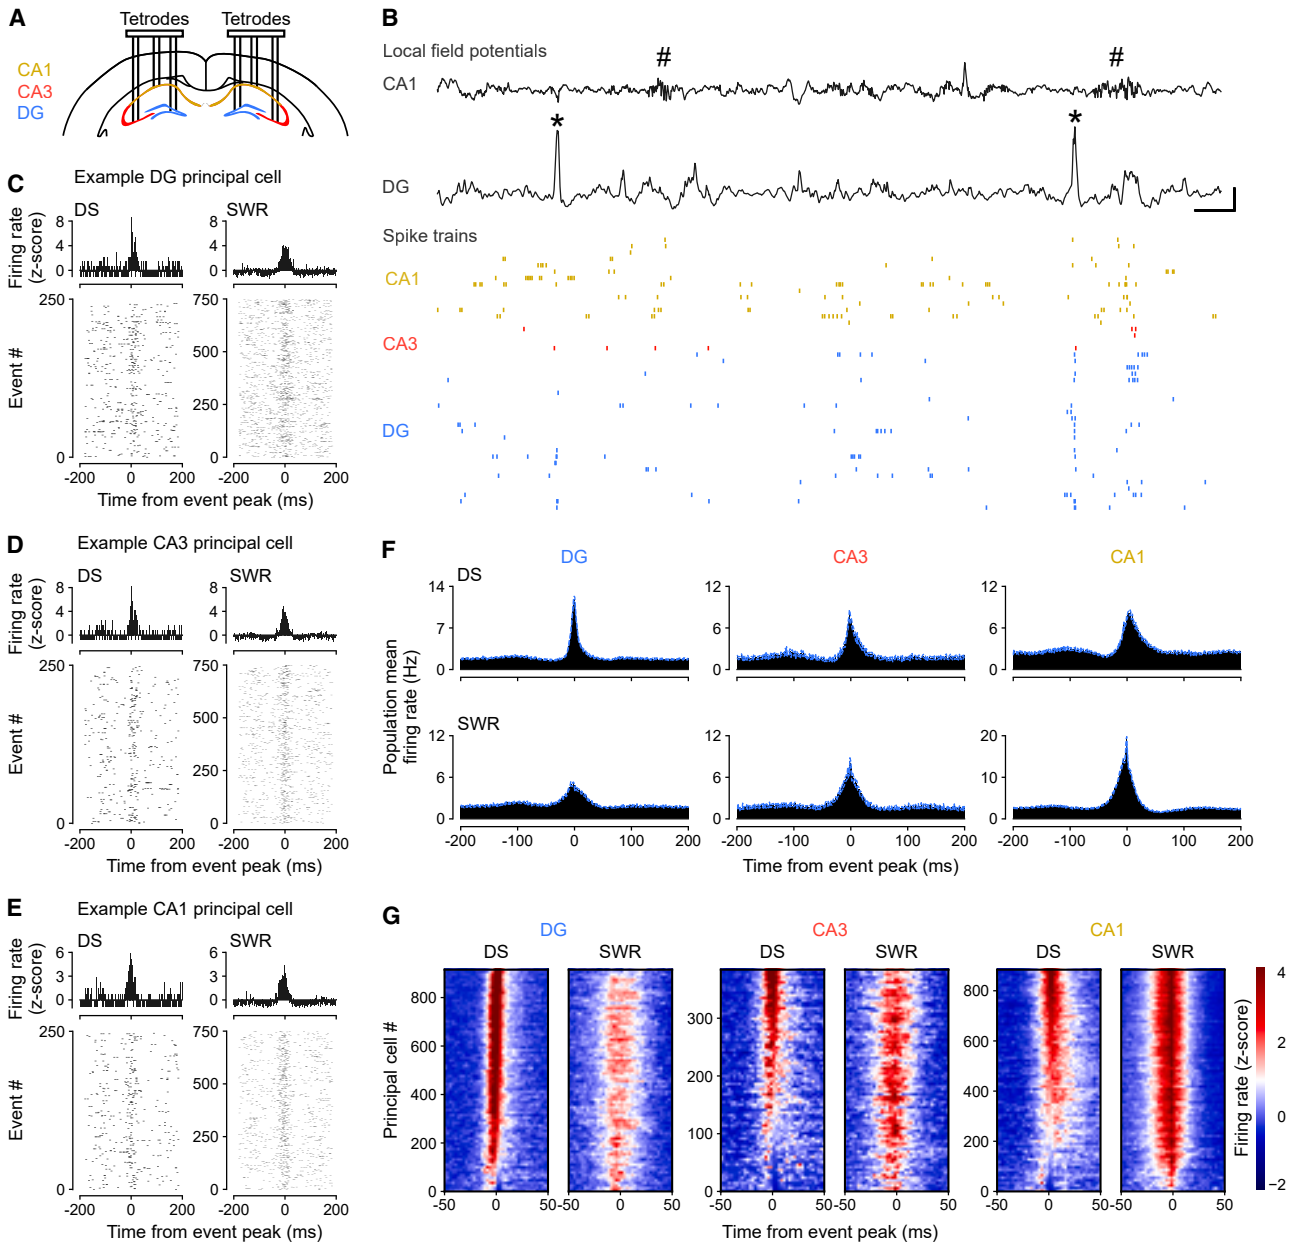
RESULTS

Firing activity of hippocampal neurons synchronizes during DS events

We first used triple-site (DG-CA3-CA1) tetrode recordings to monitor network events in the LFPs and the spike trains of neuronal ensembles from the dorsal hippocampus of mice during sleep/rest (Figures 1A, 1B, and S1A–S1C; $n = 12$ mice). From these LFPs, we detected DSs in DG and SWRs in CA1 to compare the spiking activity of principal cells between these two types of network events. Across mice and recording sessions, DS waveforms were highly consistent (Figures S1D and S1E). Both DG DSs and CA1 SWRs were of short duration (Figure S1F; median [interquartile range, IQR] duration: DSs = 42.4 [40.0–46.4] ms; SWRs = 47.5 [45.2–51.0] ms) and occurred intermittently (median [IQR] occurrence frequency: DSs = 0.40 [0.25–0.58] Hz; SWRs = 0.75 [0.26–1.26] Hz) during behavioral and LFP profiles indicative of sleep/rest (Figures S1G and S1H). These two network events rarely occurred simultaneously, consistent with previous reports,[22,29] with the vast majority of DSs not expressed within ±50 ms of a SWR (median [IQR]: 92.5% [87.9%–95.6%]; DS-SWR co-occurrence frequency: 0.03 [0.02–0.05] Hz; Figures S1I and S1J). We computed the firing responses of individual principal cells ($n = 2,196$ total recorded principal cells; CA1, 887; CA3, 388; DG, 921 cells; Figures S2A–S2F) with respect to the peak of either DSs or SWRs, excluding those temporal windows where both events co-occurred within ±50 ms. With the term "principal cells," we refer to CA pyramidal cells and DG granule cells that constitute the dominant (hence "principal") cell type in the hippocampus, exhibiting lower mean firing rates compared with local fast-spiking inhibitory cells (Figure S2E). In line with previous work, DG principal cells transiently increased their firing activity during DSs,[22,25] and the activity of CA principal cells increased during SWRs (Figures 1C–1G and S2G–S2P).[8,30] We further observed that DG principal cell firing increased during SWRs (Figures 1C, 1F, 1G, and S2G–S2P) and that CA principal cells also increased their firing rate during DSs (Figures 1D–1G and S2G–S2P), which contrasted with some earlier reports that CA principal cell firing is suppressed during DSs.[22,25]

To quantify the magnitude of neuronal activation during DSs and SWRs, we calculated the proportion of DG, CA3, and CA1 principal cells that increased their firing rate beyond a given significance threshold, using the Z scored peri-event time histograms obtained for each of these two network events (Figure 1G). During DSs, the majority of DG (91%), CA3 (56%), and CA1 (61%) principal cells increased their firing rate more than three standard deviations above baseline (Z score > 3, $p < 0.003$;

Figures S2G–S2I). During SWRs, a comparable proportion of principal cells significantly increased their firing activity beyond this threshold (Figures S2G–S2I). Hippocampal CA principal cells exhibited preferential activation during SWRs, whereas DG principal cells exhibited preferential activation during DSs (Figure S2I). DG principal cell population typically fired before CA principal cell populations during DSs (Figures S2J–S2N). While DG and CA principal cells exhibited such a temporal relationship, both DSs and SWRs were associated with an overall transient increase in hippocampal spiking activity (Figures S2O and S2P).
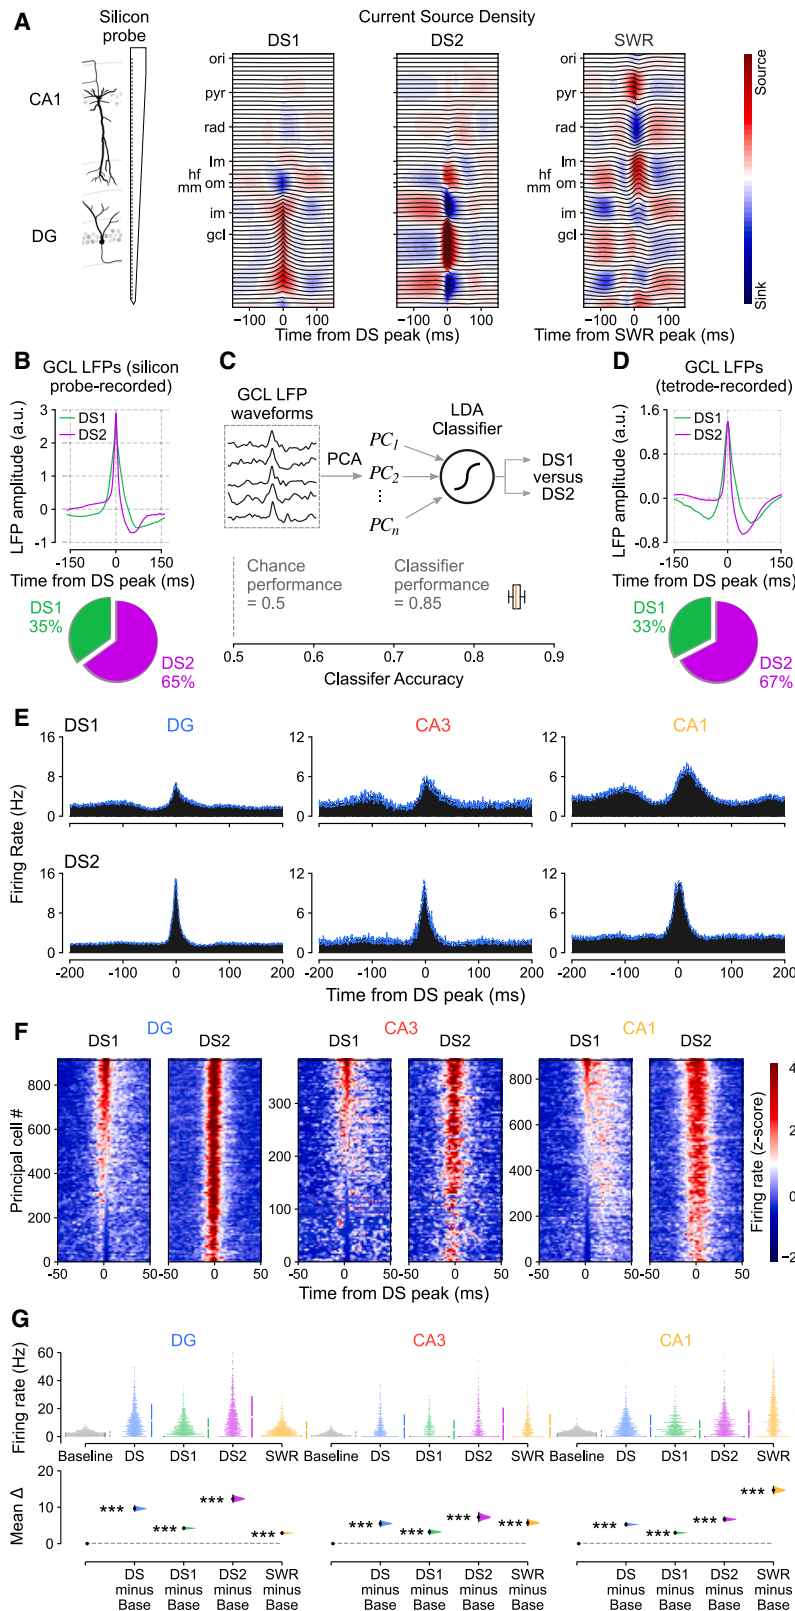
Previous studies identified two types of DS event (DS_1 and DS_2) based on the laminar profile of the transmembrane currents associated with the LFP expression of these network events.[22,24,26] Therefore, we next asked whether principal cell firing responses differed between DS_1 and DS_2. However, localizing sinks and sources of currents across hippocampal layers requires applying current source density (CSD) analysis[31] to the LFPs measured at evenly spaced sites from the CA1 oriens layer to the DG granule cell layer. Such a laminar profile is not accessible with tetrode recordings. To distinguish between DS_1 and DS_2 events, we therefore implanted linear silicon-probes spanning the somato-dendritic axis of CA1 principal cells and reaching the inferior blade of the DG in a separate group of mice ($n = 3$). Having performed silicon-probe recordings during sleep/rest, we applied CSD analysis to these LFPs measured over the radial extent of the hippocampus to identify DS_1 versus DS_2 according to their underlying profile of current sinks and sources (Figures 2A and S3A). These CSD-validated DS_1 and DS_2 events exhibited distinct DG granule cell layer LFP waveforms (Figures 2B and S3B). We then trained a linear discriminant analysis classifier to identify these CSD-validated DS_1 versus DS_2 events using only their DG granule cell layer LFP signal. When tested on the silicon-probe LFP dataset, the classifier achieved over 85% accuracy (Figure 2C). When next applied to the tetrode LFP dataset, the classifier-identified DS_1 and DS_2 events also exhibited distinct granule cell layer LFP waveforms (Figures 2D, S3C, and S3D), which were consistent with those obtained in silicon-probe recordings (Figures 2B and S3B). In both (tetrode and silicon-probe) datasets, DS_2 represented two-thirds of the DS events (median [IQR]: 66% [61%–73%]), thus constituting the dominant type. Leveraging this cross-dataset approach, we found that the firing response of DG and CA principal cells was stronger for DS_2 than DS_1 (Figures 2E–2G and S3E–S3J). A greater proportion of CA principal cells showed firing activity below baseline during DS_1 compared with DS_2 (35% versus 11%; Figures 2F, S3K, and S3L), providing insights into the previously documented DS-suppressed firing in some CA principal cells.[22,25,32] Nevertheless, the average activity of principal cells in DS_1 (and DS_2) was significantly higher than their baseline firing (calculated outside of any DS and SWR events) during sleep/rest (Figures 2G, S3G, and S3H) and their overall mean firing rate calculated over the whole recording day (Figures S3I and S3J). These results show that DS events (both DS_1 and DS_2) constitute transient network states that are qualitatively distinct from the sleep/rest epochs outside these events in terms of their capacity to increase spiking activity of individual principal cells distributed across hippocampal regions.

Figure 1. Dentate spikes recruit principal cell spiking across DG, CA3, and CA1

(A) Triple-ensemble (DG-CA3-CA1) tetrode recording allowed simultaneous monitoring of local field potentials (LFPs) and spiking activities.

(B) Upper: raw wide-band CA1 and DG LFP traces (black) showing sharp-wave ripples (SWRs, hash symbols) in CA1 and dentate spikes (DSs, asterisks) in DG. Scale bars, 100 ms (horizontal), 1.5 mV for DG, and 0.5 mV for CA1 (vertical). Lower: (color-coded) raster plot of spike trains from CA1 (orange), CA3 (red), and DG (blue) principal cells (PCs, one cell per row). Shown is a few second sample of recording for clarity.

(C–E) Spiking responses from single example DG (C), CA3 (D), and CA1 (E) principal cells. Upper: Z scored peri-event time histogram (PETH) during DSs (left) and SWRs (right). Lower: corresponding raster plot showing event-related spiking responses (one event per row).

(F) Group averaged firing rate PETHs for hippocampal PCs during DSs (top) and SWRs (bottom): DG ($n = 921$), CA3 ($n = 388$), and CA1 ($n = 887$) cells from 12 mice. Blue traces: mean ± SEM.

(G) Heatmaps showing Z scored firing rates for the DG, CA3, and CA1 PCs shown in (F). For each heatmap: one cell per row, sorted (top-to-bottom) from the most activated (highest Z score at event peak, 0 ms, red) to the least activated (lowest Z score at event peak, blue) during DSs.

Figure 2. Hippocampal principal cell firing is higher in DS_2 than DS_1

(A) Left: laminar (64-channel) silicon-probe recording allowed simultaneous monitoring of LFPs across hippocampal layers for current source density (CSD) analysis. Right: example (radially organized) mean LFP traces (gray) with superimposed CSD profile (heatmaps) for type 1 (DS_1) and type 2 (DS_2) dentate spikes and SWRs (calculated from 2,231 DS events and 8,693 SWR events in one mouse). Note the distinct CSD profiles reflecting the different transmembrane currents associated with DS_1 versus DS_2 versus SWR events. Hippocampal layers: oriens (*ori*); pyramidale (*pyr*); radiatum (*rad*); lacunosum-moleculare (*lm*); outer (*om*), middle (*mm*), and inner (*im*) moleculare; granulare (*gcl*); and hippocampal fissure (*hf*).

(B) Upper: shown for silicon-probe recorded DS_1 and DS_2 identified from their CSD profiles are example average granule cell layer LFP waveforms triggered by the peak of these events. Lower: in these recordings there was a higher proportion of DS_2 than DS_1 events (n = 15,067 events, 3 mice).

(C) Upper: we applied principal component analysis on the normalized granule cell layer LFP waveforms for all silicon-probe recorded DS events. We then used the principal components explaining 90% of the variance to train a linear discriminant classifier with the true labels (DS_1 versus DS_2) determined by the individual CSD profiles. Lower: the classifier performance (>85%) was significantly above chance level (50%) when tested on silicon-probe recorded LFP waveforms of unlabeled events. We used this classifier to next distinguish DS_1 and DS_2 from tetrode-recorded granule cell layer LFP waveforms (D).
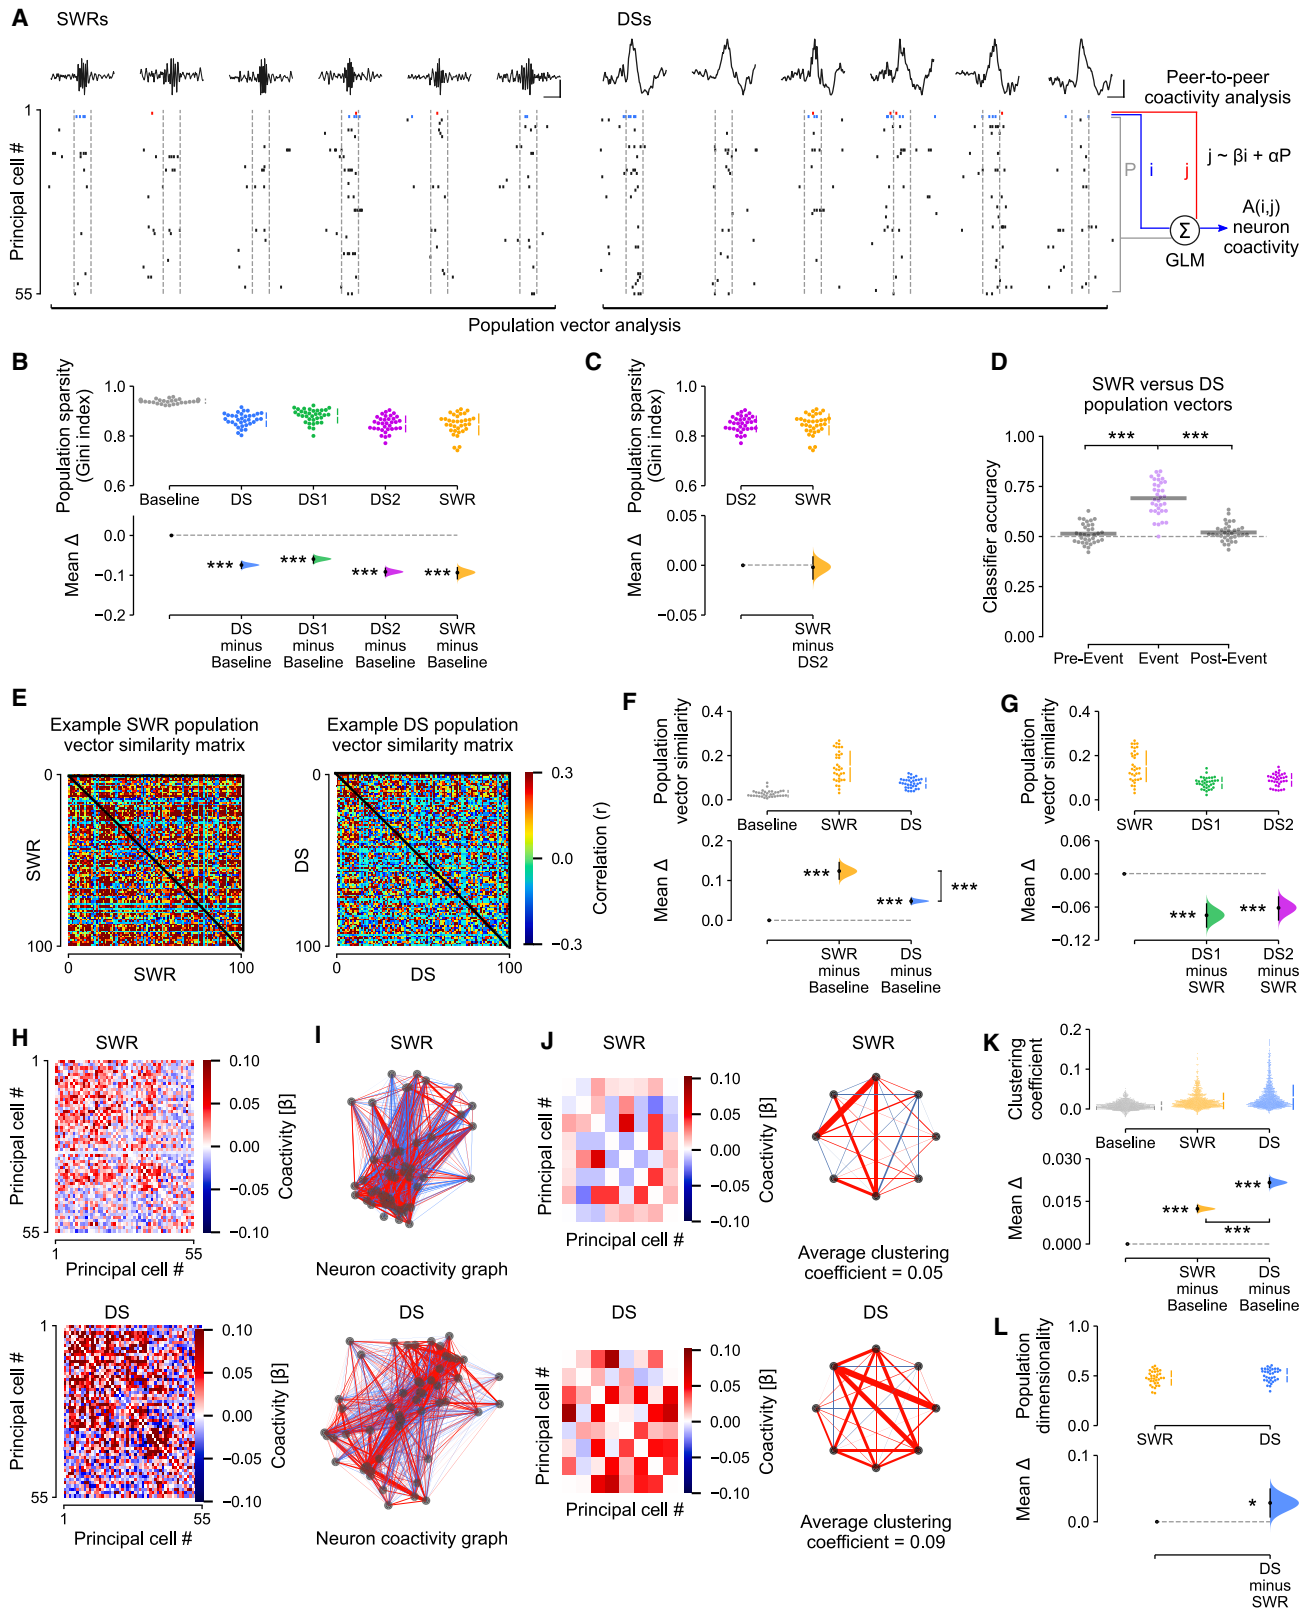
(D) Upper: shown for tetrode-recorded DS events are the average granule cell layer LFP waveforms for DS_1 and DS_2 predicted label obtained from the silicon-probe-based classifier (C). Lower: these recordings also contained a higher proportion of DS_2 than DS_1 events (n = 32,215 events, 12 mice).

(E) Group averaged firing rate PETHs for tetrode-recorded DG, CA3, and CA1 principal cells during DS_1 and DS_2 (as Figures 1F and 1G). Blue traces: mean ± SEM.

(F) Heatmaps showing Z scored firing rates for the DG, CA3, and CA1 cells shown in (E). For each heatmap: one cell per row, sorted (top-to-bottom) from the most activated (highest Z score) to least activated (lowest Z score) during DS_1 peaks.

(G) Estimation plot showing the effect size for the differences in firing rate of DG, CA3, and CA1 principal cells during all DS events, DS_1 and DS_2 events analyzed separately, and SWRs compared with equivalent (50 ms duration matched) baseline windows (base) in which no DSs or SWRs occurred. Upper: raw data points (each point represents one cell), with the gapped lines on the right as mean (gap) ± SD (vertical ends) for each event. Lower: difference (Δ) in firing rate between baseline windows and all DS, DS_1, DS_2, and SWR events computed from 5,000 bootstrapped resamples and with the difference-axis origin (dashed line) aligned to the baseline rate (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). The test statistic is the mean difference, shown on the y axis of the lower plot. p values are from paired permutation tests, event versus baseline, ***p < 0.001.

(E and G) show data from n = 2,196 hippocampal principal cells (DG: n = 921, CA3: n = 388, and CA1: n = 887) from 12 mice.

(legend on next page)

CellPress
OPEN ACCESS

Neuron
Article

DS events nest higher-dimensional patterns of population coactivity

We next investigated how the hippocampus organizes the collective activity of its principal cells both within individual DS events and across events, comparing these population-level patterns to those expressed in SWRs. To proceed, we first considered the neuron-wise vectors formed by the instantaneous spike discharge of principal cells in DSs, SWRs, or duration-matched control windows (without any DSs or SWRs) of the same sleep/rest (Figures 3A and S4A; "population vector analysis"). This was conducted for all sleep/rest epochs (both those recorded before and those after active exploratory behavior). Using the Gini index,[33,34] we noted a marked decrease of the mean population sparsity in the spiking vectors nested in DSs and SWRs compared with baseline periods of duration-matched control windows (Figure 3B), with equivalent population sparsity levels during DS_2 versus SWRs (Figures 3C, S4B, and S4C). A logistic regression classifier trained on a subset of these population firing vectors and iteratively tested on the remaining subset significantly distinguished between DS and SWR events but could not distinguish between their corresponding pre-event nor their post-event control epochs (Figure 3D). Successful classification was also obtained when using only DS_2 events, which matched SWRs in the mean population sparsity per event (Figure S4D). When evaluating the pairwise similarity of DS-nested population vectors versus those of SWR-nested vectors (Figure 3E), we further observed that DSs contained a higher diversity (i.e., lower similarity) of firing vectors compared with those

in SWRs, which in turn were more similar to one another (Figures 3F, 3G, and S4E–S4H).

This difference in population vector similarity suggested that DSs and SWRs differ with respect to their neuronal motifs of transient coactivation. By examining the topological organization of peer-to-peer firing associations, we indeed observed that DS events contain stronger motifs of coactive principal cells than SWRs. For each cell pair (i, j), we trained a generalized linear model to predict the spike discharge of neuron j from that of neuron i while accounting for the activity of the remaining peers (Figure 3A; "peer-to-peer coactivity analysis"). We performed this procedure separately for DS and SWR events, which returned for each type of network event a matrix of β regression weights that represented the coactivity structure of the population (Figure 3H). With these matrices, for both DS and SWR events, we constructed neuronal coactivity graphs (with no self-connections), where each node is a cell and the edge linking any two nodes represents the firing association of that cell pair (Figures 3I and 3J). This revealed that DS-based graphs contained stronger triads of coactive nodes compared with SWR graphs (Figures 3K, S4I, and S4J). This remained the case when directly comparing DS_2 and SWR events (Figure S4K) and when calculating the neuron-wise average coactivity strength (Figure S4L).

These findings showed that while both DS and SWR events synchronize hippocampal principal cells, population coactivity responses to DSs are more diverse. To further assess this, we applied principal component analysis to quantify and compare

Figure 3. The coactivity structure of population spiking differs between DSs and SWRs

(A) Analytical framework: the population-level coactivity structure was analyzed using population vectors of principal cell spiking transiently nested in individual SWRs, DSs, or duration matched (50 ms) baseline control windows. Scale bars show 20 ms and 0.5 mV for SWRs and 1 mV for DSs. For the analyses in (B)–(G), these population firing vectors were then binarized (for each cell: a non-zero spike count gives 1; or else 0). For the analyses in (H)–(K), we calculated the peer-to-peer coactivity, controlling for the overall population activity.
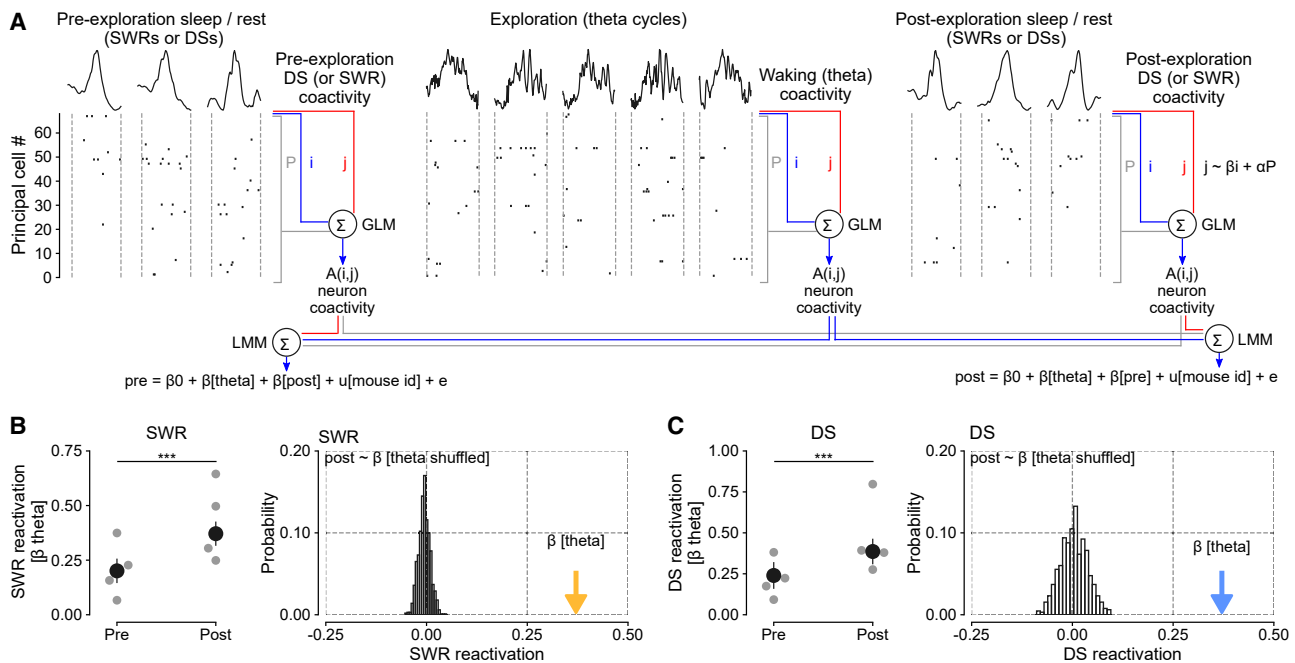
(B and C) Estimation plots showing the effect size for the differences in population sparsity (using the Gini index) during DSs (with DS_1 and DS_2 plotted altogether or separately), SWRs, and compared with equivalent (50 ms duration matched) baseline windows (baseline) in which no DSs or SWRs occurred. Upper: raw data points (each point represents one session with at least 100 of each event type and 20 principal cells), with the gapped lines on the right as mean (gap) ± SD (vertical ends) for each event. Lower: difference (Δ) in sparsity between baseline windows and all DS, DS_1, DS_2, and SWR events computed from 5,000 bootstrapped resamples and with the difference-axis origin (dashed line) aligned to the baseline sparsity (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). (C) as (B) but comparing population sparsity during SWR versus DS_2. Note that DS_2 and SWR events have equivalent sparsity, indicating they engage similar levels of neuronal activity.

(D) A logistic regression classifier trained on population vectors nested in SWR versus DS events or matched duration pre-event and post-event control windows, using a 4-fold cross-validation approach (75% of vectors for training; the remaining 25% for evaluation), significantly discriminated DSs from SWRs but could not discriminate between pre-DS versus pre-SWR and post-DS versus post-SWR vectors. Gray horizontal bars: mean classification accuracy.

(E–G) DS population firing vectors are more diverse than those in SWRs. For each sleep session, we computed the similarity (Pearson correlation coefficient) for each pair of population vectors nested in either DSs, SWRs, or duration-matched baseline windows (baseline). (E) shows example DS and SWR matrices of cross-vector similarity for one session. Cross-population vector similarity was significantly higher in SWRs compared with both DSs and control windows (F), and when compared with DS_1 and DS_2 separately (G).

(H–K) DS and SWR population firing vectors exhibit distinct topology of neuronal coactivity. The coactivity between any two (i, j) neurons was measured using a GLM that quantified their short timescale (50 ms windows centered on DS or SWR peaks) firing relationship while accounting for network-level modulation reported by the remaining principal cells in the population (A). (H) This procedure returned for both DS and SWR events an adjacency matrix of β regression weights that represented the neuron pairwise coactivity structure of the population (example matrix from one session). (I) Visualization of the corresponding matrices representing DS- and SWR-based neuronal coactivity graphs. For clarity, (J) shows an example subset (left) for each adjacency matrices shown in (H), along with its corresponding motifs of neuronal coactivity and average clustering coefficient (right). (K) Note that DS-based graphs contained stronger triads of coactive nodes compared with both SWR graphs and control graphs constructed from duration-matched baseline windows (baseline), as indicated by higher mean clustering coefficients. Each point in the upper plot of (K) represents the mean clustering coefficient for one hippocampal principal cell (n = 1,265 neurons, 8 mice).

(L) The dimensionality of population vector matrices (number of principal components required to explain 90% of the variance) was higher for DSs than SWRs. For (B)–(D), (F), (G), and (L), each data point shows one recording session (n = 34 recording sessions from 8 mice). The test statistic is the mean difference, shown on the y axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). p values are from paired permutation tests, event versus baseline (B, F, and K); event versus event (C, F, K, and L); or event versus pre-event, event versus post-event (D). *p < 0.05, **p < 0.01, ***p < 0.001.

Figure 4. Waking patterns of hippocampal coactivity reactivate in offline DSs

(A) DS and SWR reactivation of waking patterns formed by principal cell theta coactivity. For each cell pair (i,j), we predicted the spike discharge of neuron j from the activity of neuron i while regressing out the activity of the remaining population during pre-exploration sleep, exploration of open-field arenas, and post-exploration sleep (using GLMs as in Figure 3A). We separately applied this procedure for DSs and SWRs in both sleep/rest sessions (offline DS versus offline SWR coactivity) and across theta cycles in the exploration session (waking theta coactivity). This procedure returned a matrix of β regression weights that represented the neurons pairwise coactivity structure of the population in each session. We then used a linear mixed model (LMM) to compare the waking theta coactivity with post-exploration sleep coactivity (in DSs or SWRs) while controlling for pre-exploration sleep coactivity (in DSs or SWRs), and vice versa (reverse model). We included mouse identity as a random factor in each LMM.

(B) SWR reactivation (measured by the β coefficients of the LMM that predicted post-exploration SWR coactivity from waking theta coactivity, controlling for pre-exploration SWR coactivity). Left: the β coefficient for theta coactivity was significantly higher when predicting post-exploration SWR coactivity than with the reverse model (i.e., predicting pre-exploration SWR coactivity from theta coactivity, controlling for post-exploration SWR coactivity). Gray points show the β coefficient for theta coactivity for individual mice. Error bars show ± 95% confidence interval. p value from t-test comparing post versus pre β coefficients: t(7,308) = 10.29; $p < 0.0001$. Right: the histogram shows the random probability distribution of β weights for theta coactivity when cell pair identity was shuffled (i.e., a null distribution based on 1,000 random shuffles; n = 7,310 cell pairs from 4 mice). The colored arrow shows the actual β coefficient for theta coactivity.

(C) DS reactivation exhibited the same pattern of results as SWR reactivation, shown in B. p value from t-test comparing post versus pre β coefficients t(7,308) = 8.84; $p < 0.0001$.

the variance explained by the activity patterns nested in DS_1, DS_2, and SWR. This revealed higher dimensionality of population vectors in DS events compared with SWRs (Figure 3L; Table S1). This was accounted for by DS_2 firing vectors, with those nested in DS_1 requiring fewer principal components to explain equivalent variance and exhibiting lower dimensionality than those in DS_2 and SWRs (Figures S4M–S4P).

