preprint

Loss of Plasticity and Its Recovery via Replay-Driven Compression in a Spiking Model of Hippocampal CA3

Anonymous authors
Paper under double-blind review

Abstract

Artificial neural networks trained sequentially on new information are prone to interference: new learning can degrade or overwrite what was previously acquired. Research into this failure mode has recently led to a distinction between catastrophic forgetting, akin to overwriting, and loss of plasticity, in which both older and newer memories are degraded. Here we ask whether similar limitations arise in biologically-grounded spiking neural networks and study how nature might address them. In a biologically detailed recurrent circuit that models the mammalian hippocampus, multiple memory traces (modeling spatial environments) are learned sequentially by a common network and evaluated through their spontaneous and cue-evoked replay. We find that sequential learning progressively saturates recurrent connectivity, increasing interference among memories and eventually eliminating the network's ability to replay stored representations. This collapse resembles loss of plasticity more closely than classical catastrophic forgetting. We then show that alternating online learning with offline and spontaneous replay-driven synaptic compression counteracts this saturation: it reorganizes the recurrent weight distribution while preserving recall, allowing the network to store additional memory traces. Together, these results identify synaptic compression and plasticity schemes organized around spontaneous replay as complementary and neurobiologically-grounded strategies for mitigating interference in recurrent spiking memory networks. This identification connects memory capacity limits observed in biology to an active distinction in the machine learning literature between loss of plasticity and catastrophic forgetting.

1 Introduction

Continual learning in artificial neural networks remains fundamentally limited by catastrophic forgetting: when a network is trained sequentially on new tasks, the same parameters that encode new information overwrite the representations needed for earlier ones (McCloskey and Cohen, 1989; French, 1999; van de Ven et al., 2025). Regularization-based approaches such as elastic weight consolidation (Kirkpatrick et al., 2017) and synaptic intelligence (Zenke et al., 2017) mitigate this by selectively protecting important parameters, and replay-based approaches (Shin et al., 2017; Rolnick et al., 2019) interleave old and new experience during training. Catastrophic forgetting persists even at the scale of modern large language models, and incremental learning's interaction with model compositionality remains an active problem (Porrello et al., 2025).

Complementary Learning Systems (CLS) theory offers a biological account of how this problem is avoided in the brain: rapid, interference-prone encoding in the hippocampus is complemented by gradual, replay-driven transfer to neocortex, which supports stable long-term storage (McClelland et al., 1995; Kumaran et al., 2016). This theory continues to motivate new machine learning architectures directly: recent work has proposed hybrid cortico–hippocampal networks for task-agnostic continual learning, evaluated on standard

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

image-classification benchmarks (Shi et al., 2025), and hippocampus-inspired architectures for continual learning more broadly (Kapoor et al., 2026). It remains less clear how interference and capacity limits emerge mechanistically in a biophysically detailed spiking circuit, where memory is expressed not as classification accuracy but as the fidelity of spontaneous and cued replay.

A separate line of work tackles catastrophic forgetting directly in spiking neural networks. Golden et al. (2022) show that sleep-like offline reactivation constrains synaptic weights to a manifold shared by old and new tasks, mitigating forgetting through spontaneous replay and local, unsupervised plasticity. Other work uses *replay* in the rehearsal sense: storing and replaying past training examples (Proietti et al., 2023), rather than activity the network generates on its own. To our knowledge, no prior study of interference in spiking networks has validated memory retention across sequentially learned tasks using the sequential, time-compressed structure of spontaneous and cue-evoked replay itself. The scoring method itself is not new: replay is standardly validated by comparing spike-time order to the underlying behavioral trajectory (Davidson et al., 2009; Foster and Wilson, 2006; Diba and Buzsáki, 2007), as in Ecker et al. (2022) and Baron et al. (2026). What is new is applying it here: we use this readout, rather than task performance or rehearsal-style replay, to measure how memory capacity degrades as a recurrent network learns more environments.

We investigate these questions in a biophysically detailed spiking model of hippocampal area CA3, building on the network of Ecker et al. (2022), chosen for its support of empirically observed bidirectional replay (Diba and Buzsáki, 2007), which reproduces a learned trajectory in both forward and reverse order. Recurrent connectivity is further shaped by the anti-Hebbian offline plasticity rule introduced by Baron et al. (2026), chosen for its ability to reorganize connectivity during offline replay while preserving the network's capacity to generate replay. We systematically increase memory load by sequentially learning multiple spatial environments within a shared recurrent connectivity matrix, and evaluate memory integrity through spontaneous and cue-evoked replay.

Our contributions are as follows. First, we characterize a clear capacity limit under symmetric timing long-term potentiation (st-LTP), beyond which replay of previously learned environments is lost. Second, we show that replay-driven anti-Hebbian plasticity combined with homeostatic synaptic compression extends this capacity by reorganizing recurrent connectivity while preserving recall. Third, we use the sequential structure of spontaneous and cue-evoked replay as a direct measure of memory integrity under increasing load, revealing a retrieval-dependent capacity limit not captured by a single task-performance metric. Together, these results connect specific, testable synaptic and circuit-level mechanisms to open problems in continual learning, at a level of biological detail largely absent from current brain-inspired continual learning architectures

2 Results

2.1 Sequential learning saturates recurrent connectivity and collapses replay

We first replicated the results of Ecker et al. (2022) in a two-environment setting. Consistent with the original study, the network spontaneously generated replay events corresponding to both learned environments. We next asked whether the same recurrent network can continue acquiring additional environments without degrading replay of previously learned traces. We hypothesized that sequential learning under st-LTP would progressively increase interference because each new environment is written into the same CA3 PC–PC connectivity matrix, leading to synaptic saturation and loss of the structured recurrent weights required for spontaneous replay.

To test this, we ran 12 independent simulations while increasing the number of learned environments from 2 to 6 (Fig. 1a). After each training episode, we detected replay candidates and classified them using the replay validation procedure described in Methods.

Replay quality degraded sharply with memory load (Fig. 1a). Replay remained robust after learning two to three environments, but both replay candidates and validated events

dropped once more than four environments were stored. Under the standard criteria, no validated replay events were detected after learning five environments (Fig. S1a), and none were detected after six.

To test whether replay was merely masked by elevated background firing, we repeated the analysis with a more stringent spike-rate filter (70 Hz instead of 10 Hz). This revealed replay-like sequences in five-environment simulations (Fig. S1b), but replay remained absent after learning six environments, indicating that background activity alone does not explain the loss of replay.

Overall, sequential learning under st-LTP shows a clear capacity limit: replay progressively deteriorates and ultimately disappears as recurrent connectivity saturates, motivating mechanisms that compress or reorganize synapses to reduce interference.

