

# Theory of systems memory consolidation via recall-gated plasticity

**Reviewed Preprint**

Published from the original preprint after peer review and assessment by eLife.

[About eLife's process](#)

**Reviewed preprint posted**

November 24, 2023 (this version)

**Sent for peer review**

July 14, 2023

**Posted to bioRxiv**

December 10, 2022

Jack Lindsey , Ashok Litwin-Kumar 

Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY 10027, USA

 [https://en.wikipedia.org/wiki/Open\\_access](https://en.wikipedia.org/wiki/Open_access)

 Copyright information

## Abstract

In a variety of species and behavioral contexts, learning and memory formation recruits two neural systems, with initial plasticity in one system being consolidated into the other over time. Moreover, consolidation is known to be selective; that is, some experiences are more likely to be consolidated into long-term memory than others. Here, we propose and analyze a model that captures common computational principles underlying such phenomena. The key component of this model is a mechanism by which a long-term learning and memory system prioritizes the storage of synaptic changes that are consistent with prior updates to the short-term system. This mechanism, which we refer to as recall-gated consolidation, has the effect of shielding long-term memory from spurious synaptic changes, enabling it to focus on reliable signals in the environment. We describe neural circuit implementations of this model for different types of learning problems, including supervised learning, reinforcement learning, and autoassociative memory storage. These implementations involve learning rules modulated by factors such as prediction accuracy, decision confidence, or familiarity. We then develop an analytical theory of the learning and memory performance of the model, in comparison to alternatives relying only on synapse-local consolidation mechanisms. We find that recall-gated consolidation provides significant advantages, substantially amplifying the signal-to-noise ratio with which memories can be stored in noisy environments. We show that recall-gated consolidation gives rise to a number of phenomena that are present in behavioral learning paradigms, including spaced learning effects, task-dependent rates of consolidation, and differing neural representations in short- and long-term pathways.

### eLife assessment

This **fundamental** work proposes a novel mechanism for memory consolidation where short-term memory provides a gating signal for memories to be consolidated into long-term storage. The work combines extensive analytical and numerical work applied to three different scenarios and provides a **convincing** analysis of the benefits of the proposed model, although some of the analyses are limited to the type of memory consolidation the authors consider (and don't consider), which limits the impact. The work could be revised to include a more thorough comparison to existing models of memory consolidation and a discussion of its limitations, and the revision could also streamline the technical terminology. The work will be of interest to neuroscientists and many other researchers interested in the mechanistic underpinnings of memory.

## Introduction

Systems that learn and remember confront a tradeoff between memory acquisition and retention. Plasticity enables learning but can corrupt previously stored information. Consolidation mechanisms, which stabilize or render more resilient certain plasticity events associated with memory formation, are key to navigating this tradeoff (Kandel et al., 2014 [↗](#)).

Consolidation may be mediated both by molecular dynamics at the synapse level (synaptic consolidation) and dynamics at the neural population level (systems consolidation).

Theoretical studies have described how synaptic consolidation affects the strength and lifetime of memories (Fusi et al., 2005 [↗](#); Lahiri and Ganguli, 2013 [↗](#); Benna and Fusi, 2016 [↗](#)). In these studies, synapses are modeled with multiple internal variables, operating at distinct timescales, which enable individual synapses to exist in more labile or more rigid (“consolidated”) states. Such models can prolong memory lifetime and recapitulate certain memory-related phenomena, notably power-law forgetting curves. Moreover, this line of work has established theoretical limits on the memory retention capabilities of any such synaptic model, and shown that biologically realistic models can approximately achieve these limits (Lahiri and Ganguli, 2013 [↗](#); Benna and Fusi, 2016 [↗](#)). These theoretical frameworks leave open the question of what computational benefit is provided by *systems* consolidation mechanisms that take place in a coordinated fashion across populations of neurons.

The term systems consolidation most often refers to the process by which memories stored in the hippocampus are transferred to the neocortex (Fig. 1A [↗](#); Squire and Alvarez, 1995 [↗](#); Frankland and Bontempi, 2005 [↗](#); McClelland et al., 1995 [↗](#); McClelland and Goddard, 1996 [↗](#)). Prior work has described the hippocampus and the neocortex as “complementary learning systems,” emphasizing their distinct roles: the hippocampus stores information about individual experiences, and the neocortex extracts structure from many experiences (McClelland et al., 1995 [↗](#); McClelland and Goddard, 1996 [↗](#)). Related phenomena also occur in other brain systems. In rodents, distinct pathways underlie the acquisition and later execution of some motor skills, with motor cortex apparently passing off responsibility to basal ganglia structures as learning progresses (Kawai et al., 2015 [↗](#); Dhawale et al., 2021 [↗](#)). A similar consolidation process is observed during vocal learning in songbirds, where song learning is dependent on the region LMAN but song execution can, after multiple days of practice, become LMAN-independent and rely instead on the song motor pathway (Fig. 1B [↗](#); Warren et al., 2011 [↗](#)). Some insects also display a form of systems consolidation. Olfactory learning experiments in fruit flies reveal that

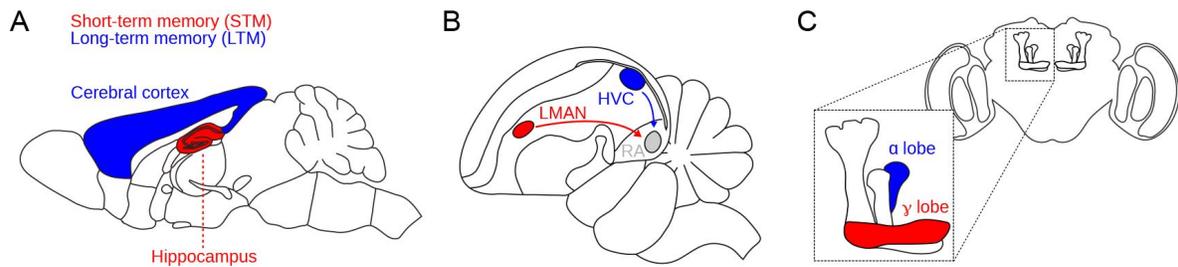
short and long-term memory retrieval recruit distinct neurons within the fly mushroom body, and the short-term pathway is necessary for long-term memory formation (Fig. 1C; Cervantes-Sandoval et al., 2013; Dubnau and Chiang, 2013).

The examples above are all characterized by two essential features: the presence of two systems involved in learning similar information and an asymmetric relationship between them, such that learning in one system (the “long-term memory”) is facilitated by the other (the “short-term memory”). Moreover, mounting evidence indicates that across all these systems, there exist mechanisms that selectively modulate or gate consolidation into long-term storage. In flies, for instance, recent work has shown that short-term olfactory memory recall gates long-term memory storage via a disinhibitory circuit, such that repeated stimulus-outcome pairings are consolidated into long-term memory but once-presented pairings are not (Awata et al., 2019). A recent study in songbirds indicates that the rate at which song learning is consolidated into the song motor pathway is modulated by performance quality (Tachibana et al., 2022). Finally, a large body of work has shown that propensity of hippocampal memories to be cortically consolidated is modulated by a variety of factors including repetition, reliability, and novelty (Terada et al., 2021; Gorriz et al., 2023; Jackson et al., 2006; Brodt et al., 2016).

The ubiquity of the systems consolidation motif across species, neural circuits, and behaviors suggests that it offers broadly useful computational advantages that are complementary to those offered by synaptic mechanisms. In this work, we propose that the ability to selectively consolidate memories is the key computational feature that distinguishes systems from synaptic memory consolidation. To formalize this idea, we generalize prior theoretical studies by considering environments in which some memories should be prioritized more than others for long-term retention. We then introduce a model of systems consolidation and show that can provide substantial performance advantages in such environments. In the model, synaptic updates are consolidated into long-term memory depending on their consistency with knowledge already stored in short-term memory. We term this mechanism “recall-gated consolidation.” This model is well-suited to prioritize storage of reliable patterns of synaptic updates which are consistently reinforced over time. We derive neural circuit implementations of this model for several tasks of interest. These involve plasticity rules modulated by globally broadcast factors such as prediction accuracy, confidence, and familiarity. We develop an analytical theory that describes the limits on learning performance achievable by synaptic consolidation mechanisms and shows that recall-gated systems consolidation exhibits qualitatively different and complementary benefits. Our theory depends on a quantitative treatment of environmental statistics, in particular the statistics with which similar events recur over time. Different model parameter choices suit different environmental statistics, and give rise to different learning behavior, including spaced training effects. The model makes predictions about the dependence of consolidation rate on the consistency of features in an environment, and the amount of time spent in it. It also predicts that short-term memory benefits from employing sparser representations compared to long-term memory. A variety of experimental data support predictions of the model, which we review in the Discussion.

## Results

Following prior work (Fusi et al., 2005; Benna and Fusi, 2016), we consider a network of neurons connected by an ensemble  $N$  synapses whose weights are described by a vector  $\mathbf{w} \in \mathbb{R}^N$ . For now, we are agnostic as to the structure of the network and its synaptic connections. The network’s synapses are subject to a stream of patterns of candidate synaptic potentiation and depression events. We refer to such a pattern as a *memory*, defined by a vector  $\Delta\mathbf{w} \in \mathbb{R}^N$ . Learning proceeds according to a specified synaptic update rule that translates candidate potentiation and depression events (memories) into synaptic changes. One simple example of a synaptic update rule is a “binary switch” model, in which synapses can exist in two states (active or inactive), and



**Figure 1**

Schematic of short- and long-term memory systems across species and brain areas. A. In mice and other mammals, hippocampal memories are consolidated into cerebral cortex. B. Zebrafish song learning initially depends on LMAN but later requires only HVC-to-RA synapses in the song motor pathway. C. In the *Drosophila* mushroom body (inset), short- and long-term memories depend on dopamine-dependent plasticity in the  $\gamma$  and  $\alpha$  lobes, respectively.

candidate synaptic updates are binary (potentiation or depression). In this model inactive synapses activate (resp. active synapses inactivate) in response to potentiation (resp. depression) events with some probability  $p$ . However, our systems consolidation model can be used with any underlying synaptic mechanisms, and we will consider a variety of synaptic plasticity rules as underlying substrates.

In our framework, the same memory can be reinforced repeatedly over time. We distinguish memories by the reliability with which they are reinforced. The notion of reliability in our framework is meant to capture the idea that the structure of events in the world which drive synaptic plasticity is in some cases consistent over time, and in other cases inconsistent. For now, we focus on a simple environment model which captures this essential distinction, in which there are two kinds of memories: “reliable” and “unreliable.” Reliable memories are consistent patterns of synaptic updates that are reinforced regularly over time, whereas unreliable memories are spurious, randomly sampled patterns of synaptic updates. Concretely, in simulations, we assume that a given reliable memory is reinforced with independent probability  $\lambda$  at each timestep, and otherwise a randomly sampled unreliable memory is encountered.

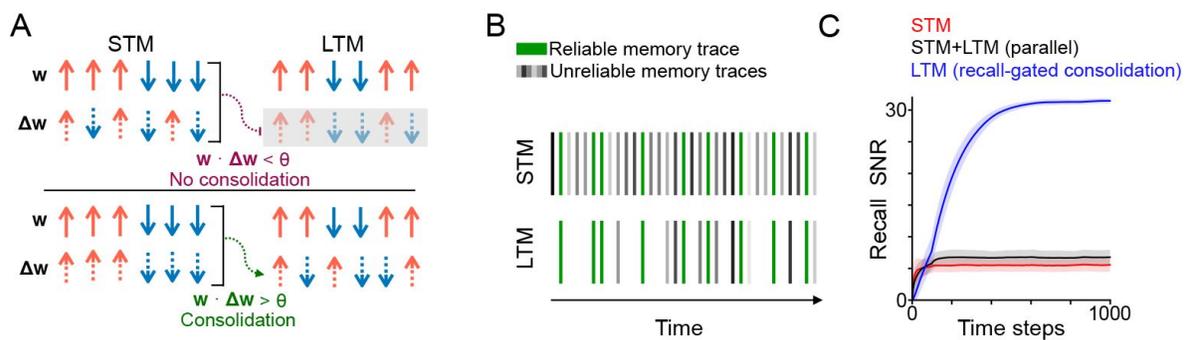
A useful measure of system performance is memory *recall*, defined as the overlap  $r = \Delta \mathbf{w} \cdot \mathbf{w}$  between a memory and the current synaptic state. Specifically, we are interested in the signal-to-noise ratio (SNR) of the recall of reliable memories, (Fusi et al., 2005 [↗](#); Benna and Fusi, 2016 [↗](#)), which normalizes the recall strength relative to the scale of fluctuations in recall of random memory vectors  $\Delta \mathbf{w}'$ :

$$\text{SNR} = \frac{\mathbf{w} \cdot \Delta \mathbf{w}}{\sqrt{\mathbb{E}[(\mathbf{w} \cdot \Delta \mathbf{w}')^2]}}. \quad (1)$$

## Recall-gated systems consolidation

In our model (Fig. 2A [↗](#)), we propose that the population of  $N$  synapses is split into two subpopulations which we call the “short-term memory” (STM) and “long-term memory” (LTM). Upon every presentation of a memory  $\Delta \mathbf{w}$ , the STM recall  $r_{\text{STM}} = \Delta \mathbf{w} \cdot \mathbf{w}_{\text{STM}}$  is computed. Learning in the LTM is modulated by a factor  $g(r_{\text{STM}})$ . We refer to  $g$  as the “gating function.” For now we assume  $g$  to be a simple threshold function, equal to 0 for  $r_{\text{STM}} < \theta$  and 1 for  $r_{\text{STM}} \geq \theta$ , for some suitable threshold  $\theta$ . This means that consolidation occurs only when a memory is reinforced at a time when it can be recalled sufficiently strongly by the STM. Later we will consider different choices of the gating function  $g$ , which may be more appropriate depending on the statistics of memory recurrence in the environment.

We refer to this mechanism as *recall-gated consolidation*. Its function is to filter out unreliable memories, preventing them from affecting LTM synaptic weights. With an appropriately chosen gating function, reliable memories will pass through the gate at a higher rate than unreliable memories. Consequently, events that trigger plasticity the LTM will consist of a higher proportion of reliable memories (Fig. 2B [↗](#)), and hence attain a higher SNR than the STM. The cost of this gating is to incur some false negatives—reliable memory presentations that fail to update the LTM. However, some false negatives can be tolerated given that we expect reliable memories to recur multiple times, and information about these events is still present in the STM. As a proof of concept of the efficacy of recall-gated consolidation, we conducted a simulation in which memories correspond to random binary synaptic update patterns and plasticity follows a binary switch rule (Fig. 2C [↗](#)). Notably, recall-dependent consolidation results in reliable memory recall with a much higher SNR than an alternative model in which LTM weight updates proceed independently of STM recall.



**Figure 2**

A. Schematic of systems consolidation model. Memories correspond to patterns of candidate potentiation and depression events (dashed arrows) applied to a synaptic population (solid arrows). The synaptic population is divided into an STM and LTM. Memories that provoke strong recall in the STM – that is, overlap strongly with the present synaptic state – enable plasticity (consolidation) in the LTM. B. Schematic of the environmental statistics. A reliable memory (green) arrives repeatedly with probability  $\lambda$  at each time step, with randomly sampled “unreliable” memories (gray) interspersed. The LTM is exposed to a filtered subset of consolidated memory traces with a higher proportion of reliable memories. C. Simulation of recall performance of a single reliable memory with time as it is presented with probability  $\lambda = 0.25$  at each time step,  $N = 2000$  synapses (1000 each in the STM and LTM). The STM and LTM learning rates (binary switching probabilities) are  $p = 0.25$  and  $p = 0.05$ , respectively, and the synaptic state is initialized randomly, each synapse initially active with probability 0.5. In the recall-gated model, the gating threshold is set at  $\theta = 2^{-3}$ . Shaded regions indicate standard deviation across 1000 simulations.

## Neural circuit implementations of recall-gated consolidation

Our model requires a computation of recall strength, which we defined as the overlap between the candidate synaptic weight changes associated with memory storage and the current state of the synaptic population. From this definition, it is not clear how recall strength can be computed biologically. The mechanisms underlying computation of recall strength will depend on the task, network architecture, and plasticity rule giving rise to candidate synaptic changes. A simple example is the case of a population of input neurons connected to a single downstream output neuron, subject to a learning rule that potentiates synapses corresponding to active inputs. In this case, the recall strength quantity corresponds exactly to the total input received by the output neuron, which acts as a familiarity detector. Below, we give the corresponding recall strength factors for other learning and memory tasks: supervised learning, reinforcement learning, and unsupervised auto-associative memory, summarized in [Fig. 3A](#) and derived in the Supplementary Information. We emphasize that our use of the term “recall” refers the familiarity of *synaptic update patterns*, and does not necessarily correspond to familiarity of stimuli or other task variables.

### Supervised learning

Suppose a population of neurons with activity  $\mathbf{x}$  representing stimuli is connected via feedforward weights to a readout population with activity  $\hat{\mathbf{y}} = \mathbf{W}\mathbf{x}$ . The goal of the system is to predict ground-truth outputs  $\mathbf{y}$ . A simple plasticity rule which will train the system appropriately is a Hebbian rule,  $\Delta W_{ij} = y_i x_j$ . The corresponding recall factor is  $\mathbf{y} \cdot \hat{\mathbf{y}}$ , corresponding to prediction accuracy ([Fig. 3B](#)).

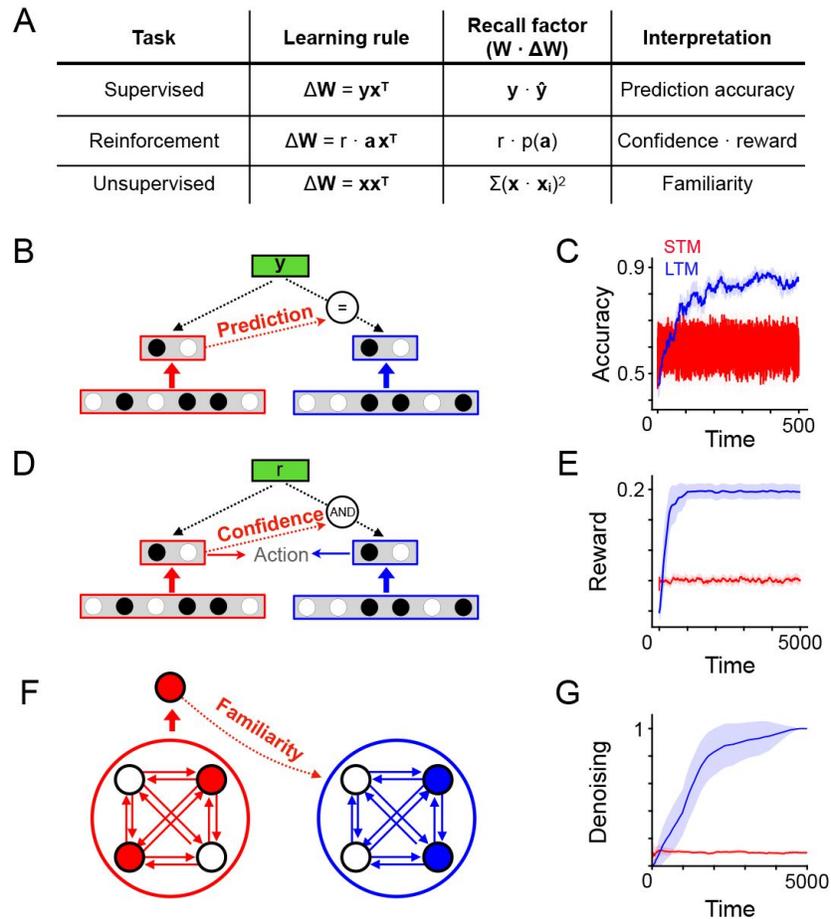
### Reinforcement learning

Suppose a population of neurons with activity  $x$  representing an animal’s state is connected to a population with activity  $\boldsymbol{\pi} = \mathbf{W}\mathbf{x}$  representing probabilities of selecting different actions, with  $\pi_i$  equal to the probability of selecting action  $a_i$ . Following action selection, the animal receives a reward  $r$ , depending on its decision. A simple approach to reinforcement learning is to use a three-factor rule  $\Delta W_{ij} = r \cdot a_i \cdot x_j$ , which reinforces actions that lead to reward. For this model, the corresponding recall factor is  $r \cdot p(a_i)$ , a multiplicative combination of reward and the animal’s confidence in its selected action. Intuitively, the recall factor will be high when a confidently chosen action leads to reward ([Fig. 3D](#)).

