

UC San Diego

UC San Diego Previously Published Works

Title

Selective Memory Generalization by Spatial Patterning of Protein Synthesis

Permalink

<https://escholarship.org/uc/item/32d9r7vw>

Journal

Neuron, 82(2)

ISSN

0896-6273

Authors

O'Donnell, Cian
Sejnowski, Terrence J

Publication Date

2014-04-01

DOI

10.1016/j.neuron.2014.02.028

Peer reviewed



Published in final edited form as:

Neuron. 2014 April 16; 82(2): 398–412. doi:10.1016/j.neuron.2014.02.028.

Selective memory generalization by spatial patterning of protein synthesis

Cian O'Donnell¹ and Terrence J. Sejnowski^{1,2}

¹Howard Hughes Medical Institute, Salk Institute for Biological Studies, La Jolla, CA 92037, USA.

²Division of Biological Sciences, University of California at San Diego, La Jolla, CA 92093, USA.

Summary

Protein synthesis is crucial for both persistent synaptic plasticity and long-term memory. De novo protein expression can be restricted to specific neurons within a population, and to specific dendrites within a single neuron. Despite its ubiquity, the functional benefits of spatial protein regulation for learning are unknown. We used computational modeling to study this problem. We found that spatially patterned protein synthesis can enable selective consolidation of some memories but forgetting of others, even for simultaneous events that are represented by the same neural population. Key factors regulating selectivity include the functional clustering of synapses on dendrites, and the sparsity and overlap of neural activity patterns at the circuit level. Based on these findings we proposed a novel two-step model for selective memory generalization during REM and slow-wave sleep. The pattern-matching framework we propose may be broadly applicable to spatial protein signaling throughout cortex and hippocampus.

Introduction

The persistence of new memories beyond a few hours requires the synthesis of new proteins at the time of learning (Davis and Squire, 1984). At the cellular level, consolidation of long-term synaptic plasticity also requires de novo protein synthesis (Kelleher et al., 2004; Krug et al., 1984). Although the molecular identities of these plasticity-related proteins (PRPs) remain unclear, their expression is tightly regulated in both time and space. In the temporal domain, a wave of protein synthesis occurs rapidly (within minutes) following the induction of synaptic plasticity, and returns to baseline less than one hour later (Kelleher et al., 2004; Otani et al., 1989). In the spatial domain, PRP expression is restricted at two distinct levels of granularity: the neural level and the dendritic level. At the neural level, protein expression following synaptic plasticity induction is specific to single cells within a given population (Mackler et al., 1992), and PRPs are presumably not shared between neurons. At the dendritic level, substantial evidence indicates that synaptic activity can drive PRP synthesis

© 2014 Elsevier Inc. All rights reserved.

Correspondence cian@salk.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

in the dendrites local to the activated synapses (Figure 1A) (Sutton and Schuman, 2006). Once synthesized, the PRPs can remain localized within or near the particular dendritic branch where they originated, on a spatial scale of ~100 μm (Govindarajan et al., 2011; Wang et al., 2009).

What functional benefits does spatial and temporal PRP expression during learning provide for the organism? In this study we address for the first time the potential functions of the spatial regulation of PRP synthesis. In contrast, previous studies have focused almost exclusively on the functions of temporally bounded PRP synthesis. One stream of research has led to the idea that time-restricted PRP expression could be used to gate which memories persist and which are forgotten, resulting in the 'synaptic tagging and capture' theory (STC) (Redondo and Morris, 2011) (schematized in Figure 1B-E). According to the STC model, there are two types of synaptic plasticity stimuli termed 'weak' and 'strong'. Weak stimuli trigger induction of long-term potentiation and activation of a molecular 'tag' at the activated synapses. However, if left unaided both the potentiation and the tag signal decay back to baseline levels over a time period of 2–3 hrs. Hence, weak stimuli alone trigger synaptic strength changes which are eventually forgotten. Strong stimuli, in contrast, trigger induction of long-term potentiation, activation of the tag, *and* the de novo synthesis of plasticity-related-proteins (PRPs) in cytosol near the synapse (Figure 1). These PRPs can be captured by tagged synapses to stabilize synaptic strength changes, which can then persist for long times (days–months). Hence, strong protein-synthesis-inducing events create a ~2 hour time window within which other nearby weak synaptic plasticity events can become consolidated.

Analogous processes to synaptic tagging have been found at the whole animal level, termed 'behavioral tagging' (Moncada and Viola, 2007). Rats exposed to a novel environment for 5 minutes showed enhanced and persistent memory for a learning task in a different familiar environment when tested 24 hours later. This novelty-induced enhancement in memory persistence required both hippocampal protein-synthesis and dopamine receptor activation (Moncada and Viola, 2007; Wang et al., 2010), similar to the STC process at the synaptic level (O'Carroll and Morris, 2004; Wang et al., 2010). These mechanisms have been postulated to underlie the 'flashbulb memory' effect in humans (Brown and Kulik, 1977), where memories for unimportant everyday events persist if they occur nearby in time to a behaviorally salient event, such as remembering our whereabouts when hearing of the 9/11 terrorist attacks.

These proposals, and several theoretical studies (Barrett et al., 2009; Clopath et al., 2008; Papper et al., 2011; Smolen et al., 2012), have suggested how STC could be used to select memories according to their alignment in time. In contrast, the potential effects of spatial restrictions of PRP expression at the dendritic and neural circuit levels remain unclear (Govindarajan et al., 2006). We built a novel framework to study this problem.

We found that spatially patterned PRP synthesis comprises a powerful mechanism for the selective consolidation of some memories over others, even for events that occur nearby in time. The effectiveness of this mechanism depends on the specificity of synaptic wiring at the dendritic level, and on the overlap of activity patterns at the neural circuit level. We

applied this framework to quantitatively link existing experimental results from rodents at the neural level to those at the behavioral level. Finally, we used the framework to develop a new model for how STC might allow selective generalization of memories during sleep.

Results

Our general goal was to develop a quantitative framework that can predict the degree of consolidation of a weak memory event as a function of its temporal and spatial overlap with a strong protein-synthesis-inducing memory event. To do this we separately studied the effects of spatial patterning of protein synthesis at the 1) dendritic and 2) neural circuit levels, in turn.

Timing and spatial overlap of synaptic inputs at dendritic level determines the degree of synaptic plasticity consolidation

What is the expected long-term synaptic change from a weak plasticity-inducing stimulus onto a single neuron? We derived an expression for this quantity based on the following simple model (Figure 2A). Consider a single postsynaptic neuron with three dendrites that is innervated by two presynaptic neurons, labeled *pre1* and *pre2*. Neuron *pre1* synapses onto the first dendrite of the postsynaptic neuron. This synapse is activated with a strong stimulus which causes LTP and the synthesis of PRPs. For simplicity we assume that the PRPs are restricted to first dendrite and do not reach the other two dendrites (Govindarajan et al., 2011). Now consider the synapse from neuron *pre2*. If it is activated with a weak stimulus, it will have access to the PRPs (hence becoming stabilized) only if it also targets the first dendrite. What determines whether the synapse from neuron *pre2* targets the same dendrite as the synapse from neuron *pre1*? We propose that this process can be parameterized by a quantity we term the dendritic correlation coefficient, c_{dend} , which may vary in most cases between 0 and 1. If synapses from neurons *pre1* and *pre2* always terminate onto the same dendrites, then they are perfectly correlated with $c_{dend} = 1$. If they select dendrites independently, then they are uncorrelated with $c_{dend} = 0$. In that situation, the probability p that the two presynaptic neurons synapse onto the same dendrite is just equal to chance. In our above example because the postsynaptic neuron has three dendrites, chance level is simply $1/3$. Intermediate levels of c_{dend} bias the probability that the two neurons to synapse onto the same dendrite without guaranteeing it. For example, in our three-dendrite example if $c_{dend} = 1/2$, then $p = 2/3$. In general, the probability that the neurons synapse onto the

same dendrite is $p = c_{dend} + (1 - c_{dend}) \frac{1}{d}$, where d is the number of dendrites on the postsynaptic neuron. Using this equation, and the assumption that the temporal window for PRP capture can be described by an exponential function (Govindarajan et al., 2011), we can write an equation for the mean long-term synaptic strength change for a synapse receiving a weak stimulus:

$$\begin{aligned} \langle \Delta w \rangle &= \alpha e^{-\frac{|\Delta t|}{\tau}} p \\ &= \alpha e^{-|\Delta t|/\tau} \left[C_{dend} + (1 - C_{dend}) \frac{1}{d} \right] \quad (1) \end{aligned}$$

where α is a constant, t is the time interval between the strong and the weak stimuli, and τ is the time constant of the tagging and PRP capture windows. In Figure 2B we plot the consolidated synaptic strength change as a function of the time interval between the strong and weak stimuli, for a range of different values of c_{dend} , both as averaged over multiple random simulations and as predicted by equation 1. When c_{dend} is large (close to 1), then consolidation is effective. However, when c_{dend} is smaller (close to 0), consolidation is much reduced. In this way, the degree of spatial correlation between strong and weakly activated synapses can enable selective consolidation of plasticity at some weakly activated synapses but not others (Figure 2B-C).

Importantly, the effectiveness of the dendritic correlation mechanism depends on the number of dendrites available (Figure 2D). For example, if a neuron had only one dendrite, then all synapses onto the neuron would have access to PRPs and there could be no selective consolidation. In contrast, a greater numbers of dendrites allow for increased discrimination between synaptic populations by reducing the baseline chance level that a weakly activated synapse will randomly terminate onto the same dendrite as a strongly activated synapse.

Overlap between strong and weak patterns at neural circuit level determines degree of consolidation of weak pattern

Above we considered plasticity at synapses onto a single neuron. However, neural representations underlying cognitive processes are distributed over many neurons simultaneously (Churchland and Sejnowski, 1994). We studied the implications of the STC theory for distributed neuronal activity patterns by considering a simple model of a two-layer feedforward neural network (Figure 3A). Although this model was generic, it could potentially be applied to many different projections in the nervous system involved in learning and memory. We defined an activity pattern as the subset of neurons in the presynaptic and postsynaptic populations that were involved in encoding a specific memory. Our aim was to calculate the mean consolidated synaptic strength change resulting from a weak activity pattern as a function of its temporal and spatial overlap with a strong activity pattern.