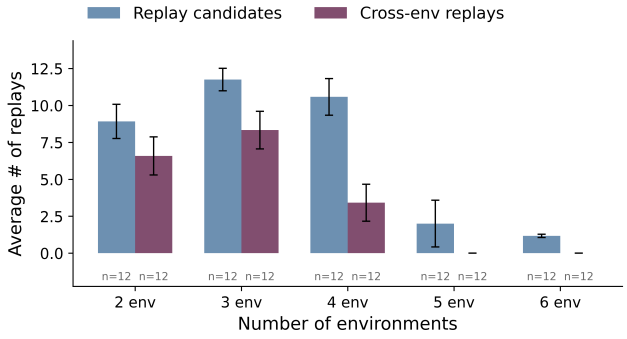

(a) Progressive degradation of replay activity with increasing memory load

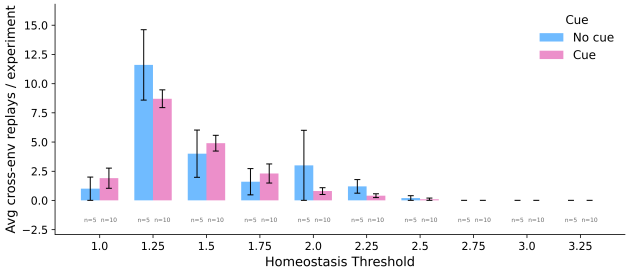

(b) Effect of progressive synaptic homeostasis on replay recovery in a five environments setup

Figure 1: Replay detection under increasing memory load, and recovery via homeostatic synaptic compression. (a) Replay statistics as the number of learned environments increases. Candidates were evaluated against the place-cell ordering of *all previously learned environments* and classified as replay if they satisfied validation criteria for any one. Both candidates and validated events decline sharply beyond four environments, with none detected after five. (b) Effect of progressive synaptic homeostasis on replay recovery after learning five environments. The x-axis shows the homeostatic threshold applied to recurrent weights; the y-axis shows average replay events per experiment (blue: 15 s spontaneous offline replay; pink: 6 s cue-evoked replay). Moderate compression increased replay, peaking at 1.0–2.0 nS thresholds, before declining at higher thresholds. Moderate compression increased both spontaneous and cue-evoked replay, with performance peaking near 1.25 nS and declining at higher thresholds. Given the limited number of spontaneous-replay runs, the present data do not establish a reliable difference between the compression levels supporting spontaneous and cue-evoked replay. Max weight was capped at 4.0–4.1 nS throughout.

2.2 Homeostatic compression restores replay after saturation

The loss of replay after learning five environments suggested the network had reached its storage capacity under st-LTP, likely due to progressive synaptic saturation: repeated learning potentiates an increasing fraction of recurrent synapses, degrading the structured connectivity required for replay generation. Consistent with this, Figure 2 shows a substantial increase in weak-to-moderately potentiated synapses (0.5–1.5 nS) after five environments. To counteract this, we introduced a homeostatic compression mechanism inspired by synaptic tagging and homeostatic plasticity (Frey and Morris, 1997; Okuno et al., 2012; Turrigiano et al., 1998; Tononi and Cirelli, 2014): following the learning of five environments, all recurrent synapses below a specified threshold were reset to baseline (0.1 nS), while stronger synapses were preserved. We varied this threshold from 1.0 to 3.25 nS in steps of 0.25 nS and evaluated replay across two independent runs at each threshold (120 simulations total). For each threshold and run, we assessed spontaneous replay (one 15 second offline simulation, evaluated via the cross environment classification procedure in Methods) and cue evoked recall (stimulating 100 place cells in the first quarter of the track and monitoring the subsequent six seconds, for each of the five learned environments, five simulations per run).

Figure 1b summarizes the results. Moderate compression substantially increased replay relative to the non homeostatic condition, peaking at a threshold of 1.25 nS for both spontaneous and cue evoked replay, then declining sharply and becoming undetectable by 2.75 nS. Spontaneous replay is currently estimated from only two simulations per threshold, so the shape of its decline, in particular the local rebound at 1.5 nS, should be treated as preliminary pending additional runs. At the sampling density collected so far, cue evoked and spontaneous replay decline over a similar range rather than showing a clear dissociation; together, these results show that homeostatic compression can restore replay in networks that otherwise fail after learning multiple environments, but do not yet establish whether cue evoked recall is supported by a sparser subset of connections than spontaneous replay.

2.3 Replay-driven anti-Hebbian plasticity reorganizes connectivity while preserving recall

Homeostatic compression alone remained sensitive to the choice of threshold, suggesting that additional mechanisms could further compress recurrent connectivity while preserving recall. Since synaptic homeostasis and replay are both thought to occur during sleep (Tononi and Cirelli, 2014; Buzsáki, 2015; Wilson and McNaughton, 1994), we tested whether replay-driven plasticity could enhance homeostatic compression. We incorporated the anti-Hebbian spike-timing-dependent plasticity (anti-hSTDP) rule, shown previously to preserve and strengthen replay dynamics (Baron et al., 2026). Following the learning of each environment, the network underwent a 25-second offline replay phase governed by anti-hSTDP before learning the next, mimicking the alternation between learning and sleep-associated replay observed in the hippocampus. To verify that this combined protocol did not impair replay, we ran a separate control experiment: after learning five environments with interleaved anti-hSTDP replay and a 1.0 nS homeostatic threshold, 10 independent 15-second offline simulations (70 Hz detection criterion) generated an average of 7 replay events each, confirming that the combined protocol preserves spontaneous replay while substantially altering synaptic connectivity.

Figure 2 summarizes the resulting weight distributions. As expected, synapses in the 0.5–1.0 nS range were eliminated by the homeostatic threshold. More surprisingly, synapses in the 1.0–1.5 nS range decreased by approximately 79%, despite exceeding the threshold and therefore not being directly targeted by the homeostatic rule – indicating that anti-Hebbian replay actively weakens a large population of moderately potentiated synapses. The strongest synapses were also redistributed non-uniformly: the 4.0–4.1 nS population decreased by approximately 7%, while the 3.5–4.0 nS population increased by approximately 17%. Together, these results indicate that anti-Hebbian replay does not merely eliminate weak synapses, but reorganizes connectivity by selectively weakening moderate and strongly potentiated connections while preserving a sparse backbone of high-strength synapses.