### Unsupervised autoassociative memory

Suppose a population of neurons with activity  $\mathbf{x}$  and recurrent weights  $\mathbf{W}$  stores memories as attractors according to an autoassociative Hebbian rule  $\Delta W_{ij} = x_i x_j$ . In this case, the recall strength  $\mathbf{W} \cdot \Delta \mathbf{W}$  can be expressed as  $\mathbf{x} \cdot (\mathbf{W}\mathbf{x})$ , a comparison between stimulus input  $\mathbf{x}$  and recurrent input  $\mathbf{W}\mathbf{x}$ . Intuitively, the recall factor measures the familiarity of the stimulus, as highly familiar stimuli will exhibit attractor behavior, making  $\mathbf{x}$  and  $\mathbf{W}\mathbf{x}$  highly correlated. Such a quantity could be computed directly by using separate dendritic compartments for the two input sources, or indirectly by comparing the network state at successive timesteps. Familiarity also be approximated using a separate novelty readout trained alongside the recurrent weights ([Fig. 3F](#)), which is the implementation we use in our simulation. In this approach, a set of readout weight  $\mathbf{u}$  receive the neural population activity as input, and follow the learning rule  $\Delta \mathbf{u} = \mathbf{x}$ . The readout  $\mathbf{u} \cdot \mathbf{x}$ , an estimate of familiarity, is used as the recall factor.

To verify that the advantages of recall-gated consolidation illustrated in [Fig. 2](#) apply in these tasks, we simulated the three architectures and learning rules described above (see Methods for simulation details). In each case, learning takes place online, with reliable task-relevant stimuli appearing a fraction  $\lambda$  of the time, interspersed among randomly sampled unreliable stimuli. In the case of supervised and reinforcement learning tasks, unreliable stimuli are paired with



**Figure 3**

A. Description of learning rules corresponding to different types of learning problems and corresponding expressions for the recall factor  $\mathbf{W} \cdot \Delta\mathbf{W}$  used in the recall-gated consolidation model.

B. Schematic indicating a possible implementation of the model in a supervised learning problem, where LTM plasticity is modulated by the consistency between STM predictions and ground-truth labels.

C. Simulation of a binary classification problem,  $N = 4000$ ,  $\lambda = 0.2$ ,  $\theta = 0.125$ ,  $p = 0.1$ . There are twenty total stimuli each associated with a random binary ( $\pm 1$ ) label. Plot shows the classification accuracy over time, given by the outputs of the STM and LTM of the consolidation model. Shaded region indicates standard deviation over 1000 simulations.

D. Like B, but for a reinforcement learning problem. LTM plasticity is gated by both STM action confidence and the presence of reward.

E. Simulation of a reinforcement learning problem,  $N = 4000$ ,  $\lambda = 0.25$ ,  $\theta = 0.75$ ,  $p = 1.0$ . There are five total stimuli and two possible actions. Each stimulus has a corresponding action that yields reward. The plot shows average reward per step over time, using the actions given by the STM or LTM.

F. Like B and D, but for an autoassociative unsupervised learning problem. LTM plasticity is gated by familiarity detection in the STM module, learned using a separate set of weights.

G. Simulation of an autoassociative learning problem.  $N = 4000$ ,  $\lambda = 0.1$ ,  $\theta = 0.25$ ,  $p = 1.0$ . Recall performance is evaluated by exposing the system to a noisy version of the reliable patterns seen during training, allowing the recurrent dynamics of the network to run for 5 timesteps, and measuring the correlation of the final state of the network with the ground-truth pattern.

random labels and random rewards, respectively. Reliable stimuli are associated with consistent labels, or action-reward contingencies. We find that recall-gated consolidation provides significant benefits in each case, illustrating that the theoretical benefits of increased SNR in memory storage translate to improved performance on meaningful tasks (Fig. 3C,E,G).

## An analytical theory of the recall of repeatedly reinforced memories

We now turn to analyzing the behavior of the recall-gated systems consolidation model more systematically, to understand the source of its computational benefits and characterize other predictions it makes. To do so, we developed an analytic theory of memory system performance, with and without recall-gated consolidation. Importantly, our framework differs from prior work (Fusi et al., 2005; Benna and Fusi, 2016) in considering environments with intermittent repeated presentations of the same memory. We adopt several assumptions for analytical tractability. First, as in previous studies, we assume that inputs have been preprocessed so that the representations of different memories are random and uncorrelated with one another (Gluck and Myers, 1993; Benna and Fusi, 2016). We also assume, for now, that each memory consists of an equal number of candidate potentiation and depression events, though later we will relax this assumption. We are interested in tracking the SNR of recall for a given reliable memory. We emphasize that this quantity is an abstract measure of system performance reflecting the degree to which a specific set of synaptic changes (a memory trace) is retained in the system, and its interpretation varies according to the task in question (Fig. 3).

The dynamics of memory storage depend strongly on the underlying synapse model and plasticity rule. Given a synaptic model, an important quantity to consider is its associated “forgetting curve”  $m(t)$ , defined as the average SNR of recall for a memory  $\Delta w$  at  $t$  timesteps following its first presentation, assuming a new randomly sampled memory has been presented at each timestep since. For example, the binary switch model with transition probability  $p$  has an associated forgetting curve  $m(t) \approx \sqrt{N}pe^{-pt}$  (Fusi et al., 2005). More sophisticated synapse models, such as the cascade model of Fusi et al. (2005) and multivariable model of Benna and Fusi (2016) achieve power-law forgetting curves (see Methods). In the limit of large system size  $N$  and under the assumption that memories are random, uncorrelated patterns, the forgetting curve is an exact description of the decay of recall strength.

Forgetting curves capture the behavior of a system in response to a single presentation of a memory, but we are concerned with the behavior of memory systems in response to multiple reinforcements of the same memory trace. Thus, another key quantity in our theory is the interarrival distribution  $p(I)$ , which describes the distribution of intervals between repeated presentations of the same memory, and its expected value  $\tau = \mathbb{E}[I]$ , the average interval length. Our simplest case of interest is the case in which a given memory recurs according to a Poisson process; that is, it is reinforced with probability  $\lambda$  at each timestep, independent of the history of recent reinforcements (as in the simulation in Fig. 2C). This case corresponds to an exponential interarrival distribution  $p(I) = \lambda e^{-\lambda x}$ , with mean interarrival time  $\tau = 1/\lambda$ .

We now quantify the recall strength for a memory that has been reinforced  $R$  times. For the synapse models we consider, this quantity can be approximated accurately (see Supplemental Information) by summing the strengths of preceding memory traces, that is:

$$\text{SNR} \approx \sum_{i=1}^R m(t_i). \quad (2)$$

where  $t_i$  is the time elapsed since the  $i$ th reinforcement of the memory. This quantity is a random variable whose value depends on the history of interarrival intervals of the memory, and the specific unreliable memories that have been stored in intervening timesteps. To more concisely

characterize a system's memory performance, we introduce a new summary metric, the *learnable timescale* of the system. For a given target SNR value and allowable probability of error  $\epsilon$ , the learnable timescale  $\tau_{\beta}^{\epsilon}$  is defined as the maximum interarrival timescale  $\tau$  for which the SNR of recall will exceed  $\beta$  with probability  $1 - \epsilon$ . We fix  $\epsilon = 0.1$  throughout this work; this choice has no qualitative effect on our results. Learnable timescale captures the system's ability to reliably recall memories that are presented intermittently. We note that there exists a close relationship between learnable timescale and the memory capacity of the system (the number of memories it can store), with the two quantities becoming linearly related in environments with a high frequency of unreliable memory presentations (see Supplementary Information and [Fig. S3](#)).

The quantifications of recall SNR and learnable timescale we present in figures are computed numerically, as deriving exact analytical expressions for learnable timescale is difficult due to the randomness of the interarrival distribution. However, to gain theoretical intuition, we find it useful to consider the following approximation, corresponding to an environment in which memories are reinforced at deterministic intervals of length  $\tau$ :

$$\text{SNR} \approx \sum_{i=1}^R m(i\tau). \quad (3)$$

This approximation is an upper bound on the true SNR in the limit of small  $\epsilon$ , and empirically provides a close match to the true dependence of SNR on  $R$  (Supplementary Information, [Fig. S9](#)). Using this approximation allows us to provide closed-form analytical estimates of the behavior of SNR and learnable timescale as a function of system and environment parameters.

## Theory of recall-gated consolidation

In the recall-gated consolidation model, the behavior of the STM is identical to that of a model without systems consolidation. The LTM, on the other hand, behaves differently, updating only in response to the subset of memory presentations that exceed a threshold level of recall in the STM. From the perspective of the LTM, this phenomenon has the effect of changing the distribution of interval lengths between repeated reinforcements of a reliable memory. For exponentially distributed interarrival times, the induced effective interarrival distribution in the LTM is also exponential with new time constant  $\tau_{\text{LTM}}$  given by

$$\tau_{\text{LTM}} \approx \frac{P(I < m^{-1}(\theta))}{1 - \Phi(\theta)} \tau. \quad (4)$$

where  $I$  is the distribution of the lengths of intervals between presentations of the same reliable memory,  $\theta$  is the consolidation threshold, and  $\Phi$  is the cumulative distribution function of the Gaussian distribution. This approximation is valid in the limit of large system sizes  $N$ , where responses to unreliable memories are nearly Gaussian. For general (non-exponential) interarrival distributions, the shape of the effective LTM interarrival distribution may change, but the above expression for  $\tau_{\text{LTM}}$  remains valid.

We note that although the consolidation threshold  $\theta$  can be chosen arbitrarily, setting it to too high a value has the effect of reducing the probability with which reliable memories are consolidated, by a factor of  $P(I < m^{-1}(\theta))$ . For large values of  $\theta$  this reduction can become unacceptably small. For a given number of memory repetitions  $R$ , we restrict ourselves to values of  $\theta$  for which the probability that no consolidation takes place after  $R$  repetitions is smaller than the allowable probability of error  $\epsilon$ .

## Recall-gated consolidation increases SNR and learnable timescale of memories

For fixed statistics of memory presentations, as the SNR of the STM increases (say, due to increasing  $N$ ), stricter thresholds can be chosen for consolidation which filter out an increasing proportion of unreliable memory presentations, without reducing the consolidation rate of reliable memories (Fig. 4A, Fig. S1A). Consequently, the SNR of the LTM can grow much larger than that of the STM, and the amplification of SNR increases with the SNR of the STM. Notably, the SNR of the LTM in the recall-gated consolidation model also exceeds that of a control model in which STM and LTM modules are both present but do not interact (Fig. 4B, Fig. S1B).

We may also view the benefits of consolidation in terms of the learnable timescale of the system. Recall-gated consolidation enables longer learnable timescales, particularly at high target SNRs (Fig. 4C, Fig. S2). We note that our definition of SNR considers only noise arising from random memory sampling and presentation order. High SNR values may be essential for adequate task performance in the face of additional sources of noise, or when the system is asked to generalize by recalling partially overlapping memory traces (Benna and Fusi, 2016).

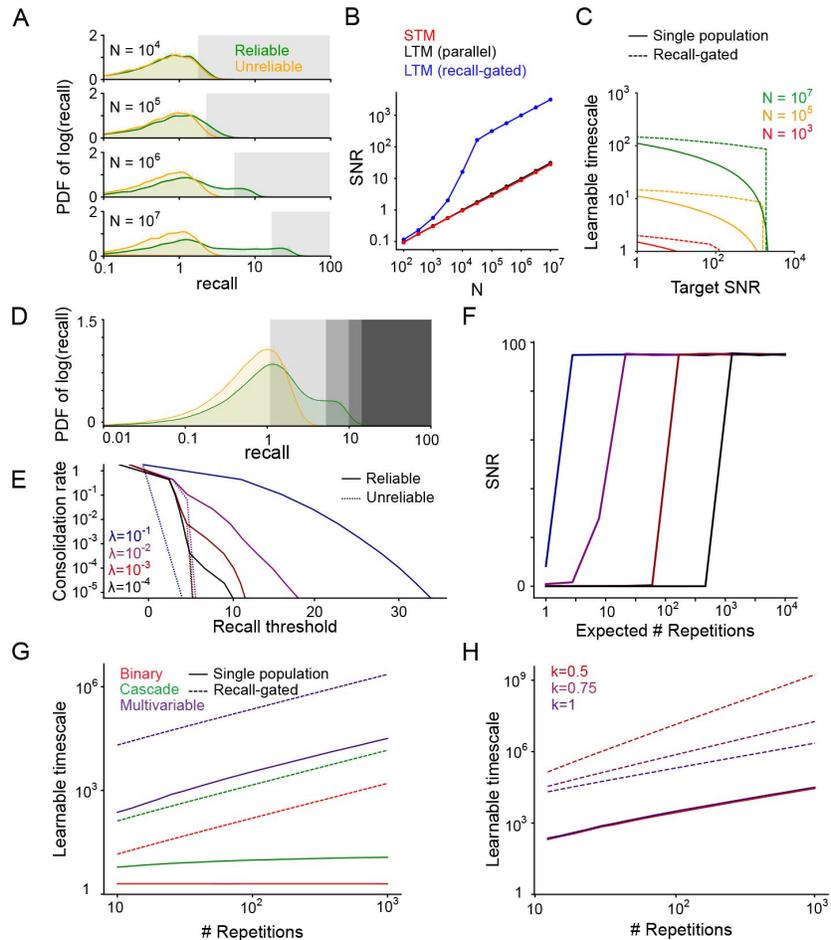
## Recall-gated consolidation enables better scaling of memory retention with repeated reinforcement

As mentioned previously, higher consolidation thresholds reduce the rate at which reliable memories are consolidated. However, the consolidation rate of *unreliable* memories decreases even more quickly as a function of the threshold (Fig. 4D,E). Hence, higher thresholds increase the fraction of consolidated memories which are reliable, at the expense of reducing the rate of consolidation into LTM. This tradeoff may be acceptable if reliable memories are reinforced a large number of times, as in this case they can still be consolidated despite infrequent LTM plasticity. In other words, as the number of anticipated repetitions  $R$  of a single reliable memory increases, higher thresholds can be used in the gating function, without preventing the eventual consolidation of that memory. Doing so allows more unreliable memory presentations to be filtered out and consequently increases the SNR in the LTM (Fig. 4F).

Assuming, as we have so far, that reliable memories are reinforced at independently sampled times at a constant rate, we show (calculations in Supplementary Information) that the dependence of learnable timescale on  $R$  is linear, regardless of the underlying synaptic model (Fig. 4G, Fig. S4). Synaptic models with a small number of states, such as binary switch or cascade models, are unable to achieve this scaling without systems consolidation (Fig. 4G). In particular, the learnable timescale is roughly invariant to  $R$  for the binary switch model, and scales approximately logarithmically with  $R$  for the cascade model (see Supplementary Information for derivation). Synaptic models employing a large number of internal states (growing exponentially with the intended timescale of memory retention), like the multivariable model of Benna and Fusi (2016), can also achieve linear scaling of learnable timescale on  $R$ . However, these models still suffer a large constant factor reduction in learnable timescale compared to models employing recall-gated consolidation (Fig. 4G).

## Consolidation dynamics depend on the statistics of memory recurrence

The benefit of recall-gated consolidation is even more notable when the reinforcement of reliable memories does not occur at independently sampled times, but rather in clusters. Such irregular interarrival times might naturally arise in real-world environments. For instance, exposure to one environmental context might induce a burst of high-frequency reinforcement of the same pattern



**Figure 4**

A. Distribution (probability density function) of reliable and unreliable memory overlaps (on a log scale), varying the number of synapses  $N$ . Shaded regions indicate consolidation thresholds that preserve 10% of reliable memory presentations. Units are standard deviations of the distribution of recall for randomly sampled memories.

B. LTM SNR induced by consolidation (with threshold set as in D) as  $N$  varies. The parallel model uses 10% slower learning rate (value of  $p$  in the binary switching synapse model) in the LTM than the STM.

C. Learnable timescale as a function of target SNR, for several values of  $N$ , using the binary switching synapse model with  $p = 100/\sqrt{N}$

D. Distribution of reliable and unreliable memory overlaps, with various potential gating thresholds indicated.

E. Fraction of memory presentations consolidated (log scale) vs recall threshold for consolidation.

F. LTM SNR induced by consolidation vs. the expected number of repetitions before consolidation occurs. Increasing the expected number of repetitions corresponds to setting a more stringent consolidation threshold which filters a higher proportion of reliable memory presentations.

G. Learnable timescale at a target SNR of 10 as a function of number of reliable memory repetitions for several underlying synapse models,  $N = 10^7$ .

H. Same as G, considering only the multivariable model as the underlying synapse model, and varying the interarrival interval regularity factor  $k$ .

of synaptic updates, followed by a long drought when the context changes. Intentional bouts of study or practice could also produce such effects. The systems consolidation model can capitalize on such bursts of reinforcement to consolidate memories while they can still be recalled.