Waking theta coactivity patterns reactivate in offline DSs and support flexible memory

The DS-nested motifs of peer-to-peer firing associations could instantiate population patterns of neuronal coactivity undergoing offline reactivation to support memory-guided behavior. Notably, the link between hippocampal SWRs and memory reactivation was initially established through the observation that the neural patterns of joint spiking activity expressed during exploratory behavior are more strongly correlated with those nested in post-exploration sleep/rest SWRs than those in SWRs before waking experience.[8,13] Accordingly, we next

determined whether DSs constitute another hippocampal time frame for offline reactivation of waking coactivity patterns. To proceed, we used our peer-to-peer coactivity analysis (Figure 3A), applying it to DS versus SWR events of sleep/rest before and after exploration of open-field arenas (Figures 4A and S5). Likewise, we obtained the waking patterns of population coactivity in theta cycles during exploration. With these, we computed DS and SWR reactivation by measuring the tendency of the peer-to-peer theta firing associations to reoccur in post-exploration sleep/rest DS (or SWR) events while controlling for prior pre-exploration DS (or SWR) coactivity and mouse identity using a linear mixed model. In line with previous work, offline patterns of SWR coactivity reflected those of theta coactivity significantly more during post-exploration than pre-exploration sleep/rest (Figures 4B, left, S5B, and S5D). This SWR reactivation was significantly higher than that obtained with a null distribution generated from models using randomly shuffled cell pair identities (Figure 4B, right). Importantly, we observed that theta coactivity patterns were also strongly reactivated in

Figure 5. DS- and SWR-informed offline suppressions of DG granule cell activity

(A) Triple-ensemble (DG-CA3-CA1) recording with LFP-informed yellow (561 nm) DG light delivery. Dentate granule cells (DGCs) transduced with ArchT-GFP (DG^Grm2::ArchT). Closed-loop light delivery to suppress DGC spiking immediately upon either DS detection (DS-sync condition) or SWR detection (SWR-sync) or their respective control conditions (DS-delay and SWR-delay, where light delivery was offset by 100 ms after event detection).

(B) ArchT-GFP-expressing DGCs in a DG^Grm2::ArchT mouse. Neuronal nuclei stained with NeuN. Scale bar, 100 μm. Granule cell layer, gcl; molecular layer, mol; pyramidal cell layer, pyr; stratum oriens, ori; radiatum, rad; and lucidum, s.l.

(C and D) Closed-loop feedback transiently silenced DGCs during either DG DS (C; "DS-sync") or CA1 SWR (D; "SWR-sync") events, illustrated with raw data examples. Scale bars, 30 ms (horizontal) and 1.5 mV (vertical).

(E) Raster plots (event-related spiking response; one light pulse per row, upper), and peri-event time histograms (lower) showing photo-silencing of two example DG cells from a DG^Grm2::ArchT mouse in DS-delay and DS-sync.

(F and G) Corresponding quantification of average DGC firing rate (Z score) for DS-delay versus DS-sync (F and G; n = 548 cells in 9 mice). In (F), the orange box shows the laser-on period for DS-sync.

(H–J) As (E)–(G) but showing DGC photo-silencing during SWR-delay and SWR-dync conditions (I and J; n = 181 cells in 3 mice). In (I), the orange box shows the laser-on period for SWR-sync.

For (F) and (I), the traces show mean ± SEM. For (G) and (J), the test statistic is the mean difference, shown on the y axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). p values are from unpaired permutation tests, delay versus sync, ***p < 0.001.

post-exploration DSs (Figures 4C, S5C, and S5E; Table S2). The clustering coefficient was higher in DSs than SWRs and also increased from pre- to post-exploration sleep in both SWRs and DSs (Figures S5F and S5G). Analyses of DS$_1$ and DS$_2$ separately showed evidence for reactivation during both types of post-exploration DS events (Figures S5H and S5I). By applying the same analytical framework to the neuronal ensembles tracked in SWRs and DSs, this result provided evidence for offline DS reactivation of hippocampal waking firing patterns.

The offline reactivation of waking population patterns in sleep/rest DSs (Figure 4C), which contain more diverse and higher-dimensional patterns of neuronal coactivation than those found

in SWRs (Figures 3F and 3L), raised the question of their network contribution to memory-guided behavior. We thus tested whether the offline population response during DSs was necessary to perform tasks that require integrating multiple items in memory to flexibly distinguish between familiar and novel stimuli. To this end, we transduced DG granule cells with the yellow (561 nm) light-driven optogenetic silencer Archaerhodopsin T (ArchT) in Grm2-Cre mice (Figures 5A and 5B). We then implanted these DG^Grm2::ArchT mice for triple-ensemble (DG-CA3-CA1) recordings combined with bilateral optic fibers for DG light delivery. In these experiments, DG light delivery was performed in a closed-loop manner during sleep/rest using the

CellPress
OPEN ACCESS

real-time detection of either DSs (in the DG LFPs) or SWRs (in the CA1 LFPs; Figures 5A, 5C, and 5D; "DS-sync" or "SWR-sync" conditions). We also used a within-subject control paradigm whereby, on different days, light was not synchronized to but instead delivered after each DS or SWR had elapsed ("DS-delay" or "SWR-delay" conditions). DS-synchronized light delivery did not affect the amplitude of the DG or CA1 LFPs, nor CA1 ripple duration, occurrence probability, or power (Figures S6A–S6I). DS-synchronized light delivery significantly reduced spiking activity in DG neurons compared with when laser-onset was DS delayed (Figures 5E–5G) and also reduced spiking activity in CA principal cells (Figures S6J and S6K). Paired analysis of the firing rates of dentate granule cells during DS events without light delivery versus those with light deliver (DS-sync) also showed significantly reduced instantaneous spiking (15.0 ± 0.7 versus 10.4 ± 0.6 Hz). SWR-synchronized light delivery also significantly suppressed DG neuronal spiking (Figures 5H–5J).

We applied these closed-loop light-delivery approaches during interposed sleep/rest sessions in three behavioral tasks (Figures 6 and S6L–S6T). The first, hippocampal-dependent task required mice to recognize previously encountered (familiar) versus novel objects.[35] In this novel-object recognition task, mice repeatedly explored a square-walled arena containing four objects (Figures 6A and 6B). In the first session ("sampling"), mice encountered four distinct novel objects, each one placed beside a wall. On the subsequent sessions ("test"), one of the initially sampled objects was replaced by a different novel object so that the mouse could explore one completely novel object along with the three "familiar" objects from the previous session that day. These exploration (sampling and test) sessions alternated with sleep/rest sessions where mice received DG-targeted light delivery, either synchronized or delayed with respect to either DS or SWR onset, thus yielding four distinct experimental conditions (DS-delay, DS-sync, SWR-delay, and SWR-sync). In each test session n, we measured novelty preference using the proportion of time mice spent investigating the novel versus the familiar objects, thereby probing recognition memory for session $n − 1$. We found that novelty detection was not impaired in test sessions following sleep with DG granule cell suppression in either DS-delayed, SWR-delayed, or SWR-synchronized conditions: mice subsequently expressed a stronger preference for novel over familiar objects under these three conditions (Figures 6C and 6D). This novel-object preference was similar to that observed in control mice without any optogenetic intervention (Figure S6L). However, novel-object preference was absent in test sessions following DS-synchronized suppression (Figures 6C and 6D). The total object exploration time, number of laser pulses, and number of SWRs did not differ between the DS-synchronized and DS-delayed conditions (Figures S6M–S6O).

We also tested the offline DS-informed suppression of DG granule cells after tone fear conditioning as a non-hippocampal-dependent task.[36] Mice were trained with five tone-shock pairings, and following DS-synchronized or DS-delayed suppression, we evaluated fear memory by measuring freezing behavior during a recall session in which tones were played but no shocks were given. Compared with baseline freezing (measured during the first tone of training, before any shocks were given), mice exhibited higher (and equivalent) freezing levels during recall regardless of whether they had received DS-synchronized or DS-delayed suppression (Figures S6P–S6S).
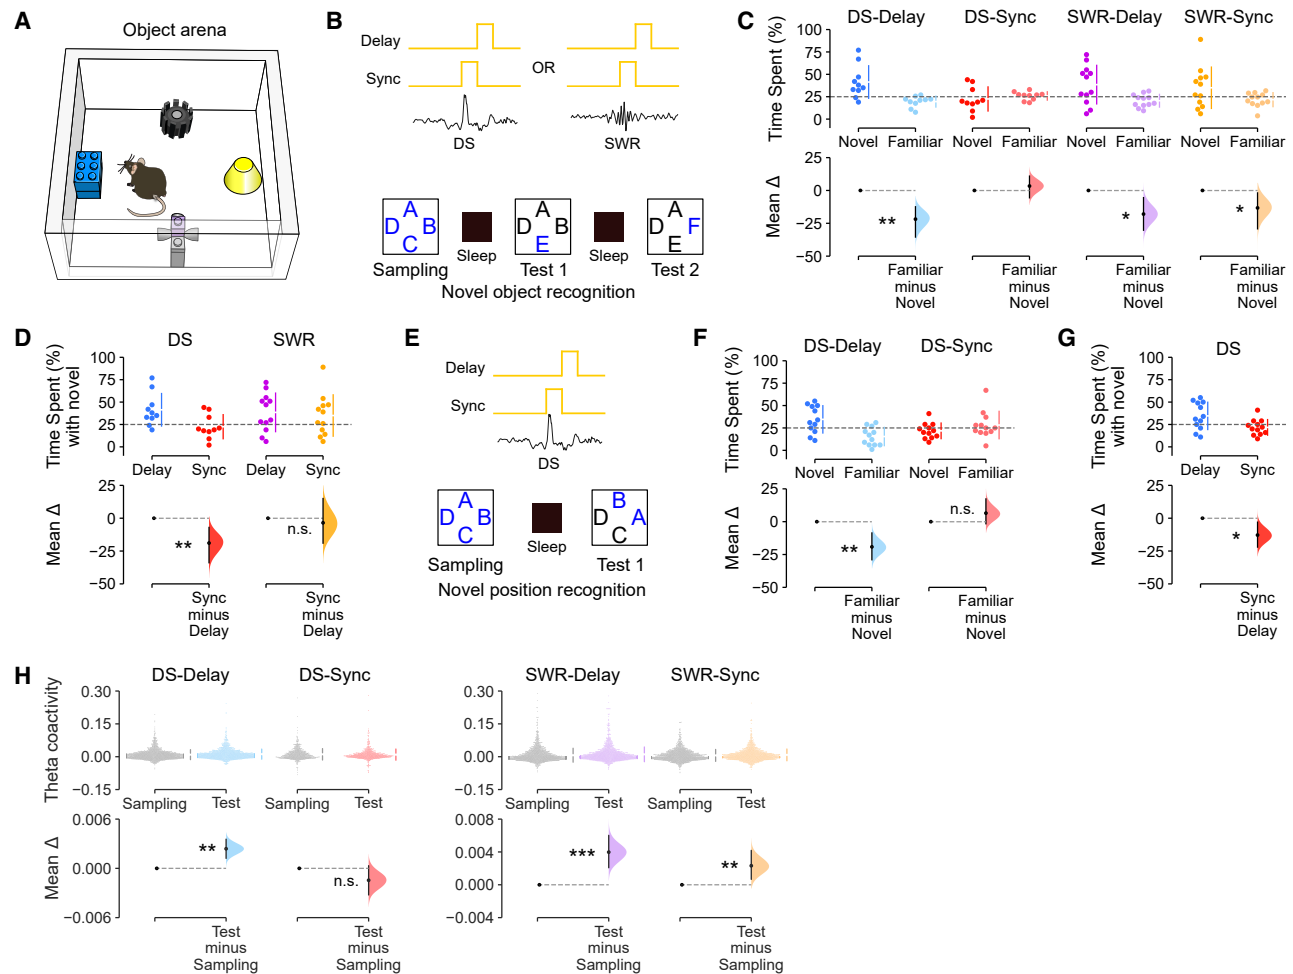
We finally tested whether DS-synchronized suppression affected performance in a novel-position recognition task that is reportedly more sensitive to DG than CA1 lesions, whereas novel-object recognition requires both DG and CA1.[37] This novel-position task is similar to the novel-object recognition task in that mice explore four novel objects during the sampling phase (Figure 6E). However, rather than introducing a new object in the test phase, the locations of two of the initially sampled objects are swapped for the subsequent session, leaving the other two objects in their original locations (Figure 6E). We found that mice preferentially explored the novel-positioned objects following DS-delayed suppression of DG granule cells but showed no such preference following DS-synchronized suppression (Figures 6F, 6G, and S6T).

Recent work has identified that the continued integration of new items in memory is associated with increased neuronal coactivity patterns nested in hippocampal theta oscillations.[38] In line with this, we found that the preserved object recognition memory observed after offline DG cell suppression in the DS-delayed, SWR-delayed, and SWR-synchronized conditions was accompanied by stronger theta coactivity (Figure 6H). This was not the case following DS-synchronized suppression (Figure 6H), indicating that DS silencing disrupts the integration of recently experienced information. Collectively, these results show that the hippocampal population response to offline DS events is required for flexible, memory-based recognition of previously encountered items and associated network gain in theta coactivity.

DISCUSSION

Our findings establish that offline DSs activate neurons across the DG and CA regions. DSs are therefore a second hippocampal network event that hosts short timescale coactivation forming population-level neural patterns, like the well-established SWRs. However, the activity structure and neuronal content are distinct in DSs. Notably, we found that DSs nest stronger motifs of coactive neurons, yielding population patterns of higher diversity and dimensionality compared with those in SWRs. Like SWRs, DSs reactivate hippocampal population patterns expressed in prior waking experience. This offline reactivation is behaviorally significant: closed-loop suppression of DG granule cell spiking selectively during offline DS events is sufficient to disrupt downstream CA principal cell activity and impair flexible recognition memory for previously encountered items, as well as the associated network gain in theta-nested neuronal coactivity. Collectively, these findings identify a core contribution for DSs to hippocampal patterns of population activity and memory-guided behavior.

We started this investigation by observing that DSs increase spiking activity in principal cells across the DG, CA3, and CA1 regions of the hippocampus. This finding is consistent with previous reports of DS-evoked spiking activity in DG granule cells but contrasts with some earlier reports of DS-suppressed CA pyramidal cell spiking.[22,24–26] Notably, Bragin and colleagues

Figure 6. Offline suppression of DS activity impairs flexible recognition memory

(A) Behavioral arena used for the recognition memory tasks.

(B–D) Offline DS events are required for novel-object recognition memory. (B) Task layout. During sleep sessions (interposed between novel-object exploration sessions), closed-loop optogenetic suppression of DG cells in DG^{Grm2}::ArchT mice was achieved using real-time monitoring of either DG or CA1 LFPs to actuate either DS synchronized (DS-sync) or delayed (DS-delay), SWR synchronized (SWR-sync), or delayed (SWR-delay) DG light delivery. Letters depict object locations in the task arena (A), with novel objects in blue. (C) Estimation plot showing the percentage of time spent by mice with the novel versus the familiar objects in each of the four LFP-informed closed-loop conditions. Upper: each data point represents the percentage time spent with the novel object versus the mean percentage time spent with the three familiar objects; chance performance is shown by the dashed line. Lower: mean difference between novel and familiar object exploration time. (D) as (C), but directly comparing novel-object preference in the delay versus sync conditions for DS and SWR events. Mice in the DS-delay, SWR-delay, and SWR-sync conditions, but not the DS-sync condition, exhibited a significant preference for novel over familiar objects (DS-delay and DS-sync: n = 10 sessions, in 3 mice; SWR-delay and SWR-sync: n = 12 sessions in 3 mice).

(E–G) Likewise, offline DS events are required for novel-position recognition memory. (E) Task layout. Letters depict object locations, with novel positions in blue. (F) Estimation plot showing the percentage of time spent by DG^{Grm2}::ArchT mice with the novel versus the familiar object locations following sleep sessions with DS-sync or DS-delay suppression of DG cells. Upper: each data point represents the percentage time spent with objects in novel locations versus objects in familiar locations; chance performance is shown by the dashed line. Lower: mean difference between novel and familiar location exploration times. (G) As (F) but directly comparing novel location preference in DS-delay versus DS-sync. Mice in the DS-delay but not the DS-sync condition exhibited a significant preference for objects in novel over familiar locations (n = 12 novel versus n = 12 familiar locations, 6 sessions, in 4 mice).

(H) In the object recognition task, the theta peer-to-peer coactivity increased from the initial object sampling to the memory test following offline DG cell suppression in the DS-delay, SWR-delay, and SWR-sync conditions, but this was not the case in the DS-sync condition (where mice exhibited no novel-object preference). Paired estimation plot showing theta coactivity during sampling versus test. Upper: each point represents a beta coefficient for the theta-nested peer-to-peer coactivity between pairs of hippocampal principal cells (n = 1,537, n = 678, n = 1,719, and n = 1,482 cell pairs, respectively, in 6 mice). Lower: black dot, mean difference between sampling and test sessions; black ticks, 95% confidence interval.

For (C) and (D) and (F)–(H), the test statistic is the mean difference, shown on the y axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). p values are from paired permutation tests, familiar versus novel (C and F); delay versus sync (D and G); or test versus sampling (H), *p < 0.05, **p < 0.01, ***p < 0.001.

reported suppressed spiking in 3/14 CA3 principal cells and suppressed CA1 multi-unit activity in 2/10 rats, showing some, rather than consistent, CA suppression. In addition, Penttonen et al. reported suppressed CA1 multi-unit activity and hyperpolarization of 4 intracellularly recorded CA1 neurons during DSs in anesthetized rats. However, DS rates are ~10-fold lower and have smaller amplitude during anesthesia compared with DSs observed during natural sleep/rest.[25] Other studies reported increased CA1 multi-unit activity[26] and increased CA3 single-unit spiking during DS events[24] in non-anesthetized mice. Here, we report a variety of firing responses across individual CA neurons, ranging from strong activation to suppression during DS events (Figures 1G and 2F). However, our systematic study (including > 3,500 principal cells; Table S3) shows that DSs do indeed drive increased mean population spiking activity in both DG and CA principal cells (Figure 2G).

Previous studies distinguished between two types of DS events (DS_1 and DS_2) based on the laminar profile of their transmembrane currents.[22,24,26] Here, we found that DS_2 are more effective than DS_1 events at recruiting hippocampal principal cells, with higher spike rates per cell and more coactive cells per event. This result is consistent with a recent report that DS_2 but not DS_1 events reliably increase spiking in DG and CA3 principal cells.[24] The same report saw only a slight increase in CA1 spiking during DS_2 and no effect of DS_1 events on either CA1 or CA3 principal cells. They also found that DG principal cell spiking was suppressed during DS_1, in contrast to our study and previous reports.[22] Here we found that both DS_1 and DS_2 events evoked significantly increased spiking activity in DG, CA3, and CA1 principal cells, but again would emphasize the diversity of CA cell responses, especially during DS_1 events (Figure 2F). The observed differences between studies might reflect differences between DS events in head-fixed awake mice versus those in sleep/rest.[22,29] They could also indicate a moment-to-moment diversity across individual DSs, similar to that highlighted for individual SWRs[39] and theta cycles.[40]

In this study, we directly compared patterns of population spiking activity in DSs versus SWRs. We found that population firing differed across these two types of network events, allowing a classifier to distinguish DSs versus SWRs based on their instantaneous vectors of principal cell spiking (Figure 3D). We also observed that population patterns in DSs are overall more diverse (less correlated) than those in SWRs, showing stronger triads of coactive neurons and higher dimensionality (Figures 3 and S4). DS_1 and DS_2 population patterns yet showed distinct trends with respect to SWRs: DS_1 firing vectors exhibited less diversity (i.e., required fewer principal components to explain most of the variance; Figure S4O) and lower dimensionality (Figure S4P) than those in SWRs; DS_2 firing vectors showed the opposite trend. Previous work has reported that DS_1 and DS_2 events relate to distinct entorhinal cortex inputs, with DS_1 relying more on the lateral entorhinal cortex while DS_2 events rely more on the medial entorhinal cortex.[24] This suggests a possible division of mnemonic labor where DS_1 population patterns would favor non-spatial information streams while DS_2 might favor spatial information.[41] We also found that waking patterns of neuronal coactivity nested in theta oscillations reactivate in DSs of post-exploration sleep/rest (Figures 4 and S5). Notably,

the distributions of coactivity values in DSs indicate both positive and negative firing associations (Figure S5E). The coexistence of correlated and anti-correlated spiking activities in DSs could reflect a Hebbian learning rule as reported in SWRs,[42] whereby positive and negative changes in hippocampal principal cell firing associations can shape offline DS reactivation as a function of recent waking experience. These findings provide important evidence for offline reactivation of hippocampal waking firing patterns outside of SWRs, stimulating new avenues for future work to explore.

To determine whether spiking activity observed during DS events was required for subsequent memory-guided behavior, we deployed a closed-loop optogenetic feedback approach to suppress DG granule cell activity selectively during DS events (Figures 5 and S6). Real-time inhibition of the DG in Grm2-Cre mice did not yield a complete suppression of the spiking activity in dentate granule cells. This also did not alter the magnitude of DSs, which powerful expression in the DG LFPs could reflect the high cellular density of the granule cell layer and its strong neural inputs. This DS-synchronized suppression of DG principal cells reduced concomitant spiking activity in CA principal cells but did not affect the expression of SWRs in CA1. When applied in sleep/rest following object-location exploration, this DS-synchronized neural suppression impaired subsequent memory performance in both novel-object and novel-position recognition tasks. Although ours is the first study to leverage a closed-loop optogenetic approach, these findings are consistent with previous behavioral studies using electrical stimulation to disrupt hippocampal activity during DSs.[32,43,44] While both approaches provide strong evidence for a central contribution of DS events in memory-guided behavior, it is important to recognize that optogenetic and electrical interventions do not recapitulate natural hippocampal activity patterns. Moreover, DS_2 represents the dominant type of DS event that exhibits, in comparison with DS_1, stronger firing rate increase of DG, CA3, and CA1 principal cells (Figure 2G), neuronal recruitment (Figures S3E and S3F), coactivity (Figure S4L), and lower dimensionality (Figures S4O and S4P). The major effect of DS activity on memory may be associated with DS_2, a hypothesis that future technological (closed-loop controller) development for differential manipulation of DS_1 versus DS_2 would be able to test. We also found that increased theta coactivity was associated with recognition memory and that this network gain in theta coactivity was absent following DS-synchronized neural suppression (Figure 6H). This finding is consistent with recent work showing that continual integration of new memory items across behavioral experiences increases neuronal coactivity.[38] Altogether, these results support the idea that neuronal activity during DS events plays an important role in subsequent memory-guided behavior, as SWRs do.

Why does the hippocampus use more than one offline network mechanism to support memory? DSs and SWRs are driven by distinct neural circuits. SWRs depend on excitatory inputs from CA3 to the CA1 stratum radiatum, generating high-frequency ripples in the CA1 pyramidal layer.[7,45–47] DSs are non-oscillatory events associated with excitatory inputs from the entorhinal cortex to the DG molecular layers.[22,24,29] Notably, entorhinal cortex lesions eliminate DSs but increase SWR incidence.[22] Our

CellPress
OPEN ACCESS

Neuron
Article

structural analysis of DS versus SWR population patterns raises the intriguing possibility that SWRs may be more suited for lower-dimensional network coactivity serving robust information flow, whereas DSs may promote higher-dimensional activity, allowing diverse mnemonic patterns to coexist offline and support flexible, pattern separation for subsequent behavior. Collectively, these findings open important new avenues for future work to explore the interplay between DS versus SWR events as two distinct timeframes for the hippocampus to optimize offline computations serving memory-guided behavior.

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the lead contact, David Dupret (david. dupret@bndu.ox.ac.uk).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- The electrophysiology dataset reported in this study is being used in on-going projects and can be accessed under a data transfer agreement. We welcome enquiries for sharing it; please contact david.dupret@bndu.ox.ac.uk.
- This paper does not report original code.
- Any additional information required to reanalyze the data reported in this work paper is available from the lead contact upon request.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.neuron.2024.08.022.

REFERENCES

1. Maquet, P. (2001). The Role of Sleep in Learning and Memory. Science 294, 1048–1052. https://doi.org/10.1126/science.1062856.

2. Walker, M.P., and Stickgold, R. (2006). Sleep, Memory, and Plasticity. Annu. Rev. Psychol. 57, 139–166. https://doi.org/10.1146/annurev.psych.56.091103.070307.

3. Klinzing, J.G., Niethard, N., and Born, J. (2019). Mechanisms of systems memory consolidation during sleep. Nat. Neurosci. 22, 1598–1610. https://doi.org/10.1038/s41593-019-0467-3.

4. Brodt, S., Inostroza, M., Niethard, N., and Born, J. (2023). Sleep—A brain-state serving systems memory consolidation. Neuron 111, 1050–1075. https://doi.org/10.1016/j.neuron.2023.03.005.

5. Buzsáki, G. (2010). Neural Syntax: Cell Assemblies, Synapsembles, and Readers. Neuron 68, 362–385. https://doi.org/10.1016/j.neuron.2010.09.023.

6. Girardeau, G., and Lopes-dos-Santos, V. (2021). Brain neural patterns and the memory function of sleep. Science 374, 560–564. https://doi.org/10.1126/science.abi8370.

7. Buzsáki, G. (1986). Hippocampal sharp waves: Their origin and significance. Brain Res. 398, 242–252. https://doi.org/10.1016/0006-8993(86)91483-6.

8. Buzsáki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. Hippocampus 25, 1073–1188. https://doi.org/10.1002/hipo.22488.

9. Axmacher, N., Elger, C.E., and Fell, J. (2008). Ripples in the medial temporal lobe are relevant for human memory consolidation. Brain 131, 1806–1817. https://doi.org/10.1093/brain/awn103.

10. Joo, H.R., and Frank, L.M. (2018). The hippocampal sharp wave–ripple in memory retrieval for immediate use and consolidation. Nat. Rev. Neurosci. 19, 744–757. https://doi.org/10.1038/s41583-018-0077-1.

11. Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., and Buzsáki, G. (1999). Fast network oscillations in the hippocampal CA1 region of the behaving rat. J. Neurosci. *19*, RC20. https://doi.org/10.1523/JNEUROSCI.19-16-j0001.1999.

12. Csicsvari, J., Hirase, H., Mamiya, A., and Buzsáki, G. (2000). Ensemble Patterns of Hippocampal CA3-CA1 Neurons during Sharp Wave–Associated Population Events. Neuron *28*, 585–594. https://doi.org/10.1016/S0896-6273(00)00135-5.

13. Wilson, M.A., and McNaughton, B.L. (1994). Reactivation of Hippocampal Ensemble Memories During Sleep. Science *265*, 676–679. https://doi.org/10.1126/science.8036517.

14. Girardeau, G., Benchenane, K., Wiener, S.I., Buzsáki, G., and Zugaro, M.B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. Nat. Neurosci. *12*, 1222–1223. https://doi.org/10.1038/nn.2384.

15. Ego-Stengel, V., and Wilson, M.A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. Hippocampus *20*, 1–10. https://doi.org/10.1002/hipo.20707.

16. van de Ven, G.M., Trouche, S., McNamara, C.G., Allen, K., and Dupret, D. (2016). Hippocampal Offline Reactivation Consolidates Recently Formed Cell Assembly Patterns during Sharp Wave-Ripples. Neuron *92*, 968–974. https://doi.org/10.1016/j.neuron.2016.10.020.

17. Maingret, N., Girardeau, G., Todorova, R., Goutierre, M., and Zugaro, M. (2016). Hippocampo-cortical coupling mediates memory consolidation during sleep. Nat. Neurosci. *19*, 959–964. https://doi.org/10.1038/nn.4304.

18. Fernández-Ruiz, A., Oliva, A., Fermino de Oliveira, E., Rocha-Almeida, F., Tingley, D., and Buzsáki, G. (2019). Long-duration hippocampal sharp wave ripples improve memory. Science *364*, 1082–1086. https://doi.org/10.1126/science.aax0758.

19. Treves, A., and Rolls, E.T. (1994). Computational analysis of the role of the hippocampus in memory. Hippocampus *4*, 374–391. https://doi.org/10.1002/hipo.450040319.

20. McHugh, T.J., Jones, M.W., Quinn, J.J., Balthasar, N., Coppari, R., Elmquist, J.K., Lowell, B.B., Fanselow, M.S., Wilson, M.A., and Tonegawa, S. (2007). Dentate Gyrus NMDA Receptors Mediate Rapid Pattern Separation in the Hippocampal Network. Science *317*, 94–99. https://doi.org/10.1126/science.1140263.

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

22. Bragin, A., Jandó, G., Nádasdy, Z., van Landeghem, M., and Buzsáki, G. (1995). Dentate EEG spikes and associated interneuronal population bursts in the hippocampal hilar region of the rat. J. Neurophysiol. *73*, 1691–1705. https://doi.org/10.1152/jn.1995.73.4.1691.

23. Senzai, Y., and Buzsáki, G. (2017). Physiological Properties and Behavioral Correlates of Hippocampal Granule Cells and Mossy Cells. Neuron *93*, 691–704.e5. https://doi.org/10.1016/j.neuron.2016.12.011.

24. Dvorak, D., Chung, A., Park, E.H., and Fenton, A.A. (2021). Dentate spikes and external control of hippocampal function. Cell Rep. *36*, 109497. https://doi.org/10.1016/j.celrep.2021.109497.

25. Penttonen, M., Kamondi, A., Sik, A., Acsády, L., and Buzsáki, G. (1997). Feed-forward and feed-back activation of the dentate gyrus in vivo during dentate spikes and sharp wave bursts. Hippocampus *7*, 437–450. https://doi.org/10.1002/(SICI)1098-1063(1997)7:4<437::AID-HIPO9>3.0.CO;2-F.

26. Meier, K., Merseburg, A., Isbrandt, D., Marguet, S.L., and Morellini, F. (2020). Dentate Gyrus Sharp Waves, a Local Field Potential Correlate of Learning in the Dentate Gyrus of Mice. J. Neurosci. *40*, 7105–7118. https://doi.org/10.1523/JNEUROSCI.2275-19.2020.

27. Sanchez-Aguilera, A., Wheeler, D.W., Jurado-Parras, T., Valero, M., Nokia, M.S., Cid, E., Fernandez-Lamo, I., Sutton, N., García-Rincón, D., de la Prida, L.M., and Ascoli, G.A. (2021). An update to Hippocampome.org

28. Farrell, J.S., Hwaun, E., Dudok, B., and Soltesz, I. (2024). Neural and behavioural state switching during hippocampal dentate spikes. Nature *628*, 590–595. https://doi.org/10.1038/s41586-024-07192-8.

29. Headley, D.B., Kanta, V., and Paré, D. (2017). Intra- and interregional cortical interactions related to sharp-wave ripples and dentate spikes. J. Neurophysiol. *117*, 556–565. https://doi.org/10.1152/jn.00644.2016.

30. Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., and Buzsáki, G. (1999). Oscillatory Coupling of Hippocampal Pyramidal Cells and Interneurons in the Behaving Rat. J. Neurosci. *19*, 274–287. https://doi.org/10.1523/JNEUROSCI.19-01-00274.1999.

31. Pettersen, K.H., Lindén, H., Dale, A.M., and Einevoll, G.T. (2012). Extracellular spikes and CSD. In Handbook of Neural Activity Measurement (Cambridge University Press), pp. 92–135.

32. Nokia, M.S., and Penttonen, M. (2022). Rhythmic Memory Consolidation in the Hippocampus. Front. Neural Circuits *16*, 885684. https://doi.org/10.3389/fncir.2022.885684.

33. Gini, C. (1921). Measurement of Inequality of Incomes. Econ. J. *31*, 124–126. https://doi.org/10.2307/2223319.

34. Hurley, N., and Rickard, S. (2009). Comparing Measures of Sparsity. IEEE Trans. Inf. Theor. *55*, 4723–4741. https://doi.org/10.1109/TIT.2009.2027527.

35. McHugh, S.B., Lopes-dos-Santos, V., Gava, G.P., Hartwich, K., Tam, S.K.E., Bannerman, D.M., and Dupret, D. (2022). Adult-born dentate granule cells promote hippocampal population sparsity. Nat. Neurosci. *25*, 1481–1491. https://doi.org/10.1038/s41593-022-01176-5.

36. Phillips, R.G., and LeDoux, J.E. (1992). Differential contribution of amygdala and hippocampus to cued and contextual fear conditioning. Behav. Neurosci. *106*, 274–285. https://doi.org/10.1037/0735-7044.106.2.274.

37. Lee, I., Hunsaker, M.R., and Kesner, R.P. (2005). The Role of Hippocampal Subregions in Detecting Spatial Novelty. Behav. Neurosci. *119*, 145–153. https://doi.org/10.1037/0735-7044.119.1.145.

38. Gava, G.P., McHugh, S.B., Lefèvre, L., Lopes-dos-Santos, V., Trouche, S., El-Gaby, M., Schultz, S.R., and Dupret, D. (2021). Integrating new memories into the hippocampal network activity space. Nat. Neurosci. *24*, 326–330. https://doi.org/10.1038/s41593-021-00804-w.

39. Navas-Olive, A., Amaducci, R., Jurado-Parras, M.-T., Sebastian, E.R., and de la Prida, L.M. (2022). Deep learning-based feature extraction for prediction and interpretation of sharp-wave ripples in the rodent hippocampus. eLife *11*, e77772. https://doi.org/10.7554/eLife.77772.

40. Lopes-dos-Santos, V., van de Ven, G.M., Morley, A., Trouche, S., Campo-Urriza, N., and Dupret, D. (2018). Parsing Hippocampal Theta Oscillations by Nested Spectral Components during Spatial Exploration and Memory-Guided Behavior. Neuron *100*, 940–952.e7. https://doi.org/10.1016/j.neuron.2018.09.031.

41. Knierim, J.J., Neunuebel, J.P., and Deshmukh, S.S. (2014). Functional correlates of the lateral and medial entorhinal cortex: objects, path integration and local–global reference frames. Philos. Trans. R. Soc. Lond. B Biol. Sci. *369*, 20130369. https://doi.org/10.1098/rstb.2013.0369.

42. O'Neill, J., Senior, T.J., Allen, K., Huxter, J.R., and Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. Nat. Neurosci. *11*, 209–215. https://doi.org/10.1038/nn2037.

43. Nokia, M.S., Gureviciene, I., Waselius, T., Tanila, H., and Penttonen, M. (2017). Hippocampal electrical stimulation disrupts associative learning when targeted at dentate spikes. J. Physiol. *595*, 4961–4971. https://doi.org/10.1113/JP274023.

44. Lensu, S., Waselius, T., Penttonen, M., and Nokia, M.S. (2019). Dentate spikes and learning: disrupting hippocampal function during memory consolidation can improve pattern separation. J. Neurophysiol. *121*, 131–139. https://doi.org/10.1152/jn.00696.2018.

45. Nakashiba, T., Buhl, D.L., McHugh, T.J., and Tonegawa, S. (2009). Hippocampal CA3 Output Is Crucial for Ripple-Associated Reactivation

by integrating single-cell phenotypes with circuit function in vivo. PLoS Biol. *19*, e3001213. https://doi.org/10.1371/journal.pbio.3001213.

and Consolidation of Memory. Neuron *62*, 781–787. https://doi.org/10.1016/j.neuron.2009.05.013.

46. Sullivan, D., Csicsvari, J., Mizuseki, K., Montgomery, S., Diba, K., and Buzsáki, G. (2011). Relationships between Hippocampal Sharp Waves, Ripples, and Fast Gamma Oscillation: Influence of Dentate and Entorhinal Cortical Activity. J. Neurosci. *31*, 8605–8616. https://doi.org/10.1523/JNEUROSCI.0294-11.2011.

47. Davoudi, H., and Foster, D.J. (2019). Acute silencing of hippocampal CA3 reveals a dominant role in place field responses. Nat. Neurosci. *22*, 337–342. https://doi.org/10.1038/s41593-018-0321-z.

48. Quinn, A.J., Lopes-dos-Santos, V., Dupret, D., Nobre, A.C., and Woolrich, M.W. (2021). EMD: Empirical Mode Decomposition and Hilbert-Huang Spectral Analyses in Python. J. Open Source Software *6*, 2977. https://doi.org/10.21105/joss.02977.

49. Magland, J.F., Jun, J.J., Lovero, E., Morley, A.J., Hurwitz, C.L., Buccino, A.P., Garcia, S., and Barnett, A.H. (2020). SpikeForest, reproducible web-facing ground-truth validation of automated neural spike sorters. Elife. *9*, e55167. https://doi.org/10.7554/eLife.55167.

50. Pachitariu, M., Steinmetz, N.A., Kadir, S.N., Carandini, M., and Harris, K.D. (2016). Fast and accurate spike sorting of high-channel count probes with KiloSort. In Adv. Neural Inf. Process. Syst., *29*, D.D. Lee, M. Sugiyama, U.V. Luxburg, I. Guyon, and R. Garnett, eds. (Curran Associates, Inc.), pp. 4448–4456.

51. Han, X., Chow, B. Y., Zhou, H., Klapoetke, N. C., Chuong, A., Rajimehr, R., Yang, A., Baratta, M. V., Winkle, J., Desimone, R., and Boyden, E. S. (2011). A high-light sensitivity optical neural silencer: development and application to optogenetic control of non-human primate cortex. Front. Syst. Neurosci. *5*, 18. https://doi.org/10.3389/fnsys.2011.00018.

52. Lopes-dos-Santos, V., Brizee, D., and Dupret, D. (2023). Spatio-temporal organization of network activity patterns in the hippocampus. Preprint at bioRxiv. https://doi.org/10.1101/2023.10.17.562689.