Neurons in the weak pattern can be split into two categories: those that are also part of the strong pattern, and those that are not. We label these two types of neuron as 'shared' and 'weak-only' respectively. Consequently, the synapses in the weak pattern can be split into four groups based on their pre and post-synaptic neuron types: 1) pre weak-only to post weak-only, 2) pre shared to post weak-only, 3) pre weak-only to post shared and 4) pre shared to post shared. Importantly, these four groups of synapse will experience different degrees of consolidation of synaptic plasticity (Table 1). Synapses that terminate on weak-only postsynaptic neurons (groups 1 and 2) will not have their changes consolidated, because they will not have access to PRPs. Hence their expected synaptic plasticity change is zero: $\langle w \rangle_1 = \langle w \rangle_2 = 0$. In contrast, synapses that arise from weak-only presynaptic neurons but terminate on shared postsynaptic neurons (group 3) may be consolidated if the dendrite they synapse onto has PRPs available. We assumed that a dendrite contains PRPs if at least one synapse from the strong pattern terminates there. The probability of this occurring, $p(N_{strong} \geq 1)$, is a function of the dendritic correlation between presynaptic

neurons, the number of neurons in the presynaptic population, and the presynaptic sparsity of the strong pattern (see Experimental Procedures and Supplemental Figure 1). In sum, the expected synaptic plasticity of a group 3 synapse is given by (see Experimental Procedures):

$\langle \Delta w \rangle_3 = \alpha e^{-\frac{|t|}{\tau}} p (N_{strong} \geq 1)$. The remaining set of synapses, those arising from shared pre- and post-synaptic neurons (group 4), are guaranteed to have access to PRPs because they were part of the strong activity pattern. However, their total expected synaptic plasticity change will be determined by the interactions between the signals from the strong and weak activity patterns, which are usually not additive (Abraham, 2008). For example, it may be that LTP induced by the first pattern causes occlusion of LTP from the second pattern (Frey et al., 1995). Because these interactions are complicated, incompletely understood, and outside the scope of this study, we simply denote the expected consolidated change at these synapses as an 'overwriting' term which can depend on the properties of the strong and weak patterns and the time interval between them: $\langle w \rangle_4 = \langle w \rangle_{over}$.

The fraction of synapses in a weak pattern that fall into each of these four groups (r_1, r_2, r_3, r_4) depends on the degree of overlap between the strong and weak patterns. The ratios between the groups are

$$r_1:r_2:r_3:r_4 = (1 - q_{pre})(1 - q_{post}) : q_{pre}(1 - q_{post}) : (1 - q_{pre})q_{post} : q_{pre}q_{post}$$

where q_{pre} and q_{post} are the fraction of presynaptic and postsynaptic neurons in the weak pattern that overlap with the strong pattern, respectively. We then calculated the total mean consolidated synaptic strength change in a weak pattern by adding the weighted contributions from the four groups of synapses (see Experimental Procedures):

$$\begin{aligned} \langle \Delta w \rangle &= r_1 \langle \Delta w \rangle_1 + r_2 \langle \Delta w \rangle_2 + r_3 \langle \Delta w \rangle_3 + r_4 \langle \Delta w \rangle_4 \\ &= q_{post} \left[(1 - q_{pre}) \alpha e^{-|\Delta t|/\tau} p + q_{pre} \langle \Delta w \rangle_{over} \right] \quad (2) \end{aligned}$$

This equation clarifies the distinct roles of presynaptic versus postsynaptic pattern overlap. Increasing the postsynaptic overlap between patterns always increases the consolidated synaptic strength change. In contrast, increasing the presynaptic overlap between patterns decreases the impact of PRPs sharing between synapses, but increases the impact of overwriting at shared synapses. Whether the net effect of increasing presynaptic overlap is to increase or decrease the consolidated synaptic strength change depends on the relative magnitudes of the PRP sharing term and the overwriting term. In summary, the degrees of pre and postsynaptic overlap between strong and weak patterns are critical determinants of the degree of consolidation of synaptic plasticity from a weak pattern.

Density of neural representations determines the effectiveness of protein sharing for weak memory consolidation

What determines the degree of overlap between neural activity patterns in the brain? There are several factors: the brain region involved, how dense or sparse its representations are, and the functional relationship between the specific items that are being represented in pre and post-synaptic populations. We first considered the simplest case where pre- and post-

synaptic patterns are spatially random and uncorrelated. In this case the fractional overlap between a weak and strong pattern, q , is simply equal to the sparsity of the strong pattern f_s , defined as the fraction of all neurons in the population that are part of the strong pattern. Because $q = f_s$, according to eq. 2 we should expect that the degree of sparsity of neuronal representations will influence the degree of consolidation of a weak pattern. Dense activity will lead to greater overlap between patterns. Although pattern sparsity in a brain region may take any value between 0 and 1, for illustration in Figure 4 we plot the mean consolidated synaptic strength change for two example sparsity levels which we label sparse ($f_s = 0.1$) and dense ($f_s = 0.5$) from both theory (curves) and simulation (circle symbols).

There are four types of presynaptic to postsynaptic projection in this scenario: sparse-to-sparse, sparse-to-dense, dense-to-sparse, and dense-to-dense. The four different types of projection showed substantially different dependencies of synaptic consolidation on dendritic protein translation. In a sparse-to-sparse projection (Figure 4A), there is little overlap between weak and strong pattern both pre and post-synaptically. Few synapses are shared between patterns, implying little contribution from overwriting. Although varying dendritic correlation c_{dend} does have a moderate impact, its effect is limited because few neurons are shared postsynaptically. As a result, memory consolidation here is selective but weak. In a sparse-to-dense projection (Figure 4B), few neurons are shared presynaptically but many are shared postsynaptically. This scenario minimizes the contribution of overwriting, and maximizes the contribution of dendritic protein sharing. As a result, in this situation regulation of c_{dend} constitutes a powerful mechanism for strong and selective memory consolidation. In a dense-to-sparse projection (Figure 4C) many neurons are shared presynaptically but few neurons are shared postsynaptically. This scenario is the least optimal for the utilization of dendritic protein sharing because few synapses in the weak pattern will have access to PRPs, and most of those that do will also be part of the strong pattern. In this case, memory consolidation is weak and not selective. Finally, in a dense-to-dense projection (Figure 4D), many neurons overlap both pre and postsynaptically. In this situation, many synapses are shared between strong and weak patterns, so the contribution of overwriting is large, and varying c_{dend} has relatively little impact. In this scenario, consolidation may be strong but not selective across memories.

These results show that the density of neural representations is a strong determinant of the effectiveness of dendritic protein sharing for selective memory consolidation. These results may prove helpful in predicting the prevalence of dendritic protein sharing from brain region to brain region. An important caveat is that these conclusions are based on the assumption that strong and weak patterns are spatially uncorrelated. We next explored scenarios with more structured activity patterns.

STC is of limited effectiveness when pre and postsynaptic populations code for the same stimuli

How does the effectiveness of the STC mechanism depend on the structure of the neural representation? To begin to address this question we considered a two-layer feedforward neural network where neurons in both the pre- and post-synaptic populations were arranged according to their preferred value of a 1-dimensional circular stimulus (Figure 5A). This

coding scheme is common in many brain regions, for example the orientation tuning of neurons in early mammalian visual cortex, or the place field tuning of a rodent's location on a linear track by hippocampal neurons. We assumed that presynaptic neurons which preferred similar stimuli were more likely to synapse onto the same postsynaptic dendrites than neurons which preferred different stimuli (Figure 5C). A strong pattern presented to the network causes PRP translation in only the innervated dendrites of the activated postsynaptic neurons. If a weak pattern is also presented to the network, its expected consolidation can be given in terms of four key parameters (see Experimental Procedures): the distance between the strong and weak stimuli values θ , the width of the pre- and postsynaptic tuning curves r_{pre} and r_{post} , and the width of dendritic spatial correlation window between presynaptic neurons λ . We highlight three main conclusions: 1) Consolidation of weak patterns through STC can only occur for stimuli within the range of the width of the postsynaptic tuning curve ($0 < \theta < r_{post}$). 2) Dendritic restriction of PRPs will only have an additional effect if the dendritic spatial correlation window also changes within this stimulus range: $\lambda < r_{post}$. 3) Increasing the width of the presynaptic tuning curve decreases the consolidation of weak patterns for all stimulus values. Some example scenarios are plotted in Figure 5D-F. If presynaptic tuning curves are narrow but postsynaptic tuning curves are wide (Figure 5D), substantial consolidation of a weak pattern can occur through STC. Dendrite-specific synaptic wiring schemes can also have a substantial impact (compare dashed and solid curves). In contrast, if presynaptic tuning curves are wide while postsynaptic tuning curves are narrow (Figure 5F), very little weak pattern consolidation occurs through STC and dendrite-specific synaptic wiring has little impact. If pre and post tuning curves are the same width, then moderate consolidation occurs (Figure 5E). These results imply that although STC can potentially act a mechanism for pattern consolidation in circuits where representations vary smoothly with a stimulus feature, very specific coding conditions must be met for this to occur.

STC can cause memory linking in the CA3-CA1 pathway

To explore the possible function of STC for a second biologically relevant coding scheme we modeled the rodent hippocampal CA3-to-CA1 Schaffer collateral pathway (Figure 6A). These subfields are key components of the hippocampal circuit, important for normal memory function (Nakashiba et al., 2008), and their synaptic connection is the locus at which STC has been best studied (Alarcon et al., 2006; Frey and Morris, 1997; Sajikumar and Frey, 2004). We retained a two-layer feedforward network as before, but instead of a 1-D variable, let activity patterns in CA3 and CA1 represent the entirety of a rodent's spatial environment. We attempted to replicate previous experiments (Leutgeb et al., 2005; 2004; Vazdarjanova and Guzowski, 2004), where an animal was allowed to explore three different environments: a reference environment *A*, a similar environment *A'*, and a substantially different environment *B*. The degree of overlap between the sets of neurons that get activated in these three types of environment differs qualitatively between CA3 and CA1 (Leutgeb et al., 2004; 2005; Vazdarjanova and Guzowski, 2004) (Figure 6B-E). CA3 is believed to perform a 'pattern completion' operation for similar environments, so that the set of neurons that are active when the rodent is in environment *A* is highly overlapping with the set of neurons active when the animal is either replaced in *A*, or allowed to explore *A'* (Figure 6B,D,E). However, for sufficiently distinct environments, CA3 performs 'pattern

separation', so that the set of neurons active in A is found to be statistically independent of those active in B (Figure 6B,C,E). In contrast, CA1 shows a much more graded shift in patterns between environments, so that A and A' representations share a large number of neurons (but less so than CA3), and A and B representations are found to show an above-chance overlap in their active neuron sets (Figure 6B-E).