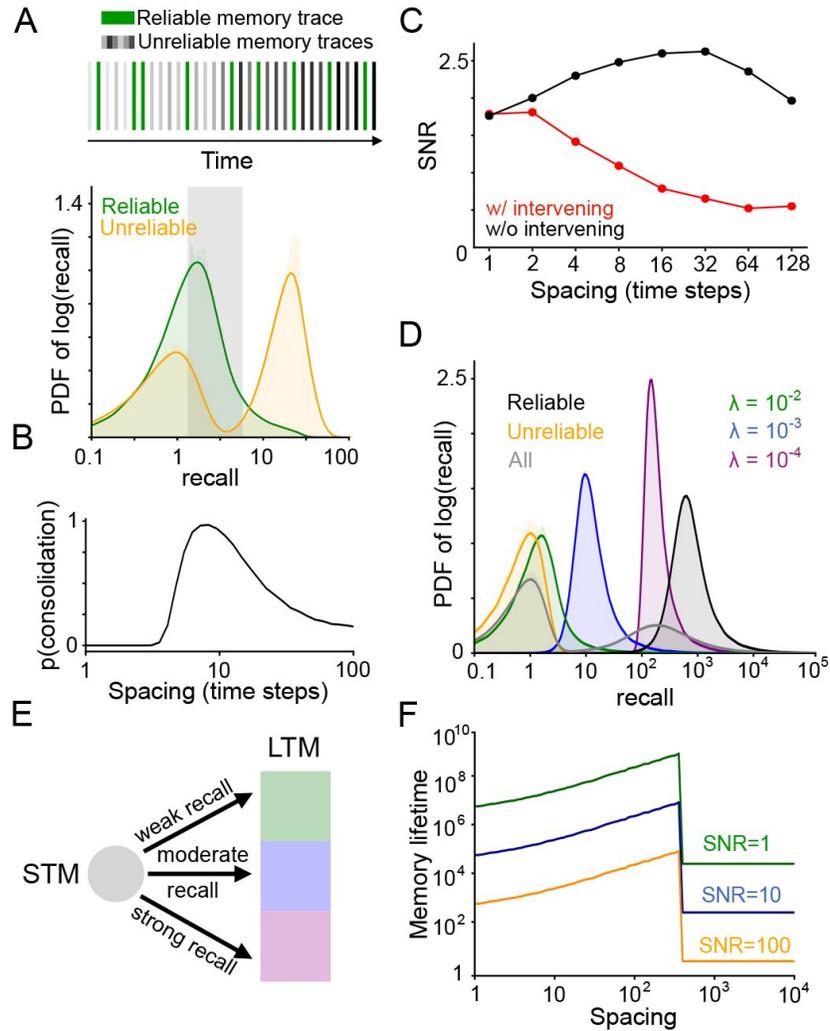
To formalize this intuition, we extend our theoretical framework to allow for more general patterns of memory recurrence. In particular, we let  $p(I)$  indicate the probability distribution of interarrival intervals  $I$  between reliable memory presentations. So far, we have considered the case of reliable memories whose occurrence times follow Poisson statistics, corresponding to an exponentially distributed interval distribution. To consider more general occurrence statistics, we consider a family of interarrival distributions known as Weibull distributions. This class allows control over an additional parameter  $k$  which modulates “burstiness” of reinforcement, and contains the exponential distribution as a special case ( $k = 1$ ). For  $k < 1$ , reliable memory presentations occur with probability that decays with time since the last presentation. In this regime, the same memory is liable to recur in bursts separated by long gaps (details in Methods).

Without systems consolidation, the most sophisticated synapse model we consider, multivariable model of [Benna and Fusi \(2016\)](#) [↗](#), achieves a scaling of learnable timescale that is linear with  $R$  regardless of the regularity factor  $k$ . In fact, we show (see Supplementary Information) that the best possible learnable timescale that can be achieved by any synaptic consolidation mechanism scales approximately linearly in  $R$ , up to logarithmic factors. However, for the recall-gated consolidation model, the learnable timescale scales as  $R^{1/k}$  when  $k \leq 1$  (**Fig. 4H** [↗](#), **Fig. S5** [↗](#)). In this sense, recall-gated consolidation outperforms any form of synaptic consolidation at learning from irregularly spaced memory reinforcement.

## Alternative gating functions suit different environmental statistics and predict spaced training effects

Thus far we have considered a threshold gating function, which is well-suited to environments in which unreliable memories are each only encountered once. We may also imagine an environment in which unreliable memories tend to recur multiple times, but over a short timescale (**Fig. 5A** [↗](#), top). In such an environment, the strongest evidence for a memory’s reliability is if it overlaps to an *intermediate* degree with the synaptic state (**Fig. 5A** [↗](#), bottom). The appropriate gating function in this case is no longer a threshold, but rather a nonmonotonic function of STM memory overlap, meaning that memories are most likely to be consolidated if reinforced at intermediate-length intervals (**Fig. 5B** [↗](#)). Such a mechanism is straightforward to implement using neurons tuned to particular ranges of recall strengths. This model behavior is consistent with spaced learning effects reported in flies ([Beck et al., 2000](#) [↗](#)), rodents ([Glas et al., 2021](#) [↗](#)), and humans ([Rovee-Collier et al., 1995](#) [↗](#); [Verkoeijen et al., 2005](#) [↗](#)), which all show a characteristic inverted U-shaped dependence of memory performance on spacing interval.

While some synapse-level models (such as the multivariable synapse model of [Benna and Fusi, 2016](#) [↗](#)) can also give rise to spaced training effects, these effects require that a synapse undergoes few additional potentiation or depression events between the spaced reinforcements (**Fig. 5C** [↗](#), **Fig. S6** [↗](#)). This is because spacing effects in such models arise when synapselocal variables are saturated, and saturation effects are disrupted when other events are interspersed between repeated presentations of the same memory. Hence, the spacing effects arising from such models are unlikely to be robust over long timescales. Recall-gated systems consolidation, on the other hand, can yield spaced training effects robustly in the presence of many intervening plasticity events.



**Figure 5**

A. Top: Example sequence of memory presentations where unreliable memories (gray) can repeat multiple times, but only within within a short timescale (note gradient of light to dark). Bottom: Distribution of reliable and unreliable memory overlaps induced by such memory presentation statistics (log scale on x axis). Shaded region indicates overlap values that are at least ten times as likely for reliable memories as for unreliable memories.

B. Probability of consolidation, with the gating function chosen such that only overlaps within the shaded region of panel A are consolidated, as a function of interarrival interval.

C. SNR at 8 timesteps following 5 spaced repetitions of a memory, with spacing interval indicated on the x axis, for the multivariable synapse model of Benna and Fusi (2016) with no systems consolidation. Spaced training effects are present at short timescales, but not if other memories are presented during the interpresentation intervals.

D. Distribution of recall strengths corresponding to different kinds of memories, in an environment with many reliable memories. In the environment model, reliable memories are reinforced with different interarrival interval distributions, and the timescales of these distributions for different memories are distributed loguniformly. The environment also has as a background rate of unreliable memory presentations appearing at fraction 0.9 of timesteps.

E. Depiction of a generalization of the model in which memories can be consolidated into different LTM sub-modules, according to gating functions tuned to different recall strengths (intended to target reliable memories with different timescales of recurrence).

F. A consequence of the model outlined in panel E is a smooth positive dependence of memory lifetime on the spacing of repetitions, up to some threshold spacing.

G. Learnable timescale calculations for this extended model still show a graceful scaling of learnable timescale with number of memory repetitions that depends on the regularity (Weibull distribution parameter  $k$ ) of interarrival intervals.

## Heterogeneous gating functions suit complex environments with multiple memories reinforced at different timescales

Thus far we have assumed a dichotomy between unreliable, one-off memories and reliable memories which recur according to particular statistics. In more realistic scenarios, there will exist many repeatedly reinforced memories, which may be reinforced at distinct timescales. We may be interested in ensuring good recall performance over a distribution of memories with varying recurrence statistics. For concreteness, we consider the specific case of an environment with a large number of distinct reliably reinforced memories, whose characteristic interarrival timescales are log-uniformly distributed. As before, unreliable memories are also presented with a constant probability per timestep.

The recall-gated plasticity model already described, using a threshold function for consolidation, still provides the benefit of filtering unreliable memory traces from the LTM. However, further improved memory recall performance is achieved with a simple extension to the model. The LTM can be subdivided into a set of subpopulations, each with distinct gating functions that specialize for different memory timescales by selecting for different recall strengths (Fig. 5D, E). That is, one subpopulation consolidates strongly recalled memories, another consolidates weakly recalled memories, and others lie on a spectrum between these extremes. The effect of this arrangement is to assign infrequently reinforced memory traces to subpopulations which experience less plasticity, allowing these traces to persist until their next reinforcement. This heterogeneity of timescales is consistent with observations in a variety of species of intermediate timescale memory traces (Rosenzweig et al., 1993; Cepeda et al., 2008; Davis, 2011).

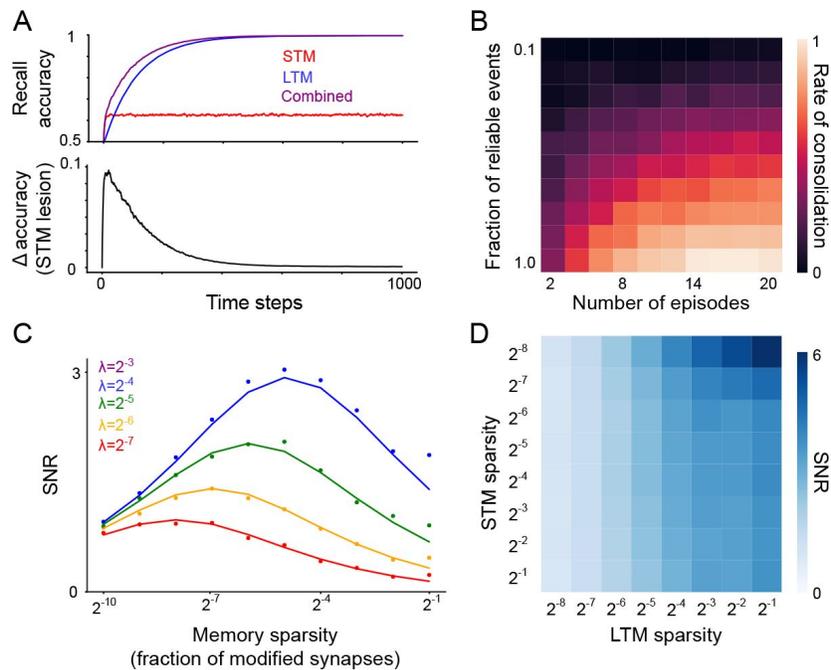
Studies of spaced training effects have found that the optimal spacing interval during training depends on the interval between training and evaluation (Cepeda et al., 2006, 2008). In particular, the timescale of memory retention is observed to increase smoothly with the spacing interval used during training. Our extended model naturally gives rise to this behavior (Fig. 5F, Fig. S7, Fig. S8), due to the fact that the lifetime of a consolidated memory scales inversely with the frequency with which memories are consolidated into its corresponding LTM subpopulation.

### Predicted features of memory representations and consolidation dynamics

The recall-gated consolidation model makes a number of key predictions. The most basic consequence of the model is that responsibility for recalling a memory will gradually shift from the STM to the LTM as consolidation progresses, rendering the recall performance of the system increasingly robust to the inactivation of the STM (Fig. 6A). A more specific prediction of the model is the *rate* of updates to the LTM increases with time, as STM recall grows stronger (Fig. 6B). The rate of LTM updates also increases with reliability of the environment (operationalized as the proportion of synaptic update events which correspond to reliable memories) (Fig. 6B).

The recall-gated consolidation model also makes predictions regarding neural representations in the STM and LTM. Until now we have assumed that memories consist of balanced potentiation and depression events distributed across the population. However, memories may involve only a sparse subset of synapses, for instance if synaptic plasticity arises from neural activity which is itself sparse. To formalize this notion, we consider memories that potentiate a fraction  $f$  of synapses, and a correspondingly modified binary switch plasticity rule such that potentiation activates synapses with probability  $p$  and depression inactivates synapses with probability  $\frac{f}{1-f}p$ .

We show analytically (see Supplementary Information) that in the limit of low  $f$ , the SNR-optimizing choice of  $f$  is proportional to the rate  $\lambda$  of reliable memory reinforcement (Fig. 6C). Other factors, such as energetic constraints and noise-robustness, may also affect the optimal coding level. In general, however, our analysis shows that environments with infrequent



**Figure 6**

A. Top: Recall performance for a single reliable memory (true positive rate, with a decision threshold set to yield a 10% false positive rate) as learning progresses. Simulation environment is the same as in Fig. 2C.  $N = 10^3$ ,  $\lambda = 0.25$ . Bottom: difference between combined recall performance and LTM-only performance. The STM makes diminishing contributions to recall over time.

B. Probability of consolidation into LTM increases with experience and with the reliability of the environment (parameterized here by the recurrence frequency  $\lambda$  of reliable memories). Simulation environment is the same as in panel A.

C. For a single population of binary synapses (no consolidation) and Poisson memory recurrence, mean SNR as a function of reliable memory recurrence frequency  $\lambda$  and memory sparsity  $f$ . Dots indicate simulation results and solid lines indicate analytical approximation.  $N = 1024$ .

D. For the systems consolidation model using binary synapses, total system SNR ( $N = 256$ ) as a function of memory sparsity in the STM and LTM.

reinforcement of a given reliable memory incentivize sparser representations. As the effective value of  $\lambda$  is amplified in the LTM module, it follows that the LTM benefits from a denser representation than the STM. Interestingly, we also find that the optimal sparsity in the STM decreases when optimizing for the overall SNR of the system—that is, the the optimal STM representation is even more sparse in the context of supporting LTM consolidation than it would be in isolation. Taken together, these two effects result in much denser representations being optimal in the LTM than in the STM (**Fig. 6D** [↗](#)). One consequence of denser representations is greater generalization in the face of input noise (Babadi and Sompolinsky, 2014 [↗](#)), implying that an optimal STM/LTM system should employ more robust and generalizable representations in the LTM.

## Discussion

We have presented a theory of systems memory consolidation via recall-gated long-term plasticity, which provides complementary benefits to synaptic consolidation mechanisms in terms of memory lifetime and retrieval accuracy. Its advantage arises from the ability to integrate over the information present in an entire neuronal population, rather than individual synapses, in order to decide which memory traces are consolidated. This capability is important in environments that induce a mixture of reliable and unreliable synaptic updates, in which a system must prioritize which updates to store long-term.

### Experimental evidence for recall-gated consolidation

The recall-gated consolidation model is by design agnostic to the underlying neural circuit and hence potentially applicable to a wide variety of species and brain regions. Here we summarize evidence consistent with recall-gated consolidation in several model organisms. As our proposal is new, the experiments we describe were not designed to directly test our model predictions, and thus provide incomplete evidence for them. We hope that future work will more directly clarify the relevance of our model to these systems as well as others, the mechanisms by which it is implemented, and the shortcomings it may have in accounting for experimental results.

### Associative learning in insects

In the *Drosophila* mushroom body, plasticity is driven by activity of dopamine neurons innervating a collection of anatomically defined compartments, some of which are associated with short or long-term memory (Aso et al., 2014 [↗](#)). These neurons receive a wide variety of inputs, including from mushroom body output neurons themselves (Li et al., 2020 [↗](#)). Such inputs provide a substrate by which long-term learning can be modulated by the outputs of short-term pathways. To implement recall-gated consolidation, the activity of dopamine neurons mediating long-term memory should be gated by learning in corresponding shortterm pathways. A recent study found an instance of this motif, in which short-term aversive learning decreases the activity of the  $\gamma 1$  output neuron, disinhibiting a dopamine neuron in the  $\alpha 2$  compartment associated with long-term aversive learning (Awata et al., 2019 [↗](#)). More work is needed to determine if other examples of this motif can be found in *Drosophila* or other insects.

### Motor learning

Several lines of work have indicated that the neural substrate of motor skills can shift with practice. In songbirds, learned changes to song performance are initially driven by a cortico-basal ganglia circuit called the anterior forebrain pathway (AFB) but eventually are consolidated into the song motor pathway (SMP) and become AFB-independent (Andalman and Fee, 2009 [↗](#); Warren et al., 2011 [↗](#)). Using transient inactivations of LMAN, a region forming part of the AFB, a recent study quantified the degree of AFB-to-SMP consolidation over time and found that it strongly

correlated with the bird's motor performance at the time (Tachibana et al., 2022). This finding is consistent with our model's prediction that the *rate* of consolidation should increase as learning progresses in the short-term pathway.

A related motor consolidation process has been observed during motor learning in rats. Experiments have shown that motor cortex disengages from heavily practiced skills (Kawai et al., 2015; Hwang et al., 2019), transferring control at least in part to the basal ganglia (Dhawale et al., 2019, 2021), and that the degree of cortical disengagement tracks motor performance, as measured by the variability of learned trajectories (Hwang et al., 2021). This finding is broadly consistent with recall-consolidation, with short-term learning being mediated by motor cortex and long-term learning being mediated by basal ganglia. However, we note that unlike in the song learning study referenced above, it neither confirms nor rejects our stronger prediction that the *rate* (rather than overall extent) of motor consolidation increases with learning.

## Spatial learning and hippocampal replay

Hippocampal replay is thought to be crucial to the consolidation of episodic memories to cortex (Carr et al., 2011; Ólafsdóttir et al., 2018). Replay has many proposed computational functions, such as enabling continual learning (van de Ven et al., 2020), or optimizing generalization performance (Sun et al., 2021), which are outside the scope of our model. However, under the assumption that replay enables long-term memory storage in cortex, the recall-gated consolidation model makes predictions about which memories should be replayed—namely, replay should disproportionately emphasize memories that are familiar to the hippocampus. That is, we would predict more frequent replay of events or associations that are frequently encountered than of those that were experienced only once, or unreliably.

Recent experimental work supports this hypothesis. A recent study found that CA3 axonal projections to CA1, those that respond visual cues associated with a fixed spatial location are recruited more readily in sharp-wave ripple events than those that respond to the randomly presented cues (Terada et al., 2021). Earlier work found that sharp-wave ripple events occur more frequently during maze navigation sessions with regular trajectories, and increase in frequency over the course of session, similar to the behavior of our model in Fig. 4B (Jackson et al., 2006). Thus, existing evidence suggests that hippocampal replay is biased toward familiar patterns of activity, consistent with a form recall-gated consolidation. Other experiments provide preliminary evidence for signatures of such a bias in cortical plasticity. For instance, fMRI study of activity in hippocampus and posterior parietal cortex (PPC) during a human virtual navigation experiment found that that the recruitment of PPC during the task, which was linked with memory performance, tended to strengthened with experience in a static environment, but did not strengthen when subjects were exposed to an constantly changing environment, consistent with consolidation of only reliable memories (Brodt et al., 2016).

## Comparison with synaptic consolidation mechanisms

Recall-gated consolidation improves memory performance regardless of the underlying synapse model (Fig. 4), indicating that its benefits are complementary to those of synaptic consolidation. Our theory quantifies these benefits in terms of the scaling behavior of the model's maximum learnable timescale with respect to other parameters. First, for any underlying synapse model, recall-gated consolidation allows learnable timescale to decay much more slowly as a function of the desired SNR of memory storage. Second, recall-gated consolidation achieves (at worst) linear scaling of learnable timescale as a function of the number of memory reinforcements  $R$ . For models with a fixed, finite number of internal states per synapse, this scaling is at best logarithmic. Our results therefore illustrate that systems-level consolidation mechanisms allow relatively simple synaptic machinery (Emes et al., 2008) to support long-term memory storage. We note that more sophisticated synaptic models, which involve a large number of internal states that scales with the memory timescale of interest (Benna and Fusi, 2016), can also achieve linear

scaling of learnable timescale with  $R$  (though recall-gated consolidation still improves their performance by a large constant factor). However, for environmental statistics characterized by concentrated bursts of repeated events separated by long gaps, recall-gated consolidation achieves superlinear power-law scaling, which we showed is not achievable by any synapse-local consolidation mechanism.