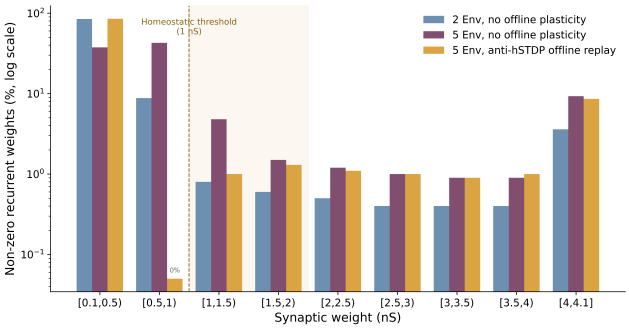

Figure 2: Distribution of non-zero recurrent synaptic weights across memory load and synaptic compression (Approx. 6.8 million synapses). Bars show networks trained on 2 or 5 environments with no plasticity during the offline phase (averaged across 12 simulations each), illustrating progressive saturation of recurrent connectivity with memory load, alongside the distribution after 5 environments when plasticity *is* applied during the offline phase, via anti-hSTDP replay combined with a 1.0 nS homeostatic threshold after each environment (averaged across 10 simulations). The dashed line marks this 1 nS threshold, which trivially depletes the $[0.5, 1)$ bin. The $[1, 1.5)$ and $[1.5, 2)$ bins fall by 79% and 17% relative to the uncompressed 5-environment network; this reduction is not a threshold effect and instead reflects the anti-hSTDP replay phase redistributing weight away from that range. This redistribution is not simply a downward shift: the $[3.5, 4)$ bin grows by 17% relative to the uncompressed network, while the strongest weights in $[4, 4.1]$ shrink by 7%, indicating that anti-hSTDP replay drives a more concentrated population of high-strength synapses just below the weight ceiling, relaxing the very strongest connections while pulling moderate weights down.

Remarkably, this extensive synaptic compression did not eliminate memory recall. Cue-evoked replay remained robust across the five learned environments despite the substantial reduction in recurrent connectivity. These findings suggest that successful recall depends on a relatively small subset of recurrent synapses, and that the subset surviving anti-Hebbian plasticity and homeostatic compression is sufficient to preserve replay despite extensive removal of moderately potentiated connectivity. Taken together, these results demonstrate that the combined anti-Hebbian replay and homeostatic protocol produces extensive reorganization of recurrent connectivity without compromising replay generation. Despite substantial reductions in the population of moderately potentiated synapses, spontaneous replay remained robust and cue-evoked recall was preserved across all five learned environments.

2.4 State-dependent compression extends memory capacity

The previous sections demonstrated that combining Anti-hSTDP replay with homeostatic compression substantially reorganizes recurrent connectivity while preserving replay generation and memory recall after learning five environments. We next investigated whether this replay-driven compression protocol could further increase the network's storage capacity by allowing additional environments to be learned without disrupting previously stored memory traces.

Starting from the compressed five-environment network, we continued training additional environments using st-LTP, followed by a 25-second offline Anti-hSTDP replay phase after each new environment, evaluated *before* homeostatic compression was applied. Recurrent synapses with strengths below a threshold in the range 1.5–3.25 nS were then reset to baseline before learning the next environment; this range was chosen based on the effective range identified in Section 2.2, and results were pooled across it as detailed in Appendix A.5. Cued-recall trials, applied after homeostatic compression of all learned environments, used the strengthened cueing protocol described in Section 4.3 (1 ms cue ISI), which was required to obtain reliable recall beyond five environments.

Table 1 summarizes the results. At the most extensively replicated homeostatic threshold (1.75 nS), cue-evoked recall of the first learned environment was observed in 7/7 simulations through six environments, 4/7 after seven environments, 1/6 after eight, and 0/6 after nine. The broader threshold sweep showed the same qualitative decline across a range of compression strengths (See A.5). Spontaneous replay degraded sharply beginning at six environments and was largely absent from seven environments onward, whereas cue-evoked recall remained robust through six environments and declined more gradually thereafter. Spontaneous and cue-evoked replay thus do not degrade in lockstep as memory load increases: autonomous replay fails first, while a sufficiently strong cue can still retrieve previously learned structure for several additional environments after full compression. This indicates that storage capacity under this protocol is retrieval-dependent rather than a single number. Replay-driven compression therefore mitigates interference relative to st-LTP alone and extends the range over which cued recall is possible, while preserving an eventual capacity limit as recurrent connectivity becomes increasingly utilized.

Table 1: Replay capacity under the replay-driven synaptic compression protocol, pooled across homeostatic thresholds from 1.5 to 3.25 nS (14 simulations; 13 for eight and nine environments; see Appendix A.5 for the per-threshold breakdown). Each entry reports the fraction (and percentage) of simulations producing at least one validated replay event. Spontaneous replay was measured during the 25-second anti-hSTDP phase immediately after training on each environment, *before* that environment's own homeostatic threshold was applied. Cued recall was measured for the first learned environment *after* homeostatic compression of all learned environments.

Environments	**Spontaneous replay** *(before homeostasis)*	**Cued recall (Env. 1)** *(after homeostasis)*
5	14/14 (100%)	13/14 (93%)
6	11/14 (79%)	13/14 (93%)
7	0/14 (0%)	7/14 (50%)
8	1/13 (8%)	3/13 (23%)
9	0/13 (0%)	0/13 (0%)

3 Limitations

Our model captures a specific, biophysically detailed learning regime built on online exploration followed by offline anti-Hebbian replay and homeostasis (Ecker et al., 2022; Baron et al., 2026), rather than a general account of online learning in artificial systems. Our results are also dependent on specific model parameters: for example, reducing the place-cell ratio to 0.35, that is, increasing representational sparsity, increased the number of environments the network could store to 9, consistent with classical theory in which sparsity itself acts as a lever on storage capacity (Tsodyks and Feigel'man, 1988). However, this manipulation required also raising the maximum synaptic weight from 4.1 nS to 5.1 nS to preserve replay, so whether the capacity gain reflects sparsity itself or the accompanying weight bound is not yet dissociated (See A.8). We do not evaluate our model on standard continual learning benchmarks (van de Ven et al., 2025), so the capacity limits and recovery mechanisms we report are not directly comparable to task-accuracy-based measures of interference used in the artificial neural network literature. Relatedly, memory in our network is represented by sequential activation across an ensemble of neurons that can be read out directly at any point in time, whereas conventional artificial neural networks typically compress information into a comparatively small output layer read out only at inference. Translating our capacity- and replay-based findings into architectures with this kind of compressed readout is not straightforward, and identifying an appropriate analogue of "replay" for such networks remains an open direction for future work.

4 Methods

4.1 Model overview

We build on the hippocampal replay model of Ecker et al. (2022) implemented in Brian2 (Stimberg et al., 2019) and extended in prior work (Baron et al., 2026; Liao et al., 2024; Liu et al., 2023). Simulations alternate between three operations: encoding new experience through online potentiation (Phase 1), compressing recurrent connectivity through replay-driven anti-Hebbian plasticity during offline replay (Phase 2), and recycling weak connections through homeostatic reset before the next environment is learned.

The CA3 network consists of 8,000 excitatory pyramidal cells (PCs) and 150 inhibitory basket cells (BCs), both modeled as single-compartment adaptive integrate-and-fire neurons, with PC→PC and BC→PC connections sampled with probability 0.1 and PC→BC and BC→BC connections sampled with probability 0.25. Plasticity is applied according to the phase-specific rules described below.