Given these data, we aimed to calculate the mean long-term synaptic strength change from a 'weak' event in either A , A' or B relative to a preceding 'strong' event in environment A . Inserting the experimentally derived pre- and post-synaptic overlap fractions (Vazdarjanova and Guzowski, 2004) into eq. 2, we plotted the expected synaptic strength change for each spatial context in Figure 6F. We found that the expected weight change for a weak event in environment A was greater than that for a weak event in environment A' , while both were greater than that for a weak event in environment B . We also considered the effect of varying the degree of dendritic correlation between weak and strong patterns, c_{dend} , and found that, as expected, decreasing c_{dend} always causes a decrease in memory strength (Figure 6F-H). However, the relative decrease in memory strength as c_{dend} varied from 1 to 0 in environment A was relatively minor (~16%) while the decrease in memory strength for environment B was almost complete (~84%) (Figure 6H). A' was intermediate to these extremes. Although predicting error bars on these estimates would be possible in principle, it would require additional assumptions for both a specific neuronal noise model and a model for the experimental measurement errors. Because little data are available at present to constrain either of these sources of variability, we instead limit our predictions to the mean outcome only. All of the above effects were due to the different degrees of pre (CA3) and post (CA1) synaptic overlap for A/A , A/A' and A/B . The high pre and post overlap for A/A implies a large number of shared synapses between the representations with relatively little role for PRP sharing between synapses. In contrast, A/B representations have almost zero pre-synaptic overlap but substantial postsynaptic overlap. This scenario is optimal for selective consolidation via STC. In summary, we found that STC is of limited use for cross-consolidation of hippocampal memories for events occurring within the same environment, but can act as a powerful mechanism for selective consolidation of hippocampal memories for events occurring in distinct environments, as observed experimentally (Ballarini et al., 2009; Moncada and Viola, 2007; Wang et al., 2010).

A two-step model for generalization of memories during sleep

The above findings show how spatial protein synthesis allows for selective memory consolidation. We next applied this same framework to a different open biological problem: how memories become generalized during sleep. Although these might appear to be unrelated problems, we show how they may be linked at the molecular level. A growing body of evidence supports the idea that human memories become reorganized during sleep (Lewis and Durrant, 2011; Rasch and Born, 2013; Stickgold and Walker, 2013). Two main effects have been observed: first, some memories are consolidated while others are forgotten. Second, the memories that are chosen for consolidation can also be generalized. Often this generalization takes the form of 'gist extraction' or abstraction from a small set of experiences to generate broader knowledge about the world (Lewis and Durrant, 2011; Stickgold and Walker, 2013; Wagner et al., 2004). For example, if a person were bitten by

their neighbor's dog, then they might generalize this information to alter their beliefs about the likelihood of getting bitten by all dogs. Importantly, this generalization process is found to be selective for items similar to the consolidated memory. The memory enhancement is not seen for items that are substantially different from the experienced items (Payne et al., 2009). To return to the biting dog example, the person should not over-generalize their new knowledge to their beliefs about cats.

At the neural level, the process of selective memory consolidation is believed to involve hippocampal-cortical interactions during slow-wave sleep (Rasch and Born, 2013; Stickgold and Walker, 2013). In contrast, little is known about how memories can become generalized during sleep. Using a simple neural network model, we explored the possibility that STC could enable selective memory generalization during sleep.

The novel generalization mechanism we propose occurs in two steps, corresponding to SWS and REM sleep respectively (Figure 7A). We considered a two-layer feedforward neural network representing the projection from one population of cortical neurons to another. During the SWS step a subset of neurons in the presynaptic population are activated by an external hippocampal 'sharp-wave ripple' input (Buzsáki, 1996; Wierzynski et al., 2009). This presynaptic pattern in turn activates a subset of the neurons in the postsynaptic population, causes potentiation at the activated synapses, heterosynaptic depression at the non-activated synapses, and the synthesis of PRPs in the activated postsynaptic neurons. We label this activity pattern 'strong'. During the following REM step, multiple 'weak' activity patterns are sequentially activated in the presynaptic population that propagate to the postsynaptic population. These patterns correspond to cortically-generated REM sleep activity patterns, as observed experimentally in both rodents (Ribeiro et al., 2004) and humans (Maquet et al., 2000). In the model, these patterns cause potentiation at their activated synapses. Crucially, this potentiation is consolidated only for the synapses onto postsynaptic neurons that have PRPs available from the earlier strong pattern. Together, we found that this two-step process was sufficient to achieve selective generalization of learned memories. In the following paragraphs we elaborate on the details of our results.

In the network model we studied, each neuron in the presynaptic population was connected to a random subset of neurons in the postsynaptic population with a low probability (see Experimental Procedures). Initially all synapses were of uniform strength. As previously found (O'Reilly and McClelland, 1994), this network configuration caused all sparse input activity patterns to undergo pattern separation: any overlap between two patterns in the presynaptic layer was reduced at the postsynaptic layer (Figure 7B). However, synaptic plasticity from the strong activity pattern changed this input-output mapping. Homosynaptic potentiation increased the probability that the targeted postsynaptic neurons were also activated for input patterns similar to the strong pattern (the region where magenta curve is above the black curve in Figure 7B). This corresponds to pattern completion. The heterosynaptic depression also played an important complementary role by decreasing the probability that the targeted postsynaptic neurons were activated for patterns dissimilar to the strong pattern (the region where the magenta curve is below the black curve in Figure 7B). This corresponds to pattern separation (O'Reilly and McClelland, 1994). This dual pattern-separation/pattern-completion property proved to be critical for selective

generalization. In Figure 7C we plot the mean strength of the synapses activated by a hypothetical pattern as a function of its overlap with the strong pattern (magenta curve and circles). As expected, the mean synaptic strength decreased with decreasing overlap with the strong pattern, because fewer synapses were shared with the strong pattern. Then, during the REM step, weak patterns were activated sequentially in random order, causing potentiation at their synapses. However this potentiation was consolidated only if the synapse had access to PRPs. After we ‘played’ several of these patterns, we replotted the mean synaptic strength of each pattern as a function of its overlap with the strong pattern (Figure 7C, blue filled circles). All patterns with high overlap (>0.5) with the strong pattern had an equally potentiated mean synaptic strength, while patterns with a low overlap (near the chance level of 0.1) had a mean synaptic strength close to the initial value. Hence, synaptic changes were consolidated only if they came from weak patterns that were similar to the strong pattern. Under the assumption that patterns with high overlap represent items or events that are similar to each other in the external world, this mechanism offers a novel neural explanation for selective generalization of memories in the brain during sleep.

To further explore how this mechanism worked, we categorized the synapses in a weak pattern according to the same scheme as before. In Figure 7D-E we plot the fraction of synapses in a weak pattern in groups 3 and 4 (see above) as a function of its overlap with the strong pattern, as calculated both analytically and from simulation results. In the initial state (Figure 7D), the fraction of synapses in group 4 (shared pre to shared post) decreased rapidly with decreasing pattern overlap, while the fraction of synapses in group 3 (weak-only pre to shared post) was only weakly dependent on overlap. Because STC can only act to consolidate synapses in group 3, in this initial state the network could not guarantee selective consolidation for patterns in any particular overlap range. However, following plasticity from the strong pattern the fraction of synapses in group 3 depended greatly on pattern overlap (Figure 7E). It was minimal for patterns of either high or low overlap, and maximal for patterns of intermediate overlap. This shift allowed weak patterns with intermediate overlap to be selectively consolidated through STC.

Further study of the model uncovered three additional findings (data not shown). First, the number of presented weak activity patterns needed to be limited to avoid over-generalization. The maximum number of possible patterns before over-generalization occurred depended on network parameters, the synaptic plasticity rule, and the statistics of the activity patterns. Second, the degree of synaptic plasticity from the strong pattern determined the extent of generalization for subsequent weak patterns. Stronger plasticity leads to broader generalization (O'Reilly and McClelland, 1994). Third, weak patterns that had an overlap with the strong pattern sufficient for substantial consolidation were extremely unlikely to arise by chance alone. Instead, the input set needed to be biased towards these patterns.

Discussion

We have introduced a novel quantitative framework for linking protein synthesis, neural circuit wiring, and activity pattern properties to selective memory consolidation. Previous research had suggested how STC could facilitate selective memory storage according to

events' separation in time (Ballarini et al., 2009; Frey and Morris, 1997). Our results quantify the additional benefits of spatial patterning of neural protein synthesis (Alarcon et al., 2006; Govindarajan et al., 2006; Sajikumar et al., 2007). Utilizing the spatial dimension lets STC further select which memories to store based on their content, even if their causal events in the outside world overlap in time.

We found that multiple key factors determine the selectivity of memory consolidation. The first factor is the spatial arrangement of synapses on a neuron's dendritic tree. If two presynaptic neurons have a tendency to synapse onto the same dendrite (perhaps based on their functional properties), then they can regulate each other through STC. On the other hand, if two neurons synapse at random locations on the dendritic tree with respect to each other, then it is unlikely that they will share the same dendrite and hence will have limited ability to affect each other's synaptic plasticity consolidation. The second factor is the degree of overlap between neural activity patterns at the circuit level. Crucially, our results stress the differing roles played by the overlap in the presynaptic and postsynaptic neural populations (examples from the mammalian hippocampus might be CA3 and CA1, respectively). A large overlap between strong and weak activity patterns in the postsynaptic population always enhances the consolidation of the weak synaptic plasticity pattern, because a large fraction of the synapses in the weak pattern will have access to PRPs. The degree of presynaptic overlap performs a different function: it determines the relative impact of protein sharing between synapses versus that of simply overwriting the same synapses twice. Whether or not increasing presynaptic overlap increases the net consolidation of a weak pattern will depend on the relative magnitudes of plasticity via these respective mechanisms, which in general may vary from one scenario to another. The important distinction is that both rely on fundamentally different processes at the molecular level and so may be differentially controlled by the cell.

Are these mechanisms specific to particular brain circuits, or widespread throughout the nervous system? STC has been experimentally demonstrated in rodent hippocampus *in vitro* (reviewed by Barco et al., 2008; Reymann and Frey, 2007), and *in vivo* (Shires et al., 2012) and invertebrate sensory neurons *in vitro* (Martin et al., 1997). The analogous behavioral tagging process has been observed for hippocampal (Ballarini et al., 2009; Moncada and Viola, 2007) and insular cortex (Ballarini et al., 2009; Merhav and Rosenblum, 2008) dependent learning tasks in rodents. These data, along with the fact that protein synthesis appears to be critical for multiple forms of persistent memory throughout the brain (Davis and Squire, 1984; Hernandez and Abel, 2008) suggests that, in principle, the mechanisms we propose could be prevalent across the nervous system.

Because PRP synthesis is triggered by electrical activity in neural circuits, any spatial patterning in PRP expression must be inherited from spatial structure already present in the electrical signals. At the circuit level it is clear that different neurons within the same population can receive distinct, structured synaptic inputs, and respond with heterogeneous gene expression (Mackler et al., 1992). However the degree of spatial structure in synaptic inputs at the sub-neuron level is less clear. At a coarse scale, layered brain structures such as the hippocampus and neocortex are wired such that the dendritic trees of larger cells (such as pyramidal neurons) may collect inputs from multiple layers. In this case, different dendritic

regions on the same cell can receive synaptic input from distinct presynaptic populations. Indeed, in certain conditions STC can obey layer-specificity in CA1 hippocampal pyramidal neurons (Alarcon et al., 2006; Pavlowsky and Alarcon, 2012; Sajikumar et al., 2007). Whether synaptic inputs from within a given layer are also structured is more controversial. There is currently evidence both for and against this possibility depending on the examined stimulus set and brain circuit (Chen et al., 2011; Jia et al., 2010; Kleindienst et al., 2011; Makino and Malinow, 2011; Takahashi et al., 2012). Although in this study we explicitly assumed that synaptic wiring is structured at the level of single dendrites, our findings will still be applicable if neural activity in a given brain region is found to be patterned at only the population or layer level.