53. Skaggs, W.E., McNaughton, B.L., and Gothard, K.M. (1993). An Information-Theoretic Approach to Deciphering the Hippocampal Code. In Advances in Neural Information Processing Systems, *5*, S.J. Hanson, J.D. Cowan, and C.L. Giles, eds. (Morgan-Kaufmann), pp. 1030–1037.

54. Mitzdorf, U. (1985). Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. Physiol. Rev. *65*, 37–100. https://doi.org/10.1152/physrev.1985.65.1.37.

55. Guest, O., and Love, B.C. (2017). What the success of brain imaging implies about the neural code. eLife *6*, e21397. https://doi.org/10.7554/eLife.21397.

56. Onnela, J.-P., Saramäki, J., Kertész, J., and Kaski, K. (2005). Intensity and coherence of motifs in weighted complex networks. Phys. Rev. E Stat. Nonlin. Soft Matter Phys. *71*, 065103. https://doi.org/10.1103/PhysRevE.71.065103.

57. Saramäki, J., Kivelä, M., Onnela, J.-P., Kaski, K., and Kertész, J. (2007). Generalizations of the clustering coefficient to weighted complex networks. Phys. Rev. E Stat. Nonlin. Soft Matter Phys. *75*, 027105. https://doi.org/10.1103/PhysRevE.75.027105.

58. Costantini, G., and Perugini, M. (2014). Generalization of Clustering Coefficients to Signed Correlation Networks. PLoS One *9*, e88669. https://doi.org/10.1371/journal.pone.0088669.

59. Estrada, E. (2019). Rethinking structural balance in signed social networks. Discrete Appl. Math. *268*, 70–90. https://doi.org/10.1016/j.dam.2019.04.019.

60. Toth, R., Zamora, M., Ottaway, J., Gillbe, T., Martin, S., Benjaber, M., Lamb, G., Noone, T., Taylor, B., Deli, A., et al. (2020). DyNeuMo Mk-2: An Investigational Circadian-Locked Neuromodulator with Responsive Stimulation for Applied Chronobiology. Conf. Proc. IEEE Int Conf Syst. Man Cybern *2020*, 3433–3440. https://doi.org/10.1109/SMC42975.2020.9283187.

61. Kavoosi, A., Toth, R., Benjaber, M., Zamora, M., Valentín, A., Sharott, A., and Denison, T. (2022). Computationally efficient neural network classifiers for next generation closed loop neuromodulation therapy - a case study in epilepsy. In 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), *2022*, pp. 288–291. https://doi.org/10.1109/EMBC48229.2022.9871793.

62. Pennington, Z.T., Dong, Z., Feng, Y., Vetere, L.M., Page-Harley, L., Shuman, T., and Cai, D.J. (2019). ezTrack: An open-source video analysis pipeline for the investigation of animal behavior. Sci. Rep. *9*, 19979. https://doi.org/10.1038/s41598-019-56408-9.

63. Ho, J., Tumkaya, T., Aryal, S., Choi, H., and Claridge-Chang, A. (2019). Moving beyond P values: data analysis with estimation graphics. Nat. Methods *16*, 565–566. https://doi.org/10.1038/s41592-019-0470-3.

64. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat. Methods *17*, 261–272. https://doi.org/10.1038/s41592-019-0686-2.

65. Harris, C.R., Millman, K.J., van der Walt, S.J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N.J., et al. (2020). Array programming with NumPy. Nature *585*, 357–362. https://doi.org/10.1038/s41586-020-2649-2.

66. Hunter, J.D. (2007). Matplotlib: A 2D Graphics Environment. Computing in Science & Engineering *9*, 90–95. https://doi.org/10.1109/MCSE.2007.55.

67. Waskom, M.L. (2021). seaborn: statistical data visualization. Journal of Open Source Software *6*, 3021. https://doi.org/10.21105/joss.03021.

68. McKinney, W.. Data Structures for Statistical Computing in Python. https://doi.org/10.25080/Majora-92bf1922-00a.

69. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine Learning in Python. J. Mach. Learn. Res. *12*, 2825–2830.

70. Seabold, S., and Perktold, J.. Statsmodels: Econometric and Statistical Modeling with Python. https://doi.org/10.25080/Majora-92bf1922-011.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Bacterial and virus strains		
AAV9-CAG-Flex-ArchT-GFP	UNC Vector Core	N/A
Experimental models: Organisms/strains		
C57BL/6J mice	Charles River	632
Grm2-Cre Tg(Grm2-cre)MR90Gsat/Mmucd	MMRRC	MMRRC_034611-UCD
Nestin-Cre B6.Cg-Tg(Nes-Cre)1Kln/J	Jackson Laboratories	IMSR_JAX:003771
Software and algorithms		
Intan RHD2000	Intan Technologies	RHD2164
Positrack	Kevin Allen	N/A
Empirical Mode Decomposition in Python	Quinn et al.[48]	N/A
Kilosort via SpikeForest	Magland et al.[49]; Pachitariu et al.[50]	N/A
Other		
12um tungsten wires	California Fine Wire	M294520
Optic fibers	Doric lenses	MFC_200/230-0.37_25mm_RM3_FLT
Head-stage amplifier	Intan Technologies	RHD2164
561nm diode-pumped solid-state laser	Laser 2000	CL561-100

EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS

Animals

These experiments used adult (4–6 months old) wild-type C57Bl6/J mice (Charles River Laboratories, Kent, UK) or hemizygous Nestin-Cre mice (Jackson Laboratories; B6.Cg-Tg(Nes-Cre)1Kln/J, stock no. 003771, RRID: IMSR_JAX:003771) for the initial investigation of principal cell spiking in DSs and SWRs using tetrodes or silicon-probe recordings (Figures 1–4). To optogenetically target DG granule cells, we used adult metabotropic-glutamate-receptor 2-Cre (Grm2-Cre) hemizygous male mice (Figures 5 and 6). This Grm2-Cre mouse strain was obtained from the Mutant Mouse Resource and Research Center (MMRRC; Tg(Grm2-cre)MR90Gsat/Mmucd; stock no. 034611-UCD, RRID:MMRRC_034611-UCD) at University of California at Davis, an NIH-funded strain repository, and was donated to the MMRRC by Nathaniel Heintz, Ph.D., The Rockefeller University, GENSAT and Charles Gerfen, Ph.D., National Institutes of Health – National Institute of Mental Health. All mice were group housed with same-sex littermates until the start of the experiment and singly housed after surgery. Mice had free access to food and water throughout, in a dedicated housing room with a 12/12 h light/dark cycle (7 a.m.–7 p.m.), 19°C–23°C ambient temperature and 40–70 % humidity. This study used mice with good health/immune status, that were not involved in previous procedures, and were drug and test naïve at the start of the experiments. Mice were adult males and the influence (or association) of age and sex, or both on the results of the study was not tested. This represents a limitation to this research's generalizability. All experiments were performed between 8 a.m.–6 p.m. during the light-on period, that is when mice sleep more. Experiments were performed 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.

METHOD DETAILS

Viral vectors

An AAV carrying a double-floxed inverse open reading frame (DIO) Cre-dependent opsin under the CAG promoter was used to deliver Archaerhodopsin (ArchT) (Han et al.[51]) into DG granule cells (AAV9-CAG-Flex-ArchT-GFP, titer: 8.3 × 10^{12} TU / mL, University of North Carolina).

Surgical procedures

Mice received viral injections and microdrive implantations under gaseous isoflurane anaesthesia (∼1% in 1 L / min O$_2$), with systemic and local analgesia administered subcutaneously (meloxicam 5 mg / kg; buprenorphine 0.1 mg / kg; bupivacaine 2 mg / kg). Viruses were injected bilaterally into the dorsal DG (3 × 200 nL per hemisphere; at the following stereotaxic coordinates from

bregma: anterior-posterior: -1.6, -2.4, -2.4; mediolateral: ±1.0, ±1.2, ±1.5; dorsoventral: -1.7, -1.7, -1.7 mm, respectively), and delivered using a pulled glass micropipette (~16 μm i.d.) at a rate of 100 nL min^{-1}, with an additional 100 nL min^{-1} diffusion time with the pipette *in situ*. In a separate surgery, mice were implanted with a microdrive containing twelve- or fourteen-independently movable tetrodes bilaterally targeting DG, CA3, and CA1, and two optic fibers (Doric Lenses Inc., Quebec, Canada) positioned bilaterally above the dorsal DG. Tetrodes were constructed by twisting together four insulated tungsten wires (12.7 μm diameter, California Fine Wire, CA, USA) 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 depth manipulation. A separate group of mice were implanted with unilateral single-shank 64-channel silicon-probe (model: ASSY-236 H3, 8 mm; Cambridge Neurotech, Cambridge, UK; stereotaxic coordinates from bregma: anterior-posterior: -2.0; mediolateral: -1.7 mm); these mice did not receive prior viral injections.

Recording procedures
Following the implantation surgery, mice recovered for at least seven days before familiarization to the recording procedure. Mice were handled daily and exposed to the sleep-box for > 0.5 h per day for at least four days. During this period, tetrodes / silicon-probes were slowly lowered to the proximity of the cell layers. Once at the correct depth, silicon-probes were left in the same position for the rest of the experiment. Tetrodes were lowered into the CA1, CA3 pyramidal or DG granule cell layers on the morning of each recording day in search of multi-unit spiking activity, using the electrophysiological profile of the local field potentials including sharp-wave ripples, gamma oscillations, and dentate spikes to further guide placement. Tetrodes were left in position for ~1.5–2 h before recordings began on that day. At the end of each recording day, tetrodes were raised (~150 μm) to protect hippocampal the cell layers from potential mechanical damage overnight. We lowered again each individual tetrode on the next morning in search of cells, making it unlikely that the recorded units are the same neurons across days. During recording sessions, mice explored open-field environments (41 cm diameter cylinder, or 41 × 41 cm square box, both with 30 cm high walls), or were placed in a sleep box containing sawdust bedding and nesting material (12 × 12 × 28 cm, length × width × height). The instantaneous speed and the theta-to-delta ratio profiles for DS and SWR events corresponded to those of sleep (Figures S1G and S1H). However, in the absence of electromyography signals or other additional signals in defining a sleep stage, we here refer to sleep/rest. Each open-field or sleep box recording session lasted ~15-30 min. Experiments were performed under dim light conditions (~20 lux) with low-level background noise (~50 dB).

Light delivery
A 561 nm diode pumped solid-state laser (Crystal Laser, model CL561-100; distributer: Laser 2000, Ringstead, UK) was used to deliver green-yellow light bilaterally to the dorsal DG (~2-4 mW) via a 2-channel rotary joint (Doric Lenses Inc.).

Multichannel data acquisition
Electrode signals were amplified, multiplexed, and digitized using a single integrated circuit (headstage) located on the head of the animal (RHD2164, Intan Technologies, USA; http://intantech.com/products_RHD2000.html). The amplified and filtered (pass band 0.09 Hz to 7.60 kHz) electrophysiological signals were digitized at 20 kHz (RHD2000 Evaluation Board) and saved to disk with the synchronization signals from the positional tracking and laser activation. To track the location of the animal, three LEDs were attached to the headstage and captured at 25 frames per second by an overhead color camera.

Spike sorting and unit isolation
Spike sorting and unit isolation were performed via automatic clustering software Kilosort[50] (https://github.com/cortex-lab/KiloSort) followed by graphically based manual recombination using cross-channel spike waveforms, auto-correlation histograms and cross-correlation histograms within the SpikeForest framework (https://github.com/flatironinstitute/spikeforest)[49]. All sessions recorded on a given day were concatenated and cluster cut together to monitor cells throughout the day. Each unit used for analyses showed consistent spike waveforms and stable firing rates throughout the entire recording day. Tetrode location in the dorsal-ventral axis for each recording day (Figure S1C) was determined using laminar LFP signatures, as described in detail in Lopes-dos-Santos et al.[52], and later confirmed in the *ex vivo* histology (Figure S1B).

Principal cell versus interneuron classification
Hippocampal principal cells were distinguished from interneurons by the trough-to-peak width of the spike waveform, as previously described.[52] Briefly, to evaluate the waveform consistency for each unit, we used the waveform with the maximum amplitude across the tetrode channels for each cluster. We compared the prominence of a unit mean waveform amplitude to the standard deviation stemming from all its spikes by computing a waveform score:

$$WV_{score} = \sqrt{\sum_{i=1}^{n} \frac{(w_i/\sigma_i)^2}{n}}$$

where w_i is the value of the mean waveform at sample i, σ_i is the standard deviation at sample i across all spikes, and n is the number of waveform samples. This metric quantifies the relative magnitude of the mean waveform amplitude against the spike-to-spike

Neuron
Article

CellPress
OPEN ACCESS

variability. Clusters with a waveform score above 0.75 and a refractory period violation below 2% (quantified as the proportion of intervals shorter than 2 ms in the ISI distribution) were included for further analyses. We categorized units as either putative interneurons or principal cells based on the width of their waveform as indicated by the trough-to-peak latency. In a prior dataset of ~4,000 neurons, we noted a bimodal distribution in trough-to-peak latency. Fitting this with a 1-dimensional, 2-component Gaussian Mixture Model (GMM), we set the classification threshold where the two Gaussian components intersect: units with latencies above were labeled as putative principal cells, and those below as putative interneurons. The same inclusion criteria and classification procedures were used for DG, CA3 and CA1 neurons. In total, this study includes 3,619 hippocampal principal cells (CA1, n = 1,322; CA3, n = 573; DG, n = 1,724; from 134 total recording days in 25 mice).

Local field potential signals
LFP signals were processed by first applying an anti-aliasing filter (8th-order Chebyshev type I filter) to the wide band signals sampled at 20 kHz. These signals were then down-sampled to 1,250 Hz using the decimate function from the signal submodule of Scipy (version 1.11.2).

Dentate spike detection
Dentate spikes were detected during sleep sessions from LFPs recorded from tetrodes located in the DG granule cell layer or silicon-probes with recording contacts in the DG granule cell layer. In silicon-probe recordings, we initially subtracted the LFP signals from all channels using a reference channel found in the stratum oriens. LFPs were band-pass filtered (1–200 Hz, using a 4th order Butterworth filter). The mean and standard deviation of the LFP amplitude were calculated across the entire sleep session and peaks that exceeded a threshold of six times the median absolute value of the filtered signals were designated as dentate spikes. The time bin with the largest local maximum was taken as the peak of the dentate spike, and this timestamp was recorded. If more than one peak appeared within a 50 ms frame, we retained only the highest amplitude peak. On recording days with several tetrodes in the DG, we used the tetrode with the largest mean DS amplitude to select DS event timestamps. Across all tetrode recordings we detected 32,215 DS events in total (mean ± SEM: 441.3 ± 29.2 per day, from 73 recording days in 12 mice); in silicon-probe recordings we detected 15,067 DS events in total (mean ± SEM: 1676.1. ± 316.5 per day, from 8 recording days in 3 mice).

Sharp-wave ripple detection
For the LFPs of each pyramidal CA1 channel, we subtracted the mean across all channels (common average reference), band-pass filtered for the ripple band (80–250 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 envelope of that channel) were regarded as candidate events. 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 twice the ripple band power in the common average reference (to eliminate common high frequency noise); (*ii*) each event had at least four ripple cycles (to eliminate events that were too brief); (*iii*) ripple band power was at least twice 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). For events passing these criteria, the local maximum of each envelope was taken as the peak of the SWR, and these timestamps were recorded. On recording days with several tetrodes in the CA1 pyramidal layer, we used the tetrode with the largest mean ripple envelope amplitude to select SWR events. In tetrode recordings we detected 65,370 SWR events (mean ± SEM: 895.0 ± 82.3 per day, from 73 recording days in 12 mice).

Place maps
To generate place maps, we divided the horizontal plane of the recording enclosure into spatial bins of 1.4 × 1.4 cm to generate the spike count map (number of spikes fired in each bin) for each neuron and the occupancy map (time spent by the animal in each spatial bin) in each task session. All maps were then smoothed by convolution with a two-dimensional Gaussian kernel (s.d. = 1.2 bin widths). Finally, spatial rate maps were generated by normalizing the smoothed spike count maps by the smoothed occupancy map.

Spatial Information
The amount of spatial information conveyed by the spike train of a given cell was calculated using the formula proposed by Skaggs et al.[53]:

$$Information\ per\ spike\ =\ \sum_{i=1}^{N} p_i \frac{\lambda_i}{\lambda} \log 2 \frac{\lambda_i}{\lambda}$$

where $i = 1, 2.. N$ represents each spatial bin of the environment, p_i is the probability of occupancy of bin i, λ_i is the mean firing rate in bin i, and λ is the mean firing rate of the cell over all spatial bins.

CellPress
OPEN ACCESS

Neuron
Article

Peri-event time histograms (PETHs)

For analysis, we excluded all DS and SWR events that occurred within 50 ms of one another. We constructed PETHs over 400 ms windows, 200 ms either side of the peak DS amplitude or the peak of the SWR envelope, using a 1 ms bin width. The mean firing rate of each neuron was calculated during each 1 ms bin over the 400 ms window for each event. Z-scored firing rates were generated (over the DS-triggered or SWR-triggered average) separately for each neuron by calculating the mean and standard deviation over the 400 ms PETH:

$$z_i = \frac{(x_i - \overline{x})}{s}$$

where z_i is the Z score at time bin i, x_i is the firing rate in time bin i, $\overline{x}$ is the mean firing rate across all time bins, and s is the standard deviation of the firing rate across all time bins. The Z-scored firing rate of each neuron was then smoothed using a 3-point moving average to eliminate spurious peaks in low firing rate neurons. For a cell to be classified as significantly activated during DS and/or SWR events, the firing rate within ± 20 ms of the event peak had to be > 3 standard deviations (s.d.) above baseline (calculated as the mean firing rate over the 400 ms window). We also calculated the proportion of activated cells as a function of activation threshold (2 < Z score < 4; Figures S2H and S3E).

Current source density analysis

Current sources and sinks were estimated from LFP recordings taken from single-shank 64-channel silicon-probes spanning the somato-dendritic axis of CA1 principal cells and reaching the inferior blade of the DG. LFP signals were first down-sampled to 1250 Hz. The current source density[54] unscaled signal at time t and electrode n, $CSD[t]n$, was estimated as:

$$CSD[t]_n = - (LFP[t]_{n-1} - 2 \times LFP[t]_n + LFP[t]_{n+1})$$

where $LFP[t]_{n-1}$, $LFP[t]_n$ and $LFP[t]_{n+1}$ are the LFP signals at time t recorded from neighboring electrodes ($n-1$ and $n+1$ are the channels immediately above and below n, respectively, with 20 μm spacing between electrodes). The silicon-probe recording site in the pyramidal layer was identified as the one with largest ripple-band power. We defined the location of radiatum and lacunosum molecular layers according to the sharp-wave and theta laminar profiles, as previously described.[52] We sorted dentate spike events into type 1 (DS$_1$) or type 2 (DS$_2$) in the following way. First, we calculated the CSD estimates for all DSs at the peak of each event and used PCA to find the first two Principal Components from the resulting CSD traces. These principal components had as many dimensions as the number of silicon-probe channels (64). We then used a 2-component Gaussian Mixture Model to classify the events based on their projection onto the first two principal components. This consistently resulted in two event classes having the strongest sinks in different areas of the molecular layer. In line with previous research,[22,24] we classified the events with the strongest sink in the outermost part of the molecular layer as DS$_1$, and events with their sink closer to the granular layer as DS$_2$. Based on CSD classification, event proportions were DS$_1$: 0.35; DS$_2$: 0.65 (5274 DS$_1$ versus 9793 DS$_2$, based on 15,067 events from 8 recording days in 3 mice).

Linear discriminant analysis classifier

To distinguish between DS$_1$ and DS$_2$ events using only the LFP traces, we trained a linear discriminant analysis (LDA) classifier using silicon-probe recorded LFPs from the granule cell layer (https://doi.org/10.5281/zenodo.10034433). LFP signals were first down-sampled to 1250 Hz and low-pass filtered at 50 Hz. We extracted 400 ms epochs centered around the peak of each DS (-200 to +200 ms, with 0.8 ms bin width), providing 500 time-based features (dimensions), one for each time bin, for each LFP trace. We then performed PCA on all silicon-probe-recorded DS LFP traces (15,067) to extract the number of components explaining 90% of the variance. This resulted in 16 principal components, which were then used to train a LDA classifier. We generated 20 models by, each time, randomly selecting 75% of the dataset, which was labelled as DS$_1$ or DS$_2$ based on the CSD classification described above, and then testing the classifier on the remaining (unlabeled) 25% of data. The classifier success rate was: median (IQR) = 85.4 (85.3–85.6) %. We then used the model with the highest accuracy to classify DS$_1$ and DS$_2$ events from LFPs recorded from the granule cell layer in our tetrode-recorded data. From tetrode-recorded LFPs, the proportions of type 1 and type 2 DS events were: median (IQR) DS$_1$ = 0.34 (0.25–0.38); DS$_2$ = 0.66 (0.62–0.75), based on 10,337 DS$_1$ versus 21,740 DS$_2$ events in 73 recording days in 12 mice.

Population spiking vectors

We generated event-based hippocampal population vectors of instantaneous principal cell spiking for every DS and SWR event using 50 ms wide windows centered on the peak of the DS or the peak envelope of the CA1 ripple (±25 ms from the peak). In addition, we calculated the spiking activity of hippocampal principal cells in equivalent 50 ms ('no event') control epochs, that contained neither DS nor SWRs. Baseline periods were selected from the same sleep sessions and excluded all epochs ± 250 ms either side of any DS or SWR events. To calculate the proportion of coactive neurons in each time window, we calculated the number of simultaneously active hippocampal principal cells (i.e., cells firing at least one spike during the 50 ms window) by the total number of simultaneously recorded hippocampal principal cells. We then calculated the mean proportion of coactive cells for each recording session. For inclusion in these analyses, each recording session required a minimum of 100 of each type of event (DS$_1$, DS$_2$, SWR) and a minimum of 20 simultaneously recorded hippocampal principal cells.

Population-level sparsity

The sparsity S of a given population firing vector x was calculated using the Gini index[33,34,55] as:

$$S = \frac{\sum_{i=1}^{N}(2i - N - 1)x_i}{N\sum_{i=1}^{N}x_i}$$

where x is the population vector containing, in ascending order, the spike counts discharged by each principal cell in a 50 ms time window (centered on the peak of the SWR, DS), N is the length of that vector (i.e. the number of simultaneously recorded principal cells), and i is the rank of spike counts in ascending order. Population vectors where the total number of spikes is more evenly distributed between neurons have a lower Gini index (lower sparsity) than population vectors where the total number of spikes is concentrated in a few neurons (higher sparsity).

Logistic regression classifier

We used a logistic regression classifier to distinguish between population vectors of hippocampal principal cell spiking activity during DS, SWR, or equivalent duration (50 ms) control vectors that were taken from 200–250 ms periods before or after the peak of either the DS or SWR events. For each recording session, we generated matrices of these population vectors (cells × epochs) for these four different event-types, and then binarized the spike counts (i.e., spike count > 0 = 1, else 0) to control for the influence of firing rate differences between neurons. For each recording day, we used the event with the lowest number of epochs to determine the training set size – for example, if there were 200 DS events, we used 75% (150 population vectors) as the DS training set, and randomly subsampled the SWR matrix for 150 SWR population vectors (with identical principal cells). This way, the training input to the classifier was balanced across event types. Similarly, the testing set consisted of the remaining (unlabeled) 25% of population vectors from the DS population vectors plus an equivalent number of SWR population vectors (e.g., 50 DS population vectors and 50 SWR population vectors, subsampled from the remaining SWR testing matrix). For each recording day, we ran three models: one to classify event epochs, one to classify pre-event epochs and one to classify post-event epochs. Model accuracy was measured as the proportion of correctly classified events (DS versus SWR, or pre-DS versus pre-SWR, respectively).

Peer-to-peer coactivity analysis

We constructed hippocampal population graphs that represent the coactivity relationships between all pairs of principal cell spike trains recorded during a given sleep or exploratory session. These coactivity graphs were computed using 50 ms time windows for DS and SWR events and theta cycles as time windows for active exploratory sessions. 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_{ij} obtained by fitting the GLM (Figure 3A):

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

where x_j, x_i are the Z-scored event-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,

$$P = \sum_{n=0}^{N-\{ij\}}x_n$$

with α_{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 $\beta_{i,i} = 0 \; \forall i \; in \; N$, 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 C_i 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.[56–58]:

$$C_i = \frac{\sum_{jq}\left(\widehat{\beta_{ij}}\,\widehat{\beta_{iq}}\,\widehat{\beta_{jq}}\right)^{1/3}}{k_i(k_i - 1)}$$

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.

CellPress
OPEN ACCESS

Neuron
Article

Note that 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. This method to assess firing relationships in the neuronal population of the hippocampus as a signed network where both positive and negative edges (i.e., correlated and anti-correlated spike trains) coexist leverages from past studies investigating community organization in social networks, indicating that triads represent the smallest motif capturing "structural balance" in patterns of peer-to-peer relationships.[59]

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. 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.

Population vector similarity

Population vectors of hippocampal principal cell spiking activity were generated for baseline, SWR, and DS events as described above, yielding separate (cell × event number) matrices of spike counts for each event-type. To remove potential biases caused by unequal numbers of events, we used the event-type with the fewest epochs to determine the final matrix size. For example, if there were 200 DS events in a given recording session, we randomly subsampled the SWR and baseline matrices to extract 200 SWR and 200 baseline population vectors (with identical principal cells) for comparison. Next we binarized these matrices (spike count > 0 = 1, else 0). Then we assessed the self-similarity for each event matrix (cells × event number) by computing the Pearson correlation coefficient for every pair of population vectors from the same event-type, and then calculating the mean across all of these correlation coefficients. As an alternative, we also calculated the Jaccard similarity coefficient (J), which measures the size of the intersection (i.e. overlap in active units) between pairs of population vectors (A, B), divided by the size of the union:

$$J(A,B) = \frac{|A \cap B|}{|A \cup B|}$$

Population dimensionality

We estimated the dimensionality of the principal cell population firing structure during SWRs and DSs from activity matrices that were matched for neuron identity and the number of DS and SWR events. We applied Principal Component Analysis (PCA) to each activity matrix, using the number of simultaneously recorded principal cells as the maximum number of components. Each matrix required at least 20 principal cells for inclusion in the analysis. We then extracted the number of components explaining 90% of the variance in these population vectors and scaled this by the total number of neurons in each matrix (Figures 3L, S4M, and S4P). We also show dimensionality for a range of explained variance values (Figures S4N and S4O). Note that the ratio of DG to CA cells in these matrices did not significantly affect the dimensionality estimate (Table S1).

Theta-cycle detection

Theta cycles were detected as described in Lopes dos Santos et al. Briefly, we used masked Empirical Mode Decomposition[48]; https://pypi.org/project/emd/) to separate CA1 LFPs into oscillatory components termed intrinsic mode functions (IMFs). We delineated individual theta cycles from their troughs and peaks, i.e. the local maxima and minima of the theta IMF. Theta cycles were defined as peak-trough-peak sequences with trough-peak and peak-trough intervals between 31-100 ms and peak-to-peak distances between 71-200 ms. Note that this method is designed to detect chains of theta cycles but to do so it identifies each cycle independently.

Reactivation of waking coactivity patterns

We leveraged our pairwise peer-to-peer coactivity measure (as described above; Figure 3A) to estimate DS and SWR reactivation. With this, we compared the tendency of principal cell pairs to co-fire in theta cycles during exploration (theta coactivity) with the tendency to co-fire in DS (or SWR) during the following post-exploration sleep/rest period (post-DS or post-SWR co-firing), controlling for their baseline co-firing in the pre-exploration sleep/rest period before (pre-DS or pre-SWR co-firing) and mouse identity, using a linear mixed model:

$$Post \sim \beta_0 + \beta_{theta} + \beta_{pre} + \upsilon_{mouseID} + e$$

where β_0 is the intercept of the regression line, β_{theta} is the regression coefficient for the theta co-firing, β_{pre} is the regression coefficient for the pre-exploration offline co-firing (in DS or SWR events), $\upsilon_{mouseID}$ is the individual mouse identity, and e the error term. Likewise, we compared the tendency of principal cell pairs to co-fire in theta cycles during exploration (theta coactivity) with the

tendency to co-fire in DS or SWR during the pre-exploration sleep/rest period (pre-DS or pre-SWR co-firing), controlling for their post-exploration co-firing in the sleep/rest period after (post-DS or post-SWR co-firing) and mouse identity, using the reverse linear mixed model:

$$Pre \sim \beta_0 + \beta_{theta} + \beta_{post} + \upsilon_{mouseID} + e$$

From these LMMs, we extracted the β coefficients predicting post-SWR or post-DS coactivity from theta coactivity (controlling for pre-SWR or pre-DS coactivity, respectively) and tested their significance in two ways. First, we performed control GLMs using the pre-DS (or pre-SWR) coactivity as the dependent variable and the theta coactivity and post-DS (or post-SWR) coactivity as the independent variables. In these models, pre-event, post-event and theta coactivity were entered as fixed-effects and mouse identity as random-effects, using the restricted maximum likelihood method (implemented using the MixedLM class, and fit() method with default parameters, from the statsmodels library (Seabold and Perktold, 2010) in Python3.10). Second, we constructed a random probability distribution of β weights for theta coactivity by shuffling the cell pair identity, thereby generating a null distribution (based on 1000 LMMs, each time randomly shuffling cell-pair identity).

Closed-loop optogenetic interventions

For DS-informed interventions, real time detection of DSs was achieved by first high pass filtering the DG LFP signals (5 Hz) using the on-board signal processing capabilities of the Intan RHD evaluation board (RHD2000, Intan Technologies, USA) and triggering a laser pulse if the LFP signal exceeded a voltage-threshold. Thresholds for DS-onset detection were set for each mouse during a sleep session at the start of each recording day so that DS events were consistently detected ($\sim$3 S.D. above mean signal amplitude). Threshold detection triggered a digital transistor-transistor logic (TTL) output pulse from the RHD interface to a Master 8 stimulation timing device (A.M.P.I., Jerusalem, Israel), which in turn sent a 100 ms duration square-wave pulse to activate the laser. In the 'DS-synchronized' condition, the laser was triggered with zero latency from DS-onset, whereas in the 'DS-delay' condition the laser was triggered 100 ms after DS detection (Figures 5 and 6). The rates of false negatives (DS not triggering laser pulse) and false positives (laser pulse emitted for LFP trace not meeting DS criteria) were 1.8±0.6% and 4.7±0.4%, respectively. The laser delivered yellow-green light (561-nm) into the dentate gyrus, which in DG^{Grm2}::ArchT mice activated the outward proton pump, Archaerhodopsin T to suppress spiking activity in DG granule cells. To investigate changes in firing rates in individual hippocampal principal cells during light-delivery, we constructed PETHs over 400 ms windows, 200 ms either side of DS-onset, using a 1 ms bin width and extracted the peak firing rate during DS-synchronized light-delivery versus DSs with no light delivery. In addition, we Z-scored the binned spike trains and calculated the mean Z score between DS-onset and 100 ms after DS-onset for each hippocampal principal cell during DS-synchronized light-delivery versus the equivalent 100 ms no-light period in the DS-delay condition.

For SWR-informed interventions, the Intan evaluation board was configured with firmware enabling additional filtering. Five operations were performed on the continuously acquired CA1 wideband LFP signal to provide a real time estimate of the instantaneous power in the ripple-band.[60,61] (1) To enable low-latency processing, the signal was first down-sampled to 2.5 kHz by averaging the raw 20 kHz data stream with a sliding window of 8 samples with no overlap. (2) This signal was then high-pass filtered (using a 1st order digital infinite impulse response filter with a corner frequency of 1.6 Hz to remove amplifier offset and electrode drift). Next, the signal was (3) band-limited to 100–200 Hz with a 4th order Butterworth filter, (4) rectified by taking its absolute value, and (5) amplified 128-fold and smoothed with an exponential moving average operation over an equivalent window size of 32 samples (12.8 ms). To detect SWR events in this band-power estimate, the threshold level for each mouse was set during a sleep session at the start of each recording day to ensure consistent ($\sim$3 S.D. above mean power) detection throughout the day. On detecting a threshold crossing, the Intan recording controller delivered a 5 ms TTL pulse to a Master 8 stimulation timing device (A.M.P.I., Jerusalem, Israel). Analogous to the DS-informed interventions, in the 'SWR-synchronized' condition the laser was triggered with zero latency from SWR-onset, whereas in the 'SWR-delay' condition the laser was triggered 100 ms after SWR detection.

Recognition memory tasks

On each day of both the novel-object and novel-position recognition tasks, mice explored a square-walled open field (Figure 6A; the 'object arena') containing four objects, each positioned midway along a given wall, $\sim$1cm from the wall edge. Objects used were $\sim$3 × 3 × 4 cm (width × depth × height) objects (e.g., Lego™ blocks or other similar items). During the first session in the object arena, mice explored four completely novel objects ('sampling' session, 10 min). After the sampling session, mice were placed into a sleep box where they received DG-targeting light delivery that was either synchronized to event detection (DS-synchronized or SWR-synchronized condition) or delayed by 100 ms from event detection (DS-delay or SWR-delay condition), as described above (sleep/rest session, 20 min). In the novel object recognition task, before the start of the next test session, one of the four objects was replaced with a different (and completely novel) object, and mice then explored the four objects again ('test 1' session, 10 min). This process was repeated, with another sleep session ($\sim$20 min, with either DS-sync or DS-delay light-delivery), followed by another object exploration session with one completely novel object and three previously encountered objects ('test 2' session, 10 min). In the novel position recognition task, the locations of two of the initially sampled objects were swapped (e.g. North and West), whereas the other two objects remained in their original positions. In the novel position task, only DS-the synchronized and DS-delay conditions and only the first test session were used. During each test session, we measured the time spent exploring each object and we calculated the percentage time spent investigating the novel object (or novel positioned objects) versus the mean percentage time spent investigating the familiar

objects (i.e., those objects seen in the previous session and/or those in the same locations). For analysis, four 'object-zones' were created by dividing the arena into nine equal sized square zones ($\sim$12 × 12 cm), such that four of these zones contained the objects. Time spent in the object zone was determined directly from the automated tracking data. Between sessions, the floor of the maze and the objects were cleaned with water. On any given day, mice received the same light-delivery condition.

Tone fear conditioning task
Fear conditioning was conducted in one of four operant chambers each with distinct visual cues (ENV-307A, Med Associates Inc., IN, USA). Mice were exposed to five auditory cues (either 2900 Hz tone or white noise, 72 dB, 30s duration), each co-terminating with a mild foot-shock (0.3 mA, 0.5 s). The mean ITI was 74s (range: 60 to 90s). Immediately after fear conditioning, mice were removed from the operant chamber and placed into the sleep box where they received DS-sync or DS-Delayed DG cell silencing for 45 minutes. For the recall session, mice were then placed into a different operant chamber than the one where they received conditioning (to reduce the impact of contextual cues on recall). Mice were exposed to the same five auditory cues but no foot-shocks were given. Fear memory was assessed by measuring freezing responses during the first two cues presented in the recall session (before extinction occurs) and comparing these responses to freezing responses to the first cue during training (before any shocks were given). Freezing was measured using automated movement detection software (ezTrack, [62]) and expressed as a % of tone duration (i.e. freezing for 15s during a 30s tone = 50% freezing).

Tissue processing and immunohistochemistry
At the completion of experiments, mice were deeply anesthetized with pentobarbital and perfused transcardially with 0.1 M PBS followed by 4% paraformaldehyde (PFA) in PBS. Brains were extracted and kept in 4% PFA for $\sim$24–72 h and then transferred to PBS (with 0.05% sodium-azide). For tetrode localization, free-floating sections (50 μm) sections were mounted on slides and imaged at × 5 using a Zeiss microscope (AxioImager M2; Zeiss, Plan-Neofluar 5× /0.16 objective). For immunostaining, free-floating sections (50 μm) were rinsed in PBS with 0.25% Triton X-100 (PBS-T) and were blocked for 1 hour at $\sim$20°C in PBS-T with 10% normal donkey serum (NDS). Sections were then incubated with primary antibodies diluted in 3% NDS blocking solution and incubated at 4°C for 72 hours (GFP anti-chicken, 1:1,000, Aves Labs, catalog no. GFP-1020; NeuN guinea pig, 1:500, Synaptic Systems, catalog no. 266 004). All sections were rinsed three times for 15 min in PBS-T and incubated for 4 hours at $\sim$20°C in secondary antibodies in the blocking solution (Cy3 donkey anti-guinea pig, 1:400, Jackson ImmunoResearch, catalog no. 706-165-148; goat anti-chicken 488, 1:1,000, Thermo Fisher Scientific, catalog no. A-11039). Sections were then rinsed three times for 15 min in PBS-T, with some sections then incubated for 1 min with DAPI (0.5 μg ml^{-1}, Sigma, D8417) diluted in PBS to label cell nuclei before three additional rinse steps of 10 min each in PBS. Sections were mounted on slides, cover-slipped with Vectashield (Vector Laboratories, catalog no. H-1000) and stored at 4°C. Sections were also used for anatomical verification of the tetrode tracks. Images were acquired using a Zeiss confocal microscope (LSM 880 Indimo, Axio Imager 2) with a Plan-Apochromat ×20/0.8 M27 objective and the ZEN (Zeiss Black 2.3) software.