4.1.1 Phase 1: online learning.

We simulate repeated traversals of a 3-m linear track at a fixed speed of 32.5 cm/s (single traversal: 9.2 s; total simulation: 400 seconds, approximately 43 traversals). Approximately half of the CA3 PCs are designated as place cells, with spatially localized Gaussian firing fields and theta modulation, while the remaining PCs exhibit low-rate background activity. To model sequential learning of distinct environments, we independently randomize the mapping between place cells and spatial locations for each environment while retaining the same underlying CA3 network, so each environment generates a distinct sequence of place-cell activity encoded into the shared CA3 PC→PC weight matrix through st-LTP (Fig. S2b). Full spike-train generation details and plasticity equations are in Appendix A.3.1.

4.1.2 Phase 2: Offline Replay

During the offline phase, the recurrent PC–PC weights learned during Phase 1 were carried forward, while the CA3 network received stochastic mossy-fiber input, modeled as independent Poisson spike trains, as its sole external drive, spontaneously generating time-compressed replay events interspersed with sparse non-replay activity. Recurrent PC–PC weights were held fixed unless otherwise specified; in experiments examining replay-driven synaptic compression, we followed Baron et al. (2026) and enabled recurrent PC–PC plasticity during the offline phase to investigate how replay-driven connectivity changes interact with sequential learning and increasing memory load. Full details of the offline input model, plasticity rules, and parameter values are in Appendix A.3.2.

4.2 Homeostatic synaptic reset

To counteract the progressive accumulation of synaptic potentiation across sequential learning sessions, we applied a threshold-based homeostatic synaptic reset inspired by the selective synaptic renormalization proposed by Frey and Morris (1997). After each training session, PC–PC synapses with conductance below a predefined threshold were reset to their initial conductance of 0.1 nS, while synapses above threshold were left unchanged, preserving strongly potentiated synapses assumed to encode stable memory traces while recycling weak ones for subsequent learning. Homeostasis was applied before testing throughout, though at a different point in each protocol: directly after training in Section 2.2, and after the anti-hSTDP compression phase in Sections 2.3 and 2.4.

4.3 Replay Detection (Cued and spontaneous)

Replay candidates were identified from population-rate transients in inhibitory basket-cell (BC) activity and validated based on the sequential structure of pyramidal-cell (PC) spiking across place fields. In simulations with five or more environments, we increased the PC spike-rate threshold used for replay validation from 10 Hz to 70 Hz, filtering elevated

background activity before the regression-based sequence analysis. Cued replay was initiated by sequentially activating 100 place cells (positions 1000–1099 in the environment-specific place-cell ordering), with each neuron receiving 150 Hz Poisson input within a 20-ms window and successive activation windows separated by 1 ms, evaluated using the same validation criteria as spontaneous replay. Complete replay detection and validation criteria are in Appendix A.4.

5 Discussion

Our results support three main claims. First, a recurrent spiking network learning multiple memory traces via st-LTP has a limited capacity to replay previously learned traces, with replay quality degrading sharply and then collapsing once storage demand exceeds the network's capacity (Fig. 1a). The Ecker et al. (2022) architecture supports spontaneous bidirectional replay, which requires structured connectivity in both directions. Because of this, replay serves as a stringent test of the network's synaptic organization. Its collapse under increasing memory load is accordingly strong evidence that sequential learning progressively degrades this structure. Second, a homeostatic regime that resets weak synapses while preserving strong ones counteracts this degradation and extends memory capacity, indicating that the information relevant to replay is concentrated in a small subset of strongly potentiated connections. Third, switching from st-LTP to Anti-hSTDP during offline replay drives further reorganization of these connections while preserving replay function. Under st-LTP, continued replay would simply reinforce the same connectivity that produces interference; under Anti-hSTDP, the same replay events instead restructure that connectivity, substantially weakening moderately potentiated synapses while redistributing the strongest connections into a sparser, more concentrated backbone.

These three claims rest on an evaluation of memory retention that differs from prior interference studies in spiking networks. Rather than validating recall through downstream task performance, as in Golden et al. (2022), or through rehearsal of stored training examples, as in Proietti et al. (2023), we treat the sequential, time-compressed structure of spontaneous and cue-evoked replay itself, evaluated directly against each environment's place-cell ordering, as the readout of whether a memory trace survives. This lets us distinguish graceful degradation from outright collapse (claim 1) and localize the synaptic changes that restore replay (claim 2) at the level of the memory representation itself, rather than inferring retention indirectly from behavior.

Together, these findings identify a specific, state-dependent plasticity mechanism, rather than a fixed architectural or algorithmic choice, as central to relieving interference in a recurrent memory system. This offers a concrete, biologically grounded target for continual learning approaches in artificial networks, where interference is typically addressed through mechanisms unrelated to the online/offline distinction identified here.

Our first claim is also worth situating relative to a distinction from recent machine learning work. Loss of plasticity—a declining ability to learn further—is distinct from catastrophic forgetting, in which new learning selectively overwrites old representations (Dohare et al., 2024). Our network does not selectively lose access to older environments while retaining newer ones; replay of *all* learned environments collapses together once connectivity saturates, a pattern more consistent with plasticity loss than classical interference.

In our model, each new environment is explicitly written into the recurrent matrix through st-LTP. Failure is therefore not an inability to update weights. Instead, it reflects a declining ability of continued synaptic modification to produce a structure that supports retrieval. Once the network saturates, neither newly nor previously encoded representations can be reliably replayed.

Our homeostatic mechanism differs from parameter-protection approaches such as elastic weight consolidation (Kirkpatrick et al., 2017), instead resembling continual backpropagation's reinitialization of underused units (Dohare et al., 2024); both recycle capacity rather than protecting task-specific parameters. Finally, unlike standard deep learning's discrete update/frozen phases (Dohare et al., 2024), our network has no frozen phase: st-LTP governs

online learning while Anti-hSTDP and homeostatic reset govern offline periods, gated by behavioral state and consistent with dopaminergic control of plasticity sign (Baron et al., 2026; Sugisaki et al., 2011; Zhang et al., 2009). Interference control and capacity extension thus emerge from a state-dependent plasticity and homeostatic cycle, a principle that may be worth exploring in artificial learning systems.

Reproducibility. Code supporting the experiments is provided in the supplementary material. The implementation builds on the publicly available CA3 model of Ecker et al. (Ecker et al., 2022) and includes the modified simulation and analysis code, replay-detection procedures, database setup, dependency requirements, and scripts used to run the principal experiments. Because the simulations include stochastic spike generation and parameter sampling, individual runs are not expected to produce identical spike trains; the supplied code is intended to reproduce the experimental procedures and statistical results reported here.

References

Baron, L., Diba, K. & Amarasingham, A. (2026) The role of plasticity in replay: Stability through anti-hebbian rules. *Hippocampus*, 36(3), e70089.