The information stored in memories may be useful for dealing with future situations (Schacter et al., 2007). However, episodic memories are defined as those for specific events that occurred in the past. Because it is unlikely that an event that occurred in an organism's past will reoccur in an identical fashion in the organism's future, the information stored in an episodic memory might be best utilized if it were combined with prior knowledge to generalize the organism's beliefs about a larger set of related events (Tenenbaum et al., 2011).

We have proposed a new two-step model for this process of selective generalization during sleep. This model may help unify several disparate pieces of data. First, sleep can enable generalization of learned information (Cai et al., 2009; Ellenbogen et al., 2007; Pace-Schott et al., 2009; Payne et al., 2009). Second, memory consolidation is found to be maximally effective when both SWS and REM occur in succession (Gais et al., 2000; Mednick et al., 2003; Stickgold et al., 2000). In the model we propose, proper function of the REM step requires a preceding SWS step. Third, most REM dreams in humans are neither veridical replays of previously experienced events nor completely unrelated, but somewhere intermediate (Fosse et al., 2003; Wamsley et al., 2010). This intermediate degree of similarity of REM activity patterns to the veridical activity patterns of the preceding SWS is a fundamental feature of the model we suggest. Fourth, rodents are found to show coordinated hippocampo-cortical 'replay' of previously experienced neural activity patterns during SWS (Ribeiro et al., 2004; Siapas and Wilson, 1998; Wierzynski et al., 2009; Wilson and McNaughton, 1994), while showing more weakly correlated, but still statistically similar, hippocampus-independent cortical activity patterns during REM sleep (Ribeiro et al., 2004; Wierzynski et al., 2009). Fifth, in rodents new episodic memories are detailed and hippocampus dependent, but over subsequent weeks become both more generalized and less hippocampus dependent (Wiltgen and Silva, 2007; Winocur et al., 2007). The dual properties of extra-hippocampal transfer and content generalization are deeply linked in the model we propose.

The framework we introduce makes several new and testable predictions. First, the degree of both postsynaptic and, crucially, presynaptic overlap of neural activity patterns at the circuit level will determine the magnitude of consolidation of a weak memory event. These measurements can now be readily made either *ex vivo* or *in vivo* using neural activity reporters such as immediate-early gene expression (Vazdarjanova et al., 2002) or fluorescent calcium indicators (Dombeck et al., 2010), respectively. Second, due to the distinct

computations performed by CA1 and CA3, our model predicts that hippocampus-dependent behavioral tagging processes will most effectively work across events occurring in distinct environments, but will only weakly influence interactions between different events in the same environment (Fig 6). Third, our novel model for generalization during sleep assigns specific roles to SWS and REM sleep phases (Fig 7). According to this model, blocking protein synthesis during SWS alone will be sufficient to block both consolidation of the initial memory and its generalization during the subsequent REM phase, while blocking tagging processes (Redondo and Morris, 2011) during the REM phase should block memory generalization while leaving the original memory intact.

In general, neuronal protein synthesis may constitute a powerful signal for encoding information in the nervous system because it is complementary to electrical signaling in both time and space. Electrical signals operate on fast timescales (ms to s) and relatively wide spatial scales (single neurons to networks), whereas PRP synthesis and degradation operates on slow timescales (minutes to hours) and small spatial scales (dendrites to single neurons). Hence, patterned PRP synthesis could be used by the brain for linking or discriminating behavioral events that would be difficult to achieve using electrical signaling alone. Problems may arise if this process goes awry. Many neurodevelopmental disorders that are associated with learning disabilities have been linked to altered neuronal protein translation (Kelleher and Bear, 2008; Zoghbi and Bear, 2012). Indeed, recent experiments have shown that hippocampal STC is altered in a mouse model of Fragile-X Syndrome, a common cause of autism (Connor et al., 2011). The framework we propose may help link these deficits at the cellular level to learning deficits at the cognitive level.

Experimental Procedures

All simulations and analysis were performed using MATLAB (Mathworks).

Generation of correlated synaptic locations on dendrites

The dendritic locations for each synapse throughout this study were generated using an algorithm previously proposed for generating correlated spike trains (Macke et al., 2009). This algorithm generates correlated binary variables with arbitrary specified means and pairwise correlations by applying suitable thresholds to an underlying correlated multivariate Gaussian. Correlated multivariate Gaussian samples are readily generated using standard software packages like MATLAB (Mathworks), for example, because they contain no higher-order correlations beyond pairwise. Our dendritic synapses had the additional constraint that we required exactly one synapse from each presynaptic neuron onto each postsynaptic neuron. We obeyed this constraint by posthoc rejection of any samples that did not meet the requirement, and verified that this correction did not alter the resulting synaptic correlations.

Expected strength change of a synapse that is 'weak only' presynaptically but 'shared' postsynaptically

A synapse from a weak pattern that is from a 'weak only' presynaptic neuron to a 'shared' postsynaptic neuron will have access to PRPs only if its dendrite receives synaptic input

from at least one synapse in the strong pattern. We sought the probability that this occurs $p(N_{strong} \geq 1 | \text{weak synapse})$, in terms of the following parameters: the number of neurons in the presynaptic population N_{pre} , the density of the strong presynaptic pattern f_{pre} , the number of dendrites per postsynaptic neuron d , and the dendritic correlation between the neurons in the presynaptic population c_{dend} (here assumed uniform). One way to find this probability is to first calculate the probability distribution for the number of synapses per dendrite $P(N_{syn})$ (Supplemental Figure 1). If synaptic locations were uncorrelated ($c_{dend} = 0$), then synapse numbers would be distributed binomially. On the other hand, if all synapses were perfectly correlated ($c_{dend} = 1$), then all synapses would terminate onto the same single dendrite, while all other dendrites on that neuron received no synapses. Varying c_{dend} from zero to one moves $P(N_{syn})$ between these two extremes (Supplemental Figure 1). We solved for this distribution numerically by simulating many realizations of synaptic wiring given values for the parameters N_{pre} , d , and c_{dend} . Although for some very restricted synaptic wiring models it would be possible to find a closed-form expression for the distribution $P(N_{syn})$ (e.g. Diniz et al., 2010), we chose to solve for it numerically because doing so allows for arbitrary arrangements of pairwise dendritic correlations between all of the presynaptic neurons. After evaluating $P(N_{syn})$, we calculated the distribution for the fraction of synapses $f(i)$ that have a given number i synapses on their dendrite from

$$f(i) = \frac{P(i) \times i \times d}{N_{pre}}$$

which is the product of the probability of getting i synapses on a dendrite, and that number of synapses i , and the number of dendrites d , all divided by the total number of synapses N_{pre} . Finally, the probability of at least one synapse from the strong pattern terminating on the same dendrite as a given weak synapse, $p(N_{strong} \geq 1 | \text{weak synapse})$, is derived using the hypergeometric distribution as follows. We start from the vantage point of a weak synapse on a given dendrite. The hypergeometric distribution describes the probability $H(k)$ of drawing k “successes” from a population of size N which contains K total “successes” when drawing n times without replacement. If we consider the population size N as the total number of remaining synapses $= N_{pre} - 1$, the total number of “successes” K as the total number of strong synapses $= N_{pre} f_{pre}$, and the drawn number n as the number of other synapses on the dendrite $= i - 1$, and k the eventual number of successes as the number of strong synapses on that dendrite N_{strong} . The total probability that at least one synapse is strong is the sum of $H(k)$ from $k = 1$ to $k = N_{pre} f_{pre}$, which is also equal to one minus the probability that there are no strong synapses on that dendrite, $H(k = 0)$. Finally, this quantity must then be averaged over all possible values of i , the probabilities of which are given by $f(i)$. In summary, $p(N_{strong} \geq 1 | \text{weak synapse}) = \sum_{i=1}^{N_{syn}} f(i) \times (1 - H(k=0))$ where $H(k)$ is the hypergeometric distribution with parameters $N = N_{pre} - 1$, $K = N_{pre} f_{pre}$ and $n = i - 1$. Although the hypergeometric distribution does not directly take into account the correlation in synaptic locations, in this case its use is valid because the effect of correlations has already been included in the earlier step of the calculation, when solving for $P(N_{syn})$.

Network simulations

Most of our STC simulations (Figures 2, 4, 5 and 6) involved two-layer feedforward neural networks with all-to-all pre to post-synaptic connectivity. The simulations consisted of two main steps. First, we generated the set of postsynaptic dendritic locations of all synapses

given the following parameters: the number of pre N_{pre} and post-synaptic N_{post} neurons, the number of dendrites per postsynaptic neuron d , and the dendritic correlation between each pair of pre-synaptic neurons c_{dend} . For Figures 2, 4 and 6, c_{dend} was the same for all pairs of neurons, but for Figure 5 the pairwise c_{dend} values depended on the difference between the preferred stimulus values of the given pair of neurons (Figure 5C). For the simulations presented in Figure 2, $N_{pre} = 2$, $N_{post} = 1$ and $d = 10$; Figure 4, $N_{pre} = 20$, $N_{post} = 100$ and $d = 15$; Figure 5 $N_{pre} = 100$, $N_{post} = 100$ and $d = 50$; Figure 6, $N_{pre} = 20$, $N_{post} = 100$ and $d = 30$. For the second step, we presented one weak and one strong activity pattern separated by a specified time interval and simulated the tag, PRP and synaptic strength dynamics. The synaptic tag $K(t)$ and dendritic PRP $P(T)$ dynamics were simulated as exponential time courses. Synaptic strength dynamics also followed the form outlined above:

$$\frac{dw}{dt} = \rho (P(t) \times K(t)).$$

The values of the scaling constants A_P , A_K and ρ were chosen so that the maximal synaptic potentiation was ~40%, similar to the magnitude of LTP observed experimentally, while the time constant of PRP and tag decay was set to $\tau = 30$ mins, to reproduce the ~90 min wide time window observed experimentally (Govindarajan et al., 2011). For Figure 2, we present the synaptic strength changes averaged over 100 realizations for each configuration of c_{dend} and t .