QUANTIFICATION AND STATISTICAL ANALYSIS

Analyses were performed in Python 3.8 (https://www.python.org/downloads/release/python-3816/) and Python 3.10 (https://www.python.org/downloads/release/python-31011/), using the Python packages DABEST (Ho et al.[63]), scipy (Virtanen et al.[64]), numpy (Harris et al.[65]), matplotlib (Hunter[66]), seaborn (Waskom[67]), pandas (McKinney[68]), scikit-learn (Pedregosa et al.[69]), statsmodels (Seabold and Perktold[70]). Error bars, mean ± S.E.M unless otherwise stated. We used throughout this study a bootstrap-coupled estimation of effect sizes, plotting the data against a mean difference between the left-most condition and one or more conditions on the right and compare this difference against zero using 5,000 bootstrapped resamples. In these estimation graphics (DABEST plots), each black dot indicates a mean difference and the associated black ticks depict error bars representing 95% confidence intervals; the shaded area represents the bootstrapped sampling-error distribution. Bandwidth estimates for the kernel density estimate were computed using the scikit-learn package. We used the DABEST package to calculate test statistics and p-values and visualize data. The test statistic is the mean difference and the p-value is the is the probability of observing the effect size (or greater), assuming the null hypothesis of zero difference is true. Paired permutation tests (or equivalent paired tests) were performed for repeated-measures analyses and unpaired tests used for independent samples. Data distributions were assumed to be normal, but this was not formally tested. Our results were replicable across mice and recording days. For the optogenetic interventions, the different closed-loop conditions (DS-sync, SWR-sync, DS-Delay, and SWR-Delay) were experienced in a randomized order across days. In the object recognition tasks, objects and their positions and the order of their replacement were randomized. 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 analyses. No statistical methods were used to pre-determine sample sizes, but our sample sizes are similar to or larger than those reported in previous publications. Inclusion criteria for well-isolated single units were used as published in previous studies and are described in the corresponding subsections of the method details. For the population vector analyses (Figures 3 and S4), each recording session required a minimum of 100 of each type of event (DS$_1$, DS$_2$, SWR) and a minimum of 20 simultaneously recorded hippocampal principal cells for inclusion.

Supplemental information

Offline hippocampal reactivation during dentate

spikes supports flexible memory

Stephen B. McHugh, Vítor Lopes-dos-Santos, Manfredi Castelli, Giuseppe P. Gava, Sophie E. Thompson, Shu K.E. Tam, Katja Hartwich, Brook Perry, Robert Toth, Timothy Denison, Andrew Sharott, and David Dupret

Supplemental Figures and Tables

Figure S1. Tetrode locations and characterization of dentate spikes, related to Figure 1.
(A) Triple-(DG-CA3-CA1) tetrode layout schematic.
(B) Example histology showing tetrode tracks in DG, CA3, and CA1, with color-coded contours of the pyramidal cell layers and the granule cell layer.
(C) Histogram of tetrode locations in the dorsal-ventral plane (CA1 tetrodes n=176; DG tetrodes n=244, CA3 tetrodes n=92, from n=73 sessions, 12 mice).
(D) Examples of mean ± SEM dentate spike (DS) local field potential (LFP) waveforms from individual mice.

(E) Group mean ± SEM DS LFP waveform (n=73 sessions, 12 mice).

(F) Kernel density estimates (KDEs) and boxplots for DS (left) and SWR (right) event durations (based on 32,215 DS events and 65,370 SWR events, n=73 sessions, 12 mice).

(G) DS and SWR events occur when mice are asleep or in quiet rest and not when they are active. The graph shows KDEs for the probability of DS (blue) and SWR (orange) occurrence versus active behavior (gray) for a range of movement speeds. Active behavior was determined from the theta-to-delta ratio (>2.4) from the CA1 LFP during open field exploration.

(H) Estimation plot showing that DS and SWR events occur when the theta-to-delta ratio is low compared to active behavior. For this analysis, we defined a minimum activity level (movement speed > 3 cm/s) based on the active behavior movement speed distribution (panel F) to include 99% of the area under the curve. We then extracted the theta-to-delta ratio for speeds above this minimum and compared this distribution to the theta-to-delta ratio distributions during DSs and SWRs. Upper: raw data points (each point shows mean theta-to-delta power during one active behavior session or one sleep session), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event. Lower: difference (Δ) in theta-to-delta ratio between active epochs versus DS and SWR epochs computed from 5,000 bootstrapped resamples and with the difference-axis origin (dashed line) aligned to active behavior (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). The test statistic is the mean difference, shown on the y-axis of the lower plot. P-values are from unpaired permutation tests, active versus event, ***$P < 0.001$.

(I) Cross-correlogram of the probability of DS occurrence with respect to SWRs. Note that there was a small increase in the probability of DS around the time of a SWR, with more than 90% of DSs not occurring within ± 50 ms of a SWR. Red dashed lines show peak of event, gray dashed lines show ± 50 ms from peak.

(J) As for panel I but for SWR occurrence with respect to DSs.

Figure S2. Dentate spikes activate hippocampal principal cells, related to Figure 1.

(A) Top: Example spike waveform showing the trough-to-peak measurement for spike width. Scale bar 100 µV and 0.5 ms. Middle: example auto-correlograms from individual DG, CA3 and CA1 principal cells. Bottom: example place maps from three individual (simultaneously recorded) DG, CA3 and CA1 principal cells illustrating the spatial distribution of spiking activity. The number in the top right corner shows the maximum firing rate of the cell in its place-field.

(B, C) Trough-to-peak width was used to classify principal cells versus interneurons. (B) Estimation plot showing that the trough-to-peak width is narrower in DG principal cells versus CA3 and CA1 principal cells. (C) Trough-to-peak width narrower in DG versus CA3 interneurons but is wider in DG versus CA1 interneurons. Note that the trough-to-peak width for DG principal cells remains wider for DG principal cells than DG interneurons.

(D) The auto-correlogram shape differs between principal cells and interneurons in DG, CA3 and CA1. Here, we used the Gini index to evaluate the sparsity in the spike probability distribution for each 1 ms bin of the auto-correlogram (between 0 and +50 ms). This distribution was more unequal for principal cells in all three regions, hence a higher Gini index; and was more equal for interneurons, hence a lower Gini index.

(E) The mean firing rate is higher in hippocampal interneurons (iDG, iCA3, iCA1: recorded cells with a trough-to-peak width < 0.45 ms) versus hippocampal principal cells.

(F) DG, CA3 and CA1 principal cells exhibit higher spatial information scores than interneurons in these respective regions.

For B-F estimation plots, Upper: raw data points (each point shows one principal cell); Lower: difference (Δ) in trough-to-peak width, firing rate, or spatial information (respectively) in DG versus CA3 and DG versus CA1. Other plot details as in Figure S1H.

(G) Examples of three individual principal cells' z-scored firing rates during DS (blue) and SWR (orange) events. The horizontal dashed line shows z-score = 3.

(H) Percentage of principal cells active during DS (left) versus SWR (right) events, as defined by crossing various z-thresholds.

(I) Percentage of significantly activated principal cells, as defined by a z-score > 3 (within ± 20 ms of the event peak, shown at time 0 in **G**, during DS (left) and SWR (right) events.

(J) Estimation plot showing the effect size for the differences in the time point of maximum neuronal spiking activity during DSs and SWRs in DG, CA3, and CA1 principal cells. Note that DG cells reached their peak firing significantly before CA1 neurons during DSs; CA3 cells reached their peak firing significantly before CA1 neurons during SWRs. Upper: raw data points (each point shows one principal cell that was significantly active (z > 3) during each event); Lower: difference (Δ) in time of peak activation in DG versus CA3 and DG versus CA1.

(K, L) Corresponding time course of principal cell instantaneous firing rate (z-score) during SWR (K) and DS (L) events.

(M, N) Using the time to cross the z-score > 3 threshold, we observed that during DS events, DG principal cells increase their firing activity significantly earlier than both CA3 and CA1 principal cells. Panel N is the corresponding estimation plot of the response latency (i.e. when each cell crossed the z > 3 threshold relative to the event peak) showing that DG cells are active before CA3 and CA1 principal cells during DSs. Upper and lower plots as in panel J.

(O, P) Estimation plots comparing the overall mean firing rate of each principal cell (calculated across the entire recording session) to its peri-event firing rate (calculated as the mean firing rate ± 5 ms around the peak of the event) during DS (O) and SWR (P) events. Upper: raw data points (each point shows one principal cell's mean rate and peri-event rate) in DG, CA3 and CA1; Lower: difference (Δ) in firing rate between mean rate and peri-event rate for DG, CA3 and CA1 separately. Other plot details as in Figure S1H.

For B-F,J,N-P, the test statistic is the mean difference, shown on the y-axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). P-values are from unpaired permutation tests, cell type versus cell type (B-F,J,N); or paired permutation tests, baseline (mean rate) versus event (O-P), *$P < 0.05$, **$P < 0.01$, ***$P < 0.001$.

Figure S3. Current source density, local field potential profiles, and hippocampal principal cell spiking for DS$_1$ and DS$_2$ events, related to Figure 2.

(A) Examples of current source density (CSD) profiles for individual Type 1 (DS$_1$) and Type 2 (DS$_2$) dentate spikes (two DS$_1$ examples, two DS$_2$ examples) recorded from a 64-channel silicon-probe. *Left panel*: the instantaneous CSD ±150 ms around the event peak. *Right panel:* the CSD amplitude at each depth (based on the mean amplitude from –25 to + 25 ms around the event peak). Stratum oriens: ori; pyramidal layer: pyr; stratum radiatum: rad; lacunosum moleculare: lm; hippocampal fissure: hf; outer molecular layer: om; middle molecular layer: mm; inner molecular layer: im; granule-cell layer: gr.

(B) Group mean ± SEM LFP waveforms for DS$_1$ and DS$_2$ events from silicon-probe recordings (8 sessions, 3 mice).

(C) Examples of mean ± SEM LFP waveforms for DS$_1$ and DS$_2$ events from tetrode recordings in individual mice.

(D) Group mean ± SEM LFP waveforms for DS$_1$ and DS$_2$ events from tetrode recordings (73 sessions, 12 mice)

(E) Percentage of principal cells active during DS$_1$ (left) versus DS$_2$ (right) events, and as defined by crossing various z-thresholds.

(F) Percentage of significantly activated principal cells, as defined by a z-score > 3 (within ± 20 ms of the event peak), during DS$_1$ (left) and DS$_2$ (right) events.

(G-J) Peri-event time histograms (G,H) showing z-scored firing rates ± 100 ms around the event peak and estimation plots (I,J) comparing overall mean firing rate (calculated across the entire recording session) to peri-event firing rate (calculated as the mean firing rate ± 5 ms around the peak of the event) for all principal cells during DS$_1$ and DS$_2$ events. DG n=921, CA3 n=388, CA1 n=887 principal cells (12 mice). Upper and lower plots as described in Figure S2O-P.

(K) Percentage of suppressed principal cells (i.e., cells with a z-score < 0 during the event peak) during DS$_1$ and DS$_2$ events.

(L) Peri-event time histograms showing z-scored firing rates ± 25 ms around the event peak for the lowest quartile of activated principal cells (i.e., the 25% least activated / suppressed principal cells) during DS$_1$ events. (DG n=230, CA3 n=97, CA1 n=221 principal cells).

For I and J, the test statistic is the mean difference, shown on the y-axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). P-values are from paired permutation tests, baseline (mean rate) versus event, ***$P < 0.001$.

Figure S4. Distinct population coactivity structures in DSs versus SWRs, related to Figure 3.

(A) Toy example illustrating construction of principal cell × event spiking activity matrices using the population firing vectors nested in 50 ms time windows centered on DSs, SWRs, or duration-matched (pre- and post-) control events. The proportion of active cells in each column vector (B-C) was defined as the proportion of cells with a non-zero spike count during the individual time window. Logistic regression classifiers (D) were trained using 75% of the population spiking activity vectors and tested with the remaining 25% of vectors, on any given session. We used the event type with the lowest number of epochs to determine the training and testing set size and then randomly subsampled the other event matrix to generate the same number of training and testing vectors for each event type, so that the classifier was balanced across event types. Separate matrices and classifiers were utilized for event, pre-event, and post-event epochs. For all of the analyses below, the inclusion criterion was that the session had to contain a minimum of 20 principal cells and 100 events of each type.

(B) Estimation plot showing the mean proportion of cells active during DS, DS_1, DS_2 and SWR events relative to equivalent duration-matched baseline periods during sleep when none of these events were present. Upper: raw data points (each point shows mean proportion of active cells in each vector across one recording session), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event type. Lower: difference (Δ) in proportion of active cells between baseline epochs versus DS, DS_1, DS_2 and SWR epochs. Note that all events contained more active cells than baseline epochs but there was no statistical difference between the proportion of active cells in DS_2 versus SWR events.

(C) Likewise, shown is the proportion of cells active during Baseline, DS_1, DS_2 and SWR events split by hippocampal subregion. Upper: raw data points (each point shows mean proportion of active cells by subregion), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each subregion. Lower: difference (Δ) in proportion of active DG cells versus CA3 and CA1 cells.

(D) Classifier performance for SWR versus DS_1, SWR versus DS_2, and DS_1 versus DS_2 population vectors. Dashed line shows chance performance.

(E) Estimation plot showing the population vector similarity for all event types compared to their control condition in which each event vector was correlated with 'shuffled' population spiking vectors, where the cell identity was randomly shuffled. Upper: raw data points (each point shows mean population similarity for one recording session or shuffled equivalent), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event. Lower: difference (Δ) in population similarity between real data for each event and its shuffled equivalent.

(F) Estimation plot showing population vector similarity for DS and SWR events compared to baseline control, as in Figure 3F, but separately for CA (left) and DG (right) principal cells. Upper: raw data points (each point shows mean population similarity by event type for one recording session), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event type. Lower: difference (Δ) in population similarity between baseline epochs versus DS and SWR epochs.

(G) Our results showing significantly higher population vector similarity for DS and SWR events versus baseline, and for SWR versus DS events (Figure 3F) were recapitulated at the level of individual mice (8/8 mice for DS and SWR events versus baseline, 7/8 mice for SWR versus DS events; $p < 0.05$ permutation tests for DS versus baseline, SWR versus baseline, DS versus SWR with n=number of population vector pairs per mouse). Black traces show individual mice ± SEM (calculated on the number of population vectors per mouse); red trace shows the group mean ± SEM (calculated on n=8 mice).

(H) Our results using the Pearson correlation to compare population vector similarity (Figure 3F-G) were replicated using the Jaccard similarity measure. Estimation plot showing the population vector similarity for all event types compared to baseline epochs. Upper: raw data points (each point shows mean population similarity by event type for one recording session), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event type. Lower: difference (Δ) in population similarity between baseline epochs versus DS, DS_1, DS_2 and SWR epochs.

(I) Estimation plot showing mean clustering coefficients (as Figure 3K) but for CA1-3 cells only (left panel) or DG cells only (right panel), which also show higher clustering coefficients for DS events.

(J) Our results showing significantly higher clustering coefficient for DS and SWR events versus baseline and DS versus SWR events (Figure 3K) were recapitulated at the level of individual mice (8/8 mice for DS and SWR events versus baseline, 7/8 mice for SWR versus DS events; $p < 0.05$ permutation tests for DS versus baseline, SWR versus baseline, DS versus SWR with n=number of neurons per mouse). Black traces show individual mice ± SEM (calculated on the number of neurons per mouse); red trace shows the group mean ± SEM (calculated on n=8 mice).

(K) Estimation plot showing that the neuronal coactivity graphs nested in both DS_1 and DS_2 events contained significantly stronger triads of coactive nodes compared to SWR graphs, as indicated by higher mean clustering coefficients. This was notably the case for DS_2.

(L) Single-neuron coactivity strength. As an alternative method to the triadic clustering coefficient parameter, we also show in this estimation plot that SWR and DS events differ in neuronal coactivity strength and from baseline epochs. We defined the single-neuron coactivity strength as the average

pairwise coactivity relation of a given neuron with its population peers. For any two neurons (i, j), we obtained the regression coefficient β from a generalized linear model predicting the spike discharge of neuron j from the activity of neuron i while regressing out the activity of the remaining population. The strength of neuron i is then the average across all the weights β_{ij}.

(M) Estimation plot showing population dimensionality required to explain 90% of the variance in DS and SWR events compared to baseline control separately for population vectors containing only CA (left) or DG (right) principal cells. Upper: raw data points (each point shows mean population dimensionality by event type for one recording session), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event type. Lower: difference (Δ) in population dimensionality between baseline epochs versus DS and SWR epochs. Importantly, considering here DG and CA regions separately for the dimensionality analysis bears the caveat that, to comply with our criterion of at least 20 simultaneously recorded principal cells for each recording session, these analyses use far fewer recording days and the resulting neuronal population vectors are much smaller (n=17 days, n=25.8±1.0 cells per vector for CA; n=15 days, n=25.5±0.9 cells per vector for DG) than the data presented in Figure 3K (n=34 days, n=37.2±1.8 cells per vector), which limits the comparison.

(N) PCA to compare the dimensionality of SWR versus DS matrices (cell × event number), matching the number of events for each event type, determining the number of components required to explain 70–95% of the variance. In each case, the dimensionality was significantly higher for DS versus SWR events at α < 0.05 (Wilcoxon test for paired samples, one-tailed).

(O-P) The number of components required to explain equivalent amounts of variance was lower in DS$_1$ versus SWR and DS$_2$, and higher in DS$_2$ versus SWR events (N); as illustrated by the estimation plot in panel O, showing that a lower number of principal components was required to explain 90% of the variance across the population vectors nested in DS$_1$, compared to SWR and DS$_2$ events.

For B-F,H-I,K-M,P, the test statistic is the mean difference, shown on the y-axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). P-values are from paired permutation tests, baseline versus event (B,F,H,I,L,M), event versus pre-event and event versus post-event (D), data versus shuffle (E), event versus event (K,P); or unpaired permutation tests, cell type versus cell type (C), *$P < 0.05$, **$P < 0.01$, ***$P < 0.001$.

Figure S5. Offline reactivation of theta population coactivity in individual mice, related to Figure 4. (A) Raw data examples of hippocampal principal cells spiking activity during active exploration marked by theta oscillations (Left) and offline sleep/rest (Right; example SWRs shown). Scale bars 0.5 mV and 150 ms (Left) or 50 ms (Right). We applied our peer-to-peer coactivity analysis (Figure 3A) to obtain the waking patterns of population coactivity nested in theta cycles during exploration, and the offline patterns of population coactivity nested in either DS or SWR events during sleep/rest before and after exploration. With these, we next computed DS and SWR reactivation by measuring the tendency of the peer-to-peer theta firing associations to reoccur in post-exploration sleep/rest DS (or SWR) events, while controlling for prior pre-exploration DS (or SWR) coactivity, using a linear mixed model (Figure 4A).

(B) SWR reactivation in individual mice (measured by the β coefficients of the linear regression that predicted post-exploration SWR coactivity from waking theta coactivity, controlling for pre-exploration SWR coactivity).

(C) DS reactivation in individual mice (measured by the β coefficients of the linear regression that predicted post-exploration DS coactivity from waking theta coactivity, controlling for pre-exploration DS coactivity).

(D-E) Distribution of peer-to-peer coactivity values (β coefficients) for SWRs (D) and DSs (E) in the pre- and post-exploration sleep/rest sessions. Significance was tested using the Wilcoxon test for matched pairs: post > pre for DSs (U=1198580, n=7310, p < 0.001); and for SWRs (U=12839635, n=7310, p < 0.001).

(F-G) Estimation plots showing that the neuronal coactivity graphs nested in DSs contained significantly stronger triads of coactive nodes compared to SWR graphs, as indicated by higher mean clustering coefficients, during both pre- and post-exploration sleep sessions (F). Also, DSs and SWRs contained significantly stronger triads of coactive nodes during post- versus pre-exploration sleep sessions (G).

(H) Mean DS_1 reactivation pooled across mice (left panel) and in individual mice (right four panels; measured by the β coefficients of the linear regression that predicted post-exploration DS_1 coactivity from waking theta coactivity, controlling for pre-exploration DS_1 coactivity)

(I) As H but for DS_2 reactivation.

For F and G, the test statistic is the mean difference, shown on the y-axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). P-values are from paired permutation tests, event versus event (F), pre versus post (G), *$P < 0.05$, **$P < 0.01$, ***$P < 0.001$.

Figure S6. Closed-loop optogenetic suppression of dentate granule cells during DSs, related to Figures 5 and 6.

(A-C) During DS-triggered DG light-delivery for the optogenetic silencing of DG cells, there was no difference in the amplitude (A, B) or duration (C) of DSs in light-on versus light-off (control) DS events.

(D-I) DG light-delivery did not affect CA1 LFPs (D), ripple duration (E), intra-ripple frequency (F), ripple power (G), nor the probability of ripple occurrence (H-I; 22 recording sessions in 8 mice).

(J) DS-triggered DG light-delivery significantly reduced firing rates in DG, CA3 and CA1 principal cells relative to DSs with no light delivery (Control), in a paired analysis (same cells under both conditions). Each panel shows mean ± SEM firing rate. Average firing rates across the entire recording session (including sleep and exploration epochs) were: DG: 1.4±0.1 Hz, CA3: 1.7±0.3 Hz, CA1: 1.5±0.1 Hz.

(K) Estimation plot showing the maximum firing rate (during DS events) in DG, CA3 and CA1 principal cells during DS-Sync or Control (DSs with no light delivery), as in panel J. Upper: raw data points (each

point shows maximum firing rate), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event type. Lower: difference (Δ) in maximum firing rates between DS-Sync and no-light control condition (paired permutation test in DG: n=216 cells, CA3: n=50 cells, and CA1: n=133 cells from n=22 recording sessions in 8 mice).

(L) Estimation plot for object preference in a 'no-laser' control group of mice, showing significantly more time spent investigating the novel object (n=20 test sessions in 5 mice).

(M-O) In the continuous novel object recognition task, the total time spent exploring the objects (M), the number of laser pulses delivered (N) and the number of SWRs detected during sleep sessions (O) did not differ between the DS-delay and DS-sync conditions.

(P-S) The tone fear task in which mice had 5 tone-shock pairings during conditioning followed by either DS-Delayed or DS-Sync stimulation, and then fear memory recall. Mice froze more during recall than the baseline (Q), but this did not differ between the DS-Delay and DS-Sync groups (R; n=8 sessions in 4 mice). Mice received an equivalent number of laser pulses in the two groups (S).

(T) In the novel position recognition task, mice in the DS-Delayed and DS-Sync groups received equivalent numbers of laser pulses.

E-G, M-O, S,T show mean ± SEM. For B,C,I,K,L,Q,R, the test statistic is the mean difference, shown on the y-axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). P-values are from paired permutation tests, Control (no laser) versus DS-Sync (B,C,I,K,R), Novel versus Familiar (L), Baseline versus DS-Delay and DS-Sync (Q), $*P < 0.05$, $**P < 0.01$, $***P < 0.001$.

Table S1. Ratio of DG to CA neurons influence on the dimensionality and similarity of population firing vectors, related to Figure 3.

Event	Dependent variable	Independent variables	Degrees of freedom	r-value	p-value
DS	Dimensionality	Ratio of DG:CA neurons	33	0.24	0.18
DS1	Dimensionality	Ratio of DG:CA neurons	33	0.12	0.51
DS2	Dimensionality	Ratio of DG:CA neurons	33	0.26	0.14
SWR	Dimensionality	Ratio of DG:CA neurons	33	0.32	0.07
DS	Similarity	Ratio of DG:CA neurons	33	-0.06	0.75
DS1	Similarity	Ratio of DG:CA neurons	33	-0.13	0.46
DS2	Similarity	Ratio of DG:CA neurons	33	-0.08	0.66
SWR	Similarity	Ratio of DG:CA neurons	33	-0.39	0.02

Table S2. Linear mixed model analysis for SWR and DS reactivation of waking theta coactivity patterns, related to Figure 4.

Event	Dependent variable	Independent variables	No. Observation
SWR	Post-exploration	Theta coactivity Pre-exploration	7310
SWR	Pre-exploration	Theta coactivity Post-exploration	7310
DS	Post-exploration	Theta coactivity Pre-exploration	7310
DS	Pre-exploration	Theta coactivity Post-exploration	7310

Event	Dependent variable	Independent variables	β coefficient	CI (95%)	z	Prob.
SWR	Post-exploration	Theta coactivity	0.37	[0.32, 0.43]	13.1	P < 0.0001
		Pre-exploration	0.11	[0.09, 0.13]	9.5	P < 0.0001
SWR	Pre-exploration	Theta coactivity	0.20	[0.15, 0.26]	7.0	P < 0.0001
		Post-exploration	0.11	[0.04, 0.08]	9.5	P < 0.0001
DS	Post-exploration	Theta coactivity	0.39	[0.31, 0.46]	9.7	P < 0.0001
		Pre-exploration	0.25	[0.23, 0.27]	23.3	P < 0.0001
DS	Pre-exploration	Theta coactivity	0.24	[0.16, 0.32]	5.8	P < 0.0001
		Post-exploration	0.27	[0.25, 0.30]	23.3	P < 0.0001

Table S3. Counts of principal cells by session and mouse, related to Figures 1-6.

Figure number / panel	# sessions	# mice	# principal cells (mean ± SEM per mouse)		
			DG	CA3	CA1
Figure. 1F,G	73	12	921 (76.5±22.2 per mouse)	388 (32.3±9.6 per mouse)	887 (73.6±17.9 per mouse)
Figure. 2A-C	8	3	n/a	n/a	n/a
Figure. 2E-G	73	12	921 (76.5±22.2 per mouse)	388 (32.3±9.6 per mouse)	887 (73.6±17.9 per mouse)
Figure. 3B-D,F-G,K,L	34	8	647 (85.6±28.9 per mouse)	169 (25.6±9.1 per mouse)	449 (77.9±23.0 per mouse)
Figure. 4B-C	9	4	114 (28.5±7.2 per mouse)	5 (1.25±1.1 per mouse)	232 (58.0±18.0 per mouse)
Figure. 5F-G	43	9	548 (60.9±11.8 per mouse)	n/a	n/a
Figure. 5I-J	13	3	181 (60.3±14.0 per mouse)	n/a	n/a
Figure. 6H (DS)	10	3	79 (26.3±10.3 per mouse)	22 (7.3±3.8 per mouse)	73 (24.3±4.0 per mouse)
Figure. 6H (SWR)	12	3	173 (57.7±15.7 per mouse)	32 (10.7±6.4 per mouse)	68 (22.7±8.1 per mouse)

Offline hippocampal reactivation during dentate spikes supports flexible memory

Highlights

- Neurons across hippocampal regions synchronize during offline dentate spike events

- Hippocampal population activity differs in dentate spikes versus sharp-wave ripples

- Dentate spikes nest higher-dimensional patterns of hippocampal population activity

- Waking coactivity patterns reactivate in offline DS events to support flexible memory

Authors

Stephen B. McHugh,
Vítor Lopes-dos-Santos,
Manfredi Castelli, ..., Timothy Denison,
Andrew Sharott, David Dupret

Correspondence

stephen.mchugh@bndu.ox.ac.uk
(S.B.M.),
david.dupret@bndu.ox.ac.uk (D.D.)

In brief

The offline reactivation of hippocampal neuronal activity during sleep is thought to be crucial for memory. We investigated the contribution of dentate spikes to offline reactivation of hippocampal activity patterns and flexible recognition memory by combining triple-ensemble (DG-CA3-CA1) recordings and closed-loop optogenetic feedback in mice undergoing sleep/rest following active exploratory behavior. This work establishes dentate spikes as a second network event that, together with sharp-wave ripples, supports offline hippocampal dynamics for memory-guided behavior.

CellPress

Neuron

Article

Offline hippocampal reactivation during dentate spikes supports flexible memory

Stephen B. McHugh,[1,*] Vítor Lopes-dos-Santos,[1] Manfredi Castelli,[1] Giuseppe P. Gava,[1] Sophie E. Thompson,[1] Shu K.E. Tam,[1] Katja Hartwich,[1] Brook Perry,[1] Robert Toth,[1] Timothy Denison,[1] Andrew Sharott,[1] and David Dupret[1,2,*]
[1]Medical Research Council Brain Network Dynamics Unit, Nuffield Department of Clinical Neurosciences, University of Oxford, Oxford OX1 3TH, UK
[2]Lead contact
*Correspondence: stephen.mchugh@bndu.ox.ac.uk (S.B.M.), david.dupret@bndu.ox.ac.uk (D.D.)
https://doi.org/10.1016/j.neuron.2024.08.022

SUMMARY

Stabilizing new memories requires coordinated neuronal spiking activity during sleep. Hippocampal sharp-wave ripples (SWRs) in the cornu ammonis (CA) region and dentate spikes (DSs) in the dentate gyrus (DG) are prime candidate network events for supporting this offline process. SWRs have been studied extensively, but the contribution of DSs remains unclear. By combining triple-ensemble (DG-CA3-CA1) recordings and closed-loop optogenetics in mice, we show that, like SWRs, DSs synchronize spiking across DG and CA principal cells to reactivate population-level patterns of neuronal coactivity expressed during prior waking experience. Notably, the population coactivity structure in DSs is more diverse and higher dimensional than that seen during SWRs. Importantly, suppressing DG granule cell spiking selectively during DSs impairs subsequent flexible memory performance during multi-object recognition tasks and associated hippocampal patterns of neuronal coactivity. We conclude that DSs constitute a second offline network event central to hippocampal population dynamics serving memory-guided behavior.

INTRODUCTION

Memories are stabilized during periods of sleep and rest.[1–4] Decades of work have provided important insights into the underlying brain network mechanisms and have identified offline hippocampal activity as essential for this process.[5,6] Central to our current understanding are hippocampal sharp-wave ripples (SWRs) that feature an intermittent, high-frequency (100–250 Hz) network event detected in the local field potentials (LFPs) of the cornu ammonis (CA)1 region.[7–10] During SWRs, the firing activity of CA1 principal cells is transiently modulated[11,12] and reactivates the population-level firing patterns expressed in previous waking experience.[13] These offline spiking correlates have behavioral significance: suppressing CA1 neurons during SWRs impairs memory recall for recently acquired information.[14–16] Conversely, prolonging SWRs or reinforcing the coordination between SWRs and neocortical activity promotes memory consolidation and subsequent behavioral performance.[17,18] Hippocampal SWRs therefore constitute an offline network event important for memory-guided behavior. However, during sleep/rest periods, the hippocampus exhibits another prominent network event: dentate spikes (DSs), which are seen in the LFPs of the dentate gyrus (DG). To date, DSs have received little attention compared with SWRs. Accordingly, here we characterize the neuronal spiking dynamics nested in DSs with respect to SWRs and evaluate whether DSs constitute a second network

event central to offline reactivation of waking firing patterns and subsequent memory-guided behavior.

The DG gates sensory information to the hippocampus, notably decorrelating these inputs into dissimilar neural patterns.[19–21] This function may be crucial for the hippocampus to integrate multiple items in memory and to flexibly distinguish between stimuli with overlapping features. Within the DG, DSs represent intermittent, large-amplitude network events recorded in the LFPs of the DG granule cell layer and are associated with increased spiking activity in dentate cells.[22–24] However, across the literature, both increased and suppressed spiking activity of CA principal cells have been reported,[22,24–28] although, notably, some of these studies were in anesthetized[25,27] or head-fixed animals.[24,28] Thus, here we further performed a systematic comparative assessment of DG and CA principal cell spiking activity during DSs versus SWRs in non-anesthetized, freely behaving mice.

To investigate the influence of DSs on hippocampal population activity and memory, we combined triple-site (DG-CA3-CA1) extracellular multichannel recordings and closed-loop optogenetic interventions in mice during active exploratory behavior and offline sleep/rest. We observed that during offline DSs, principal cell spiking transiently increased across the DG and CA regions of the hippocampus, nesting offline population-level activity patterns that are distinct from those in SWRs. Further, we report that the cell-to-cell coactivity seen during prior waking experience is reactivated during DSs (as well as SWRs). DS-nested neuronal

activity is relevant to whole-hippocampus population dynamics and memory-guided behavior: closed-loop suppression of DG granule cell spiking selectively during offline DSs, but not SWRs, impairs subsequent flexible memory performance in hippocampal-dependent, multi-object recognition tasks. We propose that DSs constitute a second hippocampus network event that plays a complementary role to that of SWRs by supporting the offline reactivation of diverse population patterns of neuronal coactivity in support of memory-guided behavior.
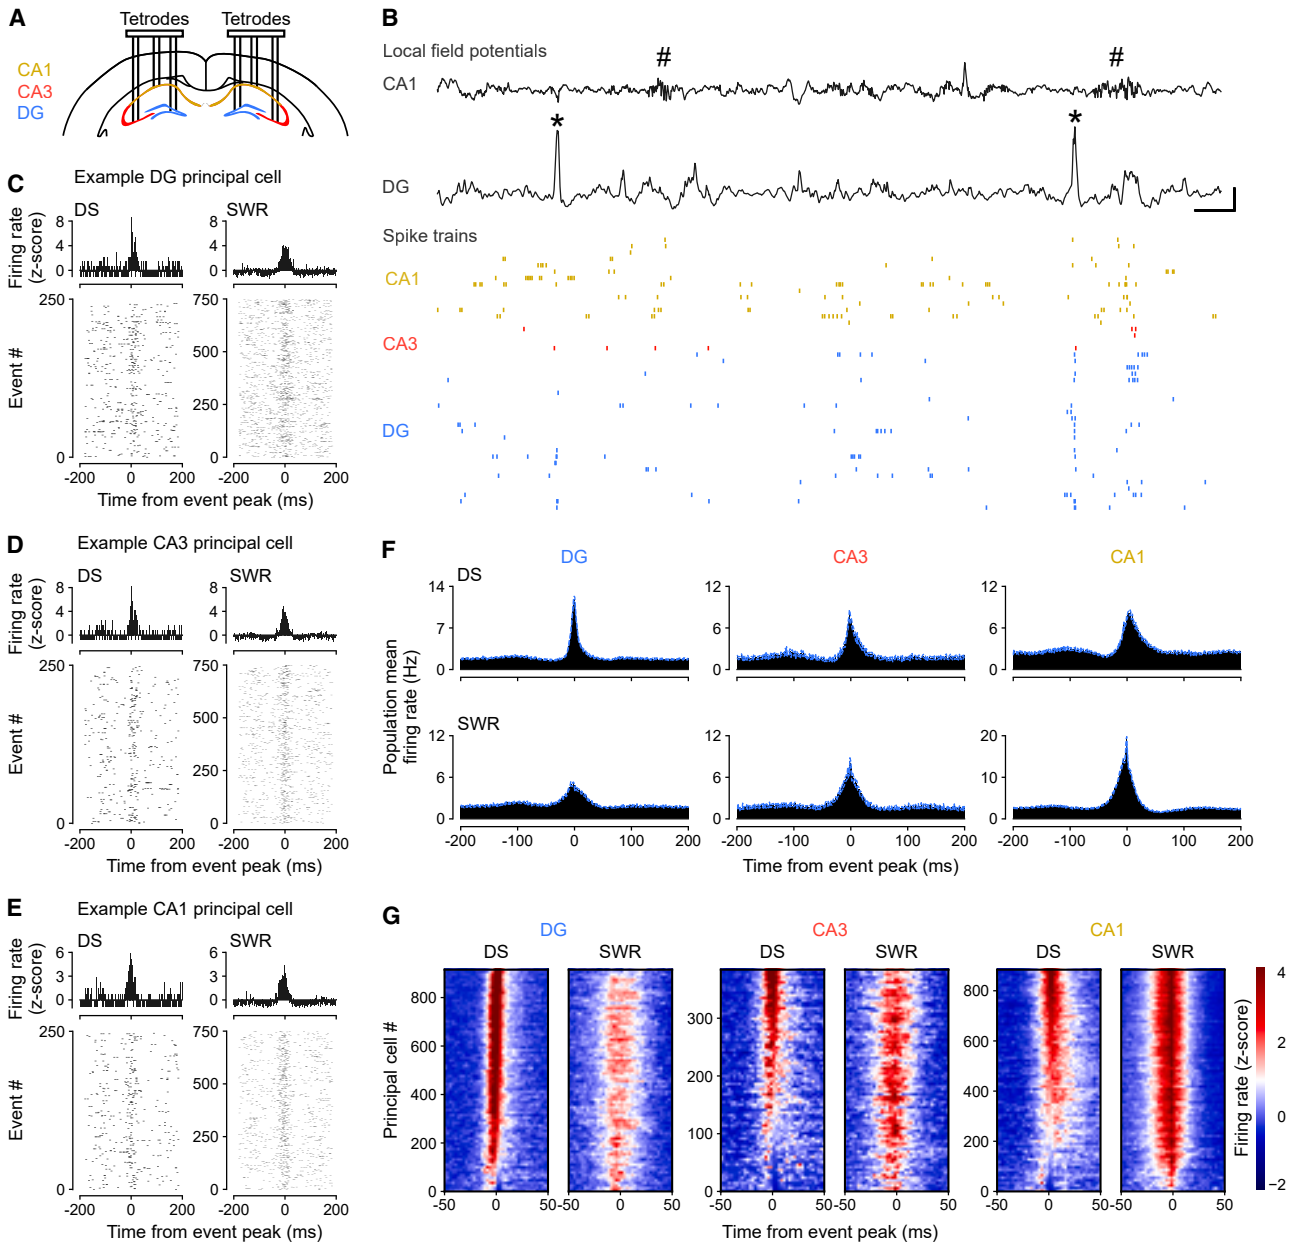
RESULTS

Firing activity of hippocampal neurons synchronizes during DS events

We first used triple-site (DG-CA3-CA1) tetrode recordings to monitor network events in the LFPs and the spike trains of neuronal ensembles from the dorsal hippocampus of mice during sleep/rest (Figures 1A, 1B, and S1A–S1C; $n = 12$ mice). From these LFPs, we detected DSs in DG and SWRs in CA1 to compare the spiking activity of principal cells between these two types of network events. Across mice and recording sessions, DS waveforms were highly consistent (Figures S1D and S1E). Both DG DSs and CA1 SWRs were of short duration (Figure S1F; median [interquartile range, IQR] duration: DSs = 42.4 [40.0–46.4] ms; SWRs = 47.5 [45.2–51.0] ms) and occurred intermittently (median [IQR] occurrence frequency: DSs = 0.40 [0.25–0.58] Hz; SWRs = 0.75 [0.26–1.26] Hz) during behavioral and LFP profiles indicative of sleep/rest (Figures S1G and S1H). These two network events rarely occurred simultaneously, consistent with previous reports,[22,29] with the vast majority of DSs not expressed within ±50 ms of a SWR (median [IQR]: 92.5% [87.9%–95.6%]; DS-SWR co-occurrence frequency: 0.03 [0.02–0.05] Hz; Figures S1I and S1J). We computed the firing responses of individual principal cells ($n = 2,196$ total recorded principal cells; CA1, 887; CA3, 388; DG, 921 cells; Figures S2A–S2F) with respect to the peak of either DSs or SWRs, excluding those temporal windows where both events co-occurred within ±50 ms. With the term "principal cells," we refer to CA pyramidal cells and DG granule cells that constitute the dominant (hence "principal") cell type in the hippocampus, exhibiting lower mean firing rates compared with local fast-spiking inhibitory cells (Figure S2E). In line with previous work, DG principal cells transiently increased their firing activity during DSs,[22,25] and the activity of CA principal cells increased during SWRs (Figures 1C–1G and S2G–S2P).[8,30] We further observed that DG principal cell firing increased during SWRs (Figures 1C, 1F, 1G, and S2G–S2P) and that CA principal cells also increased their firing rate during DSs (Figures 1D–1G and S2G–S2P), which contrasted with some earlier reports that CA principal cell firing is suppressed during DSs.[22,25]

To quantify the magnitude of neuronal activation during DSs and SWRs, we calculated the proportion of DG, CA3, and CA1 principal cells that increased their firing rate beyond a given significance threshold, using the Z scored peri-event time histograms obtained for each of these two network events (Figure 1G). During DSs, the majority of DG (91%), CA3 (56%), and CA1 (61%) principal cells increased their firing rate more than three standard deviations above baseline (Z score > 3, $p < 0.003$;

Figures S2G–S2I). During SWRs, a comparable proportion of principal cells significantly increased their firing activity beyond this threshold (Figures S2G–S2I). Hippocampal CA principal cells exhibited preferential activation during SWRs, whereas DG principal cells exhibited preferential activation during DSs (Figure S2I). DG principal cell population typically fired before CA principal cell populations during DSs (Figures S2J–S2N). While DG and CA principal cells exhibited such a temporal relationship, both DSs and SWRs were associated with an overall transient increase in hippocampal spiking activity (Figures S2O and S2P).