Our model provides an explanation for spaced training effects (Fig. 5) based on optimal gating of long-term memory consolidation depending on the recurrence statistics of reliable stimuli. It is important to note that, depending on the specific form of internal dynamics present in individual synapses, synaptic consolidation models can also reproduce spacing effects. For example, the initial improvement of memory strength with increased spacing arises in the model of Benna and Fusi (2016) due to saturation of fast synaptic variables, meaning that the timescale of these internal variables determines optimal spacing, and that intervening stimuli can block the effect by preventing saturation (Fig. 5). In contrast, in our model this timescale is set by population-level forgetting curves, rendering spacing effects robust over long timescales and in the presence of intervening events. It is likely that mechanisms at both the synaptic and systems level contribute to spacing effects; our results suggest that effects observed at longer timescales are likely to arise from memory recall mechanisms at the systems level.

## Other models of systems consolidation

Unlike previous theories, our study emphasizes the role of repeated memory reinforcement in consolidation and explicitly quantifies the consequences of such reinforcement. However, there are important connections between our work other computational models of systems consolidation that have been proposed. Much of this work focuses on consolidation via hippocampal replay. Prior work has proposed that replay (or similar mechanisms) can prolong memory lifetimes (Shaham et al., 2021; Remme et al., 2021), alleviate the problem of catastrophic forgetting of previously learned information (van de Ven et al., 2020; González et al., 2020; Shaham et al., 2021), and facilitate generalization of learned information (McClelland et al., 1995; Sun et al., 2021). One prior theoretical study (Roxin and Fusi, 2013), which uses replayed activity to consolidate synaptic changes from short to long-term modules, explored how systems consolidation extends forgetting curves. Unlike our work, this model does not involve gating of memory consolidation, and consequently provides no additional benefit in consolidating repeatedly reinforced memories. Our model is thus distinct from, but also complementary to, these prior studies. In particular, recall-gated consolidation can be implemented in real-time, without replay of old memories. However, as discussed above, selective replay of familiar memories is one possible implementation of recall-gated consolidation. Selective replay is a feature of some of the work cited above (Shaham et al., 2021; Sun et al., 2021), which suggests it can provide advantages for retention and generalization (Shaham et al., 2021; Sun et al., 2021).

Other work has proposed models of consolidation, particularly in the context of motor learning, in which one module “tutors” another to perform learned behaviors by providing it with target outputs (Murray et al., 2017; Teşileanu et al., 2017). Murray and Escola (2020) proposes a fast-learning pathway (which learns using reward or supervision) which tutors the slow-learning long-term module via a Hebbian learning rule. In machine learning, a similar concept has become popular (typically referred to “knowledge distillation”), in which the outputs of a trained neural network are used to supervise the learning of a second neural network on the same task (Hinton et al., 2015; Gou et al., 2021). Empirically, this procedure is found to improve generalization performance and enable the use of smaller networks. Our model can be interpreted as a form of partial tutoring of the LTM by the STM, as learning in the LTM is partially dictated by outputs of the STM. In this sense, our work provides a theoretical justification for the use of tutoring signals between two neural populations.

## Limitations and future work

In addition to motivating new experiments to test the predictions of a recall-gated consolidation model, our work leaves open a number of theoretical questions that future modeling could address. Our theory assumes fixed and random representations of memory traces. Subject to this assumption, we showed that short-term memory benefits from sparser representations than long-term memory. In realistic scenarios, synaptic updates are likely to be highly structured, and the optimal representations in each module could differ in more sophisticated ways. Moreover, adapting representations online—for instance, in order to decorrelate consolidated memory traces—may improve learning performance further. Addressing these questions requires extending our theory to handle memory statistics with nontrivial correlations. Another possibility we left unaddressed is that of more complex interactions between memory modules—for instance, reciprocal rather than unidirectional interactions—or the use of more than two interacting systems.

Finally, in this work we considered only a limited family of ways in which long-term consolidation may be modulated—namely, according to threshold-like functions of recall in the short-term pathway. Considering richer relationships between recall and consolidation rate may enable improved memory performance and/or better fits to experimental data. Moreover, in real neural circuits, additional factors besides recall, such as reward or salience, are likely to influence consolidation as well. Unlike our notion of recall, which can be modeled in task-agnostic fashion, the impact of such additional factors on learning likely depends strongly on the behavior in question. Our work provides a theoretical framework that will facilitate more detailed models of the rich dynamics of consolidation in specific neural systems of interest.

## Acknowledgements

We would like to thank Stefano Fusi and Samuel Muscinelli for helpful discussions and comments on the manuscript. A. L.-K. and J. L. were supported by the Gatsby Charitable Foundation, NSF award DBI-1707398. J. L. was also supported by the DOE CSGF (DE-SC0020347). A. L.-K. was also supported by the Burroughs Wellcome Foundation, the McKnight Endowment Fund, the Mathers Foundation, and NIH award R01EB029858.

## Methods

### Theoretical Framework

We consider a population of  $N$  synapses, indexed by  $i \in \{1, 2, \dots, N\}$  each with a synaptic weight  $w_i \in \mathbb{R}$ . The set of synaptic weights across the population can be denoted by the vector  $\mathbf{w} \in \mathbb{R}^N$ . The synapses may retain additional information besides strength as well; if each synapse carries  $d$ -dimensional state information in addition to its strength, the synaptic state can be written as  $\tilde{\mathbf{w}}_i \in \mathbb{R}^d$ , with the scalar synaptic strengths  $w_i \in \mathbb{R}$  defined as a function of the high-dimensional state  $\tilde{\mathbf{w}}_i$ . We define memories as patterns of candidate synaptic modifications, following prior work (Benna and Fusi, 2016 [↗](#); Fusi et al., 2005 [↗](#)). More specifically, we model each memory as a vector  $\Delta \mathbf{w} \in \mathbb{R}^N$  of candidate potentiation and depression events. By defining memories in this fashion, our analysis can remain agnostic to the network architecture and learning rule that give rise to synaptic modifications. We will typically model memories as binary potentiation/depression events for simplicity, but in principle, memories can be continuous valued. Synaptic are updated by memories according to a learning rule  $(\tilde{\mathbf{w}}_i)_{new} = \ell(\tilde{\mathbf{w}}_i, \Delta w_i)$ , which maps the synaptic state at the time of a memory event to the subsequent synaptic state.

For theoretical calculations, we assume as in prior work (Fusi et al., 2005 [↗](#); Benna and Fusi, 2016 [↗](#)), that the components of each memory  $\Delta\mathbf{w}$  are independent and uncorrelated with those of other memories (though this assumption is violated in our task learning simulations). We also assume for simplicity that memories are mean-centered so that  $\mathbb{E}[\mathbf{w} \cdot \Delta\mathbf{w}] = 0$  over randomly sampled memories  $\Delta\mathbf{w}$ .

We define the *recall strength* associated with memory as the overlap  $r = \mathbf{w} \cdot \Delta\mathbf{w}$ . This definition reflects an “ideal observer” perspective, as it requires direct and complete access to the state of the synaptic population. The ideal observer perspective provides an upper bound on the recall performance of a real system, and should be a fairly good approximation assuming that memory readout mechanisms are sophisticated enough. We are particularly interested in the normalized recall strength

$$\text{SNR} = \frac{\mathbf{w} \cdot \Delta\mathbf{w}}{\sqrt{\mathbb{E}[(\mathbf{w} \cdot \Delta\mathbf{w}')^2]}} \quad (5)$$

where the expectation is taken over randomly sampled memories  $\Delta\mathbf{w}'$ . We refer to this quantity as the signal-to-noise ratio (SNR) of memory recall.

## Synaptic models and plasticity rules

In this paper we primarily consider three synaptic models and corresponding learning rules, taken from prior work.

The first and simplest is a “binary switch” model in which synapses take on binary ( $\pm 1$ ) values and stochastically activate (resp. inactivate) in response to candidate potentiation (resp. depression) events with probability  $p$  (Amit and Fusi, 1994 [↗](#)). No auxiliary state variables are used in this model.

The second is the “cascade” model of Fusi et al. (2005) [↗](#), in which synapses are modeled as a Markov chain with a finite number  $2k$  of discrete states with transition probabilities dependent on the kind of memory event (potentiation or depression). Half the states (states  $a_1, \dots, a_k$ ) are considered potentiated (strength +1) and half (states  $b_1, \dots, b_k$ ) are depressed (strength -1). Intuitively, states of the same potentiation level correspond to different propensities for plasticity in the synapse, enabling a form of synaptic consolidation. Formally, for  $i < k$ , the potentiated state  $a_i$  (resp. depressed state  $b_i$ ) transitions to state  $a_{i+1}$  (resp.  $b_{i+1}$ ) with probability  $\frac{\alpha^i}{1-\alpha}$  following a potentiation (resp. depression) event. And for  $i < k$ , the potentiated state  $a_i$  (resp. depressed state  $b_i$ ) transitions to state  $b_1$  (resp.  $a_1$ ) with probability  $\alpha^{i-1}$  following a depression (resp. potentiation) event. For  $i = k$  this latter transition occurs with probability  $\frac{\alpha^{i-1}}{1-\alpha}$  (as described in Fusi et al. (2005) [↗](#)), this choice is made for convenience to ensure equal occupancy of the different synaptic states). We assume  $\alpha = 0.5$  throughout.

The third synaptic model is the model of Benna and Fusi (Benna and Fusi, 2016 [↗](#)), which we refer to as the “multivariable” model. In this model, synapses are described by a chain of  $m$  interacting continuous-valued variables  $u_1, \dots, u_m$ , the first of which corresponds to synaptic strength. Potentiation and depression events increment or decrement the value of the first synaptic variable, and a set of difference equations governs the evolution of the multidimensional state at each time step:

$$\frac{du_i}{dt} = n^{-2i+2}\alpha(u_{i-1} - u_i) - n^{-2i+1}\alpha(u_i - u_{i+1}), \quad (6)$$

where  $n$  and  $\alpha$  a parameter of the model (we assume  $n = 2$ ,  $\alpha = 0.5$  throughout). This model also provides the ability for synapses to store information at different timescales, due to the information retained in auxiliary variables.

## Model implementation for example tasks

**Supervised Hebbian learning** We simulated a single-layer feedforward network with a population of  $N$  input units (activity denoted by  $\mathbf{x}$ ) and a single output unit, (activity denoted by  $\hat{y}$ ), connected with a  $1 \times N$  binary weight matrix  $\mathbf{W}$ , such that  $\hat{y} = \mathbf{W}\mathbf{x}$ . In each simulation, a set of  $P = 20$  reliable stimuli were randomly generated, which corresponded to binary  $(\pm \frac{1}{N})$  random  $N$ -dimensional activity patterns in the input units. Note that due to the scaling of the inputs and use of binary synaptic weights, the activity  $\hat{y}$  is constrained to lie in the interval  $[-1, 1]$ . Each reliable stimulus was associated with a randomly chosen (but consistent across the simulation) label  $y$ , 1 or -1. At each time step one of the reliable stimuli (along with its label) was presented to the network with probability  $\lambda_i = 0.01$  for all  $i = 1, \dots, P$ . Otherwise (with probability  $1 - \sum_i \lambda_i$ ), a randomly sampled unreliable stimulus was presented with a randomly chosen label. Weights experienced candidate potentiation or depression events given by a Hebbian learning rule  $\Delta W = y\mathbf{x}^T$ , corresponding to the product of the binary input neuron activity and the corresponding label. Learning followed the binary switch rule with  $p = 0.1$ ; that is, candidate potentiation events resulted in potentiation with probability  $p$ , and likewise for depression events. At each timestep the product of the STM output and the  $\pm 1$  label was computed, and if it exceeded the consolidation threshold  $\theta = 0.125$ , plasticity was permitted in the LTM network.

## Reinforcement learning

We used the same setup as in the supervised learning task, with the following modifications. The activity of the output unit (denoted by  $\pi$  in this problem) represented the probability of taking one of two possible actions:  $p(a_1) = \pi$  and  $p(a_2) = 1 - \pi$ . Each reliable stimulus was associated with a correct action. Taking the correct action yielded a reward of  $r = 1$ , while taking the other action yielded a reward of  $r = -1$ . Plasticity events followed the following three-factor learning rule:  $\Delta W = (r \cdot a)\mathbf{x}^T$ . At each timestep the product of the STM action probability and the reward was computed, and if it exceeded the consolidation threshold  $\theta = 0.75$ , plasticity was permitted in the LTM network. For all plasticity we used the binary switch rule with  $p = 1.0$ .

## Unsupervised Hebbian learning

We simulated two recurrent neural networks with  $N = 1000$  binary units each and with binary recurrent weight matrices  $\mathbf{W}_{\text{STM}}$  and  $\mathbf{W}_{\text{LTM}}$ , respectively. Memories consisted of binary (entries equal to  $\pm \frac{1}{N}$ ) random  $N$ -dimensional vectors that provided direct input  $\mathbf{x}$  to the network units at each timestep. The network state  $\mathbf{h}$  evolved for  $T = 5$  timesteps according to the following dynamics equation

$$\mathbf{h}_{t+1} = \mathbf{W}\phi(\mathbf{h}) + \mathbf{x}, \quad (7)$$

where  $\phi$  is a binary threshold nonlinearity with threshold set so that 50 percent of units were active at each time step (corresponding to a mechanism that normalizes activity across the network). The weights  $\mathbf{W}$  of the network were binary and initialized as binary random variables with equal on/off probability. The network weights  $W_{ij}$  were subjected to potentiation events when  $x_i$  and  $x_j$  were both active at  $t = 0$ , and otherwise subjected to depression events. Synaptic updates followed the binary switch rule with probability  $p = 1.0$ .

Additionally, a set of  $N$  weights  $\mathbf{u}$  connected the STM units to a single readout unit that measured familiarity. These weights were also binary and experienced candidate potentiation/depression events when their corresponding unit was active/inactive, respectively. These weights followed the

binary switch rule with probability  $p = 1.0$ .

Plasticity in the LTM proceeded according to the same rule as in the STM but was gated by recall strength  $r = \mathbf{u} \cdot \mathbf{x}_{\text{STM}}$ , according to a threshold function with threshold equal to 0.25.

The performance of the network was determined by presenting noise-corrupted versions of the reliable memory and measuring the correlation between the network state and the uncorrupted memory after  $T = 5$  time steps. The corrupted patterns were obtained by adding Gaussian noise of variance  $\frac{1}{N} = 0.001$  to the ground-truth pattern, and binarizing the result by choosing the fraction 0.5 of units with the highest values to be active.

## Forgetting curves for different synaptic learning rules

Prior work (Fusi et al., 2005; Benna and Fusi, 2016) has considered environments in which a given memory is presented to the system only once. In this case, the performance of a single population of synapses with a given learning rule depends crucially on the memory trace function  $m(t)$ . This is defined as

$$m(t) = \frac{\mathbf{w}(t) \cdot \Delta \mathbf{w}}{\sqrt{\mathbb{E}[(\mathbf{w}(t) \cdot \Delta \mathbf{w}')^2]}}, \quad (8)$$

the recall SNR at time  $t$  for a memory  $\Delta \mathbf{w}$  presented at  $t = 0$ , assuming randomly sampled memories have been presented in the intervening timesteps. For the binary switch model,  $m(t) \approx \sqrt{N} p e^{-pt}$ . More sophisticated synaptic models, like the cascade and multivariable models, can achieve power-law scalings (Fusi et al., 2005; Benna and Fusi, 2016). The key feature of these models that enables power-law forgetting is that their synapses maintain additional information besides their weight, which encodes their propensity to change state. In this fashion, memories can be consolidated at the synaptic level into more stable, slowly decaying traces. The cascade model of Fusi et al. achieves

$$m(t) \approx \frac{\sqrt{N}}{t \log T} e^{-t/T} \quad (9)$$

for some characteristic timescale  $T$  which can be chosen as a model parameter. Hence, its performance is upper bounded by

$$m(t) \approx \frac{\sqrt{N}}{t \log t}. \quad (10)$$

The model of Benna and Fusi can achieve

$$m(t) \approx \sqrt{\frac{N}{t \log T}} e^{-t/T} \quad (11)$$

which is upper bounded by

$$\sqrt{\frac{N}{t \log t}}. \quad (12)$$

Benna and Fusi also show that  $\sqrt{N/t}$  scaling is an upper bound on the performance of any synapse model with finite dynamic range.

## Implementation of SNR and learnable timescale computations

To compute recall strengths associated with single synaptic populations, we first sampled interarrival intervals  $I$  from the environmental statistics  $p(I)$ . Given a number of repetitions  $R$ , we computed then computed recall strength samples  $r' = \sum_i^R m(I_i)$ , where  $m$  is the forgetting curve associated with the underlying synaptic learning rule, and  $I_i$  are independent samples from  $p(I)$ . We scaled recall strengths by a factor of  $\frac{1}{\sqrt{N}}$  to compute the recall SNR (an approximation that is exact in the large- $N$  limit).

To compute recall strengths associated with the recall-gated consolidation model, we repeated the above procedure using a new interarrival distribution  $p(I_{LTM})$  induced by the gating model. The induced distribution  $I_{LTM}$  is obtained by drawing as samples the lengths of intervals between consecutive interarrival interval samples for which the corresponding recall SNR in the STM exceeds the gating threshold  $\theta$  (corresponding to the interval between consolidated reliable memory presentations), and rescaling it by the fraction of unreliable memories that are consolidated. Strictly speaking, in the general case this distribution is nonstationary, as the probability of STM recall exceeding the threshold can change as synaptic updates accumulate across repetitions for sophisticated synapse models like that of [Benna and Fusi \(2016\)](#). We adopt a conservative approximation that ignores such effects and thus slightly underestimates the rate of consolidation when such synaptic models are used (and consequently underestimates the SNR and learnable timescale of the recall-gated consolidation model). With this approximation, the random variable  $I_{LTM}$  is defined as as the following mixture distribution

$$\begin{aligned} \tilde{I}_{LTM} &= \sum_{i=1}^j I_i \quad w.p. \quad q(1-q)^{j-1}, \\ I_{LTM} &= \frac{\tilde{I}_{LTM}}{\sum_{t=1}^{\tilde{I}_{LTM}} 1[\zeta_t > \theta]} \end{aligned} \quad (13)$$

where each  $I_i \sim p(I)$ ,  $\zeta_t \sim \mathcal{N}(0, 1)$ , and  $q$  indicates the probability of a reliable memory presentation inducing consolidation. The value of  $j$  corresponds to the number of reinforcements that go by between instances of consolidation. For sufficiently large  $\tau$  this distribution can be approximated by

$$I_{LTM} = \frac{\sum_{i=1}^j I_i}{1 - \Phi(\theta)} \quad w.p. \quad q(1-q)^{j-1} \quad (14)$$

where  $\Phi$  is the CDF of the standard normal distribution. For large  $N$ , the probability of consolidation  $q = P(I < m^{-1}(\theta))$ .