STC in networks where pre- and post-synaptic populations code for the same 1-D variable

We examined the case where activity patterns in two layers of feedforward network are determined by the value of a particular one-dimensional circular variable in the external world, θ . This variable may represent, for example, the orientation of a bar stimulus in the visual field. We assumed that each neuron is tuned to respond preferentially to a particular value of $\theta = \theta_{pref}$ (but also responds within a range $\theta_{pref} \pm r_{\theta}/2$), that the θ preferences of neurons are evenly distributed across the populations, and that they tile the entire range of possible values of θ . We also assume that the postsynaptic dendritic correlation between any two neurons i and j in the presynaptic population $c_{ij} = |\theta_i - \theta_j|$ is a decaying function of the distance between their preferred values (Figure 5C): $c_{ij} = e^{-|\theta_i - \theta_j|/\lambda}$ where λ is a parameter that determines the extent of dendritic correlations with distance in feature space (we set $\lambda = 0.1 \times \theta_{max}$). Hence, presynaptic neurons that prefer similar values of θ are more likely to synapse onto the same dendrites in postsynaptic neurons. We then derived an expression for the component of the mean synaptic strength change induced by a weak stimulus pattern that is due to STC, in terms of $r_{\theta pre}$, $r_{\theta post}$, θ and λ as follows. The presynaptic overlap of any two patterns with a stimulus difference of θ is:

$$q_{pre} = \begin{cases} (1 - \Delta\theta) / r_{\theta pre} & \text{if } \Delta\theta < r_{\theta pre} / 2 \\ 0 & \text{if } \Delta\theta > r_{\theta pre} / 2 \end{cases}.$$

An analogous expression exists for q_{post} . By

inserting the overlap terms into equation 2 (Results) we can calculate the mean synaptic strength change from a weak stimulus pattern:

$$\langle \Delta w \rangle = \left(\frac{1 - \Delta\theta}{r_{\theta post}} \right) \left[\left(1 - \frac{1 - \Delta\theta}{r_{\theta pre}} \right) \alpha e^{-|\Delta t|/\tau} p + \left(1 - \frac{\Delta\theta}{r_{\theta pre}} \right) \langle w \rangle_{over} \right]$$

where p is the probability that a synapse in the weak pattern terminates on a dendrite that also receives a strongly activated synapse. As above, we solved for p numerically by simulating many realizations of synaptic wiring given values for the parameters N_{pre} , d and λ .

Because representations in pre and post populations are centered on the same value of θ , then every weak pattern that overlaps postsynaptically with a strong pattern will necessarily contain a non-zero fraction of presynaptic neurons that are also part of the strong pattern. Hence, in most cases there will be a significant contribution from the overwriting term. We further analyzed the conditions for maximal STC-based enhancement of consolidation. First, the degree of consolidation is linearly correlated with the postsynaptic overlap between weak and strong patterns. Hence, $r_{\theta_{post}}$ should be large. This gives condition 1: $\theta < r_{\theta_{post}}/2$. Second, protein sharing is most important relative to the overwriting terms when the presynaptic overlap is minimized. Hence, condition 2: $\theta > r_{\theta_{pre}}/2$. These first two conditions set upper and lower limits for possible ranges of θ . In order to exploit dendritic PRP compartmentalization a further constraint for the dendritic correlation should vary in the range between these two limits: $r_{\theta_{pre}}/2 < \lambda < r_{\theta_{post}}/2$.

For the simulations presented in Figure 5 we used the following parameters: $N_{pre} = 100$, $N_{post} = 100$, $d = 50$, $\lambda = 0.1 \times \theta_{max}$.

Calculation and simulations of pattern overlap and memory generalization during sleep

We studied a novel two-step model for generalization of memories during sleep based on STC mechanisms, using both analytical calculation for the mean behavior and Monte-Carlo simulation for specific realizations of the model. For the simulations, we modeled a two-layer feedforward network of binary neurons (representing two cortical cell populations) where each presynaptic neuron ($N_{pre} = 1000$) was randomly connected to a subset of the postsynaptic neurons ($N_{post} = 1000$) with a fixed low probability ($p_{conn} = 0.1$). Each synapse had an initial weight of 1 (arbitrary units). A hippocampus-driven ‘strong’ activity pattern was initiated in the presynaptic layer by randomly choosing a subset of the neurons to be active, with a specified level of sparsity ($f_{pre} = 0.1$). A postsynaptic neuron was then activated if it received synaptic inputs from a sufficient number of activated neurons in the presynaptic layer. The threshold for activation Θ was chosen such that the expected pattern sparsity in the postsynaptic layer was close to a specified level ($f_{post} = 0.1$) by using the hypergeometric distribution, as follows (O’Reilly and McClelland, 1994). The probability for any postsynaptic neuron to receive m active synapses $p(m)$ is given by:

$$p(m) = \frac{\binom{N_{pre} f_{pre}}{m} \binom{N_{pre} - N_{pre} f_{pre}}{N_{pre} p_{conn} - m}}{\binom{N_{pre}}{N_{pre} p_{conn}}}. \quad \text{The threshold required to ensure a target level of postsynaptic layer sparsity } f_{post} \text{ can then be calculated from the cumulative of this}$$

distribution $P(m) = \sum_{m_i=0}^m p(m_i)$ by finding the smallest value of m where $P(m) \geq 1 - f_{post}$. After presenting the strong stimulus, we simulated potentiation of all of the activated synapses by increasing their weights to $\eta p_{strong} = 2$, and simulated heterosynaptic

depression at all synapses from non-activated presynaptic neurons onto the activated postsynaptic neurons by setting their weights to $\eta_D = 1/2$. We then sequentially presented weak patterns with the same presynaptic sparsity as the strong pattern ($f_{pre} = 0.1$). Each weak pattern caused potentiation of its activated synapses that were not part of the strong pattern, simulated by setting their weights to $\eta_{P_{weak}} = 2$. Synapses that were shared between the strong pattern and weak pattern were left unchanged at $\eta_{P_{strong}} = 2$. The weak patterns were not chosen by randomly selecting a subset of the presynaptic neurons, but instead were selected to have a desired overlap with the strong pattern in order to demonstrate the behavior of the model (Figure 7B-E). This selection process was necessary because for the parameters investigated a weak pattern chosen by chance was extremely unlikely to have substantial overlap with the strong pattern. The mean synaptic drive D from a pattern (Figure 7C) was calculated as the sum of the activated synaptic weights onto a single postsynaptic neuron that was part of the strong pattern, averaged across all activated

postsynaptic neurons: $D = \left\langle \sum_{i=1}^{N_{pre}} w_i x_i \right\rangle_{N_{postactive}}$ where w_i is the weight of the synapse from the i^{th} neuron and x_i is the binary state of the i^{th} neuron's activity (0 or 1). It is a measure of how likely it is that that pattern would activate a postsynaptic neuron.

For the analytical calculations of the mean behavior of the same model we followed the method of (O'Reilly and McClelland, 1994), who showed using further employment of the hypergeometric distribution how to calculate the probability that a previously activated postsynaptic neuron was reactivated by a second pattern, as a function of the presynaptic population overlap of the two patterns (Figure 7B). From this quantity we calculated the fraction of synapses in each group (weak-only pre to shared post, shared pre to shared post), plotted in Figure 7D-E.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Monika Jadi, Clare Puddifoot and Romain Veltz for comments on an earlier version of the manuscript. This research was supported by funding from the National Institute of Health grant number R01 NS059740 (CO'D, TJS), Swartz Foundation (CO'D), Howard Hughes Medical Institute (CO'D, TJS) and FRAXA Research Foundation (CO'D).

References

- Abraham WC. Metaplasticity: tuning synapses and networks for plasticity. *Nat Rev Neurosci.* 2008; 9:387. [PubMed: 18401345]
- Alarcon JM, Barco A, Kandel ER. Capture of the late phase of long-term potentiation within and across the apical and basilar dendritic compartments of CA1 pyramidal neurons: synaptic tagging is compartment restricted. *J Neurosci.* 2006; 26:256–264. [PubMed: 16399695]
- Ballarini F, Moncada D, Martinez MC, Alen N, Viola H. Behavioral tagging is a general mechanism of long-term memory formation. *Proc Natl Acad Sci USA.* 2009; 106:14599–14604. [PubMed: 19706547]

- Barco A, Lopez de Armentia M, Alarcon JM. Synapse-specific stabilization of plasticity processes: the synaptic tagging and capture hypothesis revisited 10 years later. *Neurosci Biobehav Rev.* 2008; 32:831–851. [PubMed: 18281094]
- Barrett AB, Billings GO, Morris RGM, van Rossum MCW. State based model of long-term potentiation and synaptic tagging and capture. *PLoS Comput Biol.* 2009; 5:e1000259. [PubMed: 19148264]
- Brown R, Kulik J. Flashbulb memories. *Cognition.* 1977; 5:73–99.
- Buzsáki G. The hippocampo-neocortical dialogue. *Cereb Cortex.* 1996; 6:81–92. [PubMed: 8670641]
- Cai DJ, Mednick SA, Harrison EM, Kanady JC, Mednick SC. REM, not incubation, improves creativity by priming associative networks. *Proc Natl Acad Sci USA.* 2009; 106:10130–10134. [PubMed: 19506253]
- Chen X, Leischner U, Rochefort NL, Nelken I, Konnerth A. Functional mapping of single spines in cortical neurons in vivo. *Nature.* 2011; 475:501–505. [PubMed: 21706031]
- Churchland PS, Sejnowski TJ. *The Computational Brain* (Mit Press). 1994
- Clopath C, Ziegler L, Vasilaki E, Büsing L, Gerstner W. Tag-trigger-consolidation: a model of early and late long-term-potentiation and depression. *PLoS Comput Biol.* 2008; 4:e1000248. [PubMed: 19112486]
- Connor SA, Hoeffler CA, Klann E, Nguyen PV. Fragile X mental retardation protein regulates heterosynaptic plasticity in the hippocampus. *Learn Mem.* 2011; 18:207–220. [PubMed: 21430043]
- Davis HP, Squire LR. Protein synthesis and memory: a review. *Psychol Bull.* 1984; 96:518–559. [PubMed: 6096908]
- Diniz CAR, Tutia MH, Leite JG. Bayesian analysis of a correlated binomial model. *Brazilian Journal of Probability and Statistics.* 2010; 24:68–77.
- Dombeck DA, Harvey CD, Tian L, Looger LL, Tank DW. Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. *Nat Neurosci.* 2010; 13:1433–1440. [PubMed: 20890294]
- Ellenbogen JM, Hu PT, Payne JD, Titone D, Walker MP. Human relational memory requires time and sleep. *Proc. Natl. Acad. Sci. U.S.A.* 2007; 104:7723–7728. [PubMed: 17449637]
- Fosse MJ, Fosse R, Hobson JA, Stickgold RJ. Dreaming and episodic memory: a functional dissociation? *J Cogn Neurosci.* 2003; 15:1–9. [PubMed: 12590838]
- Frey U, Morris RG. Synaptic tagging and long-term potentiation. *Nature.* 1997; 385:533–536. [PubMed: 9020359]
- Frey U, Schollmeier K, Reymann KG, Seidenbecher T. Asymptotic hippocampal long-term potentiation in rats does not preclude additional potentiation at later phases. *Neuroscience.* 1995; 67:799–807. [PubMed: 7675206]
- Gais S, Plihal W, Wagner U, Born J. Early sleep triggers memory for early visual discrimination skills. *Nat Neurosci.* 2000; 3:1335–1339. [PubMed: 11100156]
- Govindarajan A, Israely I, Huang S-Y, Tonegawa S. The dendritic branch is the preferred integrative unit for protein synthesis-dependent LTP. *Neuron.* 2011; 69:132–146. [PubMed: 21220104]
- Govindarajan A, Kelleher RJ, Tonegawa S. A clustered plasticity model of long-term memory engrams. *Nat Rev Neurosci.* 2006; 7:575–583. [PubMed: 16791146]
- Hernandez PJ, Abel T. The role of protein synthesis in memory consolidation: progress amid decades of debate. *Neurobiol Learn Mem.* 2008; 89:293–311. [PubMed: 18053752]
- Jia H, Rochefort NL, Chen X, Konnerth A. Dendritic organization of sensory input to cortical neurons in vivo. *Nature.* 2010; 464:1307–1312. [PubMed: 20428163]
- Kelleher RJ, Bear MF. The autistic neuron: troubled translation? *Cell.* 2008; 135:401–406. [PubMed: 18984149]
- Kelleher RJ, Govindarajan A, Tonegawa S. Translational regulatory mechanisms in persistent forms of synaptic plasticity. *Neuron.* 2004; 44:59–73. [PubMed: 15450160]
- Kleindienst T, Winnubst J, Roth-Alpermann C, Bonhoeffer T, Lohmann C. Activity-dependent clustering of functional synaptic inputs on developing hippocampal dendrites. *Neuron.* 2011; 72:1012–1024. [PubMed: 22196336]