Previous studies identified two types of DS event (DS_1 and DS_2) based on the laminar profile of the transmembrane currents associated with the LFP expression of these network events.[22,24,26] Therefore, we next asked whether principal cell firing responses differed between DS_1 and DS_2. However, localizing sinks and sources of currents across hippocampal layers requires applying current source density (CSD) analysis[31] to the LFPs measured at evenly spaced sites from the CA1 oriens layer to the DG granule cell layer. Such a laminar profile is not accessible with tetrode recordings. To distinguish between DS_1 and DS_2 events, we therefore implanted linear silicon-probes spanning the somato-dendritic axis of CA1 principal cells and reaching the inferior blade of the DG in a separate group of mice ($n = 3$). Having performed silicon-probe recordings during sleep/rest, we applied CSD analysis to these LFPs measured over the radial extent of the hippocampus to identify DS_1 versus DS_2 according to their underlying profile of current sinks and sources (Figures 2A and S3A). These CSD-validated DS_1 and DS_2 events exhibited distinct DG granule cell layer LFP waveforms (Figures 2B and S3B). We then trained a linear discriminant analysis classifier to identify these CSD-validated DS_1 versus DS_2 events using only their DG granule cell layer LFP signal. When tested on the silicon-probe LFP dataset, the classifier achieved over 85% accuracy (Figure 2C). When next applied to the tetrode LFP dataset, the classifier-identified DS_1 and DS_2 events also exhibited distinct granule cell layer LFP waveforms (Figures 2D, S3C, and S3D), which were consistent with those obtained in silicon-probe recordings (Figures 2B and S3B). In both (tetrode and silicon-probe) datasets, DS_2 represented two-thirds of the DS events (median [IQR]: 66% [61%–73%]), thus constituting the dominant type. Leveraging this cross-dataset approach, we found that the firing response of DG and CA principal cells was stronger for DS_2 than DS_1 (Figures 2E–2G and S3E–S3J). A greater proportion of CA principal cells showed firing activity below baseline during DS_1 compared with DS_2 (35% versus 11%; Figures 2F, S3K, and S3L), providing insights into the previously documented DS-suppressed firing in some CA principal cells.[22,25,32] Nevertheless, the average activity of principal cells in DS_1 (and DS_2) was significantly higher than their baseline firing (calculated outside of any DS and SWR events) during sleep/rest (Figures 2G, S3G, and S3H) and their overall mean firing rate calculated over the whole recording day (Figures S3I and S3J). These results show that DS events (both DS_1 and DS_2) constitute transient network states that are qualitatively distinct from the sleep/rest epochs outside these events in terms of their capacity to increase spiking activity of individual principal cells distributed across hippocampal regions.

Neuron
Article

Figure 1. Dentate spikes recruit principal cell spiking across DG, CA3, and CA1

(A) Triple-ensemble (DG-CA3-CA1) tetrode recording allowed simultaneous monitoring of local field potentials (LFPs) and spiking activities.

(B) Upper: raw wide-band CA1 and DG LFP traces (black) showing sharp-wave ripples (SWRs, hash symbols) in CA1 and dentate spikes (DSs, asterisks) in DG. Scale bars, 100 ms (horizontal), 1.5 mV for DG, and 0.5 mV for CA1 (vertical). Lower: (color-coded) raster plot of spike trains from CA1 (orange), CA3 (red), and DG (blue) principal cells (PCs, one cell per row). Shown is a few second sample of recording for clarity.

(C–E) Spiking responses from single example DG (C), CA3 (D), and CA1 (E) principal cells. Upper: Z scored peri-event time histogram (PETH) during DSs (left) and SWRs (right). Lower: corresponding raster plot showing event-related spiking responses (one event per row).

(F) Group averaged firing rate PETHs for hippocampal PCs during DSs (top) and SWRs (bottom): DG ($n = 921$), CA3 ($n = 388$), and CA1 ($n = 887$) cells from 12 mice. Blue traces: mean ± SEM.

(G) Heatmaps showing Z scored firing rates for the DG, CA3, and CA1 PCs shown in (F). For each heatmap: one cell per row, sorted (top-to-bottom) from the most activated (highest Z score at event peak, 0 ms, red) to the least activated (lowest Z score at event peak, blue) during DSs.

Figure 2. Hippocampal principal cell firing is higher in DS_2 than DS_1

(A) Left: laminar (64-channel) silicon-probe recording allowed simultaneous monitoring of LFPs across hippocampal layers for current source density (CSD) analysis. Right: example (radially organized) mean LFP traces (gray) with superimposed CSD profile (heatmaps) for type 1 (DS_1) and type 2 (DS_2) dentate spikes and SWRs (calculated from 2,231 DS events and 8,693 SWR events in one mouse). Note the distinct CSD profiles reflecting the different transmembrane currents associated with DS_1 versus DS_2 versus SWR events. Hippocampal layers: oriens (*ori*); pyramidale (*pyr*); radiatum (*rad*); lacunosum-moleculare (*lm*); outer (*om*), middle (*mm*), and inner (*im*) moleculare; granulare (*gcl*); and hippocampal fissure (*hf*).

(B) Upper: shown for silicon-probe recorded DS_1 and DS_2 identified from their CSD profiles are example average granule cell layer LFP waveforms triggered by the peak of these events. Lower: in these recordings there was a higher proportion of DS_2 than DS_1 events (n = 15,067 events, 3 mice).

(C) Upper: we applied principal component analysis on the normalized granule cell layer LFP waveforms for all silicon-probe recorded DS events. We then used the principal components explaining 90% of the variance to train a linear discriminant classifier with the true labels (DS_1 versus DS_2) determined by the individual CSD profiles. Lower: the classifier performance (>85%) was significantly above chance level (50%) when tested on silicon-probe recorded LFP waveforms of unlabeled events. We used this classifier to next distinguish DS_1 and DS_2 from tetrode-recorded granule cell layer LFP waveforms (D).

(D) Upper: shown for tetrode-recorded DS events are the average granule cell layer LFP waveforms for DS_1 and DS_2 predicted label obtained from the silicon-probe-based classifier (C). Lower: these recordings also contained a higher proportion of DS_2 than DS_1 events (n = 32,215 events, 12 mice).

(E) Group averaged firing rate PETHs for tetrode-recorded DG, CA3, and CA1 principal cells during DS_1 and DS_2 (as Figures 1F and 1G). Blue traces: mean ± SEM.

(F) Heatmaps showing Z scored firing rates for the DG, CA3, and CA1 cells shown in (E). For each heatmap: one cell per row, sorted (top-to-bottom) from the most activated (highest Z score) to least activated (lowest Z score) during DS_1 peaks.

(G) Estimation plot showing the effect size for the differences in firing rate of DG, CA3, and CA1 principal cells during all DS events, DS_1 and DS_2 events analyzed separately, and SWRs compared with equivalent (50 ms duration matched) baseline windows (base) in which no DSs or SWRs occurred. Upper: raw data points (each point represents one cell), with the gapped lines on the right as mean (gap) ± SD (vertical ends) for each event. Lower: difference (Δ) in firing rate between baseline windows and all DS, DS_1, DS_2, and SWR events computed from 5,000 bootstrapped resamples and with the difference-axis origin (dashed line) aligned to the baseline rate (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). The test statistic is the mean difference, shown on the y axis of the lower plot. p values are from paired permutation tests, event versus baseline, ***p < 0.001.

(E and G) show data from n = 2,196 hippocampal principal cells (DG: n = 921, CA3: n = 388, and CA1: n = 887) from 12 mice.

DS events nest higher-dimensional patterns of population coactivity

We next investigated how the hippocampus organizes the collective activity of its principal cells both within individual DS events and across events, comparing these population-level patterns to those expressed in SWRs. To proceed, we first considered the neuron-wise vectors formed by the instantaneous spike discharge of principal cells in DSs, SWRs, or duration-matched control windows (without any DSs or SWRs) of the same sleep/rest (Figures 3A and S4A; "population vector analysis"). This was conducted for all sleep/rest epochs (both those recorded before and those after active exploratory behavior). Using the Gini index,[33,34] we noted a marked decrease of the mean population sparsity in the spiking vectors nested in DSs and SWRs compared with baseline periods of duration-matched control windows (Figure 3B), with equivalent population sparsity levels during DS_2 versus SWRs (Figures 3C, S4B, and S4C). A logistic regression classifier trained on a subset of these population firing vectors and iteratively tested on the remaining subset significantly distinguished between DS and SWR events but could not distinguish between their corresponding pre-event nor their post-event control epochs (Figure 3D). Successful classification was also obtained when using only DS_2 events, which matched SWRs in the mean population sparsity per event (Figure S4D). When evaluating the pairwise similarity of DS-nested population vectors versus those of SWR-nested vectors (Figure 3E), we further observed that DSs contained a higher diversity (i.e., lower similarity) of firing vectors compared with those

in SWRs, which in turn were more similar to one another (Figures 3F, 3G, and S4E–S4H).

This difference in population vector similarity suggested that DSs and SWRs differ with respect to their neuronal motifs of transient coactivation. By examining the topological organization of peer-to-peer firing associations, we indeed observed that DS events contain stronger motifs of coactive principal cells than SWRs. For each cell pair (i, j), we trained a generalized linear model to predict the spike discharge of neuron j from that of neuron i while accounting for the activity of the remaining peers (Figure 3A; "peer-to-peer coactivity analysis"). We performed this procedure separately for DS and SWR events, which returned for each type of network event a matrix of β regression weights that represented the coactivity structure of the population (Figure 3H). With these matrices, for both DS and SWR events, we constructed neuronal coactivity graphs (with no self-connections), where each node is a cell and the edge linking any two nodes represents the firing association of that cell pair (Figures 3I and 3J). This revealed that DS-based graphs contained stronger triads of coactive nodes compared with SWR graphs (Figures 3K, S4I, and S4J). This remained the case when directly comparing DS_2 and SWR events (Figure S4K) and when calculating the neuron-wise average coactivity strength (Figure S4L).

These findings showed that while both DS and SWR events synchronize hippocampal principal cells, population coactivity responses to DSs are more diverse. To further assess this, we applied principal component analysis to quantify and compare

---

Figure 3. The coactivity structure of population spiking differs between DSs and SWRs

(A) Analytical framework: the population-level coactivity structure was analyzed using population vectors of principal cell spiking transiently nested in individual SWRs, DSs, or duration matched (50 ms) baseline control windows. Scale bars show 20 ms and 0.5 mV for SWRs and 1 mV for DSs. For the analyses in (B)–(G), these population firing vectors were then binarized (for each cell: a non-zero spike count gives 1; or else 0). For the analyses in (H)–(K), we calculated the peer-to-peer coactivity, controlling for the overall population activity.
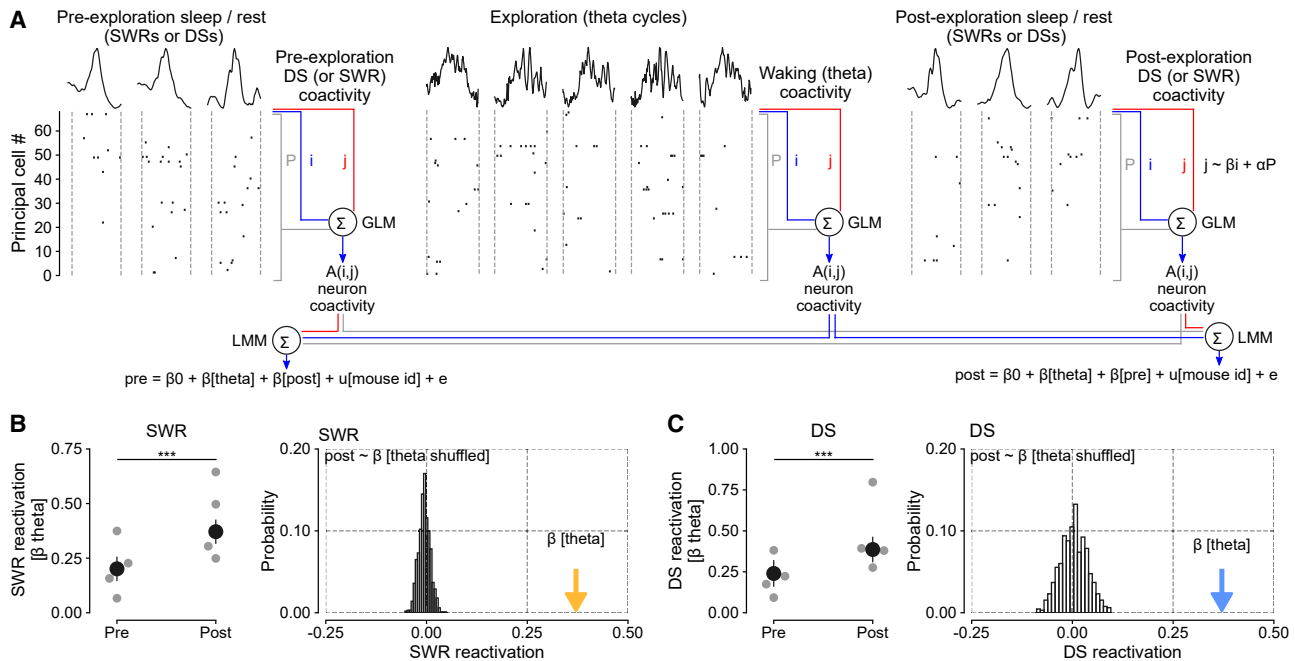
(B and C) Estimation plots showing the effect size for the differences in population sparsity (using the Gini index) during DSs (with DS_1 and DS_2 plotted altogether or separately), SWRs, and compared with equivalent (50 ms duration matched) baseline windows (baseline) in which no DSs or SWRs occurred. Upper: raw data points (each point represents one session with at least 100 of each event type and 20 principal cells), with the gapped lines on the right as mean (gap) ± SD (vertical ends) for each event. Lower: difference (Δ) in sparsity between baseline windows and all DS, DS_1, DS_2, and SWR events computed from 5,000 bootstrapped resamples and with the difference-axis origin (dashed line) aligned to the baseline sparsity (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). (C) as (B) but comparing population sparsity during SWR versus DS_2. Note that DS_2 and SWR events have equivalent sparsity, indicating they engage similar levels of neuronal activity.

(D) A logistic regression classifier trained on population vectors nested in SWR versus DS events or matched duration pre-event and post-event control windows, using a 4-fold cross-validation approach (75% of vectors for training; the remaining 25% for evaluation), significantly discriminated DSs from SWRs but could not discriminate between pre-DS versus pre-SWR and post-DS versus post-SWR vectors. Gray horizontal bars: mean classification accuracy.

(E–G) DS population firing vectors are more diverse than those in SWRs. For each sleep session, we computed the similarity (Pearson correlation coefficient) for each pair of population vectors nested in either DSs, SWRs, or duration-matched baseline windows (baseline). (E) shows example DS and SWR matrices of cross-vector similarity for one session. Cross-population vector similarity was significantly higher in SWRs compared with both DSs and control windows (F), and when compared with DS_1 and DS_2 separately (G).

(H–K) DS and SWR population firing vectors exhibit distinct topology of neuronal coactivity. The coactivity between any two (i, j) neurons was measured using a GLM that quantified their short timescale (50 ms windows centered on DS or SWR peaks) firing relationship while accounting for network-level modulation reported by the remaining principal cells in the population (A). (H) This procedure returned for both DS and SWR events an adjacency matrix of β regression weights that represented the neuron pairwise coactivity structure of the population (example matrix from one session). (I) Visualization of the corresponding matrices representing DS- and SWR-based neuronal coactivity graphs. For clarity, (J) shows an example subset (left) for each adjacency matrices shown in (H), along with its corresponding motifs of neuronal coactivity and average clustering coefficient (right). (K) Note that DS-based graphs contained stronger triads of coactive nodes compared with both SWR graphs and control graphs constructed from duration-matched baseline windows (baseline), as indicated by higher mean clustering coefficients. Each point in the upper plot of (K) represents the mean clustering coefficient for one hippocampal principal cell ($n = 1,265$ neurons, 8 mice).

(L) The dimensionality of population vector matrices (number of principal components required to explain 90% of the variance) was higher for DSs than SWRs. For (B)–(D), (F), (G), and (L), each data point shows one recording session ($n = 34$ recording sessions from 8 mice). The test statistic is the mean difference, shown on the y axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). p values are from paired permutation tests, event versus baseline (B, F, and K); event versus event (C, F, K, and L); or event versus pre-event, event versus post-event (D). *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$.

Neuron
Article

CellPress
OPEN ACCESS

Figure 4. Waking patterns of hippocampal coactivity reactivate in offline DSs

(A) DS and SWR reactivation of waking patterns formed by principal cell theta coactivity. For each cell pair (i,j), we predicted the spike discharge of neuron j from the activity of neuron i while regressing out the activity of the remaining population during pre-exploration sleep, exploration of open-field arenas, and post-exploration sleep (using GLMs as in Figure 3A). We separately applied this procedure for DSs and SWRs in both sleep/rest sessions (offline DS versus offline SWR coactivity) and across theta cycles in the exploration session (waking theta coactivity). This procedure returned a matrix of β regression weights that represented the neurons pairwise coactivity structure of the population in each session. We then used a linear mixed model (LMM) to compare the waking theta coactivity with post-exploration sleep coactivity (in DSs or SWRs) while controlling for pre-exploration sleep coactivity (in DSs or SWRs), and vice versa (reverse model). We included mouse identity as a random factor in each LMM.
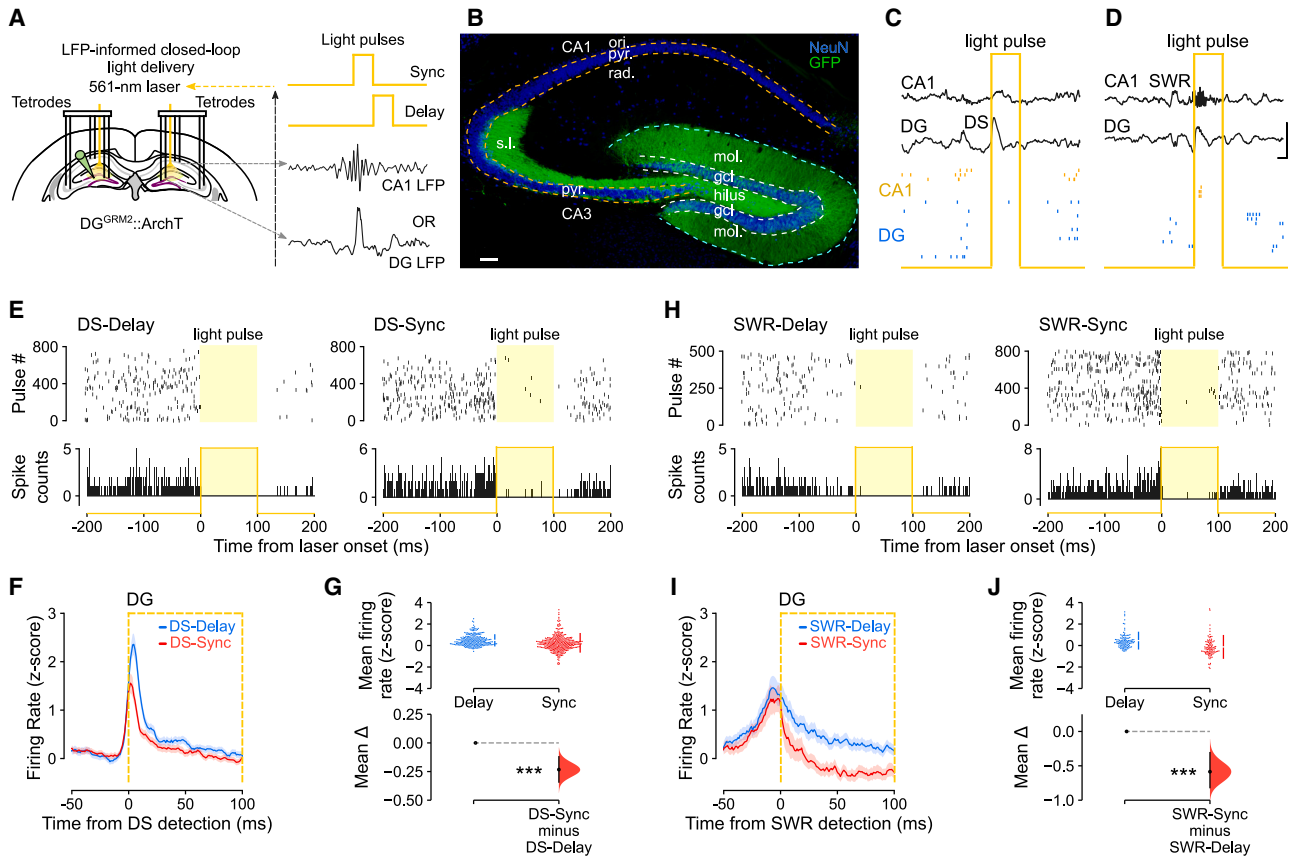
(B) SWR reactivation (measured by the β coefficients of the LMM that predicted post-exploration SWR coactivity from waking theta coactivity, controlling for pre-exploration SWR coactivity). Left: the β coefficient for theta coactivity was significantly higher when predicting post-exploration SWR coactivity than with the reverse model (i.e., predicting pre-exploration SWR coactivity from theta coactivity, controlling for post-exploration SWR coactivity). Gray points show the β coefficient for theta coactivity for individual mice. Error bars show ± 95% confidence interval. p value from t-test comparing post versus pre β coefficients: t(7,308) = 10.29; $p < 0.0001$. Right: the histogram shows the random probability distribution of β weights for theta coactivity when cell pair identity was shuffled (i.e., a null distribution based on 1,000 random shuffles; $n = 7,310$ cell pairs from 4 mice). The colored arrow shows the actual β coefficient for theta coactivity.

(C) DS reactivation exhibited the same pattern of results as SWR reactivation, shown in B. p value from t-test comparing post versus pre β coefficients t(7,308) = 8.84; $p < 0.0001$.

the variance explained by the activity patterns nested in DS_1, DS_2, and SWR. This revealed higher dimensionality of population vectors in DS events compared with SWRs (Figure 3L; Table S1). This was accounted for by DS_2 firing vectors, with those nested in DS_1 requiring fewer principal components to explain equivalent variance and exhibiting lower dimensionality than those in DS_2 and SWRs (Figures S4M–S4P).

Waking theta coactivity patterns reactivate in offline DSs and support flexible memory

The DS-nested motifs of peer-to-peer firing associations could instantiate population patterns of neuronal coactivity undergoing offline reactivation to support memory-guided behavior. Notably, the link between hippocampal SWRs and memory reactivation was initially established through the observation that the neural patterns of joint spiking activity expressed during exploratory behavior are more strongly correlated with those nested in post-exploration sleep/rest SWRs than those in SWRs before waking experience.[8,13] Accordingly, we next

determined whether DSs constitute another hippocampal time frame for offline reactivation of waking coactivity patterns. To proceed, we used our peer-to-peer coactivity analysis (Figure 3A), applying it to DS versus SWR events of sleep/rest before and after exploration of open-field arenas (Figures 4A and S5). Likewise, we obtained the waking patterns of population coactivity in theta cycles during exploration. With these, we computed DS and SWR reactivation by measuring the tendency of the peer-to-peer theta firing associations to reoccur in post-exploration sleep/rest DS (or SWR) events while controlling for prior pre-exploration DS (or SWR) coactivity and mouse identity using a linear mixed model. In line with previous work, offline patterns of SWR coactivity reflected those of theta coactivity significantly more during post-exploration than pre-exploration sleep/rest (Figures 4B, left, S5B, and S5D). This SWR reactivation was significantly higher than that obtained with a null distribution generated from models using randomly shuffled cell pair identities (Figure 4B, right). Importantly, we observed that theta coactivity patterns were also strongly reactivated in

Figure 5. DS- and SWR-informed offline suppressions of DG granule cell activity

(A) Triple-ensemble (DG-CA3-CA1) recording with LFP-informed yellow (561 nm) DG light delivery. Dentate granule cells (DGCs) transduced with ArchT-GFP (DG^{Grm2}::ArchT). Closed-loop light delivery to suppress DGC spiking immediately upon either DS detection (DS-sync condition) or SWR detection (SWR-sync) or their respective control conditions (DS-delay and SWR-delay, where light delivery was offset by 100 ms after event detection).

(B) ArchT-GFP-expressing DGCs in a DG^{Grm2}::ArchT mouse. Neuronal nuclei stained with NeuN. Scale bar, 100 μm. Granule cell layer, gcl; molecular layer, mol; pyramidal cell layer, pyr; stratum oriens, ori; radiatum, rad; and lucidum, s.l.

(C and D) Closed-loop feedback transiently silenced DGCs during either DG DS (C; "DS-sync") or CA1 SWR (D; "SWR-sync") events, illustrated with raw data examples. Scale bars, 30 ms (horizontal) and 1.5 mV (vertical).

(E) Raster plots (event-related spiking response; one light pulse per row, upper), and peri-event time histograms (lower) showing photo-silencing of two example DG cells from a DG^{Grm2}::ArchT mouse in DS-delay and DS-sync.

(F and G) Corresponding quantification of average DGC firing rate (Z score) for DS-delay versus DS-sync (F and G; $n = 548$ cells in 9 mice). In (F), the orange box shows the laser-on period for DS-sync.

(H–J) As (E)–(G) but showing DGC photo-silencing during SWR-delay and SWR-dync conditions (I and J; $n = 181$ cells in 3 mice). In (I), the orange box shows the laser-on period for SWR-sync.

For (F) and (I), the traces show mean ± SEM. For (G) and (J), the test statistic is the mean difference, shown on the y axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). p values are from unpaired permutation tests, delay versus sync, ***$p < 0.001$.

post-exploration DSs (Figures 4C, S5C, and S5E; Table S2). The clustering coefficient was higher in DSs than SWRs and also increased from pre- to post-exploration sleep in both SWRs and DSs (Figures S5F and S5G). Analyses of DS$_1$ and DS$_2$ separately showed evidence for reactivation during both types of post-exploration DS events (Figures S5H and S5I). By applying the same analytical framework to the neuronal ensembles tracked in SWRs and DSs, this result provided evidence for offline DS reactivation of hippocampal waking firing patterns.

The offline reactivation of waking population patterns in sleep/rest DSs (Figure 4C), which contain more diverse and higher-dimensional patterns of neuronal coactivation than those found

in SWRs (Figures 3F and 3L), raised the question of their network contribution to memory-guided behavior. We thus tested whether the offline population response during DSs was necessary to perform tasks that require integrating multiple items in memory to flexibly distinguish between familiar and novel stimuli. To this end, we transduced DG granule cells with the yellow (561 nm) light-driven optogenetic silencer Archaerhodopsin T (ArchT) in Grm2-Cre mice (Figures 5A and 5B). We then implanted these DG^{Grm2}::ArchT mice for triple-ensemble (DG-CA3-CA1) recordings combined with bilateral optic fibers for DG light delivery. In these experiments, DG light delivery was performed in a closed-loop manner during sleep/rest using the

real-time detection of either DSs (in the DG LFPs) or SWRs (in the CA1 LFPs; Figures 5A, 5C, and 5D; "DS-sync" or "SWR-sync" conditions). We also used a within-subject control paradigm whereby, on different days, light was not synchronized to but instead delivered after each DS or SWR had elapsed ("DS-delay" or "SWR-delay" conditions). DS-synchronized light delivery did not affect the amplitude of the DG or CA1 LFPs, nor CA1 ripple duration, occurrence probability, or power (Figures S6A–S6I). DS-synchronized light delivery significantly reduced spiking activity in DG neurons compared with when laser-onset was DS delayed (Figures 5E–5G) and also reduced spiking activity in CA principal cells (Figures S6J and S6K). Paired analysis of the firing rates of dentate granule cells during DS events without light delivery versus those with light deliver (DS-sync) also showed significantly reduced instantaneous spiking (15.0 ± 0.7 versus 10.4 ± 0.6 Hz). SWR-synchronized light delivery also significantly suppressed DG neuronal spiking (Figures 5H–5J).

We applied these closed-loop light-delivery approaches during interposed sleep/rest sessions in three behavioral tasks (Figures 6 and S6L–S6T). The first, hippocampal-dependent task required mice to recognize previously encountered (familiar) versus novel objects.[35] In this novel-object recognition task, mice repeatedly explored a square-walled arena containing four objects (Figures 6A and 6B). In the first session ("sampling"), mice encountered four distinct novel objects, each one placed beside a wall. On the subsequent sessions ("test"), one of the initially sampled objects was replaced by a different novel object so that the mouse could explore one completely novel object along with the three "familiar" objects from the previous session that day. These exploration (sampling and test) sessions alternated with sleep/rest sessions where mice received DG-targeted light delivery, either synchronized or delayed with respect to either DS or SWR onset, thus yielding four distinct experimental conditions (DS-delay, DS-sync, SWR-delay, and SWR-sync). In each test session n, we measured novelty preference using the proportion of time mice spent investigating the novel versus the familiar objects, thereby probing recognition memory for session $n − 1$. We found that novelty detection was not impaired in test sessions following sleep with DG granule cell suppression in either DS-delayed, SWR-delayed, or SWR-synchronized conditions: mice subsequently expressed a stronger preference for novel over familiar objects under these three conditions (Figures 6C and 6D). This novel-object preference was similar to that observed in control mice without any optogenetic intervention (Figure S6L). However, novel-object preference was absent in test sessions following DS-synchronized suppression (Figures 6C and 6D). The total object exploration time, number of laser pulses, and number of SWRs did not differ between the DS-synchronized and DS-delayed conditions (Figures S6M–S6O).

We also tested the offline DS-informed suppression of DG granule cells after tone fear conditioning as a non-hippocampal-dependent task.[36] Mice were trained with five tone-shock pairings, and following DS-synchronized or DS-delayed suppression, we evaluated fear memory by measuring freezing behavior during a recall session in which tones were played but no shocks were given. Compared with baseline freezing (measured during the first tone of training, before any shocks were given), mice exhibited higher (and equivalent) freezing levels during recall regardless of whether they had received DS-synchronized or DS-delayed suppression (Figures S6P–S6S).
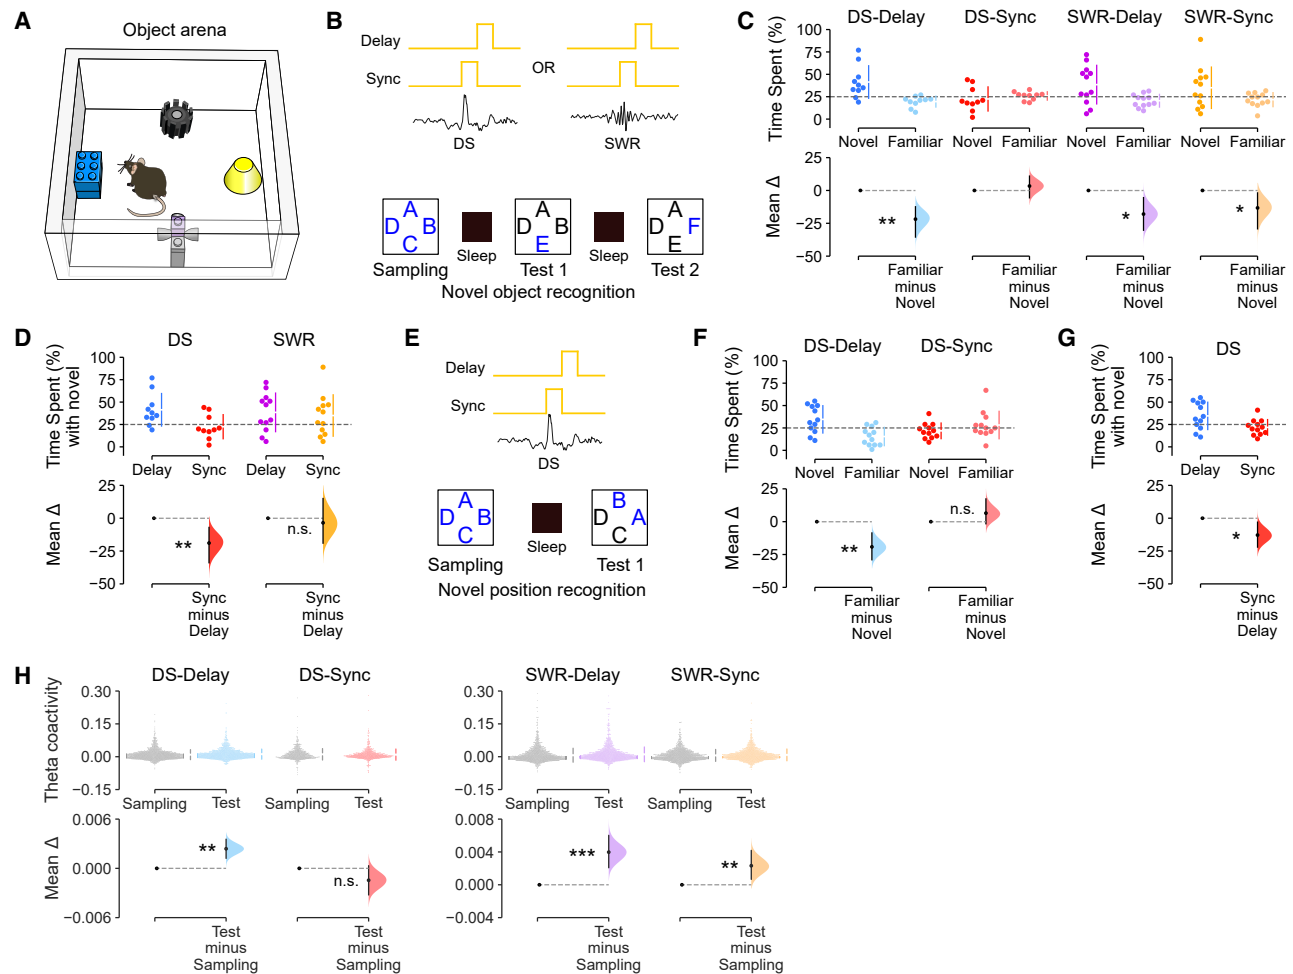
We finally tested whether DS-synchronized suppression affected performance in a novel-position recognition task that is reportedly more sensitive to DG than CA1 lesions, whereas novel-object recognition requires both DG and CA1.[37] This novel-position task is similar to the novel-object recognition task in that mice explore four novel objects during the sampling phase (Figure 6E). However, rather than introducing a new object in the test phase, the locations of two of the initially sampled objects are swapped for the subsequent session, leaving the other two objects in their original locations (Figure 6E). We found that mice preferentially explored the novel-positioned objects following DS-delayed suppression of DG granule cells but showed no such preference following DS-synchronized suppression (Figures 6F, 6G, and S6T).

Recent work has identified that the continued integration of new items in memory is associated with increased neuronal coactivity patterns nested in hippocampal theta oscillations.[38] In line with this, we found that the preserved object recognition memory observed after offline DG cell suppression in the DS-delayed, SWR-delayed, and SWR-synchronized conditions was accompanied by stronger theta coactivity (Figure 6H). This was not the case following DS-synchronized suppression (Figure 6H), indicating that DS silencing disrupts the integration of recently experienced information. Collectively, these results show that the hippocampal population response to offline DS events is required for flexible, memory-based recognition of previously encountered items and associated network gain in theta coactivity.