We note that for an exponential interarrival distribution with mean  $\tau$ , the induced distribution of  $I_{LTM}$  is also exponential, with mean  $\tau_{LTM} = \frac{P(I < m^{-1}(\theta))}{1 - \Phi(\theta)}$ . This is because the sums  $1 - \Phi(\theta)$  of  $j$  independent samples  $I_i$  are distributed according to a Gamma distribution with shape parameter  $j$ , and the mixture of such Gamma distributions with geometrically distributed mixture weights  $p(j) = q(1-q)^{j-1}$  is itself an exponential distribution with mean  $\tau/q$ .

For a given number of expected memory repetitions  $R$ , the gating threshold  $\theta$  was set such that at least one of the  $R$  repetitions would be consolidated with probability  $1 - \epsilon$ ,  $\epsilon = 0.1$ . Where  $R$  is not reported, we assume it equal to 2, the minimum number of repetitions for the notion of consolidation to be meaningful in our model.

To compute learnable timescales, we repeated the above SNR computations over a range of mean interarrival times  $\tau = \mathbb{E}[I]$ , keeping the interarrival distribution family (Weibull distributions with a fixed value of  $k$ , see below) constant. We report the maximum value of  $\tau$  for which the SNR

exceeds the designated target threshold with probability  $1 - \epsilon$ ,  $\epsilon = 0.1$ .

Throughout, for our interarrival distributions we use Weibull distributions with regularity parameter  $k$ . The corresponding cumulative distribution function is

$$P(I \leq t) = 1 - e^{-(t\Gamma(1+1/k)/\tau)^k}. \quad (15)$$

Where  $\tau = \mathbb{E}[I]$ , and  $\Gamma$  is the Gamma function. In the case  $k = 1$ , this reduces to an exponential distribution of interarrival intervals, which corresponds to memory reinforcements that occur according to a Poisson process with rate  $\lambda = 1/\tau$ . In the limit  $k \rightarrow \infty$ , it corresponds to interarrival intervals of deterministic length  $\tau$ . For  $k < 1$ , the interarrival distribution is “bursty,” with periods of dense reinforcement separated by long gaps.

## Spacing effect simulations

We simulated the multivariable synapse model of Benna and Fusi (Benna and Fusi, 2016 [↗](#)), in which each synapse is described by  $m$  continuous-valued dynamical variables  $u_1, \dots, u_m$  which evolve as follows:

$$\frac{du_i}{dt} = n^{-2i+2}\alpha(u_{i-1} - u_i) - n^{-2i+1}\alpha(u_i - u_{i+1}). \quad (16)$$

For the first variable  $u_1$ , in place of  $u_{i-1}$  we substitute components  $\Delta w_j$  of the memory traces. For the last variable  $u_m$ , in place of  $u_{i+1}$  we substitute 0. The strength of each synapse corresponds to the value of its first dynamical variable. For our simulations we chose  $m = 10$  dynamical variables,  $n = 2$ ,  $\alpha = 0.5$ , and  $N = 400$  synapses. The value of  $\alpha$  is also varied in [Fig. S6](#) [↗](#). A spacing interval  $\Delta$  was selected and a randomly drawn reliable memory was presented at  $\Delta$ -length intervals (the same pattern at each presentation). In the case without intervening memories, the dynamics of each synapse ran unimpeded between these presentations. In the case with intervening memories, new randomly drawn patterns were presented to the system at each timestep between the reliable memory presentations. Each pattern was drawn with values equal to  $\pm \frac{1}{2}$ , with equal probability.

## Generalized model with multiple memory timescales

In our generalized environment model, the environment contains a variety of distinct reliable memories  $x_i$  which recur with Poisson statistics at a variety of rates  $\lambda_i$ . Timescales  $\tau_i = 1/\lambda_i$  are distributed as  $p(\log \tau) \sim [0, A]$  where  $A$  is a large constant. This corresponds to the value of  $\log \lambda$  being uniformly distributed in  $[-A, 0]$ , or equivalently to  $p(\lambda) \sim 1/\lambda$  and bounded between  $e^{-A}$  and 1. The environment also contains an additional fraction of unreliable memories as before, sampled randomly and presented with a fixed probability at each timestep. The natural generalization of learnable timescale to this setting is the maximum interarrival interval timescale for which the lifetime of a corresponding memory (the time following last reinforcement its recall strength decays to an SNR below the target SNR) exceeds that timescale.

The distribution of interarrival intervals for memory  $i$  is

$$p(I_i) = \lambda_i e^{-\lambda_i I_i}. \quad (17)$$

Integrating across the distribution of  $\lambda$ , we get the distribution of interarrival intervals for reliable memories observed by the system:

$$\begin{aligned}
 p_{\text{reliable}}(I) &\approx \mathbb{E}_{\log \lambda \sim U([-A, 0])} [\lambda e^{-\lambda I}] \\
 &\approx \int_{\lambda=e^{-A}}^{\lambda=1} \frac{1}{\lambda} \lambda e^{-\lambda I} d\lambda \\
 &\approx \int_0^1 e^{-\lambda I} d\lambda \\
 &= \frac{1 - e^{-I}}{I} \\
 &\approx 1/I
 \end{aligned} \tag{18}$$

for large  $I$ .

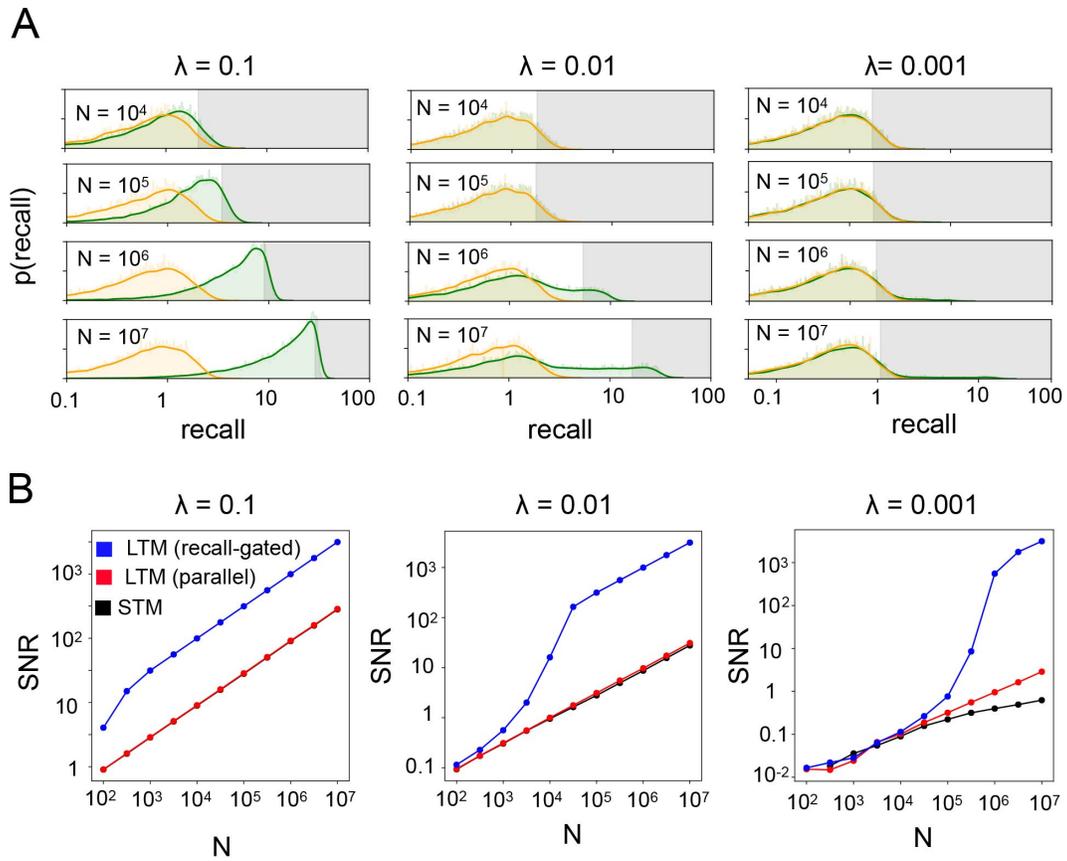
The full distribution of interval strengths (including unreliable memories) is a mixture of  $p_{\text{reliable}}$  and a delta function at  $I = \infty$ , with the latter's weight corresponding to the probability with which an unreliable memory is sampled at a given timestep (in our simulations we chose 0.9).

From here we can compute a distribution of STM recall strengths  $r$

$$p(r) = \int p(I)m(I)dI. \tag{19}$$

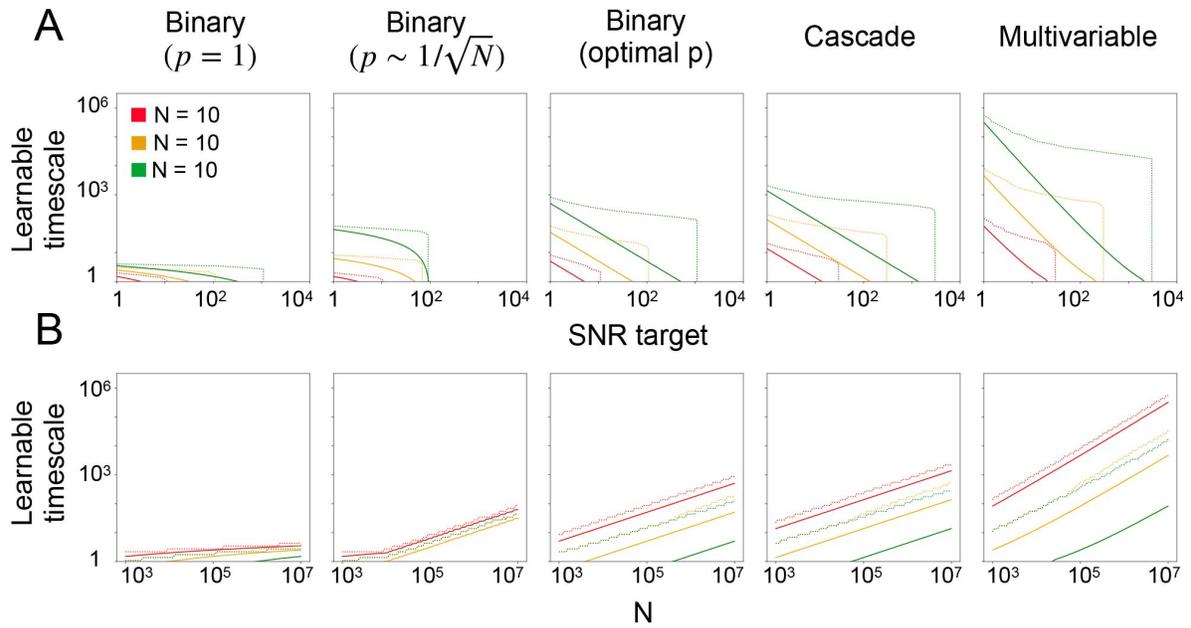
We simulated a model in which an ensemble of LTM subpopulations are assigned gating functions  $g_i(r)$  equal 1 for  $\log r \in [A_i, A_{i+1}]$  and 0 elsewhere, with the  $A_i$  spaced evenly over  $[0, \frac{1}{2} \log N]$ . The expected lifetime of a memory reinforced with a given interval  $I'$  is given approximately by the STM lifetime divided by the fraction of memory presentations for which the recall strength lies in the same interval as  $m(I')$ . This quantity reflects the proportion of memories presentations that are consolidated into the same LTM subpopulation as the memory in question.

## Supplemental Figures



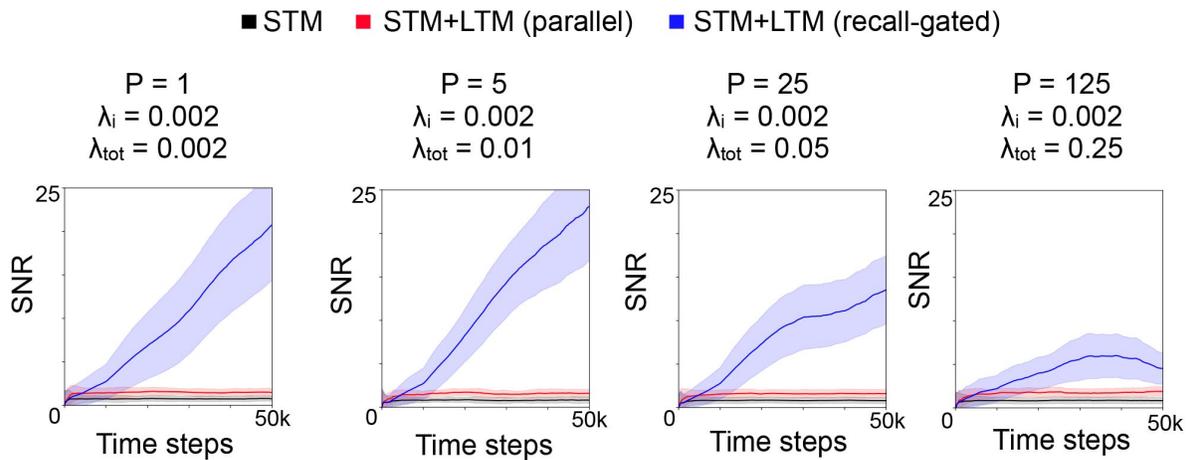
### Supplemental Figure S1

A. Same as **Fig. 4A**, also varying the presentation rate  $\lambda$  of reliable memories. B. Same as **Fig. 4B**, also varying the presentation rate  $\lambda$  of reliable memories.



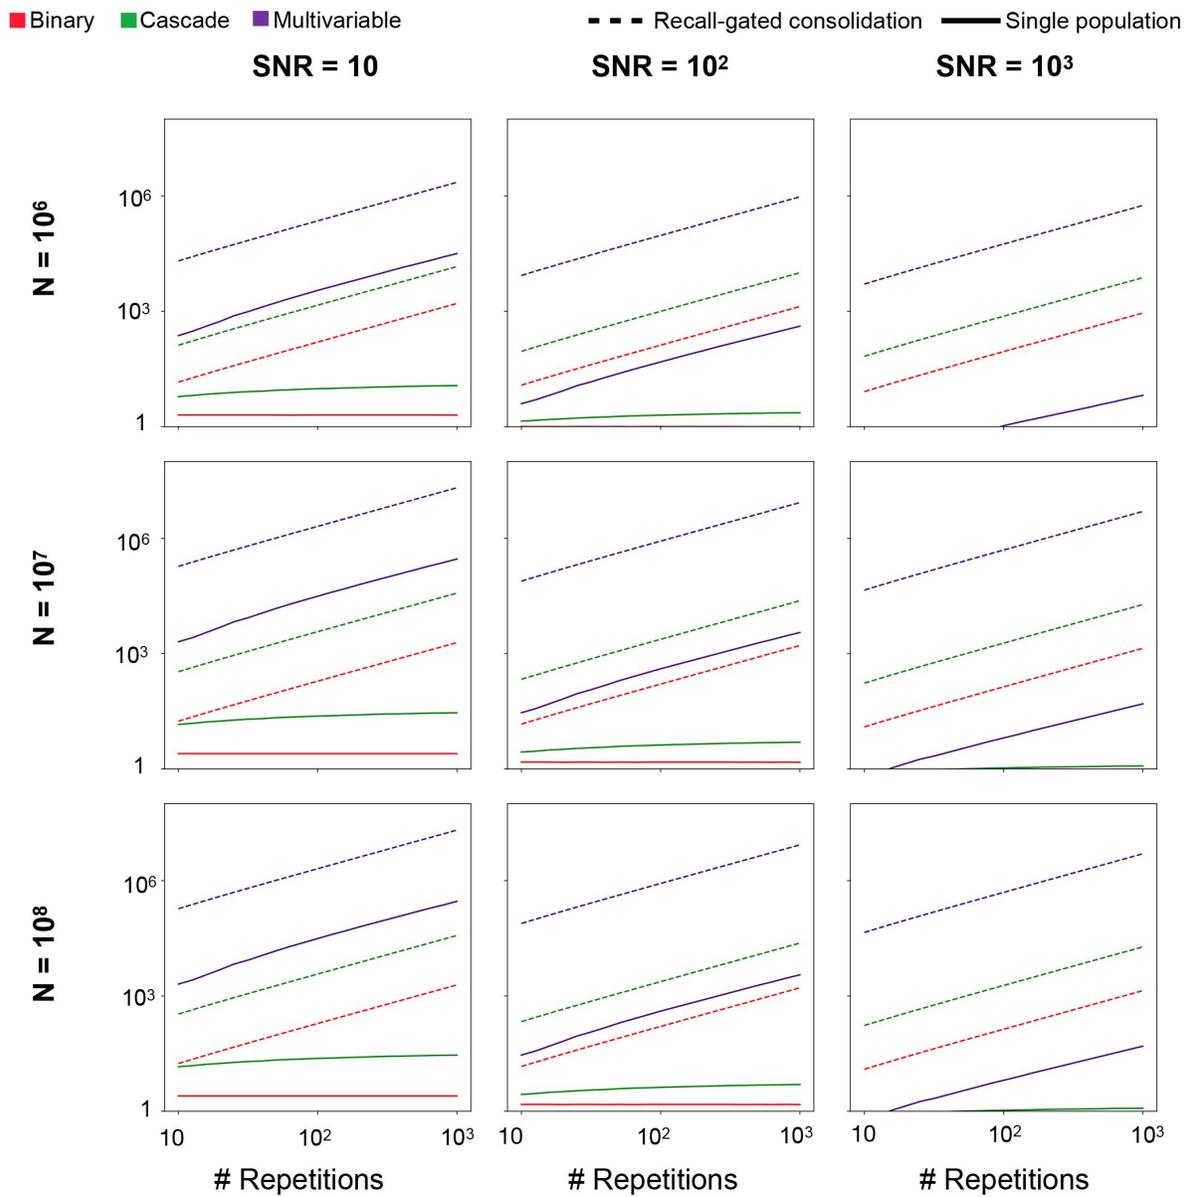
### Supplemental Figure S2

A. Same as Fig. 4C, also varying the underlying synaptic learning rule. B. Same learnable timescale information as panel A, presented as a function of the synaptic population size  $N$ .