Brette, R. & Gerstner, W. (2005) Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *Journal of neurophysiology*, 94(5), 3637–3642.

Buzsáki, G. (2015) Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, 25(10), 1073–1188.

Davidson, T.J., Kloosterman, F. & Wilson, M.A. (2009) Hippocampal replay of extended experience. *Neuron*, 63(4), 497–507.

Diba, K. & Buzsáki, G. (2007) Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, 10(10), 1241.

Dohare, S., Hernandez-Garcia, J.F., Lan, Q., Rahman, P., Mahmood, A.R. & Sutton, R.S. (2024) Loss of plasticity in deep continual learning. *Nature*, 632(8026), 768–774.

Ecker, A., Bagi, B., Vértes, E., Steinbach-Németh, O., Karlócai, M.R., Papp, O.I. et al. (2022) Hippocampal sharp wave-ripples and the associated sequence replay emerge from structured synaptic interactions in a network model of area ca3. *Elife*, 11, e71850.

Foster, D.J. & Wilson, M.A. (2006) Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440(7084), 680.

French, R.M. (1999) Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4), 128–135.

Frey, U. & Morris, R.G. (1997) Synaptic tagging and long-term potentiation. *Nature*, 385(6616), 533–536.

Golden, R., Delanois, J.E., Sanda, P. & Bazhenov, M. (2022) Sleep prevents catastrophic forgetting in spiking neural networks by forming a joint synaptic weight representation. *PLOS Computational Biology*, 18(11), e1010628. doi:10.1371/journal.pcbi.1010628.

Kapoor, K., Mackey, W., Aloimonos, Y. & Lin, X. Hicl: Hippocampal-inspired continual learning. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 40*, 2026*, pp. 22518–22526.

Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A. et al. (2017) Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521–3526.

Kumaran, D., Hassabis, D. & McClelland, J.L. (2016) What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in cognitive sciences*, 20(7), 512–534.

Liao, Z., Terada, S., Raikov, I.G., Hadjiabadi, D., Szoboszlay, M., Soltesz, I. et al. (2024) Inhibitory plasticity supports replay generalization in the hippocampus. *Nature Neuroscience*, 27(10), 1987–1998.

Liu, C., Todorova, R., Tang, W., Oliva, A. & Fernandez-Ruiz, A. (2023) Associative and predictive hippocampal codes support memory-guided behaviors. *Science*, 382(6668), eadi8237.

McClelland, J.L., McNaughton, B.L. & O'Reilly, R.C. (1995) Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3), 419–457.

McCloskey, M. & Cohen, N.J. (1989) Catastrophic interference in connectionist networks: The sequential learning problem. In: *Psychology of Learning and Motivation*. Vol. 24Academic Press, pp. 109–165.

Okuno, H., Akashi, K., Ishii, Y., Yagishita-Kyo, N., Suzuki, K., Nonaka, M. et al. (2012) Inverse synaptic tagging of inactive synapses via dynamic interaction of arc/arg3.1 with camkiiβ. *Cell*, 149(4), 886–898.

Pang, R. & Fairhall, A.L. (2019) Fast and flexible sequence induction in spiking neural networks via rapid excitability changes. *Elife*, 8, e44324.

Porrello, A., Bonicelli, L., Buzzega, P., Millunzi, M., Calderara, S. & Cucchiara, R. A second-order perspective on model compositionality and incremental learning. In: *International Conference on Learning Representations*. Vol. 2025*, 2025*, pp. 75098–75128.

Proietti, M., Ragno, A. & Capobianco, R. Memory replay for continual learning with spiking neural networks. In: *2023 IEEE 33rd International Workshop on Machine Learning for Signal Processing (MLSP), 2023*.

Ramirez-Villegas, J.F., Willeke, K.F., Logothetis, N.K. & Besserve, M. (2018) Dissecting the synapse-and frequency-dependent network mechanisms of in vivo hippocampal sharp wave-ripples. *Neuron*, 100(5), 1224–1240.

Rolnick, D., Ahuja, A., Schwarz, J., Lillicrap, T. & Wayne, G. Experience replay for continual learning. In: *Advances in Neural Information Processing Systems*. Vol. 32*, 2019*.

Shi, Q., Liu, F., Li, H., Li, G., Shi, L. & Zhao, R. (2025) Hybrid neural networks for continual learning inspired by corticohippocampal circuits. *Nature Communications*, 16(1), 1272.

Shin, H., Lee, J.K., Kim, J. & Kim, J. Continual learning with deep generative replay. In: *Advances in Neural Information Processing Systems*. Vol. 30*, 2017*.

Stimberg, M., Brette, R. & Goodman, D.F. (2019) Brian 2, an intuitive and efficient neural simulator. *Elife*, 8, e47314.

Sugisaki, E., Fukushima, Y., Tsukada, M. & Aihara, T. (2011) Cholinergic modulation on spike timing-dependent plasticity in hippocampal ca1 network. *Neuroscience*, 192, 91–101.

Tononi, G. & Cirelli, C. (2014) Sleep and the price of plasticity: from synaptic and cellular homeostasis to memory consolidation and integration. *Neuron*, 81(1), 12–34.

Tsodyks, M.V. & Feigel'man, M.V. (1988) The enhanced storage capacity in neural networks with low activity level. *Europhysics Letters*, 6(2), 101–105.

Turrigiano, G.G., Leslie, K.R., Desai, N.S., Rutherford, L.C. & Nelson, S.B. (1998) Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, 391(6670), 892–896.

van de Ven, G.M., Soures, N. & Kudithipudi, D. (2025) 1.09 - continual learning and catastrophic forgetting. In: Wixted, J. (Ed.) *Learning and Memory: A Comprehensive Reference (Third Edition)*, third edition Edition. OxfordAcademic Press, pp. 153–168.

Wilson, M.A. & McNaughton, B.L. (1994) Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172), 676–679.

Zenke, F., Poole, B. & Ganguli, S. Continual learning through synaptic intelligence. In: *Proceedings of the 34th International Conference on Machine Learning (ICML)*. Vol. 70*, 2017*, pp. 3987–3995.

Zhang, J.C., Lau, P.M. & Bi, G.Q. (2009) Gain in sensitivity and loss in temporal contrast of stdp by dopaminergic modulation at hippocampal synapses. *Proceedings of the National Academy of Sciences*, 106(31), 13028–13033.

A Appendix

A.1 AI use statement

Generative AI tools were used to assist with drafting and editing portions of the manuscript, improving clarity and presentation, identifying potentially relevant literature, and providing software coding and debugging assistance. All research questions, experimental design, simulations, analysis, interpretation of results, and scientific conclusions were developed, executed, and verified by the authors. All AI-assisted content was reviewed by the authors, who take full responsibility for the final manuscript.