- Krug M, Lössner B, Ott T. Anisomycin blocks the late phase of long-term potentiation in the dentate gyrus of freely moving rats. *Brain Res. Bull.* 1984; 13:39–42. [PubMed: 6089972]
- Lewis PA, Durrant SJ. Overlapping memory replay during sleep builds cognitive schemata. *Trends Cogn Sci.* 2011; 15:343–351. [PubMed: 21764357]
- Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser M-B. Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science.* 2005; 309:619–623. [PubMed: 16040709]
- Leutgeb S, Leutgeb JK, Treves A, Moser M-B, Moser EI. Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science.* 2004; 305:1295–1298. [PubMed: 15272123]
- Macke JH, Berens P, Ecker AS, Tolias AS, Bethge M. Generating spike trains with specified correlation coefficients. *Neural Comput.* 2009; 21:397–423. [PubMed: 19196233]
- Mackler SA, Brooks BP, Eberwine JH. Stimulus-induced coordinate changes in mRNA abundance in single postsynaptic hippocampal CA1 neurons. *Neuron.* 1992; 9:539–548. [PubMed: 1388031]
- Makino H, Malinow R. Compartmentalized versus Global Synaptic Plasticity on Dendrites Controlled by Experience. *Neuron.* 2011; 72:1001–1011. [PubMed: 22196335]
- Maquet P, Laureys S, Peigneux P, Fuchs S, Petiau C, Phillips C, Aerts J, Del Fiore G, Degueldre C, Meulemans T, et al. Experience-dependent changes in cerebral activation during human REM sleep. *Nat Neurosci.* 2000; 3:831–836. [PubMed: 10903578]
- Martin KC, Casadio A, Zhu H, Yaping E, Rose JC, Chen M, Bailey CH, Kandel ER. Synapse-specific, long-term facilitation of aplysia sensory to motor synapses: a function for local protein synthesis in memory storage. *Cell.* 1997; 91:927–938. [PubMed: 9428516]
- Mednick S, Nakayama K, Stickgold R. Sleep-dependent learning: a nap is as good as a night. *Nat Neurosci.* 2003; 6:697–698. [PubMed: 12819785]
- Merhav M, Rosenblum K. Facilitation of taste memory acquisition by experiencing previous novel taste is protein-synthesis dependent. *Learn Mem.* 2008; 15:501–507. [PubMed: 18626094]
- Moncada D, Viola H. Induction of long-term memory by exposure to novelty requires protein synthesis: evidence for a behavioral tagging. *J Neurosci.* 2007; 27:7476–7481. [PubMed: 17626208]
- Nakashiba T, Young JZ, McHugh TJ, Buhl DL, Tonegawa S. Transgenic inhibition of synaptic transmission reveals role of CA3 output in hippocampal learning. *Science.* 2008; 319:1260–1264. [PubMed: 18218862]
- O'Carroll CM, Morris RGM. Heterosynaptic co-activation of glutamatergic and dopaminergic afferents is required to induce persistent long-term potentiation. *Neuropharmacology.* 2004; 47:324–332. [PubMed: 15275821]
- O'Reilly RC, McClelland JL. Hippocampal conjunctive encoding, storage, and recall: avoiding a trade-off. *Hippocampus.* 1994; 4:661–682. [PubMed: 7704110]
- Otani S, Marshall CJ, Tate WP, Goddard GV, Abraham WC. Maintenance of long-term potentiation in rat dentate gyrus requires protein synthesis but not messenger RNA synthesis immediately post-tetanzation. *Neuroscience.* 1989; 28:519–526. [PubMed: 2710327]
- Pace-Schott EF, Milad MR, Orr SP, Rauch SL, Stickgold R, Pitman RK. Sleep promotes generalization of extinction of conditioned fear. *Sleep.* 2009; 32:19–26. [PubMed: 19189775]
- Pavlovsky A, Alarcon JM. Interaction between long-term potentiation and depression in CA1 synapses: temporal constrains, functional compartmentalization and protein synthesis. *PLoS ONE.* 2012; 7:e29865. [PubMed: 22272255]
- Payne JD, Schacter DL, Propper RE, Huang L-W, Wamsley EJ, Tucker MA, Papper M, Kempter R, Leibold C. Synaptic tagging, evaluation of memories, and the distal reward problem. *Learn Mem.* 2011; 18:58–70. [PubMed: 21191043]
- Rasch B, Born J. About Sleep's Role in Memory. *Physiol Rev.* 2013; 93:681–766. [PubMed: 23589831]
- Redondo RL, Morris RGM. Making memories last: the synaptic tagging and capture hypothesis. *Nat Rev Neurosci.* 2011; 12:17–30. [PubMed: 21170072]
- Reymann KG, Frey JU. The late maintenance of hippocampal LTP: requirements, phases, 'synaptic tagging', 'late-associativity' and implications. *Neuropharmacology.* 2007; 52:24–40. [PubMed: 16919684]

- Ribeiro S, Gervasoni D, Soares ES, Zhou Y, Lin S-C, Pantoja J, Lavine M, Nicolelis MAL. Long-lasting novelty-induced neuronal reverberation during slow-wave sleep in multiple forebrain areas. *PLoS Biol.* 2004; 2:E24. [PubMed: 14737198]
- Sajikumar S, Frey JU. Late-associativity, synaptic tagging, and the role of dopamine during LTP and LTD. *Neurobiol Learn Mem.* 2004; 82:12–25. [PubMed: 15183167]
- Sajikumar S, Navakkode S, Frey JU. Identification of compartment- and process-specific molecules required for “synaptic tagging” during long-term potentiation and long-term depression in hippocampal CA1. *J Neurosci.* 2007; 27:5068–5080. [PubMed: 17494693]
- Schacter DL, Addis DR, Buckner RL. Remembering the past to imagine the future: the prospective brain. *Nat Rev Neurosci.* 2007; 8:657–661. [PubMed: 17700624]
- Shires KL, Da Silva BM, Hawthorne JP, Morris RGM, Martin SJ. Synaptic tagging and capture in the living rat. *Nat Commun.* 2012; 3:1246. [PubMed: 23212375]
- Siapas AG, Wilson MA. Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron.* 1998; 21:1123–1128. [PubMed: 9856467]
- Smolen P, Baxter DA, Byrne JH. Molecular constraints on synaptic tagging and maintenance of long-term potentiation: a predictive model. *PLoS Comput Biol.* 2012; 8:e1002620. [PubMed: 22876169]
- Stickgold R, Whidbee D, Schirmer B, Patel V, Hobson JA. Visual discrimination task improvement: A multi-step process occurring during sleep. *J Cogn Neurosci.* 2000; 12:246–254. [PubMed: 10771409]
- Stickgold R, Walker MP. Sleep-dependent memory triage: evolving generalization through selective processing. *Nat Neurosci.* 2013; 16:139–145. [PubMed: 23354387]
- Sutton MA, Schuman EM. Dendritic protein synthesis, synaptic plasticity, and memory. *Cell.* 2006; 127:49–58. [PubMed: 17018276]
- Takahashi N, Kitamura K, Matsuo N, Mayford M, Kano M, Matsuki N, Ikegaya Y. Locally synchronized synaptic inputs. *Science.* 2012; 335:353–356. [PubMed: 22267814]
- Tenenbaum JB, Kemp C, Griffiths TL, Goodman ND. How to grow a mind: statistics, structure, and abstraction. *Science.* 2011; 331:1279–1285. [PubMed: 21393536]
- Vazdarjanova A, Guzowski JF. Differences in hippocampal neuronal population responses to modifications of an environmental context: evidence for distinct, yet complementary, functions of CA3 and CA1 ensembles. *J Neurosci.* 2004; 24:6489–6496. [PubMed: 15269259]
- Vazdarjanova A, McNaughton BL, Barnes CA, Worley PF, Guzowski JF. Experience-dependent coincident expression of the effector immediate-early genes *arc* and *Homer 1a* in hippocampal and neocortical neuronal networks. *J Neurosci.* 2002; 22:10067–10071. [PubMed: 12451105]
- Wagner U, Gais S, Haider H, Verleger R, Born J. Sleep inspires insight. *Nature.* 2004; 427:352–355. [PubMed: 14737168]
- Wamsley EJ, Perry K, Djonlagic I, Reaven LB, Stickgold R. Cognitive replay of visuomotor learning at sleep onset: temporal dynamics and relationship to task performance. *Sleep.* 2010; 33:59–68. [PubMed: 20120621]
- Wang DO, Kim SM, Zhao Y, Hwang H, Miura SK, Sossin WS, Martin KC. Synapse- and stimulus-specific local translation during long-term neuronal plasticity. *Science.* 2009; 324:1536–1540. [PubMed: 19443737]
- Wang S-H, Redondo RL, Morris RGM. Relevance of synaptic tagging and capture to the persistence of long-term potentiation and everyday spatial memory. *Proc Natl Acad Sci USA.* 2010; 107:19537–19542. [PubMed: 20962282]
- Wierzynski CM, Lubenov EV, Gu M, Siapas AG. State-dependent spike-timing relationships between hippocampal and prefrontal circuits during sleep. *Neuron.* 2009; 61:587–596. [PubMed: 19249278]
- Wilson MA, McNaughton BL. Reactivation of hippocampal ensemble memories during sleep. *Science.* 1994; 265:676–679. [PubMed: 8036517]
- Wiltgen BJ, Silva AJ. Memory for context becomes less specific with time. *Learn Mem.* 2007; 14:313–317. [PubMed: 17522020]