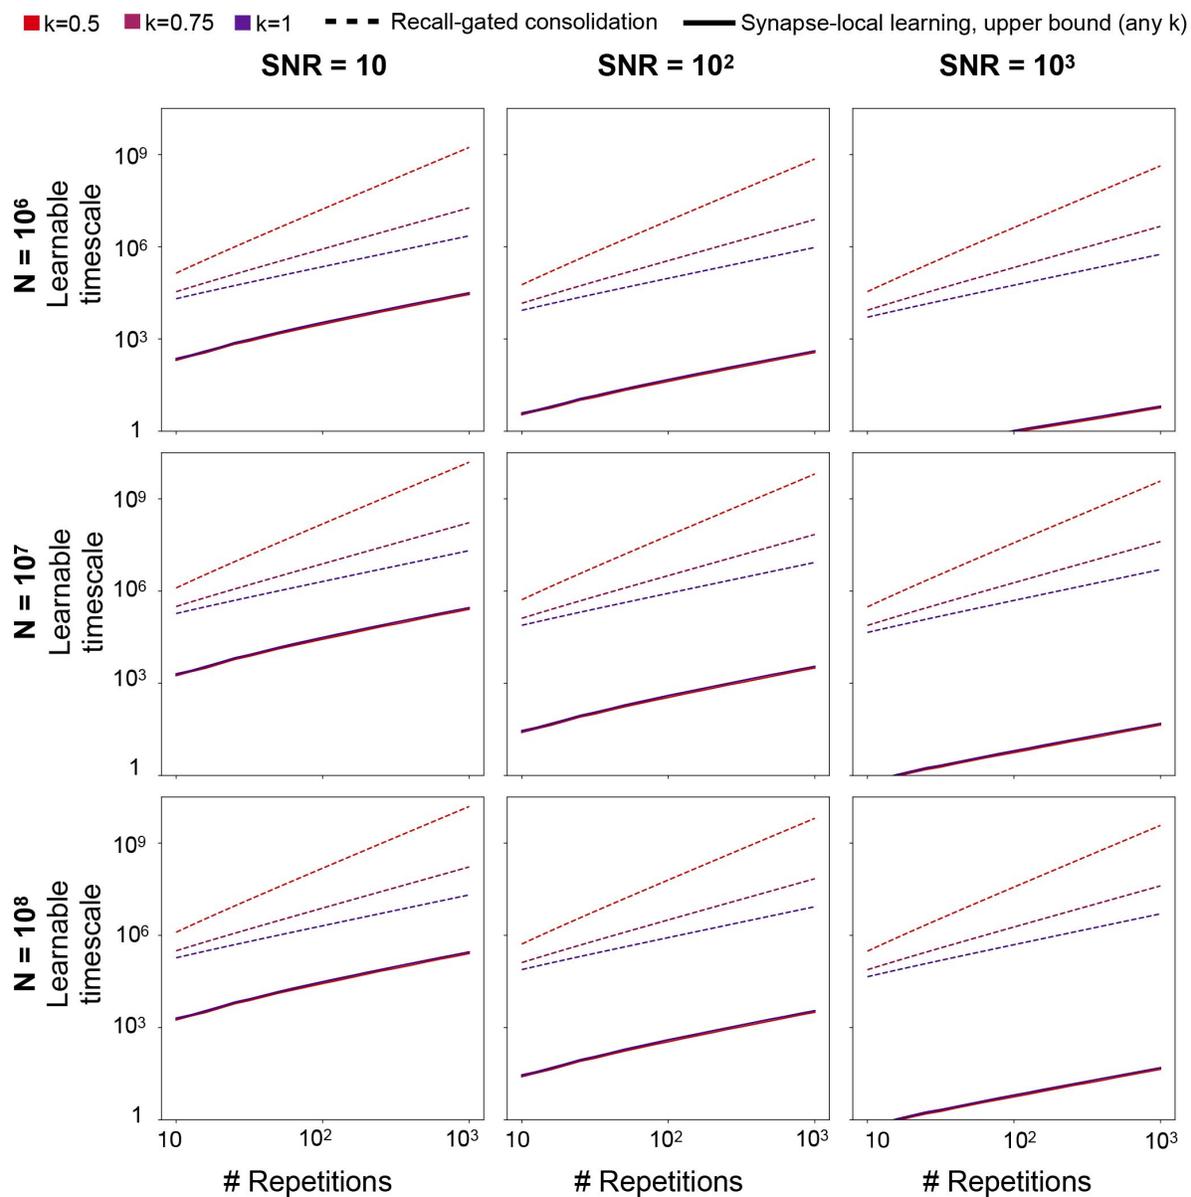
### Supplemental Figure S3

Same as Fig. 2C, but with multiple reliable memories simultaneously learned, each recurring equally often at a rate  $\lambda_i = 0.01$  for all reliable memories  $i$ . Here  $N = 10^5$  synapses, and the STM and LTM learning rates are 0.05 and 0.01, respectively. In the recall-gated model, the gating threshold is set at  $\theta = 2^{-7}$ . Each plot corresponds to a different number  $P$  of reliable memories being stored (the SNR shown is averaged across the reliable memories). The behavior of the SNR of an individual reliable memory is approximately the same as in the single-memory case for small values of  $\lambda_{tot}$  but diverges from it when  $\lambda_{tot}$  grows large.



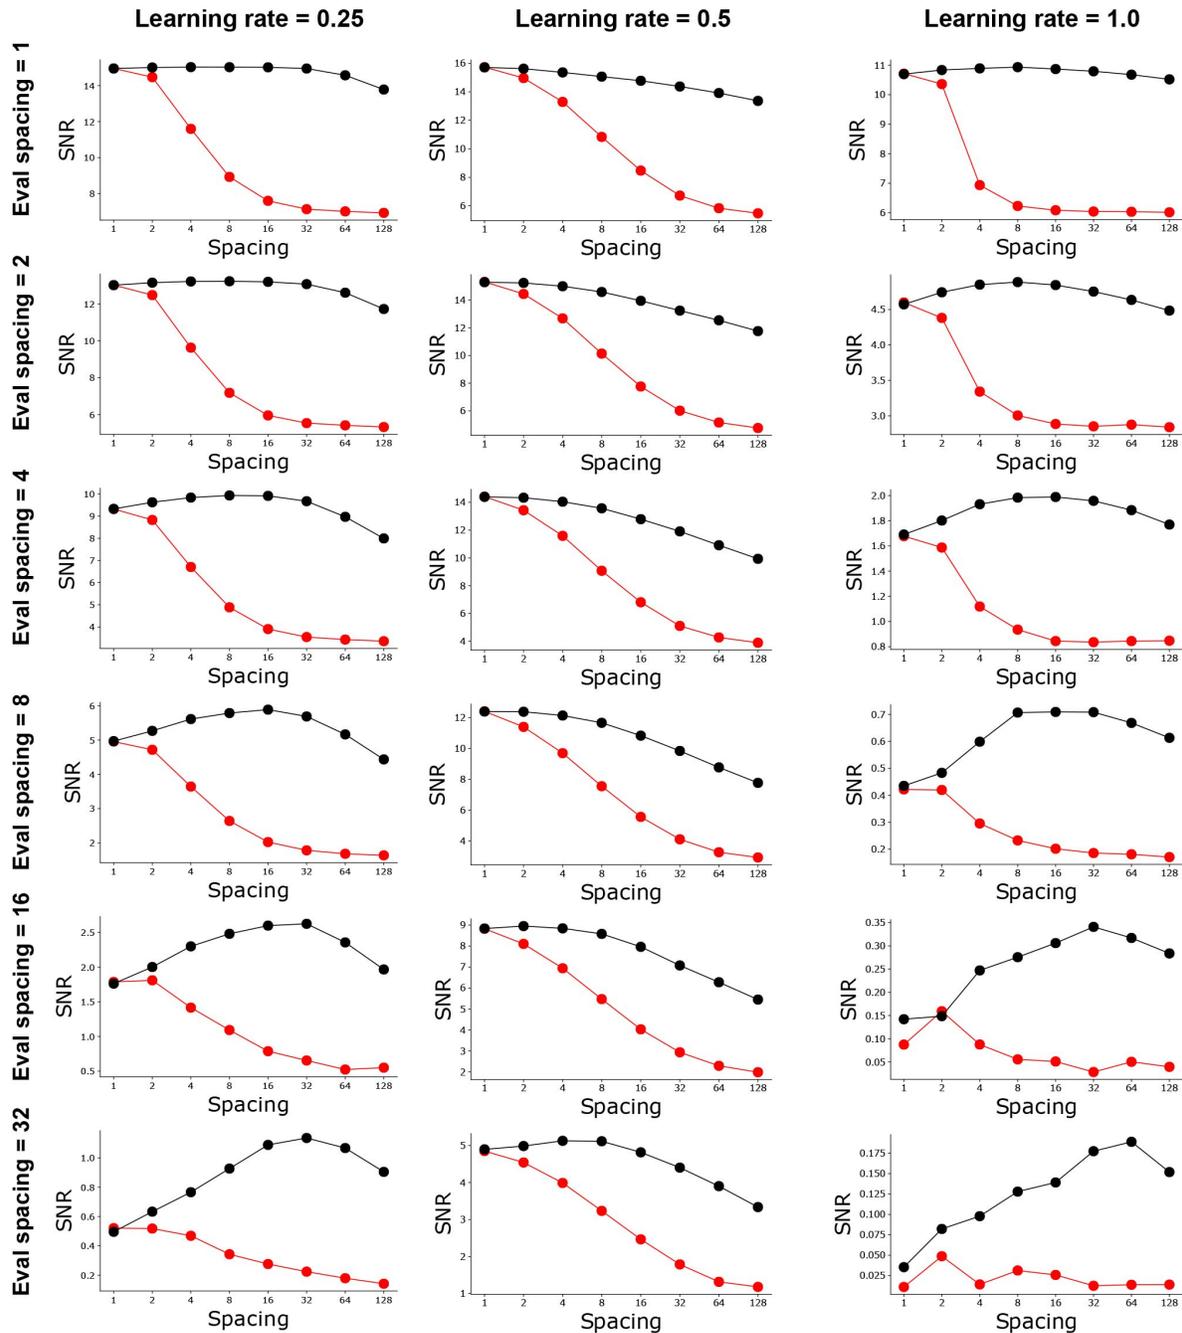
### Supplemental Figure S4

Same information as [Fig. 4G](#), varying the population size  $N$  and the desired  $SNR$ .



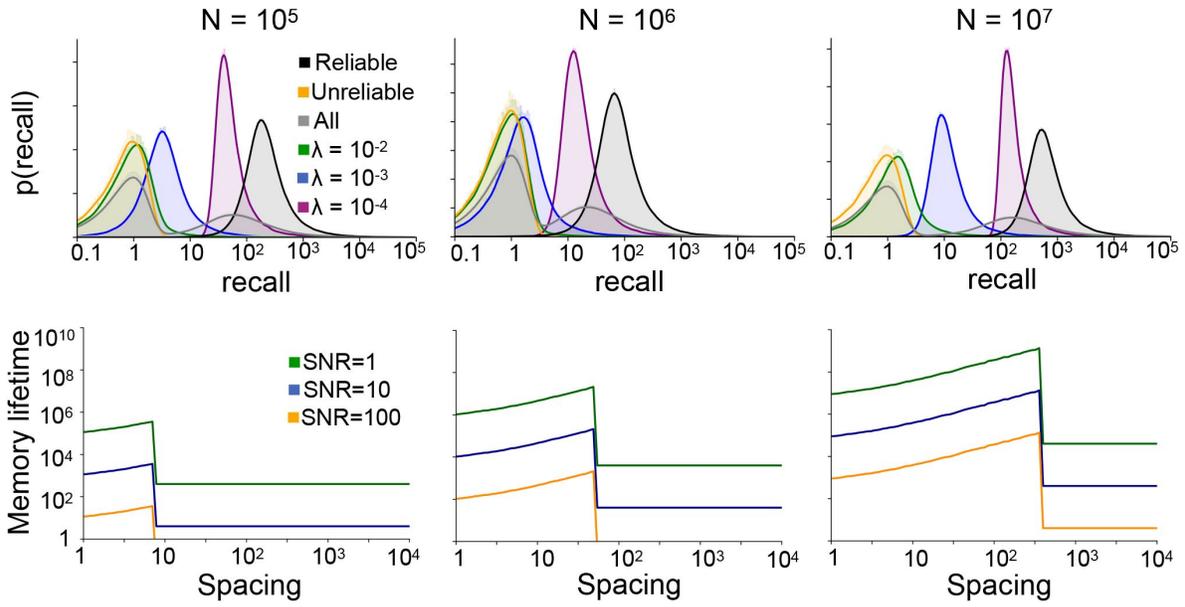
### Supplemental Figure S5

Same information as Fig. 4H, varying the population size  $N$  and the desired  $SNR$



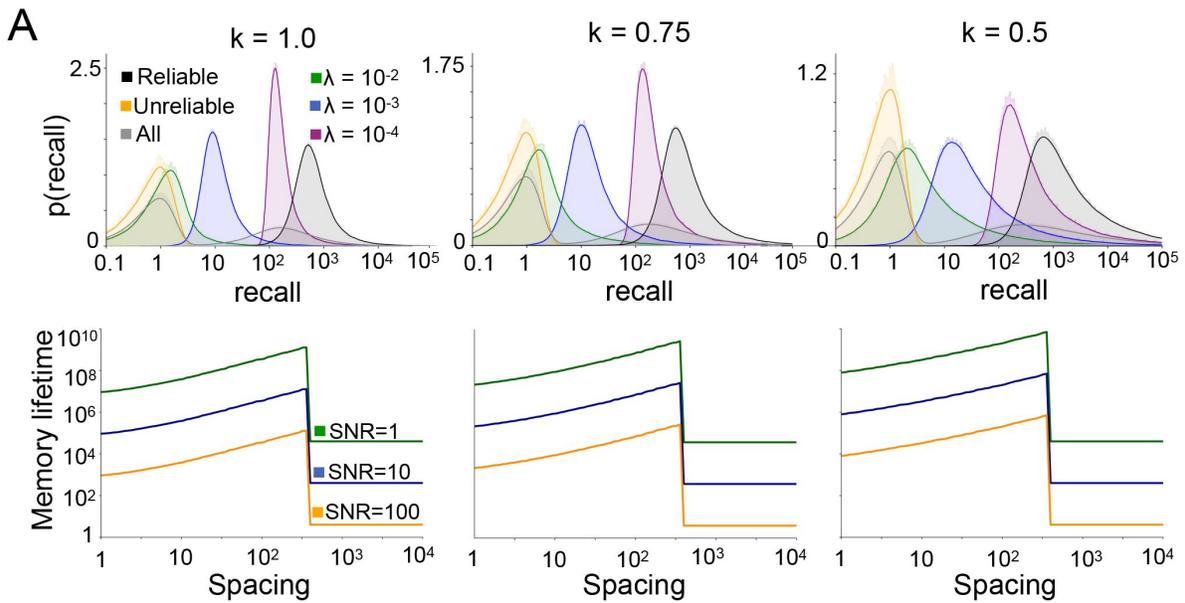
**Supplemental Figure S6**

Same information as [Fig. 5C](#), varying the learning rate (scale of potentiation/depression impulses, relative to the maximum/minimum threshold values in the model of [Benna and Fusi \(2016\)](#)), and the length of time following spaced training at which the system's recall SNR is evaluated.



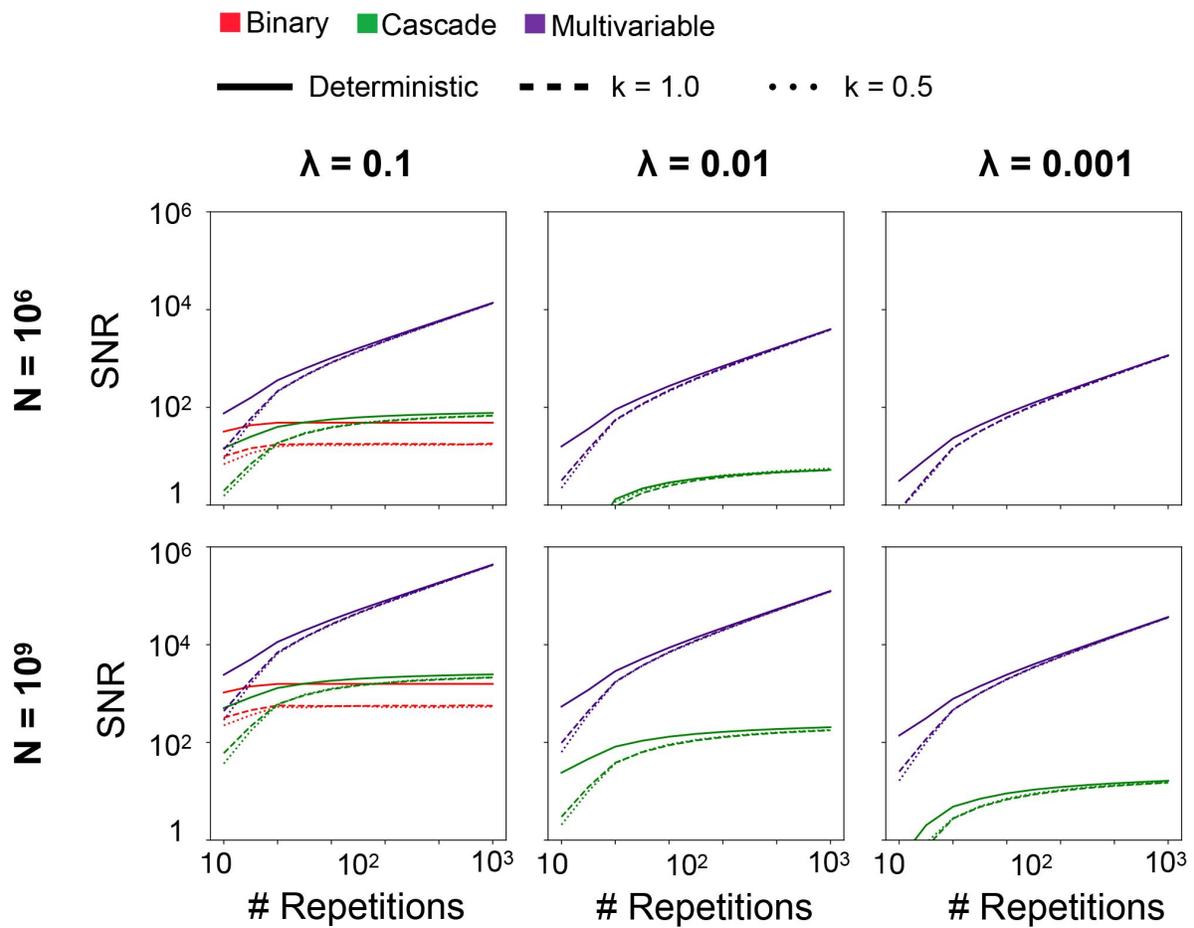
### Supplemental Figure S7

Same as Fig. 5D (top row) and Fig. 5F (bottom row), for different population sizes  $N$ .



### Supplemental Figure S8

Same as Fig. 5D (top row) and Fig. 5F (bottom row), for different memory recurrence regularity factors (Weibull distribution parameter  $k$ ).



### Supplemental Figure S9

SNR as a function of repetitions for single populations without consolidation, varying the parameter  $k$  of the Weibull distribution governing interarrival times (and defining the learnable timescale in terms of the expected interarrival time). The behavior of the systems scales similarly for diverse values of  $k \rightarrow \infty$ , justifying the use of the deterministic approximation  $k \rightarrow \infty$  for theoretical calculations. The learning rate for the binary model here is set at 0.1.

## Supplementary Information

### Derivation of recall strength quantity for specific plasticity rules

In the following derivations, we make heavy use of the fact that the elementwise dot product  $\mathbf{W}_1 \cdot \mathbf{W}_2$  between two matrices is equal to  $\text{Tr}[\mathbf{W}_1^T \mathbf{W}_2]$ .