DISCUSSION

Our findings establish that offline DSs activate neurons across the DG and CA regions. DSs are therefore a second hippocampal network event that hosts short timescale coactivation forming population-level neural patterns, like the well-established SWRs. However, the activity structure and neuronal content are distinct in DSs. Notably, we found that DSs nest stronger motifs of coactive neurons, yielding population patterns of higher diversity and dimensionality compared with those in SWRs. Like SWRs, DSs reactivate hippocampal population patterns expressed in prior waking experience. This offline reactivation is behaviorally significant: closed-loop suppression of DG granule cell spiking selectively during offline DS events is sufficient to disrupt downstream CA principal cell activity and impair flexible recognition memory for previously encountered items, as well as the associated network gain in theta-nested neuronal coactivity. Collectively, these findings identify a core contribution for DSs to hippocampal patterns of population activity and memory-guided behavior.

We started this investigation by observing that DSs increase spiking activity in principal cells across the DG, CA3, and CA1 regions of the hippocampus. This finding is consistent with previous reports of DS-evoked spiking activity in DG granule cells but contrasts with some earlier reports of DS-suppressed CA pyramidal cell spiking.[22,24–26] Notably, Bragin and colleagues

Figure 6. Offline suppression of DS activity impairs flexible recognition memory

(A) Behavioral arena used for the recognition memory tasks.

(B–D) Offline DS events are required for novel-object recognition memory. (B) Task layout. During sleep sessions (interposed between novel-object exploration sessions), closed-loop optogenetic suppression of DG cells in DG^{Grm2}::ArchT mice was achieved using real-time monitoring of either DG or CA1 LFPs to actuate either DS synchronized (DS-sync) or delayed (DS-delay), SWR synchronized (SWR-sync), or delayed (SWR-delay) DG light delivery. Letters depict object locations in the task arena (A), with novel objects in blue. (C) Estimation plot showing the percentage of time spent by mice with the novel versus the familiar objects in each of the four LFP-informed closed-loop conditions. Upper: each data point represents the percentage time spent with the novel object versus the mean percentage time spent with the three familiar objects; chance performance is shown by the dashed line. Lower: mean difference between novel and familiar object exploration time. (D) as (C), but directly comparing novel-object preference in the delay versus sync conditions for DS and SWR events. Mice in the DS-delay, SWR-delay, and SWR-sync conditions, but not the DS-sync condition, exhibited a significant preference for novel over familiar objects (DS-delay and DS-sync: $n = 10$ sessions, in 3 mice; SWR-delay and SWR-sync: $n = 12$ sessions in 3 mice).

(E–G) Likewise, offline DS events are required for novel-position recognition memory. (E) Task layout. Letters depict object locations, with novel positions in blue. (F) Estimation plot showing the percentage of time spent by DG^{Grm2}::ArchT mice with the novel versus the familiar object locations following sleep sessions with DS-sync or DS-delay suppression of DG cells. Upper: each data point represents the percentage time spent with objects in novel locations versus objects in familiar locations; chance performance is shown by the dashed line. Lower: mean difference between novel and familiar location exploration times. (G) As (F) but directly comparing novel location preference in DS-delay versus DS-sync. Mice in the DS-delay but not the DS-sync condition exhibited a significant preference for objects in novel over familiar locations ($n = 12$ novel versus $n = 12$ familiar locations, 6 sessions, in 4 mice).

(H) In the object recognition task, the theta peer-to-peer coactivity increased from the initial object sampling to the memory test following offline DG cell suppression in the DS-delay, SWR-delay, and SWR-sync conditions, but this was not the case in the DS-sync condition (where mice exhibited no novel-object preference). Paired estimation plot showing theta coactivity during sampling versus test. Upper: each point represents a beta coefficient for the theta-nested peer-to-peer coactivity between pairs of hippocampal principal cells ($n = 1,537$, $n = 678$, $n = 1,719$, and $n = 1,482$ cell pairs, respectively, in 6 mice). Lower: black dot, mean difference between sampling and test sessions; black ticks, 95% confidence interval.

For (C) and (D) and (F)–(H), the test statistic is the mean difference, shown on the y axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). p values are from paired permutation tests, familiar versus novel (C and F); delay versus sync (D and G); or test versus sampling (H), $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

Neuron
Article

CellPress
OPEN ACCESS

reported suppressed spiking in 3/14 CA3 principal cells and suppressed CA1 multi-unit activity in 2/10 rats, showing some, rather than consistent, CA suppression. In addition, Penttonen et al. reported suppressed CA1 multi-unit activity and hyperpolarization of 4 intracellularly recorded CA1 neurons during DSs in anesthetized rats. However, DS rates are ~10-fold lower and have smaller amplitude during anesthesia compared with DSs observed during natural sleep/rest.[25] Other studies reported increased CA1 multi-unit activity[26] and increased CA3 single-unit spiking during DS events[24] in non-anesthetized mice. Here, we report a variety of firing responses across individual CA neurons, ranging from strong activation to suppression during DS events (Figures 1G and 2F). However, our systematic study (including > 3,500 principal cells; Table S3) shows that DSs do indeed drive increased mean population spiking activity in both DG and CA principal cells (Figure 2G).

Previous studies distinguished between two types of DS events (DS_1 and DS_2) based on the laminar profile of their transmembrane currents.[22,24,26] Here, we found that DS_2 are more effective than DS_1 events at recruiting hippocampal principal cells, with higher spike rates per cell and more coactive cells per event. This result is consistent with a recent report that DS_2 but not DS_1 events reliably increase spiking in DG and CA3 principal cells.[24] The same report saw only a slight increase in CA1 spiking during DS_2 and no effect of DS_1 events on either CA1 or CA3 principal cells. They also found that DG principal cell spiking was suppressed during DS_1, in contrast to our study and previous reports.[22] Here we found that both DS_1 and DS_2 events evoked significantly increased spiking activity in DG, CA3, and CA1 principal cells, but again would emphasize the diversity of CA cell responses, especially during DS_1 events (Figure 2F). The observed differences between studies might reflect differences between DS events in head-fixed awake mice versus those in sleep/rest.[22,29] They could also indicate a moment-to-moment diversity across individual DSs, similar to that highlighted for individual SWRs[39] and theta cycles.[40]

In this study, we directly compared patterns of population spiking activity in DSs versus SWRs. We found that population firing differed across these two types of network events, allowing a classifier to distinguish DSs versus SWRs based on their instantaneous vectors of principal cell spiking (Figure 3D). We also observed that population patterns in DSs are overall more diverse (less correlated) than those in SWRs, showing stronger triads of coactive neurons and higher dimensionality (Figures 3 and S4). DS_1 and DS_2 population patterns yet showed distinct trends with respect to SWRs: DS_1 firing vectors exhibited less diversity (i.e., required fewer principal components to explain most of the variance; Figure S4O) and lower dimensionality (Figure S4P) than those in SWRs; DS_2 firing vectors showed the opposite trend. Previous work has reported that DS_1 and DS_2 events relate to distinct entorhinal cortex inputs, with DS_1 relying more on the lateral entorhinal cortex while DS_2 events rely more on the medial entorhinal cortex.[24] This suggests a possible division of mnemonic labor where DS_1 population patterns would favor non-spatial information streams while DS_2 might favor spatial information.[41] We also found that waking patterns of neuronal coactivity nested in theta oscillations reactivate in DSs of post-exploration sleep/rest (Figures 4 and S5). Notably,

the distributions of coactivity values in DSs indicate both positive and negative firing associations (Figure S5E). The coexistence of correlated and anti-correlated spiking activities in DSs could reflect a Hebbian learning rule as reported in SWRs,[42] whereby positive and negative changes in hippocampal principal cell firing associations can shape offline DS reactivation as a function of recent waking experience. These findings provide important evidence for offline reactivation of hippocampal waking firing patterns outside of SWRs, stimulating new avenues for future work to explore.

To determine whether spiking activity observed during DS events was required for subsequent memory-guided behavior, we deployed a closed-loop optogenetic feedback approach to suppress DG granule cell activity selectively during DS events (Figures 5 and S6). Real-time inhibition of the DG in Grm2-Cre mice did not yield a complete suppression of the spiking activity in dentate granule cells. This also did not alter the magnitude of DSs, which powerful expression in the DG LFPs could reflect the high cellular density of the granule cell layer and its strong neural inputs. This DS-synchronized suppression of DG principal cells reduced concomitant spiking activity in CA principal cells but did not affect the expression of SWRs in CA1. When applied in sleep/rest following object-location exploration, this DS-synchronized neural suppression impaired subsequent memory performance in both novel-object and novel-position recognition tasks. Although ours is the first study to leverage a closed-loop optogenetic approach, these findings are consistent with previous behavioral studies using electrical stimulation to disrupt hippocampal activity during DSs.[32,43,44] While both approaches provide strong evidence for a central contribution of DS events in memory-guided behavior, it is important to recognize that optogenetic and electrical interventions do not recapitulate natural hippocampal activity patterns. Moreover, DS_2 represents the dominant type of DS event that exhibits, in comparison with DS_1, stronger firing rate increase of DG, CA3, and CA1 principal cells (Figure 2G), neuronal recruitment (Figures S3E and S3F), coactivity (Figure S4L), and lower dimensionality (Figures S4O and S4P). The major effect of DS activity on memory may be associated with DS_2, a hypothesis that future technological (closed-loop controller) development for differential manipulation of DS_1 versus DS_2 would be able to test. We also found that increased theta coactivity was associated with recognition memory and that this network gain in theta coactivity was absent following DS-synchronized neural suppression (Figure 6H). This finding is consistent with recent work showing that continual integration of new memory items across behavioral experiences increases neuronal coactivity.[38] Altogether, these results support the idea that neuronal activity during DS events plays an important role in subsequent memory-guided behavior, as SWRs do.

Why does the hippocampus use more than one offline network mechanism to support memory? DSs and SWRs are driven by distinct neural circuits. SWRs depend on excitatory inputs from CA3 to the CA1 stratum radiatum, generating high-frequency ripples in the CA1 pyramidal layer.[7,45–47] DSs are non-oscillatory events associated with excitatory inputs from the entorhinal cortex to the DG molecular layers.[22,24,29] Notably, entorhinal cortex lesions eliminate DSs but increase SWR incidence.[22] Our

CellPress
OPEN ACCESS

Neuron
Article

structural analysis of DS versus SWR population patterns raises the intriguing possibility that SWRs may be more suited for lower-dimensional network coactivity serving robust information flow, whereas DSs may promote higher-dimensional activity, allowing diverse mnemonic patterns to coexist offline and support flexible, pattern separation for subsequent behavior. Collectively, these findings open important new avenues for future work to explore the interplay between DS versus SWR events as two distinct timeframes for the hippocampus to optimize offline computations serving memory-guided behavior.

RESOURCE AVAILABILITY

Lead contact
Further information and requests for resources and reagents should be directed to and will be fulfilled by the lead contact, David Dupret (david.dupret@bndu.ox.ac.uk).

Materials availability
This study did not generate new unique reagents.

Data and code availability
- The electrophysiology dataset reported in this study is being used in on-going projects and can be accessed under a data transfer agreement. We welcome enquiries for sharing it; please contact david.dupret@bndu.ox.ac.uk.
- This paper does not report original code.
- Any additional information required to reanalyze the data reported in this work paper is available from the lead contact upon request.

AUTHOR CONTRIBUTIONS

Conceptualization, S.B.M. and D.D.; methodology, S.B.M., V.L-.d.-S., M.C., G.P.G., S.E.T., and D.D.; formal analysis, S.B.M.; investigation, S.B.M. and S.K.E.T.; resources, S.B.M., V.L-.d.-S., G.P.G., K.H., B.P., R.T., T.D., A.S., and D.D.; writing—original draft, S.B.M. and D.D.; writing—reviewing and editing, S.B.M., V.L-.d.-S., G.P.G., M.C., K.H., S.K.E.T., B.P., R.T., T.D., A.S., and D.D.; visualization, S.B.M. and D.D.; supervision, D.D.; funding acquisition, S.B.M. and D.D.

DECLARATION OF INTERESTS

The authors declare no competing interests.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS
  - Animals
- METHOD DETAILS
  - Viral vectors
  - Surgical procedures
  - Recording procedures
  - Light delivery
  - Multichannel data acquisition
  - Spike sorting and unit isolation
  - Principal cell versus interneuron classification
  - Local field potential signals
  - Dentate spike detection
  - Sharp-wave ripple detection
  - Place maps
  - Spatial Information
  - Peri-event time histograms (PETHs)
  - Linear discriminant analysis classifier
  - Clustering coefficient
  - Single-neuron coactivity strength
  - Population vector similarity
  - Population dimensionality
  - Theta-cycle detection
  - Reactivation of waking coactivity patterns
  - Closed-loop optogenetic interventions
  - Recognition memory tasks
  - Tone fear conditioning task
  - Tissue processing and immunohistochemistry
- QUANTIFICATION AND STATISTICAL ANALYSIS

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.neuron.2024.08.022.

REFERENCES

1. Maquet, P. (2001). The Role of Sleep in Learning and Memory. Science *294*, 1048–1052. https://doi.org/10.1126/science.1062856.

2. Walker, M.P., and Stickgold, R. (2006). Sleep, Memory, and Plasticity. Annu. Rev. Psychol. *57*, 139–166. https://doi.org/10.1146/annurev.psych.56.091103.070307.

3. Klinzing, J.G., Niethard, N., and Born, J. (2019). Mechanisms of systems memory consolidation during sleep. Nat. Neurosci. *22*, 1598–1610. https://doi.org/10.1038/s41593-019-0467-3.

4. Brodt, S., Inostroza, M., Niethard, N., and Born, J. (2023). Sleep—A brain-state serving systems memory consolidation. Neuron *111*, 1050–1075. https://doi.org/10.1016/j.neuron.2023.03.005.

5. Buzsáki, G. (2010). Neural Syntax: Cell Assemblies, Synapsembles, and Readers. Neuron *68*, 362–385. https://doi.org/10.1016/j.neuron.2010.09.023.

6. Girardeau, G., and Lopes-dos-Santos, V. (2021). Brain neural patterns and the memory function of sleep. Science *374*, 560–564. https://doi.org/10.1126/science.abi8370.

7. Buzsáki, G. (1986). Hippocampal sharp waves: Their origin and significance. Brain Res. *398*, 242–252. https://doi.org/10.1016/0006-8993(86)91483-6.

8. Buzsáki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. Hippocampus *25*, 1073–1188. https://doi.org/10.1002/hipo.22488.

9. Axmacher, N., Elger, C.E., and Fell, J. (2008). Ripples in the medial temporal lobe are relevant for human memory consolidation. Brain *131*, 1806–1817. https://doi.org/10.1093/brain/awn103.

10. Joo, H.R., and Frank, L.M. (2018). The hippocampal sharp wave–ripple in memory retrieval for immediate use and consolidation. Nat. Rev. Neurosci. *19*, 744–757. https://doi.org/10.1038/s41583-018-0077-1.

Neuron
Article

CellPress
OPEN ACCESS

11. Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., and Buzsáki, G. (1999). Fast network oscillations in the hippocampal CA1 region of the behaving rat. J. Neurosci. *19*, RC20. https://doi.org/10.1523/JNEUROSCI.19-16-j0001.1999.

12. Csicsvari, J., Hirase, H., Mamiya, A., and Buzsáki, G. (2000). Ensemble Patterns of Hippocampal CA3-CA1 Neurons during Sharp Wave–Associated Population Events. Neuron *28*, 585–594. https://doi.org/10.1016/S0896-6273(00)00135-5.

13. Wilson, M.A., and McNaughton, B.L. (1994). Reactivation of Hippocampal Ensemble Memories During Sleep. Science *265*, 676–679. https://doi.org/10.1126/science.8036517.

14. Girardeau, G., Benchenane, K., Wiener, S.I., Buzsáki, G., and Zugaro, M.B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. Nat. Neurosci. *12*, 1222–1223. https://doi.org/10.1038/nn.2384.

15. Ego-Stengel, V., and Wilson, M.A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. Hippocampus *20*, 1–10. https://doi.org/10.1002/hipo.20707.

16. van de Ven, G.M., Trouche, S., McNamara, C.G., Allen, K., and Dupret, D. (2016). Hippocampal Offline Reactivation Consolidates Recently Formed Cell Assembly Patterns during Sharp Wave-Ripples. Neuron *92*, 968–974. https://doi.org/10.1016/j.neuron.2016.10.020.

17. Maingret, N., Girardeau, G., Todorova, R., Goutierre, M., and Zugaro, M. (2016). Hippocampo-cortical coupling mediates memory consolidation during sleep. Nat. Neurosci. *19*, 959–964. https://doi.org/10.1038/nn.4304.

18. Fernández-Ruiz, A., Oliva, A., Fermino de Oliveira, E., Rocha-Almeida, F., Tingley, D., and Buzsáki, G. (2019). Long-duration hippocampal sharp wave ripples improve memory. Science *364*, 1082–1086. https://doi.org/10.1126/science.aax0758.

19. Treves, A., and Rolls, E.T. (1994). Computational analysis of the role of the hippocampus in memory. Hippocampus *4*, 374–391. https://doi.org/10.1002/hipo.450040319.

20. McHugh, T.J., Jones, M.W., Quinn, J.J., Balthasar, N., Coppari, R., Elmquist, J.K., Lowell, B.B., Fanselow, M.S., Wilson, M.A., and Tonegawa, S. (2007). Dentate Gyrus NMDA Receptors Mediate Rapid Pattern Separation in the Hippocampal Network. Science *317*, 94–99. https://doi.org/10.1126/science.1140263.

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

22. Bragin, A., Jandó, G., Nádasdy, Z., van Landeghem, M., and Buzsáki, G. (1995). Dentate EEG spikes and associated interneuronal population bursts in the hippocampal hilar region of the rat. J. Neurophysiol. *73*, 1691–1705. https://doi.org/10.1152/jn.1995.73.4.1691.

23. Senzai, Y., and Buzsáki, G. (2017). Physiological Properties and Behavioral Correlates of Hippocampal Granule Cells and Mossy Cells. Neuron *93*, 691–704.e5. https://doi.org/10.1016/j.neuron.2016.12.011.

24. Dvorak, D., Chung, A., Park, E.H., and Fenton, A.A. (2021). Dentate spikes and external control of hippocampal function. Cell Rep. *36*, 109497. https://doi.org/10.1016/j.celrep.2021.109497.

25. Penttonen, M., Kamondi, A., Sik, A., Acsády, L., and Buzsáki, G. (1997). Feed-forward and feed-back activation of the dentate gyrus in vivo during dentate spikes and sharp wave bursts. Hippocampus *7*, 437–450. https://doi.org/10.1002/(SICI)1098-1063(1997)7:4<437::AID-HIPO9>3.0.CO;2-F.

26. Meier, K., Merseburg, A., Isbrandt, D., Marguet, S.L., and Morellini, F. (2020). Dentate Gyrus Sharp Waves, a Local Field Potential Correlate of Learning in the Dentate Gyrus of Mice. J. Neurosci. *40*, 7105–7118. https://doi.org/10.1523/JNEUROSCI.2275-19.2020.

27. Sanchez-Aguilera, A., Wheeler, D.W., Jurado-Parras, T., Valero, M., Nokia, M.S., Cid, E., Fernandez-Lamo, I., Sutton, N., García-Rincón, D., de la Prida, L.M., and Ascoli, G.A. (2021). An update to Hippocampome.org

28. Farrell, J.S., Hwaun, E., Dudok, B., and Soltesz, I. (2024). Neural and behavioural state switching during hippocampal dentate spikes. Nature *628*, 590–595. https://doi.org/10.1038/s41586-024-07192-8.

29. Headley, D.B., Kanta, V., and Paré, D. (2017). Intra- and interregional cortical interactions related to sharp-wave ripples and dentate spikes. J. Neurophysiol. *117*, 556–565. https://doi.org/10.1152/jn.00644.2016.

30. Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., and Buzsáki, G. (1999). Oscillatory Coupling of Hippocampal Pyramidal Cells and Interneurons in the Behaving Rat. J. Neurosci. *19*, 274–287. https://doi.org/10.1523/JNEUROSCI.19-01-00274.1999.

31. Pettersen, K.H., Lindén, H., Dale, A.M., and Einevoll, G.T. (2012). Extracellular spikes and CSD. In Handbook of Neural Activity Measurement (Cambridge University Press), pp. 92–135.

32. Nokia, M.S., and Penttonen, M. (2022). Rhythmic Memory Consolidation in the Hippocampus. Front. Neural Circuits *16*, 885684. https://doi.org/10.3389/fncir.2022.885684.

33. Gini, C. (1921). Measurement of Inequality of Incomes. Econ. J. *31*, 124–126. https://doi.org/10.2307/2223319.

34. Hurley, N., and Rickard, S. (2009). Comparing Measures of Sparsity. IEEE Trans. Inf. Theor. *55*, 4723–4741. https://doi.org/10.1109/TIT.2009.2027527.

35. McHugh, S.B., Lopes-dos-Santos, V., Gava, G.P., Hartwich, K., Tam, S.K.E., Bannerman, D.M., and Dupret, D. (2022). Adult-born dentate granule cells promote hippocampal population sparsity. Nat. Neurosci. *25*, 1481–1491. https://doi.org/10.1038/s41593-022-01176-5.

36. Phillips, R.G., and LeDoux, J.E. (1992). Differential contribution of amygdala and hippocampus to cued and contextual fear conditioning. Behav. Neurosci. *106*, 274–285. https://doi.org/10.1037/0735-7044.106.2.274.

37. Lee, I., Hunsaker, M.R., and Kesner, R.P. (2005). The Role of Hippocampal Subregions in Detecting Spatial Novelty. Behav. Neurosci. *119*, 145–153. https://doi.org/10.1037/0735-7044.119.1.145.

38. Gava, G.P., McHugh, S.B., Lefèvre, L., Lopes-dos-Santos, V., Trouche, S., El-Gaby, M., Schultz, S.R., and Dupret, D. (2021). Integrating new memories into the hippocampal network activity space. Nat. Neurosci. *24*, 326–330. https://doi.org/10.1038/s41593-021-00804-w.

39. Navas-Olive, A., Amaducci, R., Jurado-Parras, M.-T., Sebastian, E.R., and de la Prida, L.M. (2022). Deep learning-based feature extraction for prediction and interpretation of sharp-wave ripples in the rodent hippocampus. eLife *11*, e77772. https://doi.org/10.7554/eLife.77772.

40. Lopes-dos-Santos, V., van de Ven, G.M., Morley, A., Trouche, S., Campo-Urriza, N., and Dupret, D. (2018). Parsing Hippocampal Theta Oscillations by Nested Spectral Components during Spatial Exploration and Memory-Guided Behavior. Neuron *100*, 940–952.e7. https://doi.org/10.1016/j.neuron.2018.09.031.

41. Knierim, J.J., Neunuebel, J.P., and Deshmukh, S.S. (2014). Functional correlates of the lateral and medial entorhinal cortex: objects, path integration and local–global reference frames. Philos. Trans. R. Soc. Lond. B Biol. Sci. *369*, 20130369. https://doi.org/10.1098/rstb.2013.0369.

42. O'Neill, J., Senior, T.J., Allen, K., Huxter, J.R., and Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. Nat. Neurosci. *11*, 209–215. https://doi.org/10.1038/nn2037.

43. Nokia, M.S., Gureviciene, I., Waselius, T., Tanila, H., and Penttonen, M. (2017). Hippocampal electrical stimulation disrupts associative learning when targeted at dentate spikes. J. Physiol. *595*, 4961–4971. https://doi.org/10.1113/JP274023.

44. Lensu, S., Waselius, T., Penttonen, M., and Nokia, M.S. (2019). Dentate spikes and learning: disrupting hippocampal function during memory consolidation can improve pattern separation. J. Neurophysiol. *121*, 131–139. https://doi.org/10.1152/jn.00696.2018.

45. Nakashiba, T., Buhl, D.L., McHugh, T.J., and Tonegawa, S. (2009). Hippocampal CA3 Output Is Crucial for Ripple-Associated Reactivation

CellPress
OPEN ACCESS

Neuron
Article

and Consolidation of Memory. Neuron *62*, 781–787. https://doi.org/10.1016/j.neuron.2009.05.013.

46. Sullivan, D., Csicsvari, J., Mizuseki, K., Montgomery, S., Diba, K., and Buzsáki, G. (2011). Relationships between Hippocampal Sharp Waves, Ripples, and Fast Gamma Oscillation: Influence of Dentate and Entorhinal Cortical Activity. J. Neurosci. *31*, 8605–8616. https://doi.org/10.1523/JNEUROSCI.0294-11.2011.

47. Davoudi, H., and Foster, D.J. (2019). Acute silencing of hippocampal CA3 reveals a dominant role in place field responses. Nat. Neurosci. *22*, 337–342. https://doi.org/10.1038/s41593-018-0321-z.

48. Quinn, A.J., Lopes-dos-Santos, V., Dupret, D., Nobre, A.C., and Woolrich, M.W. (2021). EMD: Empirical Mode Decomposition and Hilbert-Huang Spectral Analyses in Python. J. Open Source Software *6*, 2977. https://doi.org/10.21105/joss.02977.

49. Magland, J.F., Jun, J.J., Lovero, E., Morley, A.J., Hurwitz, C.L., Buccino, A.P., Garcia, S., and Barnett, A.H. (2020). SpikeForest, reproducible web-facing ground-truth validation of automated neural spike sorters. Elife. *9*, e55167. https://doi.org/10.7554/eLife.55167.

50. Pachitariu, M., Steinmetz, N.A., Kadir, S.N., Carandini, M., and Harris, K.D. (2016). Fast and accurate spike sorting of high-channel count probes with KiloSort. In Adv. Neural Inf. Process. Syst., *29*, D.D. Lee, M. Sugiyama, U.V. Luxburg, I. Guyon, and R. Garnett, eds. (Curran Associates, Inc.), pp. 4448–4456.

51. Han, X., Chow, B. Y., Zhou, H., Klapoetke, N. C., Chuong, A., Rajimehr, R., Yang, A., Baratta, M. V., Winkle, J., Desimone, R., and Boyden, E. S. (2011). A high-light sensitivity optical neural silencer: development and application to optogenetic control of non-human primate cortex. Front. Syst. Neurosci. *5*, 18. https://doi.org/10.3389/fnsys.2011.00018.

52. Lopes-dos-Santos, V., Brizee, D., and Dupret, D. (2023). Spatio-temporal organization of network activity patterns in the hippocampus. Preprint at bioRxiv. https://doi.org/10.1101/2023.10.17.562689.

53. Skaggs, W.E., McNaughton, B.L., and Gothard, K.M. (1993). An Information-Theoretic Approach to Deciphering the Hippocampal Code. In Advances in Neural Information Processing Systems, *5*, S.J. Hanson, J.D. Cowan, and C.L. Giles, eds. (Morgan-Kaufmann), pp. 1030–1037.

54. Mitzdorf, U. (1985). Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. Physiol. Rev. *65*, 37–100. https://doi.org/10.1152/physrev.1985.65.1.37.

55. Guest, O., and Love, B.C. (2017). What the success of brain imaging implies about the neural code. eLife *6*, e21397. https://doi.org/10.7554/eLife.21397.

56. Onnela, J.-P., Saramäki, J., Kertész, J., and Kaski, K. (2005). Intensity and coherence of motifs in weighted complex networks. Phys. Rev. E Stat. Nonlin. Soft Matter Phys. *71*, 065103. https://doi.org/10.1103/PhysRevE.71.065103.

57. Saramäki, J., Kivelä, M., Onnela, J.-P., Kaski, K., and Kertész, J. (2007). Generalizations of the clustering coefficient to weighted complex networks. Phys. Rev. E Stat. Nonlin. Soft Matter Phys. *75*, 027105. https://doi.org/10.1103/PhysRevE.75.027105.

58. Costantini, G., and Perugini, M. (2014). Generalization of Clustering Coefficients to Signed Correlation Networks. PLoS One *9*, e88669. https://doi.org/10.1371/journal.pone.0088669.

59. Estrada, E. (2019). Rethinking structural balance in signed social networks. Discrete Appl. Math. *268*, 70–90. https://doi.org/10.1016/j.dam.2019.04.019.

60. Toth, R., Zamora, M., Ottaway, J., Gillbe, T., Martin, S., Benjaber, M., Lamb, G., Noone, T., Taylor, B., Deli, A., et al. (2020). DyNeuMo Mk-2: An Investigational Circadian-Locked Neuromodulator with Responsive Stimulation for Applied Chronobiology. Conf. Proc. IEEE Int Conf Syst. Man Cybern *2020*, 3433–3440. https://doi.org/10.1109/SMC42975.2020.9283187.

61. Kavoosi, A., Toth, R., Benjaber, M., Zamora, M., Valentín, A., Sharott, A., and Denison, T. (2022). Computationally efficient neural network classifiers for next generation closed loop neuromodulation therapy - a case study in epilepsy. In 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), *2022*, pp. 288–291. https://doi.org/10.1109/EMBC48229.2022.9871793.

62. Pennington, Z.T., Dong, Z., Feng, Y., Vetere, L.M., Page-Harley, L., Shuman, T., and Cai, D.J. (2019). ezTrack: An open-source video analysis pipeline for the investigation of animal behavior. Sci. Rep. *9*, 19979. https://doi.org/10.1038/s41598-019-56408-9.

63. Ho, J., Tumkaya, T., Aryal, S., Choi, H., and Claridge-Chang, A. (2019). Moving beyond P values: data analysis with estimation graphics. Nat. Methods *16*, 565–566. https://doi.org/10.1038/s41592-019-0470-3.

64. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat. Methods *17*, 261–272. https://doi.org/10.1038/s41592-019-0686-2.

65. Harris, C.R., Millman, K.J., van der Walt, S.J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N.J., et al. (2020). Array programming with NumPy. Nature *585*, 357–362. https://doi.org/10.1038/s41586-020-2649-2.

66. Hunter, J.D. (2007). Matplotlib: A 2D Graphics Environment. Computing in Science & Engineering *9*, 90–95. https://doi.org/10.1109/MCSE.2007.55.

67. Waskom, M.L. (2021). seaborn: statistical data visualization. Journal of Open Source Software *6*, 3021. https://doi.org/10.21105/joss.03021.

68. McKinney, W.. Data Structures for Statistical Computing in Python. https://doi.org/10.25080/Majora-92bf1922-00a.

69. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine Learning in Python. J. Mach. Learn. Res. *12*, 2825–2830.

70. Seabold, S., and Perktold, J.. Statsmodels: Econometric and Statistical Modeling with Python. https://doi.org/10.25080/Majora-92bf1922-011.

CellPress
OPEN ACCESS

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Bacterial and virus strains		
AAV9-CAG-Flex-ArchT-GFP	UNC Vector Core	N/A
Experimental models: Organisms/strains		
C57BL/6J mice	Charles River	632
Grm2-Cre Tg(Grm2-cre)MR90Gsat/Mmucd	MMRRC	MMRRC_034611-UCD
Nestin-Cre B6.Cg-Tg(Nes-Cre)1Kln/J	Jackson Laboratories	IMSR_JAX:003771
Software and algorithms		
Intan RHD2000	Intan Technologies	RHD2164
Positrack	Kevin Allen	N/A
Empirical Mode Decomposition in Python	Quinn et al.[48]	N/A
Kilosort via SpikeForest	Magland et al.[49]; Pachitariu et al.[50]	N/A
Other		
12um tungsten wires	California Fine Wire	M294520
Optic fibers	Doric lenses	MFC_200/230-0.37_25mm_RM3_FLT
Head-stage amplifier	Intan Technologies	RHD2164
561nm diode-pumped solid-state laser	Laser 2000	CL561-100

EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS

Animals

These experiments used adult (4–6 months old) wild-type C57Bl6/J mice (Charles River Laboratories, Kent, UK) or hemizygous Nestin-Cre mice (Jackson Laboratories; B6.Cg-Tg(Nes-Cre)1Kln/J, stock no. 003771, RRID: IMSR_JAX:003771) for the initial investigation of principal cell spiking in DSs and SWRs using tetrodes or silicon-probe recordings (Figures 1–4). To optogenetically target DG granule cells, we used adult metabotropic-glutamate-receptor 2-Cre (Grm2-Cre) hemizygous male mice (Figures 5 and 6). This Grm2-Cre mouse strain was obtained from the Mutant Mouse Resource and Research Center (MMRRC; Tg(Grm2-cre)MR90Gsat/Mmucd; stock no. 034611-UCD, RRID:MMRRC_034611-UCD) at University of California at Davis, an NIH-funded strain repository, and was donated to the MMRRC by Nathaniel Heintz, Ph.D., The Rockefeller University, GENSAT and Charles Gerfen, Ph.D., National Institutes of Health – National Institute of Mental Health. All mice were group housed with same-sex littermates until the start of the experiment and singly housed after surgery. Mice had free access to food and water throughout, in a dedicated housing room with a 12/12 h light/dark cycle (7 a.m.–7 p.m.), 19°C–23°C ambient temperature and 40–70 % humidity. This study used mice with good health/immune status, that were not involved in previous procedures, and were drug and test naïve at the start of the experiments. Mice were adult males and the influence (or association) of age and sex, or both on the results of the study was not tested. This represents a limitation to this research's generalizability. All experiments were performed between 8 a.m.–6 p.m. during the light-on period, that is when mice sleep more. Experiments were performed 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.

METHOD DETAILS

Viral vectors

An AAV carrying a double-floxed inverse open reading frame (DIO) Cre-dependent opsin under the CAG promoter was used to deliver Archaerhodopsin (ArchT) (Han et al.[51]) into DG granule cells (AAV9-CAG-Flex-ArchT-GFP, titer: 8.3 × 10^{12} TU / mL, University of North Carolina).

Surgical procedures

Mice received viral injections and microdrive implantations under gaseous isoflurane anaesthesia (∼1% in 1 L / min O₂), with systemic and local analgesia administered subcutaneously (meloxicam 5 mg / kg; buprenorphine 0.1 mg / kg; bupivacaine 2 mg / kg). Viruses were injected bilaterally into the dorsal DG (3 × 200 nL per hemisphere; at the following stereotaxic coordinates from

bregma: anterior-posterior: -1.6, -2.4, -2.4; mediolateral: ±1.0, ±1.2, ±1.5; dorsoventral: -1.7, -1.7, -1.7 mm, respectively), and delivered using a pulled glass micropipette ($\sim$16 μm i.d.) at a rate of 100 nL min^{-1}, with an additional 100 nL min^{-1} diffusion time with the pipette *in situ*. In a separate surgery, mice were implanted with a microdrive containing twelve- or fourteen-independently movable tetrodes bilaterally targeting DG, CA3, and CA1, and two optic fibers (Doric Lenses Inc., Quebec, Canada) positioned bilaterally above the dorsal DG. Tetrodes were constructed by twisting together four insulated tungsten wires (12.7 μm diameter, California Fine Wire, CA, USA) 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 depth manipulation. A separate group of mice were implanted with unilateral single-shank 64-channel silicon-probe (model: ASSY-236 H3, 8 mm; Cambridge Neurotech, Cambridge, UK; stereotaxic coordinates from bregma: anterior-posterior: -2.0; mediolateral: -1.7 mm); these mice did not receive prior viral injections.

Recording procedures

Following the implantation surgery, mice recovered for at least seven days before familiarization to the recording procedure. Mice were handled daily and exposed to the sleep-box for > 0.5 h per day for at least four days. During this period, tetrodes / silicon-probes were slowly lowered to the proximity of the cell layers. Once at the correct depth, silicon-probes were left in the same position for the rest of the experiment. Tetrodes were lowered into the CA1, CA3 pyramidal or DG granule cell layers on the morning of each recording day in search of multi-unit spiking activity, using the electrophysiological profile of the local field potentials including sharp-wave ripples, gamma oscillations, and dentate spikes to further guide placement. Tetrodes were left in position for $\sim$1.5–2 h before recordings began on that day. At the end of each recording day, tetrodes were raised ($\sim$150 μm) to protect hippocampal the cell layers from potential mechanical damage overnight. We lowered again each individual tetrode on the next morning in search of cells, making it unlikely that the recorded units are the same neurons across days. During recording sessions, mice explored open-field environments (41 cm diameter cylinder, or 41 × 41 cm square box, both with 30 cm high walls), or were placed in a sleep box containing sawdust bedding and nesting material (12 × 12 × 28 cm, length × width × height). The instantaneous speed and the theta-to-delta ratio profiles for DS and SWR events corresponded to those of sleep (Figures S1G and S1H). However, in the absence of electromyography signals or other additional signals in defining a sleep stage, we here refer to sleep/rest. Each open-field or sleep box recording session lasted $\sim$15-30 min. Experiments were performed under dim light conditions ($\sim$20 lux) with low-level background noise ($\sim$50 dB).

Light delivery

A 561 nm diode pumped solid-state laser (Crystal Laser, model CL561-100; distributer: Laser 2000, Ringstead, UK) was used to deliver green-yellow light bilaterally to the dorsal DG ($\sim$2-4 mW) via a 2-channel rotary joint (Doric Lenses Inc.).

Multichannel data acquisition

Electrode signals were amplified, multiplexed, and digitized using a single integrated circuit (headstage) located on the head of the animal (RHD2164, Intan Technologies, USA; http://intantech.com/products_RHD2000.html). The amplified and filtered (pass band 0.09 Hz to 7.60 kHz) electrophysiological signals were digitized at 20 kHz (RHD2000 Evaluation Board) and saved to disk with the synchronization signals from the positional tracking and laser activation. To track the location of the animal, three LEDs were attached to the headstage and captured at 25 frames per second by an overhead color camera.