- Winocur G, Moscovitch M, Sekeres M. Memory consolidation or transformation: context manipulation and hippocampal representations of memory. *Nat Neurosci.* 2007; 10:555–557. [PubMed: 17396121]
- Zoghbi HY, Bear MF. Synaptic Dysfunction in Neurodevelopmental Disorders Associated with Autism and Intellectual Disabilities. *Cold Spring Harb Perspect Biol.* 2012

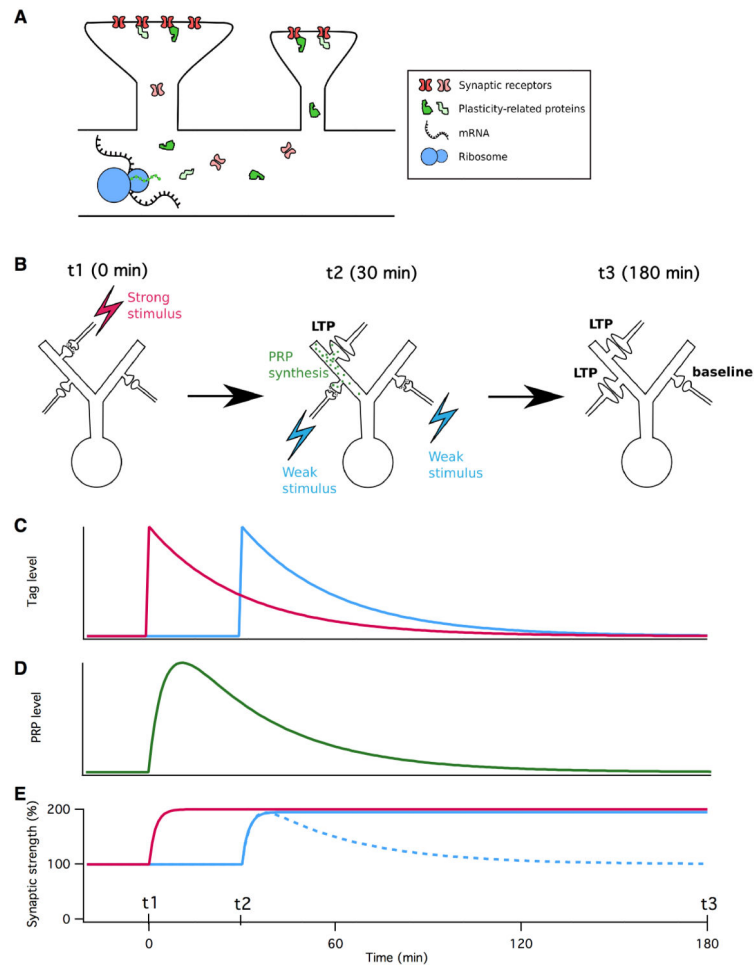
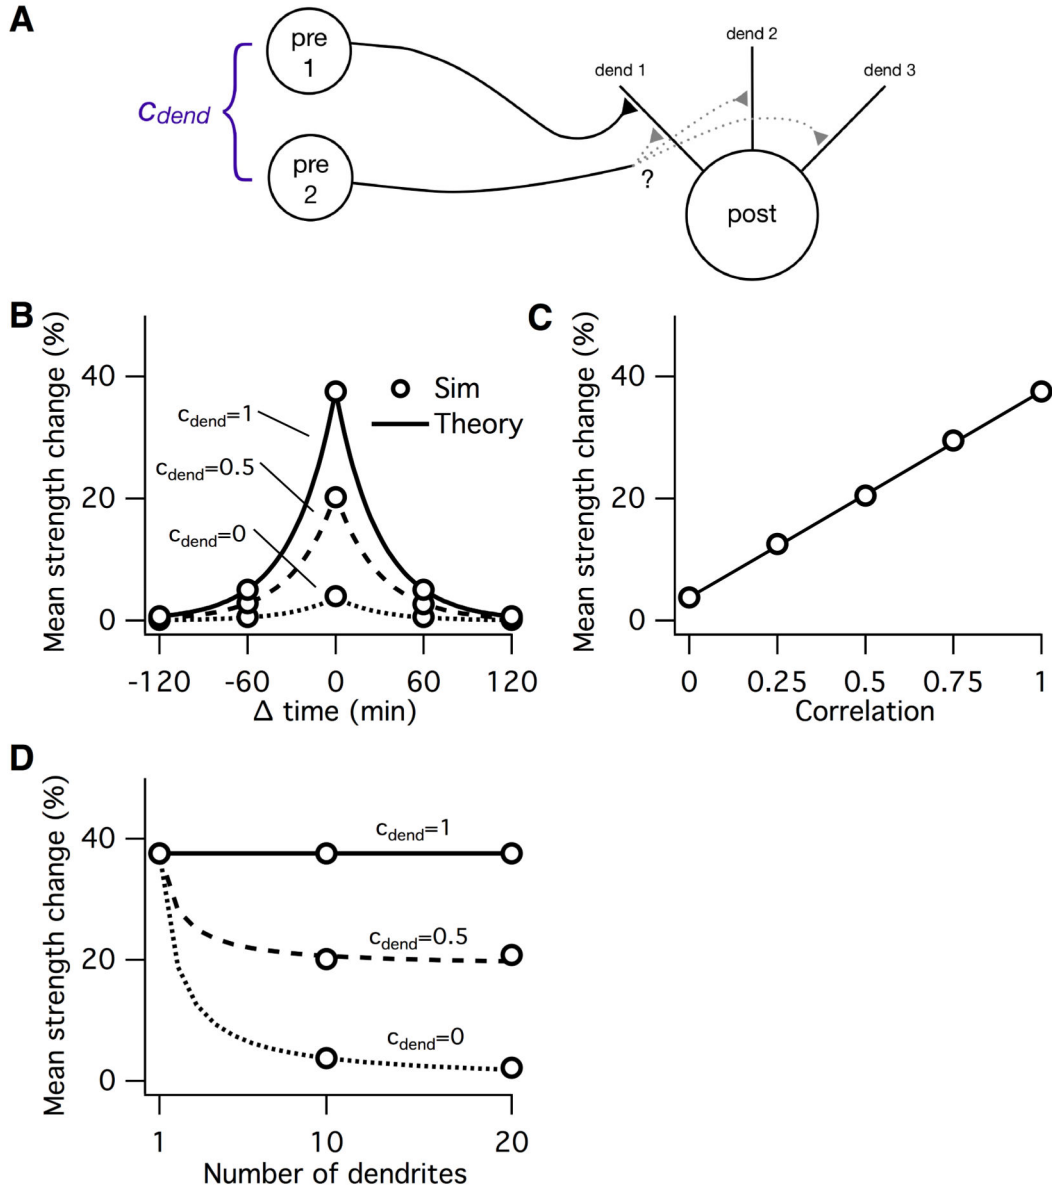


Figure 1.

Dendritic protein translation and synaptic tagging. **A:** Schematic of local protein translation in dendrites. Dendritic ribosomes (blue) transcribes mRNA (black) to synthesize the assorted synaptic receptors (red), plasticity-related proteins (green) and other proteins (not shown) that comprise and regulate nearby synapses. **B:** Cartoon of molecular events during STC in a neuron with two dendrites and three synapses. **C:** Activity level of molecular ‘tag’ at the strongly (magenta) and weakly (blue) activated synapses vs time. **D:** PRP level in the dendrite of the strongly activated synapse vs time. **E:** Synaptic strength vs time for strongly activated synapse (magenta), weakly activated synapse on same dendrite (solid blue) and weakly activated synapse on different dendrite (dashed blue). Labeled points on the time axis (t1, t2, t3) correspond to three illustrations in A.

**Figure 2.**

A: A single postsynaptic neuron innervated by two presynaptic neurons. The degree of bias for the two neurons to synapse onto the same dendrite can be captured by a parameter c_{dend} . **B:** Mean consolidated strength change of a weakly stimulated synapse as a function of the time interval with the strong stimulus. Circle symbols are results from simulations, curves are predictions from theory (see Experimental Procedures). Different curves denote different values of c_{dend} . **C:** Mean consolidated strength change of a weakly stimulated synapse as a function of the weak presynaptic neuron's correlation with the strongly activated presynaptic neuron, for a time interval of zero min. **D:** Mean consolidated strength change of a weakly stimulated synapse as a function of the number of dendrites on the postsynaptic neuron. Different curves denote different values of c_{dend} , as indicated in figure.

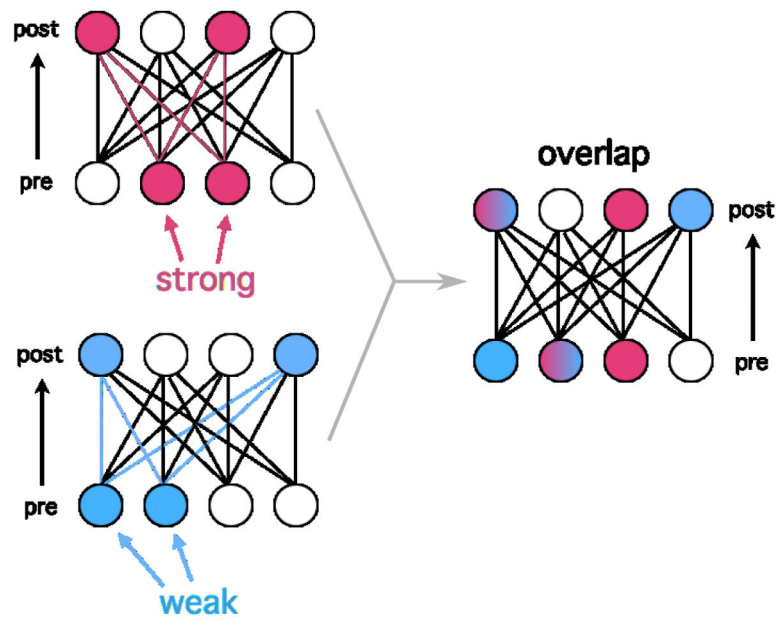


Figure 3.

Two-layer feedforward neuronal network. A strong (top left) or weak (bottom left) activity pattern covers a specific subset the pre and postsynaptic populations. The statistics of pre and postsynaptic overlap (right) determines the efficacy of STC (see Text).