A.2 Algorithmic summary of the state-dependent plasticity cycle

Algorithm 1 State-dependent plasticity cycle for environment k

Encode: learn environment k via st-LTP (online)
Compress: 25 s offline replay under anti-hSTDP
Recycle: reset synapses below threshold θ to baseline
Proceed to environment $k+1$

A.3 Adaptive integrate and fire cell modeling

We used the cell modeling from Ecker et al. (2022) model for both PC and BC population:

$$C_m \frac{dV(t)}{dt} = -g_L(V(t) - V_{\text{rest}} - \Delta_T \cdot e^{\frac{(V(t)-\theta)}{\Delta_T}}) + I_{\text{syn}}(t) + ad(t) \tag{1}$$

$$\frac{d\,ad(t)}{dt} = \frac{c\big(V(t) - V_{\text{rest}}\big) - ad(t)}{\tau_{ad}} + b\sum_k \delta(t - t_k) \tag{2}$$

where: C_m is the membrane capacitance; g_L is the leak conductance; V_{rest} is the reversal potential; θ is the intrinsic spike threshold; Δ_T is the slope factor (See (Brette and Gerstner, 2005)); I_{syn} is the synaptic current (See eq. 3); here $ad(t)$ is an adaptation current, t_k denotes the k-th spike time, and the Dirac delta term $b\sum_k \delta(t - t_k)$ produces an instantaneous increment $ad \to ad + b$ at each spike.

For pyramidal cells (PCs), parameters were

$$b = 206.9 \text{ pA}, \quad c = -0.27 \text{ nS}, \quad \tau_{ad} = 85 \text{ ms}.$$

For basket cells (BCs), parameters were

$$b = 0.91 \text{ pA}, \quad c = 3.05 \text{ nS}, \quad \tau_{ad} = 178.6 \text{ ms}.$$

The negative value of c in PCs reflects dominance of the spike-triggered adaptation component, whereas in BCs the positive c introduces stronger voltage-dependent adaptation. These parameter regimes were adopted from Ecker et al. (2022), in which they were calibrated to

reproduce physiologically realistic spike-frequency adaptation and stable replay dynamics in CA3 network simulations.

The total synaptic current $I_{\text{syn}}(t)$ received by a neuron is modeled as the sum of excitatory and inhibitory components:

$$I_{\text{syn}}(t) = g_{\text{PC}}(t)\,(V(t) - E_{\text{exc}}) + g_{\text{BC}}(t)\,(V(t) - E_{\text{inh}}) \tag{3}$$

where $E_{\text{exc}} = 0$ mV and $E_{\text{inh}} = -70$ mV are the reversal potentials for excitatory and inhibitory synapses, respectively. The total conductance from each input population (e.g., PCs for excitation, BCs for inhibition) is computed by summing the postsynaptic response to each presynaptic spike:

$$g_X(t) = \sum_{j \in \mathcal{P}_X} \sum_k \hat{g}_j \cdot Q \cdot \left[\exp\left(-\frac{t - t_j^{(k)} - d_j}{\tau_d} \right) - \exp\left(-\frac{t - t_j^{(k)} - d_j}{\tau_r} \right) \right] \cdot \mathbb{1}\{t - t_j^{(k)} - d_j \geq 0\} \tag{4}$$

Here: $X \in \{\text{PC}, \text{BC}\}$ denotes the presynaptic population (excitatory or inhibitory), $\mathcal{P}_X$ is the set of presynaptic neurons of type X, $t_j^{(k)}$ is the time of the kth spike from neuron j, d_j is the synaptic delay from neuron j, $\hat{g}_j$ is the peak conductance of the synapse from neuron j, τ_r and τ_d are the rise and decay time constants, - $Q = \left(\exp\left(-\frac{t_{\text{peak}}}{\tau_d} \right) - \exp\left(-\frac{t_{\text{peak}}}{\tau_r} \right) \right)^{-1}$ normalizes the peak conductance to $\hat{g}_j$, with t_{peak} given by:

$$t_{\text{peak}} = \frac{\tau_r \tau_d}{\tau_d - \tau_r} \ln\left(\frac{\tau_d}{\tau_r} \right). \tag{5}$$

The indicator function component $\mathbb{1}\{t - t_j^{(k)} - d_j \geq 0\}$ ensures causality in the synaptic response: the conductance change becomes nonzero only after the corresponding presynaptic spike has occurred and its associated synaptic delay has elapsed.

A.3.1 Phase 1 Online Learning: Additional Details

The online learning phase serves to embed a spatial environment into the recurrent PC–PC connectivity matrix that is subsequently used during offline replay. Connectivity is not imposed through predefined topological rules (Ramirez-Villegas et al., 2018; Pang and Fairhall, 2019); instead, recurrent weights develop from simulated spiking activity via spike-timing-dependent plasticity (STDP), following Ecker et al. (2022). All network parameters matched those of Ecker et al. (2022) unless otherwise noted.

As in Ecker et al. (2022), we modeled CA3 activity as a rodent traverses a 3-meter linear track in one direction. At the start of each simulation, every pyramidal cell was independently designated a place cell with probability 0.5, yielding approximately 4,000 place cells. Each designated neuron received a Gaussian place field centered at a location drawn uniformly along the track. Whereas the original model (Ecker et al., 2022) assigned place-field centers so that neuron identity increased monotonically with position, here the mapping between place-field centers and neuron identity was independently randomized for each environment, giving each environment a distinct place-cell ordering while leaving the underlying network unchanged.

Formally, non-place-cell spiking followed a homogeneous Poisson process with rate $\lambda = 0.1$ Hz, while place-cell spiking followed a conditionally independent point process with a peak rate of 20 Hz, governed by the conditional intensity function $\lambda_i(t; H_t)$ (Eq. 6), which combines a Gaussian spatial tuning curve with theta-band modulation (H_t denotes spiking history up to time t). The tuning curve's standard deviation σ spanned 10% of the track length, and a 5-ms refractory period was enforced by discarding spikes falling within 5 ms of the preceding one.

Following Ecker et al. (2022), the animal's simulated running speed was fixed at 32.5 cm/s, so that a single traversal of the track lasted approximately 9.2 seconds; position reset instantaneously to the start upon reaching the end. Spike trains were generated over a total of 400 seconds, corresponding to roughly 43 track traversals.