Spike sorting and unit isolation

Spike sorting and unit isolation were performed via automatic clustering software Kilosort[50] (https://github.com/cortex-lab/KiloSort) followed by graphically based manual recombination using cross-channel spike waveforms, auto-correlation histograms and cross-correlation histograms within the SpikeForest framework (https://github.com/flatironinstitute/spikeforest)[49]. All sessions recorded on a given day were concatenated and cluster cut together to monitor cells throughout the day. Each unit used for analyses showed consistent spike waveforms and stable firing rates throughout the entire recording day. Tetrode location in the dorsal-ventral axis for each recording day (Figure S1C) was determined using laminar LFP signatures, as described in detail in Lopes-dos-Santos et al.[52], and later confirmed in the *ex vivo* histology (Figure S1B).

Principal cell versus interneuron classification

Hippocampal principal cells were distinguished from interneurons by the trough-to-peak width of the spike waveform, as previously described.[52] Briefly, to evaluate the waveform consistency for each unit, we used the waveform with the maximum amplitude across the tetrode channels for each cluster. We compared the prominence of a unit mean waveform amplitude to the standard deviation stemming from all its spikes by computing a waveform score:

$$wv_{score} = \sqrt{\sum_{i=1}^{n} \frac{(w_i/\sigma_i)^2}{n}}$$

where w_i is the value of the mean waveform at sample i, σ_i is the standard deviation at sample i across all spikes, and n is the number of waveform samples. This metric quantifies the relative magnitude of the mean waveform amplitude against the spike-to-spike

variability. Clusters with a waveform score above 0.75 and a refractory period violation below 2% (quantified as the proportion of intervals shorter than 2 ms in the ISI distribution) were included for further analyses. We categorized units as either putative interneurons or principal cells based on the width of their waveform as indicated by the trough-to-peak latency. In a prior dataset of $\sim$4,000 neurons, we noted a bimodal distribution in trough-to-peak latency. Fitting this with a 1-dimensional, 2-component Gaussian Mixture Model (GMM), we set the classification threshold where the two Gaussian components intersect: units with latencies above were labeled as putative principal cells, and those below as putative interneurons. The same inclusion criteria and classification procedures were used for DG, CA3 and CA1 neurons. In total, this study includes 3,619 hippocampal principal cells (CA1, n = 1,322; CA3, n = 573; DG, n = 1,724; from 134 total recording days in 25 mice).

Local field potential signals

LFP signals were processed by first applying an anti-aliasing filter (8[th]-order Chebyshev type I filter) to the wide band signals sampled at 20 kHz. These signals were then down-sampled to 1,250 Hz using the decimate function from the signal submodule of Scipy (version 1.11.2).

Dentate spike detection

Dentate spikes were detected during sleep sessions from LFPs recorded from tetrodes located in the DG granule cell layer or silicon-probes with recording contacts in the DG granule cell layer. In silicon-probe recordings, we initially subtracted the LFP signals from all channels using a reference channel found in the stratum oriens. LFPs were band-pass filtered (1–200 Hz, using a 4th order Butterworth filter). The mean and standard deviation of the LFP amplitude were calculated across the entire sleep session and peaks that exceeded a threshold of six times the median absolute value of the filtered signals were designated as dentate spikes. The time bin with the largest local maximum was taken as the peak of the dentate spike, and this timestamp was recorded. If more than one peak appeared within a 50 ms frame, we retained only the highest amplitude peak. On recording days with several tetrodes in the DG, we used the tetrode with the largest mean DS amplitude to select DS event timestamps. Across all tetrode recordings we detected 32,215 DS events in total (mean ± SEM: 441.3 ± 29.2 per day, from 73 recording days in 12 mice); in silicon-probe recordings we detected 15,067 DS events in total (mean ± SEM: 1676.1. ± 316.5 per day, from 8 recording days in 3 mice).

Sharp-wave ripple detection

For the LFPs of each pyramidal CA1 channel, we subtracted the mean across all channels (common average reference), band-pass filtered for the ripple band (80–250 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 envelope of that channel) were regarded as candidate events. 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 twice the ripple band power in the common average reference (to eliminate common high frequency noise); (*ii*) each event had at least four ripple cycles (to eliminate events that were too brief); (*iii*) ripple band power was at least twice 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). For events passing these criteria, the local maximum of each envelope was taken as the peak of the SWR, and these timestamps were recorded. On recording days with several tetrodes in the CA1 pyramidal layer, we used the tetrode with the largest mean ripple envelope amplitude to select SWR events. In tetrode recordings we detected 65,370 SWR events (mean ± SEM: 895.0 ± 82.3 per day, from 73 recording days in 12 mice).

Place maps

To generate place maps, we divided the horizontal plane of the recording enclosure into spatial bins of 1.4 × 1.4 cm to generate the spike count map (number of spikes fired in each bin) for each neuron and the occupancy map (time spent by the animal in each spatial bin) in each task session. All maps were then smoothed by convolution with a two-dimensional Gaussian kernel (s.d. = 1.2 bin widths). Finally, spatial rate maps were generated by normalizing the smoothed spike count maps by the smoothed occupancy map.

Spatial Information

The amount of spatial information conveyed by the spike train of a given cell was calculated using the formula proposed by Skaggs et al.[53]:

$$Information\ per\ spike\ =\ \sum_{i=1}^{N} p_i \frac{\lambda_i}{\lambda} \log 2 \frac{\lambda_i}{\lambda}$$

where i = 1, 2.. N represents each spatial bin of the environment, p_i is the probability of occupancy of bin i, λ_i is the mean firing rate in bin i, and λ is the mean firing rate of the cell over all spatial bins.

Peri-event time histograms (PETHs)

For analysis, we excluded all DS and SWR events that occurred within 50 ms of one another. We constructed PETHs over 400 ms windows, 200 ms either side of the peak DS amplitude or the peak of the SWR envelope, using a 1 ms bin width. The mean firing rate of each neuron was calculated during each 1 ms bin over the 400 ms window for each event. Z-scored firing rates were generated (over the DS-triggered or SWR-triggered average) separately for each neuron by calculating the mean and standard deviation over the 400 ms PETH:

$$z_i = \frac{(x_i - \overline{x})}{s}$$

where z_i is the Z score at time bin i, x_i is the firing rate in time bin i, $\overline{x}$ is the mean firing rate across all time bins, and s is the standard deviation of the firing rate across all time bins. The Z-scored firing rate of each neuron was then smoothed using a 3-point moving average to eliminate spurious peaks in low firing rate neurons. For a cell to be classified as significantly activated during DS and/or SWR events, the firing rate within ± 20 ms of the event peak had to be > 3 standard deviations (s.d.) above baseline (calculated as the mean firing rate over the 400 ms window). We also calculated the proportion of activated cells as a function of activation threshold ($2 < Z$ score < 4; Figures S2H and S3E).

Current source density analysis

Current sources and sinks were estimated from LFP recordings taken from single-shank 64-channel silicon-probes spanning the so-mato-dendritic axis of CA1 principal cells and reaching the inferior blade of the DG. LFP signals were first down-sampled to 1250 Hz. The current source density[54] unscaled signal at time t and electrode n, $CSD[t]n$, was estimated as:

$$CSD[t]_n = -\left(LFP[t]_{n-1} - 2 \times LFP[t]_n + LFP[t]_{n+1}\right)$$

where $LFP[t]_{n-1}$, $LFP[t]_n$ and $LFP[t]_{n+1}$ are the LFP signals at time t recorded from neighboring electrodes ($n-1$ and $n+1$ are the channels immediately above and below n, respectively, with 20 μm spacing between electrodes). The silicon-probe recording site in the pyramidal layer was identified as the one with largest ripple-band power. We defined the location of radiatum and lacunosum mo-lecular layers according to the sharp-wave and theta laminar profiles, as previously described.[52] We sorted dentate spike events into type 1 (DS_1) or type 2 (DS_2) in the following way. First, we calculated the CSD estimates for all DSs at the peak of each event and used PCA to find the first two Principal Components from the resulting CSD traces. These principal components had as many dimensions as the number of silicon-probe channels (64). We then used a 2-component Gaussian Mixture Model to classify the events based on their projection onto the first two principal components. This consistently resulted in two event classes having the strongest sinks in different areas of the molecular layer. In line with previous research,[22,24] we classified the events with the strongest sink in the outermost part of the molecular layer as DS_1, and events with their sink closer to the granular layer as DS_2. Based on CSD classification, event proportions were DS_1: 0.35; DS_2: 0.65 (5274 DS_1 versus 9793 DS_2, based on 15,067 events from 8 recording days in 3 mice).

Linear discriminant analysis classifier

To distinguish between DS_1 and DS_2 events using only the LFP traces, we trained a linear discriminant analysis (LDA) classifier using silicon-probe recorded LFPs from the granule cell layer (https://doi.org/10.5281/zenodo.10034433). LFP signals were first down-sampled to 1250 Hz and low-pass filtered at 50 Hz. We extracted 400 ms epochs centered around the peak of each DS (-200 to +200 ms, with 0.8 ms bin width), providing 500 time-based features (dimensions), one for each time bin, for each LFP trace. We then performed PCA on all silicon-probe-recorded DS LFP traces (15,067) to extract the number of components explaining 90% of the variance. This resulted in 16 principal components, which were then used to train a LDA classifier. We generated 20 models by, each time, randomly selecting 75% of the dataset, which was labelled as DS_1 or DS_2 based on the CSD classification described above, and then testing the classifier on the remaining (unlabeled) 25% of data. The classifier success rate was: median (IQR) = 85.4 (85.3–85.6) %. We then used the model with the highest accuracy to classify DS_1 and DS_2 events from LFPs recorded from the granule cell layer in our tetrode-recorded data. From tetrode-recorded LFPs, the proportions of type 1 and type 2 DS events were: median (IQR) DS_1 = 0.34 (0.25–0.38); DS_2 = 0.66 (0.62–0.75), based on 10,337 DS_1 versus 21,740 DS_2 events in 73 recording days in 12 mice.

Population spiking vectors

We generated event-based hippocampal population vectors of instantaneous principal cell spiking for every DS and SWR event us-ing 50 ms wide windows centered on the peak of the DS or the peak envelope of the CA1 ripple (±25 ms from the peak). In addition, we calculated the spiking activity of hippocampal principal cells in equivalent 50 ms ('no event') control epochs, that contained neither DS nor SWRs. Baseline periods were selected from the same sleep sessions and excluded all epochs ± 250 ms either side of any DS or SWR events. To calculate the proportion of coactive neurons in each time window, we calculated the number of simultaneously active hippocampal principal cells (i.e., cells firing at least one spike during the 50 ms window) by the total number of simultaneously recorded hippocampal principal cells. We then calculated the mean proportion of coactive cells for each recording session. For in-clusion in these analyses, each recording session required a minimum of 100 of each type of event (DS_1, DS_2, SWR) and a minimum of 20 simultaneously recorded hippocampal principal cells.

Population-level sparsity

The sparsity S of a given population firing vector x was calculated using the Gini index[33,34,55] as:

$$S = \frac{\sum_{i=1}^{N}(2i - N - 1)x_i}{N\sum_{i=1}^{N}x_i}$$

where x is the population vector containing, in ascending order, the spike counts discharged by each principal cell in a 50 ms time window (centered on the peak of the SWR, DS), N is the length of that vector (i.e. the number of simultaneously recorded principal cells), and i is the rank of spike counts in ascending order. Population vectors where the total number of spikes is more evenly distributed between neurons have a lower Gini index (lower sparsity) than population vectors where the total number of spikes is concentrated in a few neurons (higher sparsity).

Logistic regression classifier

We used a logistic regression classifier to distinguish between population vectors of hippocampal principal cell spiking activity during DS, SWR, or equivalent duration (50 ms) control vectors that were taken from 200–250 ms periods before or after the peak of either the DS or SWR events. For each recording session, we generated matrices of these population vectors (cells × epochs) for these four different event-types, and then binarized the spike counts (i.e., spike count > 0 = 1, else 0) to control for the influence of firing rate differences between neurons. For each recording day, we used the event with the lowest number of epochs to determine the training set size – for example, if there were 200 DS events, we used 75% (150 population vectors) as the DS training set, and randomly sub-sampled the SWR matrix for 150 SWR population vectors (with identical principal cells). This way, the training input to the classifier was balanced across event types. Similarly, the testing set consisted of the remaining (unlabeled) 25% of population vectors from the DS population vectors plus an equivalent number of SWR population vectors (e.g., 50 DS population vectors and 50 SWR population vectors, subsampled from the remaining SWR testing matrix). For each recording day, we ran three models: one to classify event epochs, one to classify pre-event epochs and one to classify post-event epochs. Model accuracy was measured as the proportion of correctly classified events (DS versus SWR, or pre-DS versus pre-SWR, respectively).

Peer-to-peer coactivity analysis

We constructed hippocampal population graphs that represent the coactivity relationships between all pairs of principal cell spike trains recorded during a given sleep or exploratory session. These coactivity graphs were computed using 50 ms time windows for DS and SWR events and theta cycles as time windows for active exploratory sessions. 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_{ij} obtained by fitting the GLM (Figure 3A):

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

where x_j, x_i are the Z-scored event-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,

$$P = \sum_{n=0}^{N-\{ij\}} x_n$$

with α_{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 $\beta_{i,i} = 0 \; \forall i \, in \, N$, 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 C_i 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.[56–58]:

$$C_i = \frac{\sum_{jq}\left(\widehat{\beta_{ij}}\,\widehat{\beta_{iq}}\,\widehat{\beta_{jq}}\right)^{1/3}}{k_i(k_i - 1)}$$

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.

Note that 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. This method to assess firing relationships in the neuronal population of the hippocampus as a signed network where both positive and negative edges (i.e., correlated and anti-correlated spike trains) coexist leverages from past studies investigating community organization in social networks, indicating that triads represent the smallest motif capturing "structural balance" in patterns of peer-to-peer relationships.[59]

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. 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.

Population vector similarity

Population vectors of hippocampal principal cell spiking activity were generated for baseline, SWR, and DS events as described above, yielding separate (cell × event number) matrices of spike counts for each event-type. To remove potential biases caused by unequal numbers of events, we used the event-type with the fewest epochs to determine the final matrix size. For example, if there were 200 DS events in a given recording session, we randomly subsampled the SWR and baseline matrices to extract 200 SWR and 200 baseline population vectors (with identical principal cells) for comparison. Next we binarized these matrices (spike count > 0 = 1, else 0). Then we assessed the self-similarity for each event matrix (cells × event number) by computing the Pearson correlation coefficient for every pair of population vectors from the same event-type, and then calculating the mean across all of these correlation coefficients. As an alternative, we also calculated the Jaccard similarity coefficient (J), which measures the size of the intersection (i.e. overlap in active units) between pairs of population vectors (A, B), divided by the size of the union:

$$J(A,B) = \frac{|A \cap B|}{|A \cup B|}$$

Population dimensionality

We estimated the dimensionality of the principal cell population firing structure during SWRs and DSs from activity matrices that were matched for neuron identity and the number of DS and SWR events. We applied Principal Component Analysis (PCA) to each activity matrix, using the number of simultaneously recorded principal cells as the maximum number of components. Each matrix required at least 20 principal cells for inclusion in the analysis. We then extracted the number of components explaining 90% of the variance in these population vectors and scaled this by the total number of neurons in each matrix (Figures 3L, S4M, and S4P). We also show dimensionality for a range of explained variance values (Figures S4N and S4O). Note that the ratio of DG to CA cells in these matrices did not significantly affect the dimensionality estimate (Table S1).

Theta-cycle detection

Theta cycles were detected as described in Lopes dos Santos et al. Briefly, we used masked Empirical Mode Decomposition[48]; https://pypi.org/project/emd/) to separate CA1 LFPs into oscillatory components termed intrinsic mode functions (IMFs). We delineated individual theta cycles from their troughs and peaks, i.e. the local maxima and minima of the theta IMF. Theta cycles were defined as peak-trough-peak sequences with trough-peak and peak-trough intervals between 31-100 ms and peak-to-peak distances between 71-200 ms. Note that this method is designed to detect chains of theta cycles but to do so it identifies each cycle independently.

Reactivation of waking coactivity patterns

We leveraged our pairwise peer-to-peer coactivity measure (as described above; Figure 3A) to estimate DS and SWR reactivation. With this, we compared the tendency of principal cell pairs to co-fire in theta cycles during exploration (theta coactivity) with the tendency to co-fire in DS (or SWR) during the following post-exploration sleep/rest period (post-DS or post-SWR co-firing), controlling for their baseline co-firing in the pre-exploration sleep/rest period before (pre-DS or pre-SWR co-firing) and mouse identity, using a linear mixed model:

$$\text{Post} \sim \beta_0 + \beta_{theta} + \beta_{pre} + \upsilon_{mouseID} + e$$

where β_0 is the intercept of the regression line, β_{theta} is the regression coefficient for the theta co-firing, β_{pre} is the regression coefficient for the pre-exploration offline co-firing (in DS or SWR events), $\upsilon_{mouseID}$ is the individual mouse identity, and e the error term. Likewise, we compared the tendency of principal cell pairs to co-fire in theta cycles during exploration (theta coactivity) with the

tendency to co-fire in DS or SWR during the pre-exploration sleep/rest period (pre-DS or pre-SWR co-firing), controlling for their post-exploration co-firing in the sleep/rest period after (post-DS or post-SWR co-firing) and mouse identity, using the reverse linear mixed model:

$$\text{Pre} \sim \beta_0 + \beta_{theta} + \beta_{post} + \upsilon_{mouseID} + e$$

From these LMMs, we extracted the β coefficients predicting post-SWR or post-DS coactivity from theta coactivity (controlling for pre-SWR or pre-DS coactivity, respectively) and tested their significance in two ways. First, we performed control GLMs using the pre-DS (or pre-SWR) coactivity as the dependent variable and the theta coactivity and post-DS (or post-SWR) coactivity as the independent variables. In these models, pre-event, post-event and theta coactivity were entered as fixed-effects and mouse identity as random-effects, using the restricted maximum likelihood method (implemented using the MixedLM class, and fit() method with default parameters, from the statsmodels library (Seabold and Perktold, 2010) in Python3.10). Second, we constructed a random probability distribution of β weights for theta coactivity by shuffling the cell pair identity, thereby generating a null distribution (based on 1000 LMMs, each time randomly shuffling cell-pair identity).

Closed-loop optogenetic interventions

For DS-informed interventions, real time detection of DSs was achieved by first high pass filtering the DG LFP signals (5 Hz) using the on-board signal processing capabilities of the Intan RHD evaluation board (RHD2000, Intan Technologies, USA) and triggering a laser pulse if the LFP signal exceeded a voltage-threshold. Thresholds for DS-onset detection were set for each mouse during a sleep session at the start of each recording day so that DS events were consistently detected (~3 S.D. above mean signal amplitude). Threshold detection triggered a digital transistor-transistor logic (TTL) output pulse from the RHD interface to a Master 8 stimulation timing device (A.M.P.I., Jerusalem, Israel), which in turn sent a 100 ms duration square-wave pulse to activate the laser. In the 'DS-synchronized' condition, the laser was triggered with zero latency from DS-onset, whereas in the 'DS-delay' condition the laser was triggered 100 ms after DS detection (Figures 5 and 6). The rates of false negatives (DS not triggering laser pulse) and false positives (laser pulse emitted for LFP trace not meeting DS criteria) were 1.8±0.6% and 4.7±0.4%, respectively. The laser delivered yellow-green light (561-nm) into the dentate gyrus, which in DG^{Grm2}::ArchT mice activated the outward proton pump, Archaerhodopsin T to suppress spiking activity in DG granule cells. To investigate changes in firing rates in individual hippocampal principal cells during light-delivery, we constructed PETHs over 400 ms windows, 200 ms either side of DS-onset, using a 1 ms bin width and extracted the peak firing rate during DS-synchronized light-delivery versus DSs with no light delivery. In addition, we Z-scored the binned spike trains and calculated the mean Z score between DS-onset and 100 ms after DS-onset for each hippocampal principal cell during DS-synchronized light-delivery versus the equivalent 100 ms no-light period in the DS-delay condition.

For SWR-informed interventions, the Intan evaluation board was configured with firmware enabling additional filtering. Five operations were performed on the continuously acquired CA1 wideband LFP signal to provide a real time estimate of the instantaneous power in the ripple-band.[60,61] (1) To enable low-latency processing, the signal was first down-sampled to 2.5 kHz by averaging the raw 20 kHz data stream with a sliding window of 8 samples with no overlap. (2) This signal was then high-pass filtered (using a 1st order digital infinite impulse response filter with a corner frequency of 1.6 Hz to remove amplifier offset and electrode drift). Next, the signal was (3) band-limited to 100–200 Hz with a 4th order Butterworth filter, (4) rectified by taking its absolute value, and (5) amplified 128-fold and smoothed with an exponential moving average operation over an equivalent window size of 32 samples (12.8 ms). To detect SWR events in this band-power estimate, the threshold level for each mouse was set during a sleep session at the start of each recording day to ensure consistent (~3 S.D. above mean power) detection throughout the day. On detecting a threshold crossing, the Intan recording controller delivered a 5 ms TTL pulse to a Master 8 stimulation timing device (A.M.P.I., Jerusalem, Israel). Analogous to the DS-informed interventions, in the 'SWR-synchronized' condition the laser was triggered with zero latency from SWR-onset, whereas in the 'SWR-delay' condition the laser was triggered 100 ms after SWR detection.

Recognition memory tasks

On each day of both the novel-object and novel-position recognition tasks, mice explored a square-walled open field (Figure 6A; the 'object arena') containing four objects, each positioned midway along a given wall, ~1cm from the wall edge. Objects used were ~3 × 3 × 4 cm (width × depth × height) objects (e.g., Lego™ blocks or other similar items). During the first session in the object arena, mice explored four completely novel objects ('sampling' session, 10 min). After the sampling session, mice were placed into a sleep box where they received DG-targeting light delivery that was either synchronized to event detection (DS-synchronized or SWR-synchronized condition) or delayed by 100 ms from event detection (DS-delay or SWR-delay condition), as described above (sleep/rest session, 20 min). In the novel object recognition task, before the start of the next test session, one of the four objects was replaced with a different (and completely novel) object, and mice then explored the four objects again ('test 1' session, 10 min). This process was repeated, with another sleep session (~20 min, with either DS-sync or DS-delay light-delivery), followed by another object exploration session with one completely novel object and three previously encountered objects ('test 2' session, 10 min). In the novel position recognition task, the locations of two of the initially sampled objects were swapped (e.g. North and West), whereas the other two objects remained in their original positions. In the novel position task, only DS-the synchronized and DS-delay conditions and only the first test session were used. During each test session, we measured the time spent exploring each object and we calculated the percentage time spent investigating the novel object (or novel positioned objects) versus the mean percentage time spent investigating the familiar

objects (i.e., those objects seen in the previous session and/or those in the same locations). For analysis, four 'object-zones' were created by dividing the arena into nine equal sized square zones ($\sim$12 × 12 cm), such that four of these zones contained the objects. Time spent in the object zone was determined directly from the automated tracking data. Between sessions, the floor of the maze and the objects were cleaned with water. On any given day, mice received the same light-delivery condition.

Tone fear conditioning task

Fear conditioning was conducted in one of four operant chambers each with distinct visual cues (ENV-307A, Med Associates Inc., IN, USA). Mice were exposed to five auditory cues (either 2900 Hz tone or white noise, 72 dB, 30s duration), each co-terminating with a mild foot-shock (0.3 mA, 0.5 s). The mean ITI was 74s (range: 60 to 90s). Immediately after fear conditioning, mice were removed from the operant chamber and placed into the sleep box where they received DS-sync or DS-Delayed DG cell silencing for 45 minutes. For the recall session, mice were then placed into a different operant chamber than the one where they received conditioning (to reduce the impact of contextual cues on recall). Mice were exposed to the same five auditory cues but no foot-shocks were given. Fear memory was assessed by measuring freezing responses during the first two cues presented in the recall session (before extinction occurs) and comparing these responses to freezing responses to the first cue during training (before any shocks were given). Freezing was measured using automated movement detection software (ezTrack,[62]) and expressed as a % of tone duration (i.e. freezing for 15s during a 30s tone = 50% freezing).

Tissue processing and immunohistochemistry

At the completion of experiments, mice were deeply anesthetized with pentobarbital and perfused transcardially with 0.1 M PBS followed by 4% paraformaldehyde (PFA) in PBS. Brains were extracted and kept in 4% PFA for $\sim$24–72 h and then transferred to PBS (with 0.05% sodium-azide). For tetrode localization, free-floating sections (50 μm) sections were mounted on slides and imaged at × 5 using a Zeiss microscope (AxioImager M2; Zeiss, Plan-Neofluar 5× /0.16 objective). For immunostaining, free-floating sections (50 μm) were rinsed in PBS with 0.25% Triton X-100 (PBS-T) and were blocked for 1 hour at $\sim$20°C in PBS-T with 10% normal donkey serum (NDS). Sections were then incubated with primary antibodies diluted in 3% NDS blocking solution and incubated at 4°C for 72 hours (GFP anti-chicken, 1:1,000, Aves Labs, catalog no. GFP-1020; NeuN guinea pig, 1:500, Synaptic Systems, catalog no. 266 004). All sections were rinsed three times for 15 min in PBS-T and incubated for 4 hours at $\sim$20°C in secondary antibodies in the blocking solution (Cy3 donkey anti-guinea pig, 1:400, Jackson ImmunoResearch, catalog no. 706-165-148; goat anti-chicken 488, 1:1,000, Thermo Fisher Scientific, catalog no. A-11039). Sections were then rinsed three times for 15 min in PBS-T, with some sections then incubated for 1 min with DAPI (0.5 μg ml^{-1}, Sigma, D8417) diluted in PBS to label cell nuclei before three additional rinse steps of 10 min each in PBS. Sections were mounted on slides, cover-slipped with Vectashield (Vector Laboratories, catalog no. H-1000) and stored at 4°C. Sections were also used for anatomical verification of the tetrode tracks. Images were acquired using a Zeiss confocal microscope (LSM 880 Indimo, Axio Imager 2) with a Plan-Apochromat ×20/0.8 M27 objective and the ZEN (Zeiss Black 2.3) software.

QUANTIFICATION AND STATISTICAL ANALYSIS

Analyses were performed in Python 3.8 (https://www.python.org/downloads/release/python-3816/) and Python 3.10 (https://www.python.org/downloads/release/python-31011/), using the Python packages DABEST (Ho et al.[63]), scipy (Virtanen et al.[64]), numpy (Harris et al.[65]), matplotlib (Hunter[66]), seaborn (Waskom[67]), pandas (McKinney[68]), scikit-learn (Pedregosa et al.[69]), statsmodels (Seabold and Perktold[70]). Error bars, mean ± S.E.M unless otherwise stated. We used throughout this study a bootstrap-coupled estimation of effect sizes, plotting the data against a mean difference between the left-most condition and one or more conditions on the right and compare this difference against zero using 5,000 bootstrapped resamples. In these estimation graphics (DABEST plots), each black dot indicates a mean difference and the associated black ticks depict error bars representing 95% confidence intervals; the shaded area represents the bootstrapped sampling-error distribution. Bandwidth estimates for the kernel density estimate were computed using the scikit-learn package. We used the DABEST package to calculate test statistics and p-values and visualize data. The test statistic is the mean difference and the p-value is the is the probability of observing the effect size (or greater), assuming the null hypothesis of zero difference is true. Paired permutation tests (or equivalent paired tests) were performed for repeated-measures analyses and unpaired tests used for independent samples. Data distributions were assumed to be normal, but this was not formally tested. Our results were replicable across mice and recording days. For the optogenetic interventions, the different closed-loop conditions (DS-sync, SWR-sync, DS-Delay, and SWR-Delay) were experienced in a randomized order across days. In the object recognition tasks, objects and their positions and the order of their replacement were randomized. 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 analyses. No statistical methods were used to pre-determine sample sizes, but our sample sizes are similar to or larger than those reported in previous publications. Inclusion criteria for well-isolated single units were used as published in previous studies and are described in the corresponding subsections of the method details. For the population vector analyses (Figures 3 and S4), each recording session required a minimum of 100 of each type of event (DS$_1$, DS$_2$, SWR) and a minimum of 20 simultaneously recorded hippocampal principal cells for inclusion.

Supplemental information

Offline hippocampal reactivation during dentate

spikes supports flexible memory

Stephen B. McHugh, Vítor Lopes-dos-Santos, Manfredi Castelli, Giuseppe P. Gava, Sophie E. Thompson, Shu K.E. Tam, Katja Hartwich, Brook Perry, Robert Toth, Timothy Denison, Andrew Sharott, and David Dupret

Supplemental Figures and Tables

Figure S1. Tetrode locations and characterization of dentate spikes, related to Figure 1.
(A) Triple-(DG-CA3-CA1) tetrode layout schematic.
(B) Example histology showing tetrode tracks in DG, CA3, and CA1, with color-coded contours of the pyramidal cell layers and the granule cell layer.
(C) Histogram of tetrode locations in the dorsal-ventral plane (CA1 tetrodes n=176; DG tetrodes n=244, CA3 tetrodes n=92, from n=73 sessions, 12 mice).
(D) Examples of mean ± SEM dentate spike (DS) local field potential (LFP) waveforms from individual mice.

(E) Group mean ± SEM DS LFP waveform (n=73 sessions, 12 mice).

(F) Kernel density estimates (KDEs) and boxplots for DS (left) and SWR (right) event durations (based on 32,215 DS events and 65,370 SWR events, n=73 sessions, 12 mice).

(G) DS and SWR events occur when mice are asleep or in quiet rest and not when they are active. The graph shows KDEs for the probability of DS (blue) and SWR (orange) occurrence versus active behavior (gray) for a range of movement speeds. Active behavior was determined from the theta-to-delta ratio (>2.4) from the CA1 LFP during open field exploration.

(H) Estimation plot showing that DS and SWR events occur when the theta-to-delta ratio is low compared to active behavior. For this analysis, we defined a minimum activity level (movement speed > 3 cm/s) based on the active behavior movement speed distribution (panel F) to include 99% of the area under the curve. We then extracted the theta-to-delta ratio for speeds above this minimum and compared this distribution to the theta-to-delta ratio distributions during DSs and SWRs. Upper: raw data points (each point shows mean theta-to-delta power during one active behavior session or one sleep session), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event. Lower: difference (Δ) in theta-to-delta ratio between active epochs versus DS and SWR epochs computed from 5,000 bootstrapped resamples and with the difference-axis origin (dashed line) aligned to active behavior (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). The test statistic is the mean difference, shown on the y-axis of the lower plot. P-values are from unpaired permutation tests, active versus event, ***$P < 0.001$.

(I) Cross-correlogram of the probability of DS occurrence with respect to SWRs. Note that there was a small increase in the probability of DS around the time of a SWR, with more than 90% of DSs not occurring within ± 50 ms of a SWR. Red dashed lines show peak of event, gray dashed lines show ± 50 ms from peak.

(J) As for panel I but for SWR occurrence with respect to DSs.

A

Spike waveform
Trough-to-peak width

Example auto-correlellograms

DG   CA3   CA1

Counts

-50 0 50 -50 0 50 -50 0 50
Time (ms)

Example placemaps

DG   2.5   CA3   19.9   CA1   11.5

0   Firing Rate (Hz)   max

B   Principal cells

Trough-to-peak width (ms)

DG   CA3   CA1

Mean Δ

***   ***

CA3 minus DG   CA1 minus DG

C   Interneurons

iDG   iCA3   CA1

Mean Δ

*   *

iCA3 minus iDG   iCA1 minus iDG

D

AC shape (Gini index)

iDG   iCA3   CA1   iCA3   CA1   iCA1

Mean Δ

***   ***   ***

iDG minus DG   iCA3 minus CA3   iCA1 minus CA1

E

Mean firing rate (Hz)

DG   iDG   CA3   iCA3   CA1   iCA1

Mean Δ

***   ***   ***

iDG minus DG   iCA3 minus CA3   iCA1 minus CA1

F

Information per spike (bits)

DG   iDG   CA3   iCA3   CA1   iCA1

Mean Δ

***   ***   ***

iDG minus DG   iCA3 minus CA3   iCA1 minus CA1

G

DG   DS   SWR

Firing rate (z-score)

-200 -100 0 100 200
Time from event peak (ms)

CA3   DS   SWR

Firing rate (z-score)

-200 -100 0 100 200
Time from event peak (ms)

CA1   DS   SWR

Firing rate (z-score)

-200 -100 0 100 200
Time from event peak (ms)

H   DS

Cells activated (%)

DG   CA3   CA1

4.0 3.5 3.0 2.5 2.0
Activation threshold (z-score)

SWR

Cells activated (%)

DG   CA3   CA1

4.0 3.5 3.0 2.5 2.0
Activation threshold (z-score)

I   DS   SWR

Significant activation (%)

DG   CA3   CA1   DG   CA3   CA1
Principal cells

J   DS   SWR

Time at maximum activation (ms)

DG   CA3   CA1   CA1   CA3   DG

Mean Δ

***   **

CA3 minus DG   CA1 minus DG   CA3 minus DG   CA1 minus DG

K   SWR

Firing rate (z-score)

DG   CA3   CA1

-100 -50 0 50 100
Time from event peak (ms)

L   DS

Firing rate (z-score)

DG   CA3   CA1

-100 -50 0 50 100
Time from event peak (ms)

M   DS

CA1

CA3

DG

-20 -10 0 10 20
Latency (ms)

N   DS

Response latency (ms)