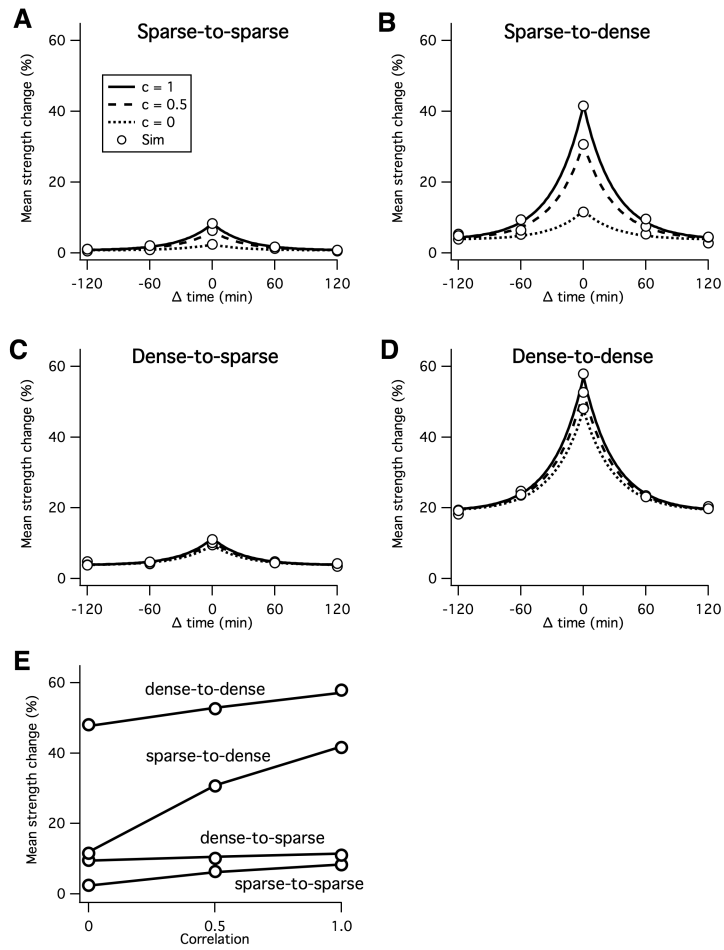


Figure 4.

Mean consolidated synaptic strength of a weak pattern as a function of its temporal interval with a strong pattern for different levels of pre- and post-synaptic pattern sparsity (**A-D**) and three levels of dendritic correlation between synapses (different curves in each subfigure). **E:** The mean synaptic strength change as a function of dendritic correlation for a time interval of zero (data replotted from A-D).

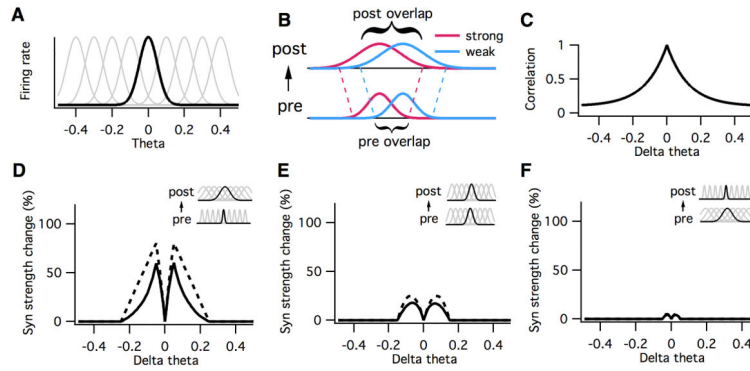


Figure 5.

Weak pattern consolidation in a two-layer network where pre- and post-synaptic populations code for the same variable. **A:** Each neuron in a layer responds maximally for its preferred value of the stimulus. Neurons are ordered according to their preference and tile the entire 1-D stimulus space (gray curves). **B:** The degree of overlap in the set of neurons activated by a strong (magenta) and weak (blue) stimulus is determined by the distance between the stimulus values. **C:** The dendritic correlation (y-axis) between any two neurons decays as a function of the difference between their preferred stimulus values (x-axis). **D-F:** The mean consolidated synaptic strength change from a weak stimulus (y-axis) as a function of the difference between its value and the value of the strong stimulus (x-axis). Solid curves are from case where dendritic correlation follows figure C, dashed curves are from case where there is only one postsynaptic dendrite. DF represent different relative widths for the pre and postsynaptic tuning curves (insets).

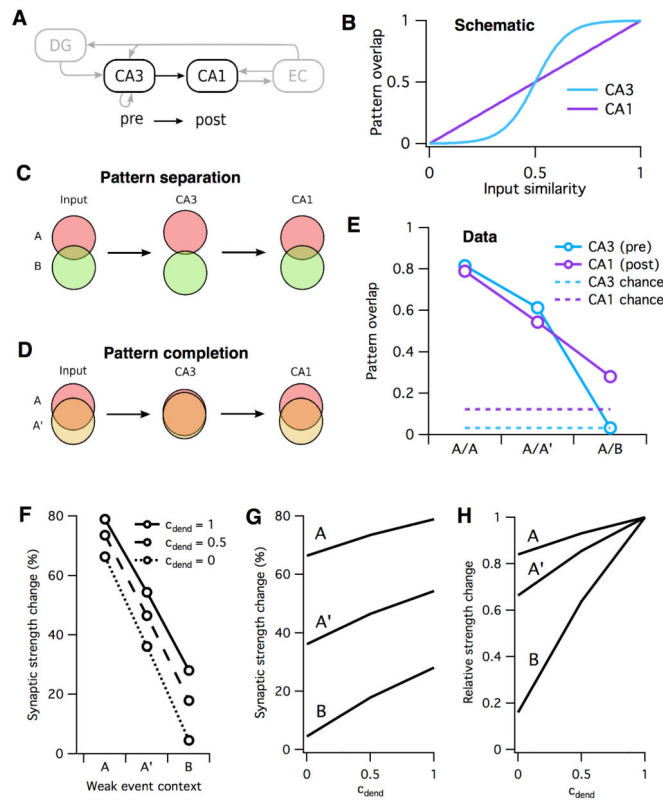


Figure 6.

Weak pattern consolidation in a model of rodent hippocampal Schaffer collateral projection (CA3-CA1).

A: Schematic diagram of the hippocampal circuit (DG: dentate gyrus, EC: entorhinal cortex). We focused on modeling the CA3-CA1 (pre-post) projection. **B:** Schematic of pattern separation (for dissimilar inputs) and pattern completion (for similar inputs) by CA3, but preservation of pattern similarity by CA1. **C:** Pattern separation for dissimilar environments by CA3, but environment differences are preserved in CA1 representation. **D:** Pattern completion for similar environments by CA3, but environment similarities are preserved in CA1. **E:** Experimentally measured (Vazdarjanova and Guzowski, 2004) activity pattern overlap in CA3 and CA1 for three different spatial context comparisons: same (A/A), minor change (A/A') or major change (A/B). **F:** Expected synaptic strength change of a weak pattern representing context 2 prior to a strong pattern representing context 1. Three curves refer to varying levels of dendritic correlation of weak pattern CA3 neurons with strong pattern CA3 neurons. **G-H:** Expected synaptic strength change for a weak event in either context A, A' or B (different curves) as a function of mean dendritic correlation of weak pattern neurons with strong pattern neurons. Panel G plots synaptic strength change relative to baseline, while H plots synaptic strength change relative to maximum at $c_{dend}=1$. Data replotted from F.

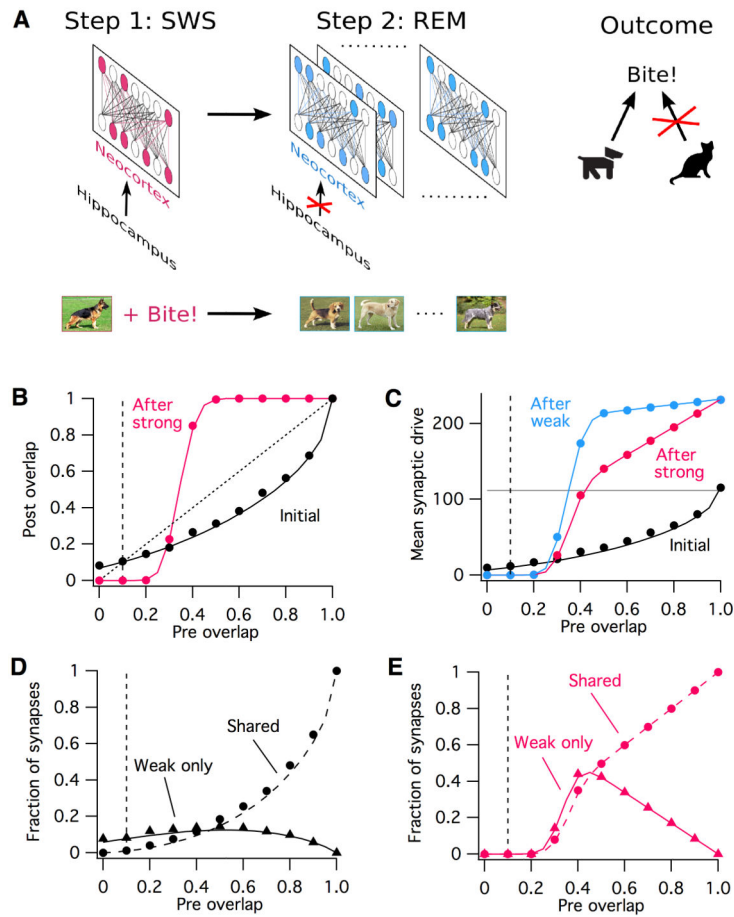


Figure 7.

STC and noisy replay as a mechanism for generalization of synaptic learning.

A: Schematic diagram of the model of memory generalization during sleep (see Results).





B: Probability of postsynaptic neuron from strong pattern being active as a function of weak pattern overlap. Black: from initial network. Magenta: from network following strong pattern plasticity. In this and subsequent subfigures, solid curves indicate predictions from theory, while symbols indicate simulation results. Diagonal dashed line indicates identity. Vertical dashed line indicates chance overlap level.

C: Mean strength of activated synapses onto strong pattern postsynaptic neurons as a function of overlap. Black: from initial network. Magenta: from network after strong pattern plasticity. Blue: from network after strong and weak pattern plasticity. Horizontal gray line indicates mean neuron spike threshold.

D-E: Fraction of synapses in weak-only (solid curve, triangles) and shared (dashed curve, circles) presynaptic categories as a function of weak pattern overlap with strong pattern. Data in D (black) calculated from initial network, data in E (magenta) from network after strong pattern plasticity.

Table 1

Synapses in a weak pattern can be separated into four groups depending on their overlap with those in a strong pattern.

<i>Symbols</i>	<i>Connection</i>	$\langle w \rangle$
	<i>pre weak-only</i> → <i>post weak-only</i>	0
	<i>pre shared</i> → <i>post weak-only</i>	0
	<i>pre weak-only</i> → <i>post shared</i>	$\alpha e^{-\frac{ t }{\tau}} p(N_{strong} - 1)$
	<i>pre shared</i> → <i>post shared</i>	$\langle w \rangle_{over}$

Each row corresponds to one of the four groups of synapses that make up a weak memory pattern, depending on their overlap with a strong pattern. The first column shows each group pictorially according to the color scheme from Figure 3. The pair of filled circles represents a pair of neurons from the presynaptic and postsynaptic populations, with blue representing neurons that are only in the weak pattern, and magenta-blue those that are part of both the strong and weak patterns. The second column describes the synapse group in words, while the third column gives the mean synaptic weight change, $\langle w \rangle$, of that group. See Results for details.