### Supervised Hebbian learning

Let  $\mathbf{x}$  be the input population activity,  $\mathbf{W}$  be the prediction weights,  $\hat{\mathbf{y}} = \mathbf{W}\mathbf{x}$  the output population activity (predicted probabilities), and  $\mathbf{y}$  indicate ground-truth target values. For a plasticity rule  $\Delta\mathbf{W} = \mathbf{y}\mathbf{x}^T$ , the recall strength  $r = \mathbf{W} \cdot \Delta\mathbf{W}$  can be written as

$$\begin{aligned}
 \mathbf{W} \cdot \Delta\mathbf{W} &= \mathbf{W} \cdot (\mathbf{y}\mathbf{x}^T) \\
 &= \text{Tr}[(\mathbf{y}\mathbf{x}^T)^T \mathbf{W}] \\
 &= \text{Tr}[\mathbf{x}\mathbf{y}^T \mathbf{W}] \\
 &= \mathbf{x} \cdot (\mathbf{y}^T \mathbf{W})^T \\
 &= \mathbf{x}^T \mathbf{W}^T \mathbf{y} \\
 &= (\mathbf{W}\mathbf{x})^T \mathbf{y} \\
 &= \hat{\mathbf{y}} \cdot \mathbf{y},
 \end{aligned} \tag{20}$$

corresponding to the accuracy of the prediction  $\hat{\mathbf{y}}$ .

If instead plasticity is driven by the delta rule,  $\Delta W_{ij} = (y_i - \hat{y}_i)x_j$ , the recall factor becomes  $(\mathbf{y} \cdot \hat{\mathbf{y}}) - \|\hat{\mathbf{y}}\|^2$ , which assuming normalized activity is simply an offset measure of prediction accuracy. In either case, the computation of the recall factor requires an explicit comparison of predictions to ground truth labels.

### Reinforcement learning

Let  $\mathbf{x}$  be the input population activity representing state information,  $\mathbf{W}$  be the output weights,  $\boldsymbol{\pi} = \mathbf{W}\mathbf{x}$  be the output population activity representing log probabilities of taking a given action,  $\mathbf{a}$  be a vector indicating the sampled action, and  $r = 1$  be the scalar reward that results. For a plasticity rule  $\Delta\mathbf{W} = r \cdot (\mathbf{a}\mathbf{x}^T)$ , the recall strength  $r = \mathbf{W} \cdot \Delta\mathbf{W}$  can be written as

$$\begin{aligned}
 \mathbf{W} \cdot \Delta\mathbf{W} &= \mathbf{W} \cdot (r \cdot \mathbf{a}\mathbf{x}^T) \\
 &= r \cdot (\mathbf{a}\mathbf{x}^T)^T \mathbf{W}.
 \end{aligned} \tag{21}$$

Following the same steps as the derivation for the supervised learning case, with  $\mathbf{a}$  in place of  $\mathbf{y}$ , gives

$$\begin{aligned}
 \mathbf{W} \cdot \Delta\mathbf{W} &= r \cdot ((\mathbf{W}\mathbf{x})^T \mathbf{a}) \\
 &= r \cdot (\boldsymbol{\pi}^T \mathbf{a}) \\
 &= r \cdot \log p_{\boldsymbol{\pi}}(\mathbf{a}),
 \end{aligned} \tag{22}$$

corresponding to the confidence with which the action  $\mathbf{a}$  was selected, modulated by reward.

Computing this factor requires preserving the network's action probability distribution, extracting from it the probability of the sampled action, and multiplicatively scaling the result by the obtained reward.

## Autoassociative memory

Let  $\mathbf{x}$  be the population activity and  $\mathbf{W}$  be the recurrent weight matrix. If  $\Delta\mathbf{W} = \mathbf{x}\mathbf{x}^T$ , and the weight matrix  $\mathbf{W}$  can be written as a sum  $\sum_i \mathbf{x}_i \mathbf{x}_i^T$  over prior plasticity-driven updates, then the recall strength  $r = \mathbf{W} \cdot \Delta\mathbf{W}$  can be written as

$$\begin{aligned}
 \mathbf{W} \cdot \Delta\mathbf{W} &= \mathbf{W} \cdot (\mathbf{x}\mathbf{x}^T) \\
 &= \left( \sum_i (\mathbf{x}_i \mathbf{x}_i^T) \right) \cdot (\mathbf{x}\mathbf{x}^T) \\
 &= \left( \sum_i (\mathbf{x}_i \mathbf{x}_i^T) \cdot (\mathbf{x}\mathbf{x}^T) \right) \\
 &= \sum_i \text{Tr} [(\mathbf{x}_i \mathbf{x}_i^T)^T (\mathbf{x}\mathbf{x}^T)] \\
 &= \sum_i \text{Tr} [\mathbf{x}_i \mathbf{x}_i^T \mathbf{x}\mathbf{x}^T] \\
 &= \sum_i \mathbf{x}_i^T \mathbf{x} \text{Tr} [\mathbf{x}_i \mathbf{x}_i^T] \\
 &= \sum_i (\mathbf{x}_i^T \mathbf{x})^2, \tag{23}
 \end{aligned}$$

corresponding the familiarity of the current pattern  $\mathbf{x}$  relative to all previously seen patterns  $\mathbf{x}_i$ .

Familiarity also be computed with a separate familiarity readout trained alongside the recurrent weights. If the familiarity readout employs a Hebbian rule, the resulting estimate of familiarity will be equal to

$$\sum_i (\mathbf{x}_i^T \mathbf{x}). \tag{24}$$

For uncorrelated patterns in a network below capacity, this strategy corresponds exactly to the true recall factor in the limit of large network size.

## Relationship between learnable timescale and capacity

We note that theoretical work on memory systems often focuses on memory *capacity*, the number of memories that can be reliably stored in the system (Gardner, 1988 [↗](#); Fusi et al., 2005 [↗](#); Benna and Fusi, 2016 [↗](#)). Our learnable timescale metric is distinct from capacity. However, the two are closely linked in a particular regime. Suppose  $P$  distinct reliable memories are reinforced independently at rates  $\lambda_1, \dots, \lambda_P$ . In the regime in which the overall rate of reliable memory presentation  $\lambda_{\text{tot}} = \sum_i \lambda_i$  is small, the SNR of memory recall for memory  $i$  will be the same as in the case of a single reliable memory with  $\lambda = \lambda_i$  (**Fig. S3** [↗](#)). Hence, for a fixed  $\lambda_{\text{tot}}$ , and for simplicity assuming that distinct reliable memories are presented at equal rates  $\lambda_i = \frac{1}{P} \lambda_{\text{tot}}$  for all  $i$ , the learnable timescale  $\tau^*$  of the system dictates its capacity, equal to  $\tau^* \lambda_{\text{tot}}$ . We note that this correspondence does not hold in the case where most observed memories are reliable. In this work, however, we are interested primarily in the regime of scarce reliability, where recall-gated consolidation provides the most benefit. In this regime, we regard the learnable timescale as the most natural measure of system performance, as the primary obstacle to memory storage is the presence of long gaps between reinforcements of reliable memories.

## The effect of repeated reinforcement on memory dynamics without recall-gated consolidation

When memories can recur multiple times, the memory trace function  $m(t)$  is no longer an adequate description of system behavior, as the synaptic updates from multiple presentations can combine. For the synaptic learning rules we consider here – the binary switch, the cascade model of Fusi et al., and the multivariable model of Benna & Fusi, this combination is approximately additive (Benna and Fusi, 2016). This is because for each of these learning rules, the change in distribution of synaptic states following the presentation of a memory is approximately independent of the existing synaptic state. The only dependencies are saturation effects – synapses which have reached the edge of their dynamic range – which can only lead to sub-additive behavior. Saturation effects can be avoided by making the dynamic range of synapses sufficiently large. Thus for these learning rules of interest we may consider additive memory trace combination to represent a close approximation (and a tight upper bound) on the combined memory trace strength.

For a reliable memory presented at times  $t_1, \dots, t_R$ , and a population of synapses using additive learning rules, the current SNR at time  $t$  can therefore be approximated as

$$\text{SNR}(t) = \sum_{i=1}^R m(t - t_i). \quad (25)$$

If memory presentations occur separated by regular intervals of length  $\tau = \frac{1}{\lambda}$ , we have

$$\text{SNR}(t) = \sum_{i=0}^{R-1} m(t - t_R + i\tau). \quad (26)$$

For the binary switch model,  $m(t)$  decays exponentially with time constant  $1/p$ , and so the second term is negligible compared to the first. Hence the learnable timescale of the system is the same as the memory lifetime, approximately  $1/p$ . For target SNR threshold  $\delta$ , we require  $p \geq \delta/\sqrt{N}$ , so the best possible learnable timescale, optimizing over  $p$ , is  $O(\sqrt{N}/\delta)$ .

For the cascade model,  $m(t)$  decays as  $\frac{\sqrt{N}}{t \log T} e^{-t/T}$ . For  $t \gg T$  the exponential factor dominates, resulting in the same behavior as the binary switch model. For  $t \ll T$ , the exponential term approximately vanishes, so the following expression for  $\text{SNR}(t)$  is a close approximation and tight upper bound:

$$\begin{aligned} \text{SNR}(t) &\approx \frac{\sqrt{N}}{(t - t_R) \log T} + \int_1^{\min(R, T/\tau)} \frac{\sqrt{N}}{\tau \cdot t \log T} dt \\ &= \frac{\sqrt{N}}{(t - t_R) \log T} + \frac{\sqrt{N}}{\tau \log T} \log[\min(T/\tau, R)] \\ &= \frac{\sqrt{N}}{(t - t_R)(\log R + \log \tau)} + \frac{\sqrt{N}}{\tau(\log R + \log \tau)} \log R. \end{aligned} \quad (27)$$

Again, for computing learnable timescale we are interested in when  $t - t_R \approx \tau$ , in which case:

$$\text{SNR}(t) \approx \frac{\sqrt{N}}{\tau} \frac{\log R}{\log R + \log \tau}. \quad (28)$$

For the multivariable model,  $m(t)$  decays as  $\sqrt{\frac{N}{t \log T}} e^{-t/T}$ . Again we are primarily interested in the  $t \ll T$  regime, in which the expression for  $\text{SNR}(t)$  is approximately

$$\begin{aligned} \text{SNR}(t) &\approx \frac{\sqrt{N}}{\sqrt{t - t_R} \sqrt{\log T}} + \int_1^{\min(R, T/\tau)} \frac{\sqrt{N}}{\sqrt{\tau \cdot t} \sqrt{\log T}} dt \\ &= \frac{\sqrt{N}}{\sqrt{t - t_R} \sqrt{\log T}} + \frac{2\sqrt{N}}{\sqrt{\tau} \sqrt{\log T}} (\sqrt{\min(T/\tau, R)} - 1). \end{aligned} \quad (29)$$

This SNR is maximized for  $T \approx R \cdot \tau$ . And for computing learnable timescale we are interested in when  $t - t_R \approx \tau$ . So we have

$$\text{SNR}(t) = \frac{\sqrt{N}}{\sqrt{\tau \log(R \cdot \tau)}} (2\sqrt{R} - 1). \quad (30)$$

To compute the learnable timescale at target SNR  $\delta$  for  $1 \ll R \ll \tau$ , we have  $4NR \geq \tau \log(R\tau)\delta^2$ , the solution of which is within logarithmic factors of  $O(RN)$ .

The above calculations assume deterministic interarrival intervals of length  $\tau$ . In general, we are interested in an interarrival distribution  $p(I)$  with mean  $\tau$ . However, we show numerically that for Weibull distributions with reasonable values of  $k$  (not too close to zero), the true learnable timescale figures are very bounded very closely to our results above (**Fig. S9**). Moreover, for the purpose of computing learnable timescale  $\tau_\epsilon^\beta$  with error probability tolerance  $\epsilon$ , for sufficiently small  $\epsilon$  the deterministic approximation represents an upper bound on the SNR for distributions with mean  $\tau$ . This is because to ensure high SNR with very high probability, deterministic intervals are a best case scenario, as stochastic interval lengths will with some nonzero probability deviate far above the mean.

## Bounds on an ideal synaptic consolidation model

In this section we show that no realistic synapse-local mechanism can achieve significantly better learnable timescale than ( $RN$ ), and hence that the ability of recall-gated systems consolidation to achieve learnable timescale scaling superlinearly with  $R$  in some environments (see previous section) represents a qualitative advantage.

We consider a very general class of synaptic learning rules. In particular, we suppose a synapse can maintain a history of sequences of potentiation and depression events for arbitrarily long time windows and track the number of windows for which  $\Delta$ , the difference in number of potentiation and depression events, exceeds a threshold  $\delta$ . Let  $p_{\text{reliable}}(\Delta; \tau)$  refer to the probability distribution of values of  $\Delta$  after  $\tau$  timesteps, given that a synapse is potentiated by the reliable memory of interest – and  $p_{\text{unreliable}}$  refers to the analogous distribution for synapses subject only to potentiation by unreliable memories. After  $\frac{T}{\tau}$  intervals of length  $\tau$ , for a synapse potentiated by the reliable memory, we have that

$$\mathbb{E} \left[ \log \left( \frac{p(\text{reliable}|\text{data})}{p(\text{unreliable}|\text{data})} \right) \right] = \frac{T}{\tau} D_{\text{KL}} (p_{\text{reliable}}(\Delta; \tau) || p_{\text{unreliable}}(\Delta; \tau)). \quad (31)$$

The memory can be considered retrievable with SNR of order  $\mathcal{O}(1)$  once the expression above exceeds  $(\frac{1}{N})$  (since evidence can be accumulated across the  $N$  synapses) for any choice of  $\tau$  (since we are interested in the best achievable performance).

Now, for large enough  $\tau$ ,  $p_{\text{unreliable}}(\Delta; \tau)$  is approximately Gaussian with mean 0 and standard deviation  $\sqrt{\tau}$ . Conditioned on the reliable memory being presented  $r$  times in  $\tau$  timesteps,  $p_{\text{reliable}}(\Delta)$  is approximately Gaussian with mean  $r$  and standard deviation  $\sqrt{\tau}$ . The KL divergence between these two distributions is  $\frac{r^2}{2\tau}$ . Now consider the distribution  $p(r)$  of number of repetitions  $r$  that occur in a time window  $\tau$ . We want to find a value  $r^{\text{max}}$  such that  $P_{\tau}(r \geq r^{\text{max}}) \leq \frac{\tau}{T}$ . From there we can assume that  $r < r^{\text{max}}$  in any of the  $\tau$ -length intervals, since after  $T$  timesteps we cannot reliably count on  $r$  exceeding  $r^{\text{max}}$  in any of the intervals.

For  $\tau \leq M [I]$ , the median of the distribution  $p(I)$ , note that  $r^{\text{max}} \leq \log(T/\tau) \leq \log T$ . For  $\tau \leq c \cdot M [I]$ , if  $R > c \log T$  then at least one interval of less than length  $M [I]$  contains at least  $\log T$  repetitions. Hence  $r^{\text{max}} \leq c \log T$ . So conservatively we can take  $r_{\text{max}} = \frac{\tau}{M[I]} \log T$ .

Thus, our log probability expression above is bounded as follows

$$\begin{aligned} \mathbb{E} \left[ \log \left( \frac{p(\text{reliable}|\text{data})}{p(\text{unreliable}|\text{data})} \right) \right] &= \frac{T}{\tau} D_{\text{KL}} (p_{\text{reliable}}(\Delta; \tau) || p_{\text{unreliable}}(\Delta; \tau)) \\ &\leq \frac{T}{\tau} \frac{r_{\text{max}}^2}{2\tau} \\ &= \frac{T}{\tau} \frac{\tau^2}{M[I]^2} \frac{\log^2(T)}{2\tau} \\ &= T \frac{\log^2(T)}{2M[I]^2}. \end{aligned}$$

Hence the KL divergence criterion becomes

$$T \frac{\log^2(T)}{2M[I]^2} \geq \mathcal{O} \left( \frac{1}{N} \right), \quad (32)$$

or equivalently,

$$TN \log^2 T \geq 2M[I]^2. \quad (33)$$

The number of repetitions is  $R \approx T/\mathbb{E}[I]$ , giving

$$RN \log^2(R\mathbb{E}[I]) \geq \frac{2M[I]^2}{\mathbb{E}[I]}, \quad (34)$$

$$RN \log^2(R\mathbb{E}[I]) \geq 2\mathbb{E}[I], \quad (35)$$

assuming  $M[I] \sim O(\mathbb{E}[I])$ . For the interarrival distributions we consider (of the Weibull family),  $M[I] < \mathbb{E}[I]$  so this is a conservative assumption. Hence the learnable timescale of any population using only synapse-local learning rules is no greater than the solution for  $\mathbb{E}[I]$  of the equation above. We have

$$RN \geq 2\mathbb{E}[I] / \log^2(R\mathbb{E}[I]), \quad (36)$$

the solution of which is within logarithmic factors of  $\mathcal{O}(RN)$ .

## Scaling behavior of the STM/LTM model

In the recall-gated consolidation model, the overlap  $r = \mathbf{w}_{\text{STM}} \cdot \Delta \mathbf{w}_{\text{STM}}$  indicates the recall strength of memory  $x$  given the current synaptic state of the STM. LTM plasticity is modulated by a factor  $g(r)$  – we refer to  $g$  as the “gating function” and  $r$  as the STM recall strength. We assume for now that the gating function  $g(r)$  is chosen to be a threshold function,  $g(r) = H(r - \theta)$ , where  $r$  is the SNR of the memory overlap,  $H$  is the Heaviside step function, and  $\theta$  is referred to as the “consolidation threshold.” With this choice, unreliable memories will be consolidated at a rate of  $1 - \Phi(\theta)$ , where  $\Phi$  is the CDF of the normal distribution, in the limit of large system size  $N$ .

Suppose a memory  $\Delta \mathbf{w}$  is presented twice with interval  $I$ . Then the SNR at the second presentation will be lower-bounded by  $m(I)$ , in expectation. It follows that the rate at which reliable memories will be consolidated at for the gating function above is lower bounded by  $P(I < m^{-1}(\theta))$ . After  $R$  repetitions of the reliable memory, the probability that consolidation has occurred will be at least

$$1 - (1 - P(I < m^{-1}(\theta)))^R. \quad (37)$$

We are interested in the maximum  $\theta$  for which this expression exceeds  $1 - \epsilon$  – this is the most stringent consolidation threshold we can set while still ensuring consolidation of the reliable memory with high probability. This value of  $\theta$  is given by

$$R \log(1 - P(I < m^{-1}(\theta))) = \log \epsilon \quad (38)$$

If  $R$  is large then the solution will be such that  $P(I < m^{-1}(\theta))$  is small, enabling the approximation:

$$P(I < m^{-1}(\theta)) = -\log \epsilon / R \quad (39)$$

For tractability we consider, as our family of interarrival distributions, Weibull distributions with regularity parameter  $k$ . The cumulative distribution function is

$$P(I \leq t) = 1 - e^{-(t\Gamma(1+1/k)/\tau)^k}. \quad (40)$$

For  $t \ll \tau$ , this is approximated as

$$P(I \leq t) \approx (t\Gamma(1 + 1/k)/\tau)^k. \quad (41)$$

Importantly,  $P(I \leq t)$  decays as  $t^k$ . Thus, increasing the number of repetitions  $R$  has the effect of scaling the  $\tau$  that satisfies **Equation 39** by  $R^{1/k}$ . That is, for a fixed  $\theta$ , and hence a fixed degree of amplification  $\tau_{\text{LTM}}/\tau$  of the effective rate of reliable memories in the LTM, the maximum  $\tau$  achieving that SNR with probability  $1 - \epsilon$  (i.e. the learnable timescale  $\tau_{\beta}^{\epsilon}$ ) scales as  $\mathcal{O}(R^{1/k})$ .