DG   CA3   CA1

Mean Δ

**   ***

CA3 minus DG   CA1 minus DG

O   DS

Firing rate (Hz)   DG   CA3   CA1

Mean   Event   Mean   Event   Mean   Event

Mean Δ

***   ***   ***

Event minus Mean   Event minus Mean   Event minus Mean

P   SWR

Firing rate (Hz)   DG   CA3   CA1

Mean   Event   Mean   Event   Mean   Event

Mean Δ

***   ***   ***

Event minus Mean   Event minus Mean   Event minus Mean

Figure S2. Dentate spikes activate hippocampal principal cells, related to Figure 1.
(A) Top: Example spike waveform showing the trough-to-peak measurement for spike width. Scale bar 100 µV and 0.5 ms. Middle: example auto-correlograms from individual DG, CA3 and CA1 principal cells. Bottom: example place maps from three individual (simultaneously recorded) DG, CA3 and CA1 principal cells illustrating the spatial distribution of spiking activity. The number in the top right corner shows the maximum firing rate of the cell in its place-field.
(B, C) Trough-to-peak width was used to classify principal cells versus interneurons. (B) Estimation plot showing that the trough-to-peak width is narrower in DG principal cells versus CA3 and CA1 principal cells. (C) Trough-to-peak width narrower in DG versus CA3 interneurons but is wider in DG versus CA1 interneurons. Note that the trough-to-peak width for DG principal cells remains wider for DG principal cells than DG interneurons.
(D) The auto-correlogram shape differs between principal cells and interneurons in DG, CA3 and CA1. Here, we used the Gini index to evaluate the sparsity in the spike probability distribution for each 1 ms bin of the auto-correlogram (between 0 and +50 ms). This distribution was more unequal for principal cells in all three regions, hence a higher Gini index; and was more equal for interneurons, hence a lower Gini index.
(E) The mean firing rate is higher in hippocampal interneurons (iDG, iCA3, iCA1: recorded cells with a trough-to-peak width < 0.45 ms) versus hippocampal principal cells.
(F) DG, CA3 and CA1 principal cells exhibit higher spatial information scores than interneurons in these respective regions.
For B-F estimation plots, Upper: raw data points (each point shows one principal cell); Lower: difference (Δ) in trough-to-peak width, firing rate, or spatial information (respectively) in DG versus CA3 and DG versus CA1. Other plot details as in Figure S1H.
(G) Examples of three individual principal cells' z-scored firing rates during DS (blue) and SWR (orange) events. The horizontal dashed line shows z-score = 3.
(H) Percentage of principal cells active during DS (left) versus SWR (right) events, as defined by crossing various z-thresholds.
(I) Percentage of significantly activated principal cells, as defined by a z-score > 3 (within ± 20 ms of the event peak, shown at time 0 in **G**, during DS (left) and SWR (right) events.
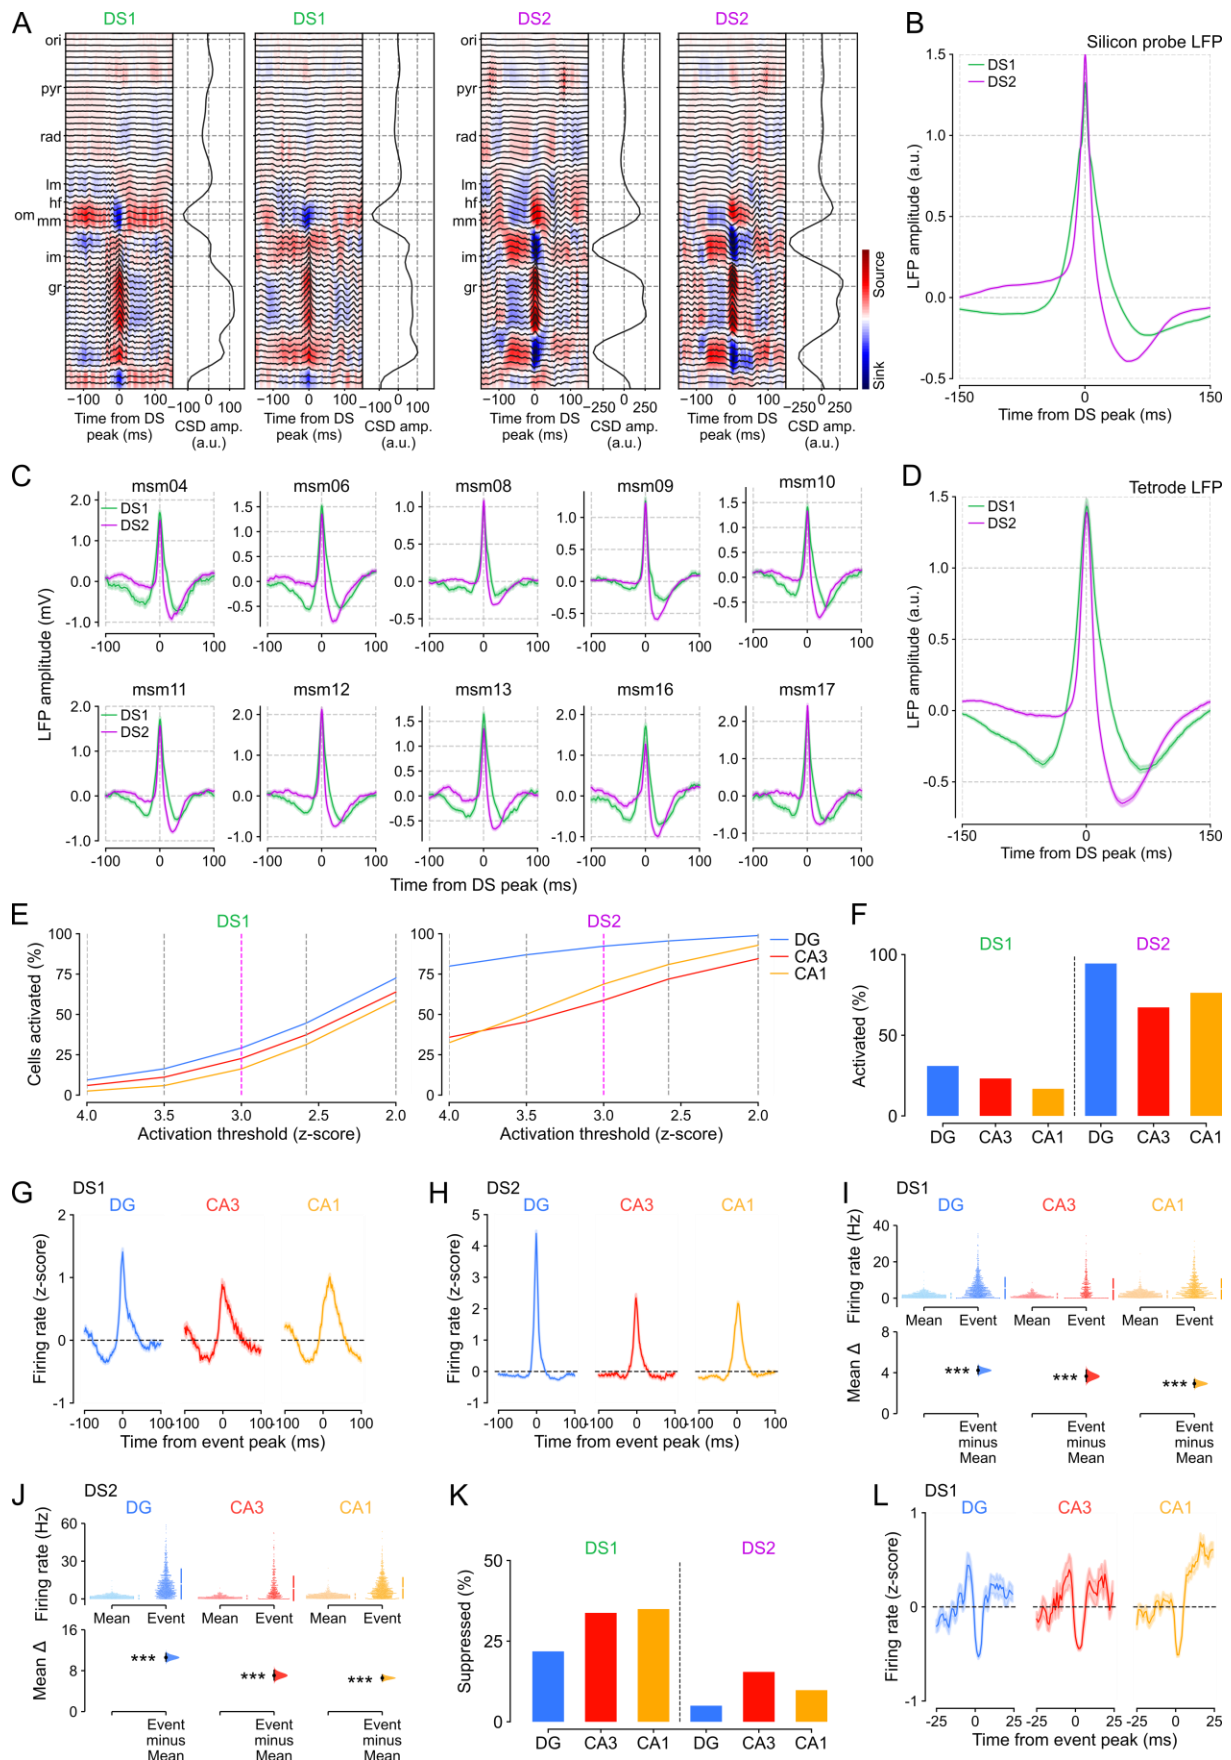**(J)** Estimation plot showing the effect size for the differences in the time point of maximum neuronal spiking activity during DSs and SWRs in DG, CA3, and CA1 principal cells. Note that DG cells reached their peak firing significantly before CA1 neurons during DSs; CA3 cells reached their peak firing significantly before CA1 neurons during SWRs. Upper: raw data points (each point shows one principal cell that was significantly active (z > 3) during each event); Lower: difference (Δ) in time of peak activation in DG versus CA3 and DG versus CA1.
(K, L) Corresponding time course of principal cell instantaneous firing rate (z-score) during SWR (K) and DS (L) events.
(M, N) Using the time to cross the z-score > 3 threshold, we observed that during DS events, DG principal cells increase their firing activity significantly earlier than both CA3 and CA1 principal cells. Panel N is the corresponding estimation plot of the response latency (i.e. when each cell crossed the z > 3 threshold relative to the event peak) showing that DG cells are active before CA3 and CA1 principal cells during DSs. Upper and lower plots as in panel J.
(O, P) Estimation plots comparing the overall mean firing rate of each principal cell (calculated across the entire recording session) to its peri-event firing rate (calculated as the mean firing rate ± 5 ms around the peak of the event) during DS (O) and SWR (P) events. Upper: raw data points (each point shows one principal cell's mean rate and peri-event rate) in DG, CA3 and CA1; Lower: difference (Δ) in firing rate between mean rate and peri-event rate for DG, CA3 and CA1 separately. Other plot details as in Figure S1H.
For B-F,J,N-P, the test statistic is the mean difference, shown on the y-axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). P-values are from unpaired permutation tests, cell type versus cell type (B-F,J,N); or paired permutation tests, baseline (mean rate) versus event (O-P), *$P < 0.05$, **$P < 0.01$, ***$P < 0.001$.

Figure S3. Current source density, local field potential profiles, and hippocampal principal cell spiking for DS$_1$ and DS$_2$ events, related to Figure 2.

(A) Examples of current source density (CSD) profiles for individual Type 1 (DS$_1$) and Type 2 (DS$_2$) dentate spikes (two DS$_1$ examples, two DS$_2$ examples) recorded from a 64-channel silicon-probe. *Left panel*: the instantaneous CSD ±150 ms around the event peak. *Right panel:* the CSD amplitude at each depth (based on the mean amplitude from –25 to + 25 ms around the event peak). Stratum oriens: ori; pyramidal layer: pyr; stratum radiatum: rad; lacunosum moleculare: lm; hippocampal fissure: hf; outer molecular layer: om; middle molecular layer: mm; inner molecular layer: im; granule-cell layer: gr.

(B) Group mean ± SEM LFP waveforms for DS$_1$ and DS$_2$ events from silicon-probe recordings (8 sessions, 3 mice).

(C) Examples of mean ± SEM LFP waveforms for DS$_1$ and DS$_2$ events from tetrode recordings in individual mice.

(D) Group mean ± SEM LFP waveforms for DS$_1$ and DS$_2$ events from tetrode recordings (73 sessions, 12 mice)

(E) Percentage of principal cells active during DS$_1$ (left) versus DS$_2$ (right) events, and as defined by crossing various z-thresholds.
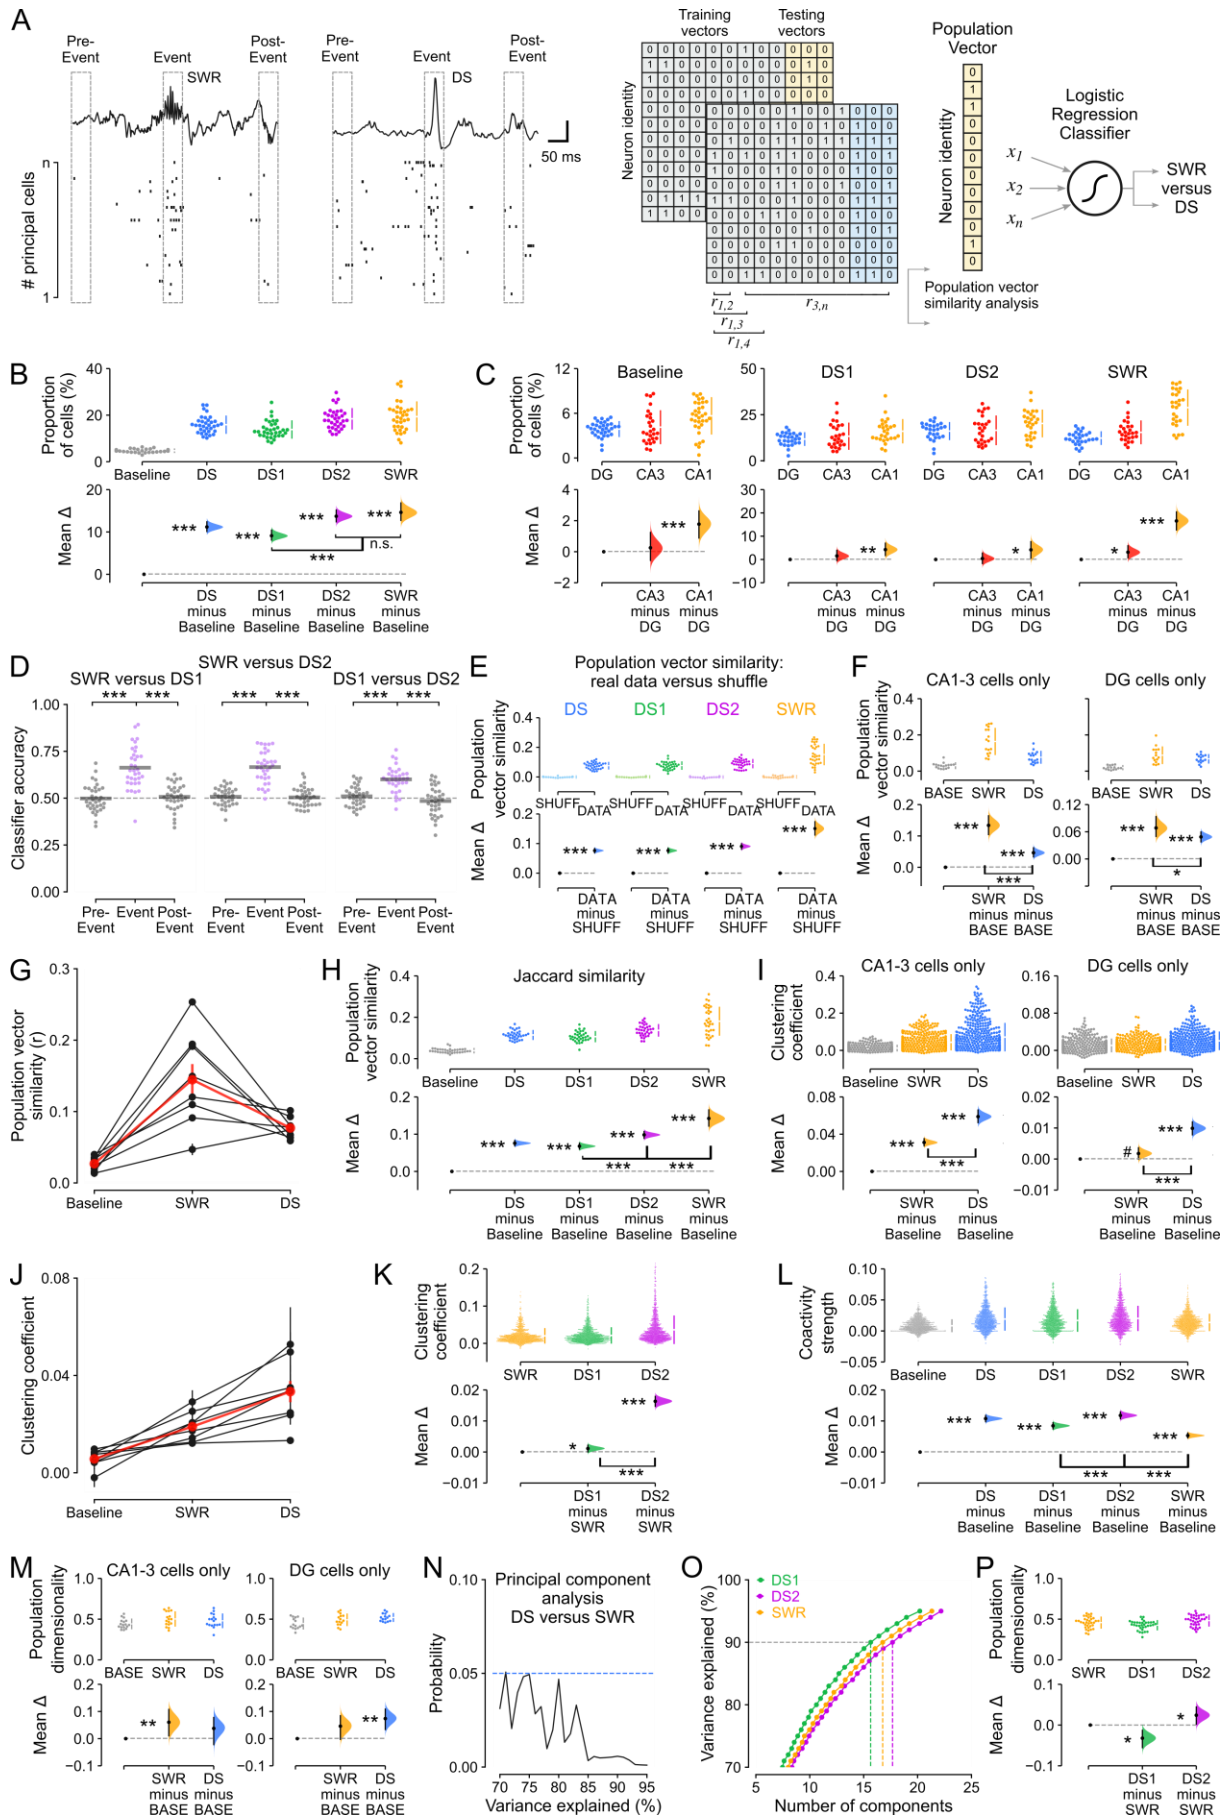
(F) Percentage of significantly activated principal cells, as defined by a z-score > 3 (within ± 20 ms of the event peak), during DS$_1$ (left) and DS$_2$ (right) events.

(G-J) Peri-event time histograms (G,H) showing z-scored firing rates ± 100 ms around the event peak and estimation plots (I,J) comparing overall mean firing rate (calculated across the entire recording session) to peri-event firing rate (calculated as the mean firing rate ± 5 ms around the peak of the event) for all principal cells during DS$_1$ and DS$_2$ events. DG n=921, CA3 n=388, CA1 n=887 principal cells (12 mice). Upper and lower plots as described in Figure S2O-P.

(K) Percentage of suppressed principal cells (i.e., cells with a z-score < 0 during the event peak) during DS$_1$ and DS$_2$ events.

(L) Peri-event time histograms showing z-scored firing rates ± 25 ms around the event peak for the lowest quartile of activated principal cells (i.e., the 25% least activated / suppressed principal cells) during DS$_1$ events. (DG n=230, CA3 n=97, CA1 n=221 principal cells).

For I and J, the test statistic is the mean difference, shown on the y-axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). P-values are from paired permutation tests, baseline (mean rate) versus event, $^{***}P < 0.001$.

Figure S4. Distinct population coactivity structures in DSs versus SWRs, related to Figure 3.

(A) Toy example illustrating construction of principal cell × event spiking activity matrices using the population firing vectors nested in 50 ms time windows centered on DSs, SWRs, or duration-matched (pre- and post-) control events. The proportion of active cells in each column vector (B-C) was defined as the proportion of cells with a non-zero spike count during the individual time window. Logistic regression classifiers (D) were trained using 75% of the population spiking activity vectors and tested with the remaining 25% of vectors, on any given session. We used the event type with the lowest number of epochs to determine the training and testing set size and then randomly subsampled the other event matrix to generate the same number of training and testing vectors for each event type, so that the classifier was balanced across event types. Separate matrices and classifiers were utilized for event, pre-event, and post-event epochs. For all of the analyses below, the inclusion criterion was that the session had to contain a minimum of 20 principal cells and 100 events of each type.

(B) Estimation plot showing the mean proportion of cells active during DS, DS_1, DS_2 and SWR events relative to equivalent duration-matched baseline periods during sleep when none of these events were present. Upper: raw data points (each point shows mean proportion of active cells in each vector across one recording session), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event type. Lower: difference (Δ) in proportion of active cells between baseline epochs versus DS, DS_1, DS_2 and SWR epochs. Note that all events contained more active cells than baseline epochs but there was no statistical difference between the proportion of active cells in DS_2 versus SWR events.

(C) Likewise, shown is the proportion of cells active during Baseline, DS_1, DS_2 and SWR events split by hippocampal subregion. Upper: raw data points (each point shows mean proportion of active cells by subregion), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each subregion. Lower: difference (Δ) in proportion of active DG cells versus CA3 and CA1 cells.

(D) Classifier performance for SWR versus DS_1, SWR versus DS_2, and DS_1 versus DS_2 population vectors. Dashed line shows chance performance.

(E) Estimation plot showing the population vector similarity for all event types compared to their control condition in which each event vector was correlated with 'shuffled' population spiking vectors, where the cell identity was randomly shuffled. Upper: raw data points (each point shows mean population similarity for one recording session or shuffled equivalent), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event. Lower: difference (Δ) in population similarity between real data for each event and its shuffled equivalent.

(F) Estimation plot showing population vector similarity for DS and SWR events compared to baseline control, as in Figure 3F, but separately for CA (left) and DG (right) principal cells. Upper: raw data points (each point shows mean population similarity by event type for one recording session), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event type. Lower: difference (Δ) in population similarity between baseline epochs versus DS and SWR epochs.

(G) Our results showing significantly higher population vector similarity for DS and SWR events versus baseline, and for SWR versus DS events (Figure 3F) were recapitulated at the level of individual mice (8/8 mice for DS and SWR events versus baseline, 7/8 mice for SWR versus DS events; $p < 0.05$ permutation tests for DS versus baseline, SWR versus baseline, DS versus SWR with n=number of population vector pairs per mouse). Black traces show individual mice ± SEM (calculated on the number of population vectors per mouse); red trace shows the group mean ± SEM (calculated on n=8 mice).

(H) Our results using the Pearson correlation to compare population vector similarity (Figure 3F-G) were replicated using the Jaccard similarity measure. Estimation plot showing the population vector similarity for all event types compared to baseline epochs. Upper: raw data points (each point shows mean population similarity by event type for one recording session), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event type. Lower: difference (Δ) in population similarity between baseline epochs versus DS, DS_1, DS_2 and SWR epochs.

(I) Estimation plot showing mean clustering coefficients (as Figure 3K) but for CA1-3 cells only (left panel) or DG cells only (right panel), which also show higher clustering coefficients for DS events.

(J) Our results showing significantly higher clustering coefficient for DS and SWR events versus baseline and DS versus SWR events (Figure 3K) were recapitulated at the level of individual mice (8/8 mice for DS and SWR events versus baseline, 7/8 mice for SWR versus DS events; $p < 0.05$ permutation tests for DS versus baseline, SWR versus baseline, DS versus SWR with n=number of neurons per mouse). Black traces show individual mice ± SEM (calculated on the number of neurons per mouse); red trace shows the group mean ± SEM (calculated on n=8 mice).

(K) Estimation plot showing that the neuronal coactivity graphs nested in both DS_1 and DS_2 events contained significantly stronger triads of coactive nodes compared to SWR graphs, as indicated by higher mean clustering coefficients. This was notably the case for DS_2.
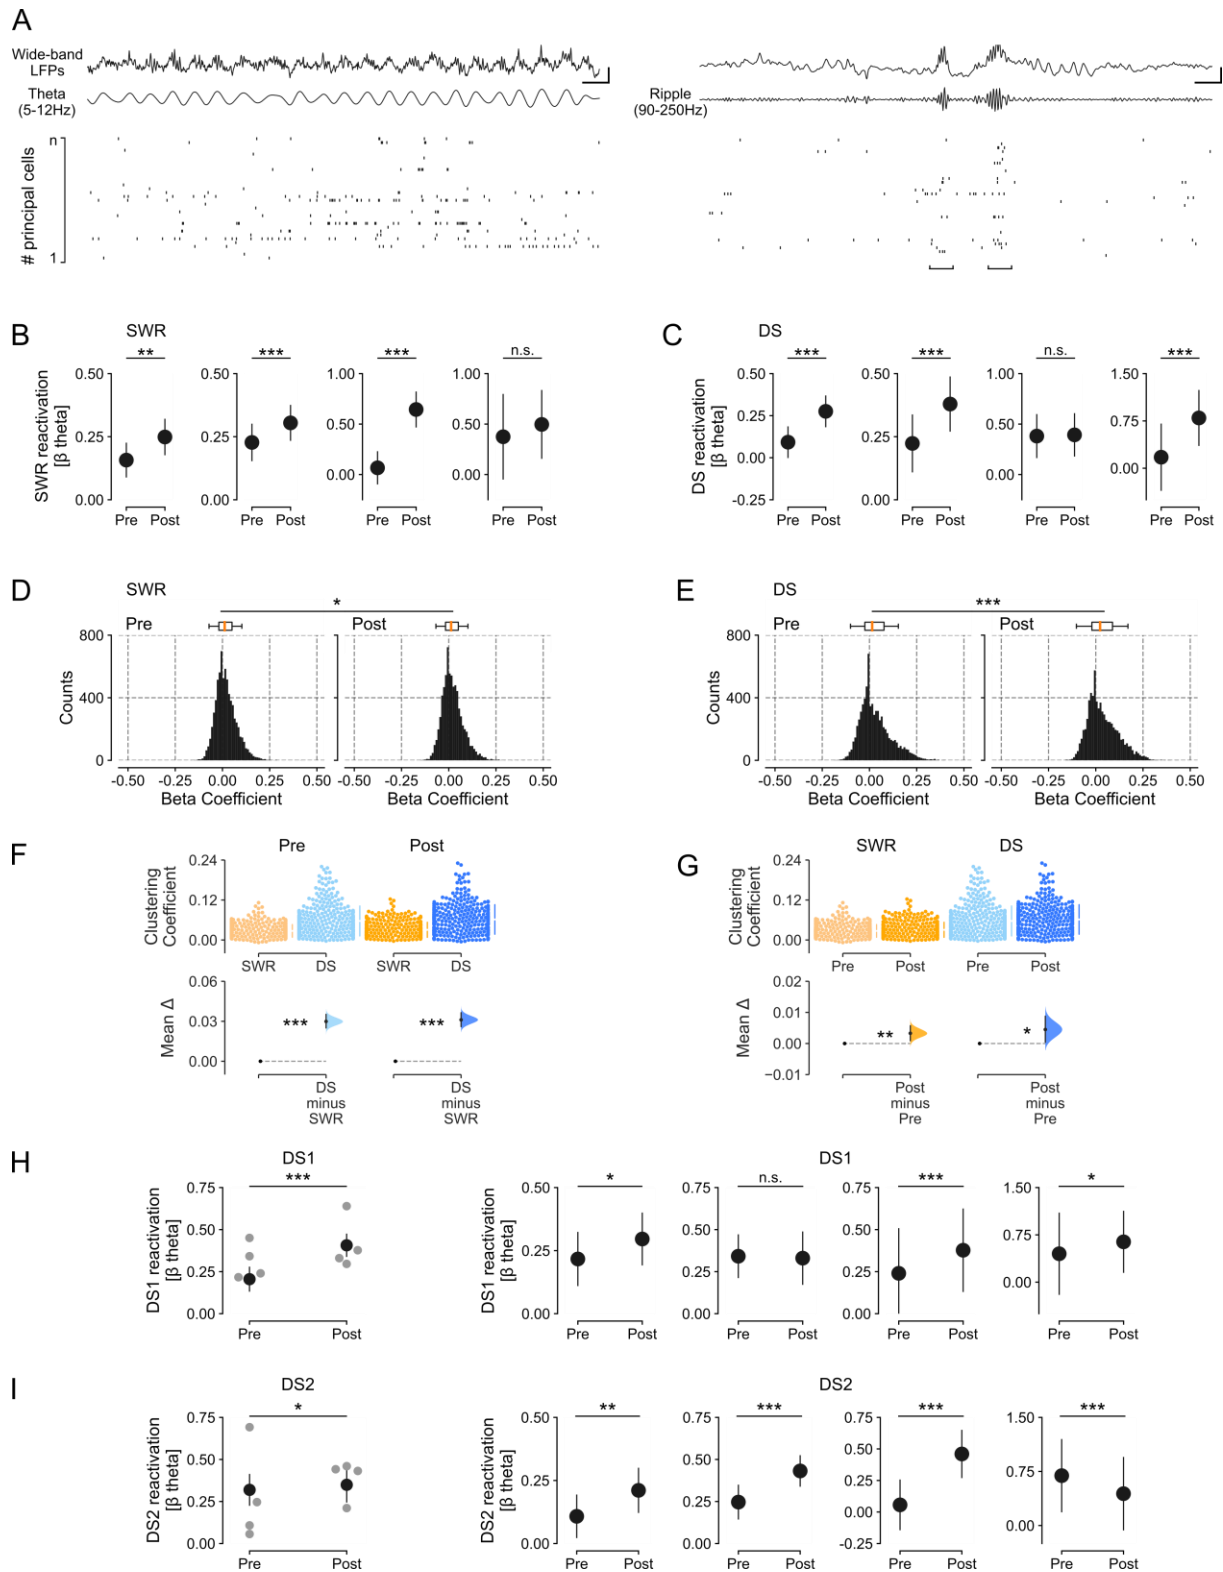
(L) Single-neuron coactivity strength. As an alternative method to the triadic clustering coefficient parameter, we also show in this estimation plot that SWR and DS events differ in neuronal coactivity strength and from baseline epochs. We defined the single-neuron coactivity strength as the average

pairwise coactivity relation of a given neuron with its population peers. For any two neurons (i, j), we obtained the regression coefficient β from a generalized linear model predicting the spike discharge of neuron j from the activity of neuron i while regressing out the activity of the remaining population. The strength of neuron i is then the average across all the weights β_{ij}.

(M) Estimation plot showing population dimensionality required to explain 90% of the variance in DS and SWR events compared to baseline control separately for population vectors containing only CA (left) or DG (right) principal cells. Upper: raw data points (each point shows mean population dimensionality by event type for one recording session), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event type. Lower: difference (Δ) in population dimensionality between baseline epochs versus DS and SWR epochs. Importantly, considering here DG and CA regions separately for the dimensionality analysis bears the caveat that, to comply with our criterion of at least 20 simultaneously recorded principal cells for each recording session, these analyses use far fewer recording days and the resulting neuronal population vectors are much smaller (n=17 days, n=25.8±1.0 cells per vector for CA; n=15 days, n=25.5±0.9 cells per vector for DG) than the data presented in Figure 3K (n=34 days, n=37.2±1.8 cells per vector), which limits the comparison.

(N) PCA to compare the dimensionality of SWR versus DS matrices (cell × event number), matching the number of events for each event type, determining the number of components required to explain 70–95% of the variance. In each case, the dimensionality was significantly higher for DS versus SWR events at α < 0.05 (Wilcoxon test for paired samples, one-tailed).

(O-P) The number of components required to explain equivalent amounts of variance was lower in DS$_1$ versus SWR and DS$_2$, and higher in DS$_2$ versus SWR events (N); as illustrated by the estimation plot in panel O, showing that a lower number of principal components was required to explain 90% of the variance across the population vectors nested in DS$_1$, compared to SWR and DS$_2$ events.

For B-F,H-I,K-M,P, the test statistic is the mean difference, shown on the y-axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). P-values are from paired permutation tests, baseline versus event (B,F,H,I,L,M), event versus pre-event and event versus post-event (D), data versus shuffle (E), event versus event (K,P); or unpaired permutation tests, cell type versus cell type (C), *$P < 0.05$, **$P < 0.01$, ***$P < 0.001$.

Figure S5. Offline reactivation of theta population coactivity in individual mice, related to Figure 4.
(A) Raw data examples of hippocampal principal cells spiking activity during active exploration marked by theta oscillations (Left) and offline sleep/rest (Right; example SWRs shown). Scale bars 0.5 mV and 150 ms (Left) or 50 ms (Right). We applied our peer-to-peer coactivity analysis (Figure 3A) to obtain the waking patterns of population coactivity nested in theta cycles during exploration, and the offline patterns of population coactivity nested in either DS or SWR events during sleep/rest before and after exploration. With these, we next computed DS and SWR reactivation by measuring the tendency of the peer-to-peer theta firing associations to reoccur in post-exploration sleep/rest DS (or SWR) events, while controlling for prior pre-exploration DS (or SWR) coactivity, using a linear mixed model (Figure 4A).

(B) SWR reactivation in individual mice (measured by the β coefficients of the linear regression that predicted post-exploration SWR coactivity from waking theta coactivity, controlling for pre-exploration SWR coactivity).

(C) DS reactivation in individual mice (measured by the β coefficients of the linear regression that predicted post-exploration DS coactivity from waking theta coactivity, controlling for pre-exploration DS coactivity).

(D-E) Distribution of peer-to-peer coactivity values (β coefficients) for SWRs (D) and DSs (E) in the pre- and post-exploration sleep/rest sessions. Significance was tested using the Wilcoxon test for matched pairs: post > pre for DSs (U=1198580, n=7310, p < 0.001); and for SWRs (U=12839635, n=7310, p < 0.001).

(F-G) Estimation plots showing that the neuronal coactivity graphs nested in DSs contained significantly stronger triads of coactive nodes compared to SWR graphs, as indicated by higher mean clustering coefficients, during both pre- and post-exploration sleep sessions (F). Also, DSs and SWRs contained significantly stronger triads of coactive nodes during post- versus pre-exploration sleep sessions (G).

(H) Mean DS_1 reactivation pooled across mice (left panel) and in individual mice (right four panels; measured by the β coefficients of the linear regression that predicted post-exploration DS_1 coactivity from waking theta coactivity, controlling for pre-exploration DS_1 coactivity)

(I) As H but for DS_2 reactivation.

For F and G, the test statistic is the mean difference, shown on the y-axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). P-values are from paired permutation tests, event versus event (F), pre versus post (G), *$P < 0.05$, **$P < 0.01$, ***$P < 0.001$.

Figure S6. Closed-loop optogenetic suppression of dentate granule cells during DSs, related to Figures 5 and 6.

(A-C) During DS-triggered DG light-delivery for the optogenetic silencing of DG cells, there was no difference in the amplitude (A, B) or duration (C) of DSs in light-on versus light-off (control) DS events.
(D-I) DG light-delivery did not affect CA1 LFPs (D), ripple duration (E), intra-ripple frequency (F), ripple power (G), nor the probability of ripple occurrence (H-I; 22 recording sessions in 8 mice).
(J) DS-triggered DG light-delivery significantly reduced firing rates in DG, CA3 and CA1 principal cells relative to DSs with no light delivery (Control), in a paired analysis (same cells under both conditions). Each panel shows mean ± SEM firing rate. Average firing rates across the entire recording session (including sleep and exploration epochs) were: DG: 1.4±0.1 Hz, CA3: 1.7±0.3 Hz, CA1: 1.5±0.1 Hz.
(K) Estimation plot showing the maximum firing rate (during DS events) in DG, CA3 and CA1 principal cells during DS-Sync or Control (DSs with no light delivery), as in panel J. Upper: raw data points (each

point shows maximum firing rate), with the gapped lines on the right as mean (gap) ± s.d. (vertical ends) for each event type. Lower: difference (Δ) in maximum firing rates between DS-Sync and no-light control condition (paired permutation test in DG: n=216 cells, CA3: n=50 cells, and CA1: n=133 cells from n=22 recording sessions in 8 mice).

(L) Estimation plot for object preference in a 'no-laser' control group of mice, showing significantly more time spent investigating the novel object (n=20 test sessions in 5 mice).

(M-O) In the continuous novel object recognition task, the total time spent exploring the objects (M), the number of laser pulses delivered (N) and the number of SWRs detected during sleep sessions (O) did not differ between the DS-delay and DS-sync conditions.

(P-S) The tone fear task in which mice had 5 tone-shock pairings during conditioning followed by either DS-Delayed or DS-Sync stimulation, and then fear memory recall. Mice froze more during recall than the baseline (Q), but this did not differ between the DS-Delay and DS-Sync groups (R; n=8 sessions in 4 mice). Mice received an equivalent number of laser pulses in the two groups (S).

(T) In the novel position recognition task, mice in the DS-Delayed and DS-Sync groups received equivalent numbers of laser pulses.

E-G, M-O, S,T show mean ± SEM. For B,C,I,K,L,Q,R, the test statistic is the mean difference, shown on the y-axis of each lower plot (black dot, mean; black ticks, 95% confidence interval; filled curve, sampling-error distribution). P-values are from paired permutation tests, Control (no laser) versus DS-Sync (B,C,I,K,R), Novel versus Familiar (L), Baseline versus DS-Delay and DS-Sync (Q), *$P < 0.05$, **$P < 0.01$, ***$P < 0.001$.

Table S1. Ratio of DG to CA neurons influence on the dimensionality and similarity of population firing vectors, related to Figure 3.

Event	Dependent variable	Independent variables	Degrees of freedom	r-value	p-value
DS	Dimensionality	Ratio of DG:CA neurons	33	0.24	0.18
DS1	Dimensionality	Ratio of DG:CA neurons	33	0.12	0.51
DS2	Dimensionality	Ratio of DG:CA neurons	33	0.26	0.14
SWR	Dimensionality	Ratio of DG:CA neurons	33	0.32	0.07
DS	Similarity	Ratio of DG:CA neurons	33	-0.06	0.75
DS1	Similarity	Ratio of DG:CA neurons	33	-0.13	0.46
DS2	Similarity	Ratio of DG:CA neurons	33	-0.08	0.66
SWR	Similarity	Ratio of DG:CA neurons	33	-0.39	0.02

Table S2. Linear mixed model analysis for SWR and DS reactivation of waking theta coactivity patterns, related to Figure 4.

Event	Dependent variable	Independent variables	No. Observation
SWR	Post-exploration	Theta coactivity Pre-exploration	7310
SWR	Pre-exploration	Theta coactivity Post-exploration	7310
DS	Post-exploration	Theta coactivity Pre-exploration	7310
DS	Pre-exploration	Theta coactivity Post-exploration	7310

Event	Dependent variable	Independent variables	β coefficient	CI (95%)	z	Prob.
SWR	Post-exploration	Theta coactivity	0.37	[0.32, 0.43]	13.1	P < 0.0001
		Pre-exploration	0.11	[0.09, 0.13]	9.5	P < 0.0001
SWR	Pre-exploration	Theta coactivity	0.20	[0.15, 0.26]	7.0	P < 0.0001
		Post-exploration	0.11	[0.04, 0.08]	9.5	P < 0.0001
DS	Post-exploration	Theta coactivity	0.39	[0.31, 0.46]	9.7	P < 0.0001
		Pre-exploration	0.25	[0.23, 0.27]	23.3	P < 0.0001
DS	Pre-exploration	Theta coactivity	0.24	[0.16, 0.32]	5.8	P < 0.0001
		Post-exploration	0.27	[0.25, 0.30]	23.3	P < 0.0001

Table S3. Counts of principal cells by session and mouse, related to Figures 1-6.

Figure number / panel	# sessions	# mice	# principal cells (mean ± SEM per mouse)		
			DG	**CA3**	**CA1**
Figure. 1F,G	73	12	921 (76.5±22.2 per mouse)	388 (32.3±9.6 per mouse)	887 (73.6±17.9 per mouse)
Figure. 2A-C	8	3	n/a	n/a	n/a
Figure. 2E-G	73	12	921 (76.5±22.2 per mouse)	388 (32.3±9.6 per mouse)	887 (73.6±17.9 per mouse)
Figure. 3B-D,F-G,K,L	34	8	647 (85.6±28.9 per mouse)	169 (25.6±9.1 per mouse)	449 (77.9±23.0 per mouse)
Figure. 4B-C	9	4	114 (28.5±7.2 per mouse)	5 (1.25±1.1 per mouse)	232 (58.0±18.0 per mouse)
Figure. 5F-G	43	9	548 (60.9±11.8 per mouse)	n/a	n/a
Figure. 5I-J	13	3	181 (60.3±14.0 per mouse)	n/a	n/a
Figure. 6H (DS) Figure. 6H (SWR)	10 12	3 3	79 (26.3±10.3 per mouse) 173 (57.7±15.7 per mouse)	22 (7.3±3.8 per mouse) 32 (10.7±6.4 per mouse)	73 (24.3±4.0 per mouse) 68 (22.7±8.1 per mouse)