The conditional intensity function of place cell i is specified via:

$$\lambda_i(t; H_t) = \lambda_i(t, t - s_i^*(t)) = \lambda_{\max} \times \tau_i(x(t)) \times \cos\left(2\pi f_\theta t + \frac{\pi}{l^{PF}}(x(t) - s_i^{PF})\right) \times \mathbb{1}\{t - s_i^*(t) > t_{\text{ref}}\} \tag{6}$$

$$\tau_i(x) = \exp\left(-\frac{(x - m_i^{PF})^2}{2\sigma^2}\right) \tag{7}$$

where t is time, $s_i^*(t)$ is the last spike of the i'th neuron before time t, $\mathbb{1}\{\cdot\}$ is the indicator function, $\tau_i(x)$ is the spatial tuning curve of the i'th neuron, $\sigma = 7.5cm$ is the standard deviation of the Gaussian tuning curve, $x(t)$ is the position of the animal, m_i^{PF} is the center of the place field, $l^{PF} = 0.3\text{m}(= 4\sigma)$, and $s_i^{PF}(= m_i^{PF} - 2\sigma)$ are the length and start of the given place field,

respectively, and $\lambda_{\max} = 20\,\text{Hz}$ is the maximum in-field firing rate, f_θ is the background 7Hz theta activity. $t - s_i^*(t)$ is the time between two consecutive spikes, $t_{\text{ref}} = 5ms$ is the refractory period.

During the phase 1 simulations, synaptic connection strength was modified based on a symmetric STDP rule:

$$\Delta w = A \cdot \exp\left(\frac{-|\Delta t|}{\tau}\right) \tag{8}$$

where Δw is the change in synaptic weights in nS, $\Delta t = t_{\text{post}} - t_{\text{pre}}$ is the time difference in ms between action potentials, and A describes the weight update in nS, which decayed exponentially with time constants τ in ms, respectively. We used the optimized parameters from Ecker et al. (2022) for our phase 1 learning ($\tau = 62.5$ ms and $A = 0.08$ nS), which were shown to produce robust bi-directional replays during phase 2.

A.3.2 Phase 2 Offline Replay - Additional Details

The PC-to-PC connections, and weights learned during Phase 1, were carried to Phase 2 for the offline simulations . Following Ecker et al. (2022), the network architecture was further expanded during the offline phase to include 8000 stochastic mossy fiber (MF) cells. The MF cells provide excitatory input to the PC population during Phase 2 and comprised the sole external input to the network during this phase. These MF cells were modeled as independent Poisson neurons spiking at rate $\lambda_{MF} = 0.1$. The MF population was connected 1:1 to the PC population with strong "detonator" synaptic weights (19.15 nS). BC and MF cells were not connected.

A.3.3 Plasticity During Phase 2

In the Ecker et al. (2022) model, spontaneous bi-directional replays are observed during the offline state, interspersed with periods of sparse and unstructured activity (non-replay periods). In contrast to the absence of offline plasticity in Ecker et al. (2022), we incorporated different synaptic plasticity rules specifically in the PC-to-PC recurrent synapses. To enable selective activation of different plasticity kernels, we introduced the variables p_+ and p_- into the Phase 1 plasticity rule of Ecker et al. (2022). Consequently, our synaptic weight adaptation rule is defined as:

$$\Delta w_+ = p_+ \cdot A_+ \exp\left(\frac{-\Delta t}{\tau^+}\right) \quad \text{at } t_{\text{post}} \text{ if } (t_{\text{pre}} + d) < t_{\text{post}} \tag{9}$$

$$\Delta w_- = p_- \cdot A_- \exp\left(\frac{\Delta t}{\tau^-}\right) \quad \text{at } t_{\text{pre}} \text{ if } (t_{\text{pre}} + d) > t_{\text{post}} \tag{10}$$

where $p_{+/-}$ is the multiplier used to activate the selected plasticity kernel, ranging from -1 to 1; $\tau_{+/-}$ is the decay time constant; d is the synaptic delay (2.2 ms for PC-to-PC synapses); t_{pre} and t_{post} are the times of the pre- and postsynaptic action potentials, respectively; $\Delta t = t_{\text{post}} - (t_{\text{pre}} + d)$ is the spike-time difference after accounting for the synaptic delay; and $A_{+/-}$ is the amplitude of the synaptic weight update (nS), which decays exponentially with the corresponding time constant $\tau_{+/-}$.

To prevent uncontrolled synaptic growth during sequential learning, the maximum PC-to-PC synaptic conductance was capped at 4.0–4.1 nS in all simulations. This range was selected empirically because it consistently preserved stable replay activity while avoiding the emergence of continuous "zigzag" replay patterns associated with excessive synaptic saturation as described by Baron et al. (2026).

Except for the maximum synaptic weight cap described above, the offline network model followed Ecker et al. (2022) and Baron et al. (2026). Model parameters, including plasticity constants and synaptic delays, were independently sampled from uniform distributions centered on the published values, with bounds of approximately ±5%. In contrast to Ecker et al. (2022), place cells were selected using a uniform sampling procedure rather than the original track-edge selection method.

A.4 Replay detection

We used the replay visualization tool developed by Baron et al. (2026) to identify an initial set of replay events and to establish replay candidate detection and replay validation parameters consistent with the methodology described in Baron et al. (2026).

Replay candidates were first identified using the BC population. A replay candidate was detected whenever the BC population spiking rate, calculated as the number of spikes within a 10 ms time window, exceeded twice the average BC spiking rate over the entire simulation. The average BC spiking rate was computed as the total number of BC spikes divided by the simulation duration (ms) and scaled to a 10 ms window.

Candidate events were subsequently validated using the PC population. PC spikes were first filtered to include only neurons with a firing rate of at least 10 Hz during the candidate event, after which the active PCs were sorted according to their place field order. For each candidate event, we analyzed a window of 500 PCs centered on the median active PC (±250 PCs). Within this window, we performed a linear regression using spike time as the independent variable (x-axis) and place field order as the dependent variable (y-axis). Candidate events with an absolute regression slope between 1.5 and 20 and containing more than 100 PC spikes were classified as replay events.

For simulations containing more than five memory traces, progressive synaptic saturation produced substantially higher background firing rates and high-frequency spiking activity, which increased the number of false replay candidates. To maintain comparable replay detection specificity across conditions, we increased the minimum PC firing-rate threshold to 70 Hz for these highly saturated networks while leaving all other replay validation criteria unchanged.

A.5 Homeostatic threshold sweep for extended memory capacity

The results in Section 2.4 pool simulations run across a range of homeostatic reset thresholds (1.5–3.25 nS in 0.25 nS increments), all using the strengthened cueing protocol (§4.3), to assess whether the network's extended-capacity behavior depends on the exact threshold value. Each threshold was tested with at least one simulation across five to nine sequentially learned environments; one threshold value received additional replicates, collected over the course of the study, to assess run-to-run variability. Table S1 reports the per-threshold breakdown underlying the pooled counts in Table 1.

Cue-evoked recall of the first learned environment was obtained at every threshold tested through six environments except one (2.75 nS, which failed at all environment counts and appears to be an outlier given the surrounding thresholds' behavior). Recall was intermittent

at seven environments and inconsistent at eight and nine, without a clear monotonic relationship to threshold value in this range. This indicates that the extended-capacity result reported in Section 2.4 reflects a broad range of effective homeostatic thresholds rather than a property of any single value.

Table S1: Cue-evoked recall of the first learned environment, by homeostatic threshold, underlying the pooled counts in Table 1. Each entry reports the fraction (and percentage) of simulations producing at least one validated replay event.