For a gating function threshold  $\theta$ , the corresponding SNR in the LTM will be the SNR induced by an interarrival distribution  $I_{\text{LTM}}$  with mean value

$$\mathbb{E}[I_{\text{LTM}}] = \frac{P(I < m^{-1}(\theta))}{1 - \Phi(\theta)} \mathbb{E}[I]. \quad (42)$$

Since  $1 - \Phi(\theta)$  decays much more rapidly than any power of  $m^{-1}(\theta)$ , it follows that  $\mathbb{E}[I_{\text{LTM}}]$  can be made  $\mathcal{O}(1)$ , and thus the SNR of the LTM can become  $\mathcal{O}(\sqrt{N})$ , for relatively small values of  $\theta$  (and hence a small number of required repetitions). In other words, for a fixed number of expected memory repetitions, the learnable timescale of the LTM decreases only slightly as the target SNR is raised from  $\mathcal{O}(1)$  to  $\mathcal{O}(\sqrt{N})$ .

Note that if  $P$  different reliable memories are present, then  $\mathbb{E}[I_{LTM}]$  for any given reliable memory will be lower-bounded by  $\mathcal{O}(P)$  instead of  $\mathcal{O}(1)$ . The induced SNR for any given reliable memory in the LTM will in this case be of order  $m(P)$ , rather than  $O(\sqrt{N})$ .

## Optimal sparsity calculations

We now consider the case of sparse memories – those which potentiate a fraction  $f$  of synapses. Consider the behavior of a single population of binary synapses employing the binary switch plasticity rule. We modify the plasticity rule slightly so that potentiation flips the state of a synapse with probability  $p$  and depression with probability  $\frac{f}{1-f}p$ , to ensure that the fractions of potentiated and depressed synapses remain balanced.

We consider an environment with a single reliable memory that is presented with probability  $\lambda$  at each time step (otherwise, a randomly sampled unreliable memory is presented). We can compute the behavior analytically by tracking how the distributions of  $u$  (the output neuron response to true stimuli) and  $v$  (the output neuron response to noise) evolve over time. We assume that the coding level  $f$  is sufficiently small that terms of order  $O(f^2)$  may be ignored.

Due to the balanced plasticity rule,  $\frac{1}{2}$  of synapses are strong at any given time, so the mean response  $v^*$  to a randomly sampled noise pattern is  $\frac{1}{2}$ . The variance of  $v$  is also constant and equal to  $\frac{1}{4Nf}$ .

The evolution of  $u$  is a stochastic process that, in the limit of large  $Nf$  (i.e. a large number of active neurons for each stimulus), can be described as an Ornstein-Uhlenbeck (OU) process:

$$\Delta u = \theta(u^* - u) + \epsilon. \quad (43)$$

where  $\epsilon \sim \mathcal{N}(0, \sigma^2)$

In the limit of small  $f$  we have:

$$\theta = p(2f + \lambda) \quad (44)$$

$$u^* = \frac{f + \lambda}{2f + \lambda} = \frac{1}{2} + \frac{\lambda}{4f + 2\lambda}. \quad (45)$$

The quantity  $u^*$  determines the asymptotic mean of  $u$  and the quantity  $\theta$  determines the rate at which  $u$  converges to this mean. Immediately we see that  $u^*$  scales with the frequency  $\lambda$  with which the true stimulus is presented, and that the rate of convergence (speed of learning) is proportional to  $p$ .

By well-known properties of OU processes, the asymptotic variance of  $u$  is equal to  $\frac{\sigma^2}{2\theta}$ . In the small- $p$  limit, this quantity comes out to

$$\text{Var}(u) = \frac{1}{4Nf}. \quad (46)$$

Note that in the low- $p$  limit (slow learning rate) this is the same as the variance of  $v$ . Thus in the limit of slow learning, we have that

$$\mathbb{E}[u - v] = \frac{\lambda}{4f + 2\lambda}, \quad (47)$$

$$\text{Var}(u - v) = \frac{1}{2Nf}. \quad (48)$$

And thus

$$\text{SNR} = \frac{N}{2} \frac{f\lambda^2}{(2f + \lambda)^2}. \quad (49)$$

From this expression we can see that for a given  $f$ , the asymptotic SNR always increases with  $\lambda$  and  $N$ . For a given  $\lambda$ , we would like to maximize this expression with respect to  $f$ .

$$\nabla_f \text{SNR} = \frac{N}{2} \frac{(2f + \lambda)^2 \lambda^2 - f\lambda^2 \cdot 4(2f + \lambda)}{(2f + \lambda)^4}. \quad (50)$$

This expression equals zero when

$$(2f + \lambda)^2 \lambda^2 - f\lambda^2 \cdot 4(2f + \lambda) = 0 \quad (51)$$

$$(2f + \lambda) = 4f \quad (52)$$

$$f = \frac{1}{2}\lambda, \quad (53)$$

so the asymptotic SNR is maximized for  $f = \frac{1}{2}\lambda$ . That is, the optimal coding level is proportional to the frequency with which reliable (as opposed to unreliable) stimuli are observed in the environment.

## References

- Amit D. J., Fusi S. (1994) **Learning in neural networks with material synapses** *Neural Computation* **6**:957–982
- Andalman A. S., Fee M. S. (2009) **A basal ganglia-forebrain circuit in the songbird biases motor output to avoid vocal errors** *Proceedings of the National Academy of Sciences* **106**:12518–12523
- Aso Y. *et al.* (2014) **The neuronal architecture of the mushroom body provides a logic for associative learning** *elife* **3**
- Awata H., Takakura M., Kimura Y., Iwata I., Masuda T., Hirano Y. (2019) **The neural circuit linking mushroom body parallel circuits induces memory consolidation in drosophila** *Proceedings of the National Academy of Sciences* **116**:16080–16085
- Babadi B., Sompolinsky H. (2014) **Sparseness and expansion in sensory representations** *Neuron* **83**:1213–1226
- Beck C., Schroeder B., Davis R. L. (2000) **Learning performance of normal and mutant drosophila after repeated conditioning trials with discrete stimuli** *Journal of Neuroscience* **20**:2944–2953
- Benna M. K., Fusi S. (2016) **Computational principles of synaptic memory consolidation** *Nature neuroscience* **19**:1697–1706
- Brodts S., Pöhlchen D., Flanagan V. L., Glasauer S., Gais S., Schönauer M. (2016) **Rapid and independent memory formation in the parietal cortex** *Proceedings of the National Academy of Sciences* **113**:13251–13256
- Carr M. F., Jadhav S. P., Frank L. M. (2011) **Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval** *Nature neuroscience* **14**:147–153
- Cepeda N. J., Pashler H., Vul E., Wixted J. T., Rohrer D. (2006) **Distributed practice in verbal recall tasks: A review and quantitative synthesis** *Psychological bulletin* **132**
- Cepeda N. J., Vul E., Rohrer D., Wixted J. T., Pashler H. (2008) **Spacing effects in learning: A temporal ridge line of optimal retention** *Psychological science* **19**:1095–1102
- Cervantes-Sandoval I., Martin-Peña A., Berry J. A., Davis R. L. (2013) **System-like consolidation of olfactory memories in drosophila** *Journal of Neuroscience* **33**:9846–9854
- Davis R. L. (2011) **Traces of drosophila memory** *Neuron* **70**:8–19
- Dhawale A. K., Wolff S. B., Ko R., O’Ölveczky B. P. (2019) **The basal ganglia can control learned motor sequences independently of motor cortex** *BioRxiv*
- Dhawale A. K., Wolff S. B., Ko R., O’Ölveczky B. P. (2021) **The basal ganglia control the detailed kinematics of learned motor skills** *Nature Neuroscience* **24**:1256–1269

Dubnau J., Chiang A.-S. (2013) **Systems memory consolidation in drosophila** *Current opinion in neurobiology* **23**:84–91

Frankland P. W., Bontempi B. (2005) **The organization of recent and remote memories** *Nature reviews neuroscience* **6**:119–130

Fusi S., Drew P. J., Abbott L. F. (2005) **Cascade models of synaptically stored memories** *Neuron* **45**:599–611

Gardner E. (1988) **The space of interactions in neural network models** *Journal of physics A: Mathematical and general* **21**

Glas A., Hübener M., Bonhoeffer T., Goltstein P. M. (2021) **Spaced training enhances memory and prefrontal ensemble stability in mice** *Current Biology* **31**:4052–4061

Gluck M. A., Myers C. E. (1993) **Hippocampal mediation of stimulus representation: A computational theory** *Hippocampus* **3**:491–516

González O. C., Sokolov Y., Krishnan G. P., Delanois J. E., Bazhenov M. (2020) **Can sleep protect memories from catastrophic forgetting?** *Elife* **9**

Gorriz M. H., Takigawa M., Bendor D. (2023) **The role of experience in prioritizing hippocampal replay** *bioRxiv* :2023–3

Gou J., Yu B., Maybank S. J., Tao D. (2021) **Knowledge distillation: A survey** *International Journal of Computer Vision* **129**:1789–1819

Hinton G., Vinyals O., Dean J. (2015) **Distilling the knowledge in a neural network**

Hwang E. J., Dahlen J. E., Hu Y. Y., Aguilar K., Yu B., Mukundan M., Mitani A., Komiyama T. (2019) **Disengagement of motor cortex from movement control during long-term learning** *Science advances* **5**

Hwang E. J., Dahlen J. E., Mukundan M., Komiyama T. (2021) **Disengagement of motor cortex during long-term learning tracks the performance level of learned movements** *Journal of Neuroscience* **41**:7029–7047

Jackson J. C., Johnson A., Redish A. D. (2006) **Hippocampal sharp waves and reactivation during awake states depend on repeated sequential experience** *Journal of Neuroscience* **26**:12415–12426

Kandel E. R., Dudai Y., Mayford M. R. (2014) **The molecular and systems biology of memory** *Cell* **157**:163–186

Kawai R., Markman T., Poddar R., Ko R., Fantana A. L., Dhawale A. K., Kampff A. R., O'Ólveczky B. P. (2015) **Motor cortex is required for learning but not for executing a motor skill** *Neuron* **86**:800–812

Lahiri S., Ganguli S. (2013) **A memory frontier for complex synapses** *Advances in neural information processing systems* **26**:1034–1042

Li F. *et al.* (2020) **The connectome of the adult drosophila mushroom body provides insights into function** *Elife* **9**

- McClelland J. L., Goddard N. H. (1996) **Considerations arising from a complementary learning systems perspective on hippocampus and neocortex** *Hippocampus* **6**:654–665
- 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**
- Murray J. M., Escola G. S. (2020) **Remembrance of things practiced with fast and slow learning in cortical and subcortical pathways** *Nature Communications* **11**:1–12
- Murray J. M., et al. (2017) **Learning multiple variable-speed sequences in striatum via cortical tutoring** *Elife* **6**
- Ólafsdóttir H. F., Bush D., Barry C. (2018) **The role of hippocampal replay in memory and planning** *Current Biology* **28**:R37–R50
- Remme M. W., Bergmann U., Alevi D., Schreiber S., Sprekeler H., Kempster R. (2021) **Hebbian plasticity in parallel synaptic pathways: A circuit mechanism for systems memory consolidation** *PLOS Computational Biology* **17**
- Rosenzweig M. R., Bennett E. L., Colombo P. J., Lee D. W., Serrano P. A. (1993) **Short-term, intermediate-term, and long-term memories** *Behavioural brain research* **57**:193–198
- Rovee-Collier C., Evancio S., Earley L. A. (1995) **The time window hypothesis: Spacing effects** *Infant Behavior and Development* **18**:69–78
- Roxin A., Fusi S. (2013) **Efficient partitioning of memory systems and its importance for memory consolidation** *PLoS computational biology* **9**
- Shaham N., Chandra J., Kreiman G., Sompolinsky H. (2021) **Stochastic consolidation of lifelong memory** *bioRxiv*
- Squire L. R., Alvarez P. (1995) **Retrograde amnesia and memory consolidation: a neurobiological perspective** *Current opinion in neurobiology* **5**:169–177
- Sun W., Advani M., Spruston N., Saxe A., Fitzgerald J. E. (2021) **Organizing memories for generalization in complementary learning systems** *BioRxiv*
- Tachibana R. O., Lee D., Kai K., Kojima S. (2022) **Performance-dependent consolidation of learned vocal changes in adult songbirds** *The Journal of Neuroscience: the Official Journal of the Society for Neuroscience*
- Terada S., Geiller T., Liao Z., O'Hare J., Vancura B., Losonczy A. (2021) **Adaptive stimulus selection for consolidation in the hippocampus** *Nature* :1–5
- Teşileanu T., Ölveczky B., Balasubramanian V. (2017) **Rules and mechanisms for efficient two-stage learning in neural circuits** *Elife* **6**
- van de Ven G. M., Siegelmann H. T., Tolias A. S. (2020) **Brain-inspired replay for continual learning with artificial neural networks** *Nature communications* **11**:1–14
- Verkoeijen P. P., Rikers R. M., Schmidt H. G. (2005) **Limitations to the spacing effect: Demonstration of an inverted u-shaped relationship between interrepetition spacing and free recall** *Experimental Psychology* **52**:257–263

## Article and author information

### Jack Lindsey

Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY 10027, USA

**For correspondence:** jackwlindsey@gmail.com

ORCID ID: [0000-0003-0930-7327](https://orcid.org/0000-0003-0930-7327)

### Ashok Litwin-Kumar

Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY 10027, USA

**For correspondence:** a.litwin-kumar@columbia.edu

ORCID ID: [0000-0003-2422-6576](https://orcid.org/0000-0003-2422-6576)

## Copyright

© 2023, Jack Lindsey & Ashok Litwin-Kumar

This article is distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use and redistribution provided that the original author and source are credited.

## Editors

Reviewing Editor

### Julijana Gjorgjieva

Technical University of Munich, Germany

Senior Editor

### Michael Frank

Brown University, United States of America

### Reviewer #1 (Public Review):

Summary:

The authors develop a memory consolidation theory utilizing the recall quality in the short-term memory system to decide what to consolidate in the long-term memory (LTM). The theory is based on a set of previously proposed models identifying memories and synaptic weights (without neuronal activity) with an addition of the second set of weights responsible for long-term storage. The rigorous analysis and numerical experiments show that under some assumptions, the long-term system achieves a high signal-to-noise ratio, particularly much higher than concurrently learning or localized in the same synapses LTM.

Strengths:

The authors take on an important problem of designing robust memory consolidation that fits the numerous experimental observations and, to a large extent, they succeed. The proposed solution is general and generalized to multiple contexts. The mathematical treatment is solid and convincing.

#### Weaknesses:

The presented model seems to be tuned for learning repetitive events. However, single-shot learning, for example, under fear conditioning or if a presented stimulus is astonishing, seems to contradict the proposed framework. I would assume that part of the load could be taken by a reply system that could vigorously replay more surprising events, but it seems to still not exactly match the proposed scheme.

For context, I would like to see the comparison/discussion of the wide range of models on synaptic tagging for consolidation by various types of signals. Notably, studies from Wulfram Gerstner's group (e.g., Brea, J., Clayton, N. S., & Gerstner, W. (2023). Computational models of episodic-like memory in food-caching birds. *Nature Communications*, 14(1); and studies on surprise).

The models that are taken for comparison with the slow but otherwise identical to STM LTM could be incapable per design. Reducing the probability of switching independently of the previous presentation does not make the system "slow"; instead, it should integrate previous signals (and thus slowly remove independent noise).

The usage of terms and streamlining of writing could be improved for better understanding.

#### **Reviewer #2 (Public Review):**

##### Summary:

In the manuscript "Recall-Gated Consolidation: A Model for Learning and Memory in Neural Systems," the authors suggest a computational mechanism called recall-gated consolidation, which prioritizes the storage of previously experienced synaptic updates in memory. The authors investigate the mechanism with different types of learning problems including supervised learning, reinforcement learning, and unsupervised auto-associative memory. They rigorously analyse the general mechanism and provide valuable insights into its benefits.

##### Strengths:

The authors establish a general theoretical framework, which they translate into three concrete learning problems. For each, they define an individual mathematical formulation. Finally, they extensively analyse the suggested mechanism in terms of memory recall, consolidation dynamics, and learnable timescales.

The presented model of recall-gated consolidation covers various aspects of synaptic plasticity, memory recall, and the influence of gating functions on memory storage and retrieval. The model's predictions align with observed spaced learning effects.

The authors conduct simulations to validate the recall-gated consolidation model's predictions, and their simulated results align with theoretical predictions. These simulations demonstrate the model's advantages over consolidating any memory and showcase its potential application to various learning tasks.

The suggestion of a novel consolidation mechanism provides a good starting point to investigate memory consolidation in diverse neural systems and may inspire artificial learning algorithms.

##### Weaknesses:

I appreciate that the authors devoted a specific section to the model's predictions, and point out how the model connects to experimental findings in various model organisms. However, the connection is rather weak and the model needs to make more specific predictions to be

distinguishable from other theories of memory consolidation (e.g. those that the authors discuss) and verifiable by experimental data.

While the article extensively discusses the strengths and advantages of the recall-gated consolidation model, it provides a limited discussion of potential limitations or shortcomings of the model, such as the missing feature of generalization, which is part of previous consolidation models. The model is not compared to other consolidation models in terms of performance and how much it increases the signal-to-noise ratio. It is only compared to a simple STM or a parallel LTM, which I understand to be essentially the same as the STM but with a different timescale (so not really an alternative consolidation model). It would be nice to compare the model to an actual or more sophisticated existing consolidation model to allow for a fairer comparison.

The article is lengthy and dense and it could be clearer. Some sections are highly technical and may be challenging to follow. It could benefit from more concise summaries and visual aids to help convey key points.

**Reviewer #3 (Public Review):**

Summary:

In their article "Theory of systems memory consolidation via recall-gated plasticity", Jack Lindsey and Ashok Litwin-Kumar describe a new model for systems memory consolidation. Their idea is that a short-term memory acts not as a teacher for a long-term memory - as is common in most complementary learning systems - but as a selection module that determines which memories are eligible for long-term storage. The criterion for the consolidation of a given memory is a sufficient strength of recall in the short-term memory.

The authors provide an in-depth analysis of the suggested mechanism. They demonstrate that it allows substantially higher SNRs than previous synaptic consolidation models, provide an extensive mathematical treatment of the suggested mechanism, show that the required recall strength can be computed in a biologically plausible way for three different learning paradigms, and illustrate how the mechanism can explain spaced training effects.

Strengths:

The suggested consolidation mechanism is novel and provides a very interesting alternative to the classical view of complementary learning systems. The analysis is thorough and convincing.

Weaknesses:

The main weakness of the paper is the equation of recall strength with the synaptic changes brought about by the presentation of a stimulus. In most models of learning, synaptic changes are driven by an error signal and hence cease once the task has been learned. The suggested consolidation mechanism would stop at that point, although recall is still fine. The authors should discuss other notions of recall strength that would allow memory consolidation to continue after the initial learning phase. Aside from that, I have only a few technical comments that I'm sure the authors can address with a reasonable amount of work.