Threshold (nS)	**5 env**	**6 env**	**7 env**	**8 env**	**9 env**
1.50	1/1 (100%)	1/1 (100%)	1/1 (100%)	0/1 (0%)	0/1 (0%)
1.75	7/7 (100%)	7/7 (100%)	4/7 (57%)	1/6 (17%)	0/6 (0%)
2.00	1/1 (100%)	1/1 (100%)	0/1 (0%)	0/1 (0%)	0/1 (0%)
2.25	1/1 (100%)	1/1 (100%)	0/1 (0%)	1/1 (100%)	0/1 (0%)
2.50	1/1 (100%)	1/1 (100%)	1/1 (100%)	0/1 (0%)	0/1 (0%)
2.75	0/1 (0%)	0/1 (0%)	0/1 (0%)	0/1 (0%)	0/1 (0%)
3.00	1/1 (100%)	1/1 (100%)	1/1 (100%)	0/1 (0%)	0/1 (0%)
3.25	1/1 (100%)	1/1 (100%)	0/1 (0%)	1/1 (100%)	0/1 (0%)

A.6 Additional Raster Plots

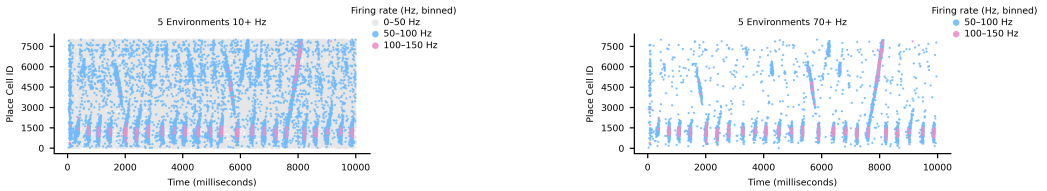

(a) Five environments (10 Hz filter; standard replay criteria)

(b) Five environments (70 Hz filter; suppresses background activity)

Figure S1: Additional raster plots. x axis - Time (in milliseconds) from simulation starts. Y Axis - Place cell ID (a-b)

A.7 Plasticity Kernels used in our analysis

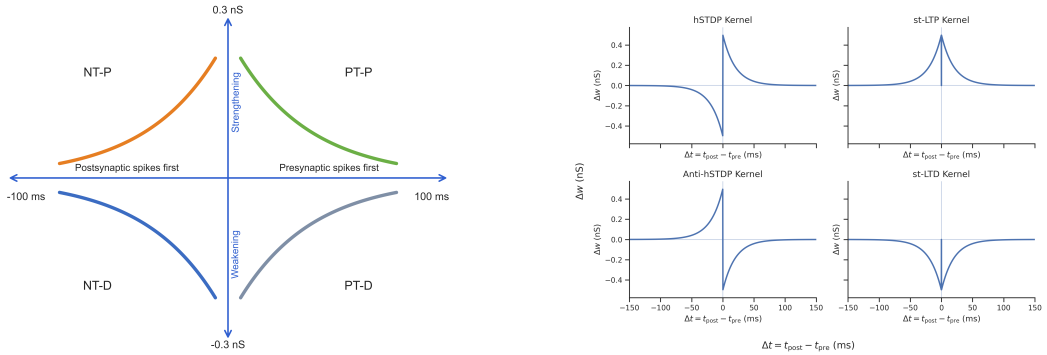

(a) Plasticity Kernel Components

(b) Kernels Types

Figure S2: Adapted from (Baron et al., 2026): a - Plasticity Kernel Component (PKC) classification. X-axis: Spike timing between pre- and post-synaptic PCs, defined as (post - pre spike time) in milliseconds - Positive Timing (PT) when x $>$ 0 and Negative Timing (NT) when x $<$ 0. Y-axis: Change in synaptic conductance (nS); Long Term Potentiation (LTP) for values $>$ 0, Long Term Depression (LTD) for values $<$ 0. PKC labels: PT-P – potentiation for positive timing; PT-D – depression for positive timing; NT-P – potentiation for negative timing; NT-D – depression for negative timing. b - Kernel Types: hSTDP: Co-activation of PT-P & NT-D; symmetric timing LTP (st-LTP): Co-activation of PT-P & NT-P; anti-hSTDP: Co-activation of PT-D & NT-P; symmetric-timing LTD (st-LTD): Co-activation of PT-D & NT-D.

A.8 Selection of a Minimum $w_{\max}$ for Reduced Place-Cell Ratios

Reducing the fraction of CA3 pyramidal cells assigned as place cells decreases the number of neurons participating in each learned representation and may therefore require stronger recurrent synaptic weights to preserve spontaneous replay. We therefore evaluated the minimum maximum synaptic weight ($w_{\max}$) sufficient to maintain replay as the place-cell ratio was reduced.

For each combination of place-cell ratio and $w_{\max}$, the network was trained during the online learning phase and subsequently simulated for 15 s under offline conditions. Two independent simulations were performed for each parameter combination, yielding 48 simulations in total. Replay events were detected using the same criteria employed throughout the study.

Table S2 shows that reducing the place-cell ratio required a corresponding increase in $w_{\max}$ to preserve replay. We therefore selected $w_{\max} = 5.0$ nS, the minimum tested value that preserved replay at a PC ratio of 0.35, for subsequent simulations using this configuration.

Table S2: Average number of replay events as a function of the fraction of CA3 pyramidal cells assigned as place cells (PC ratio) and the maximum recurrent PC–PC synaptic weight ($w_{\max}$). Each parameter combination was evaluated using two independent simulations, for a total of 48 simulations. Following Phase 1 training, replay activity was assessed during a 15-s offline simulation. Zero values indicate configurations in which no replay events were detected. The analysis was used to identify the minimum $w_{\max}$ required to preserve replay as the place-cell ratio was reduced.

	$w_{\max}$ (nS)					
PC Ratio	**3.75**	**4.00**	**4.25**	**4.50**	**4.75**	**5.00**
0.35	0.0	0.0	0.0	0.0	0.0	18.0
0.40	0.0	0.0	10.0	9.5	19.5	19.0
0.45	0.0	0.0	13.5	22.0	22.5	22.0
0.50	3.3	16.7	19.2	23.5	23.4	23.1

As the place-cell ratio decreased, larger values of $w_{\max}$ were required to maintain replay activity. At a PC ratio of 0.50, replay was observed at $w_{\max} = 3.75$ nS, whereas ratios of 0.40 and 0.45 required at least the lowest tested value of 4.25 nS to produce robust replay. For a PC ratio of 0.35, $w_{\max} = 5.0$ nS was sufficient to preserve replay, producing an average of 18 replay events during the 15-s offline simulation. Because lower values were not evaluated at this ratio, 5.0 nS should be interpreted as a validated sufficient value rather than a confirmed minimum.