

Original Article

The penumbra of learning: A statistical theory of synaptic tagging and capture

SAMUEL J. GERSHMAN

Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA

(Received 28 August 2013; revised 15 October 2013; accepted 2 November 2013)

Abstract

Learning in humans and animals is accompanied by a penumbra: Learning one task benefits from learning an unrelated task shortly before or after. At the cellular level, the penumbra of learning appears when weak potentiation of one synapse is amplified by strong potentiation of another synapse on the same neuron during a critical time window. Weak potentiation sets a molecular tag that enables the synapse to capture plasticity-related proteins synthesized in response to strong potentiation at another synapse. This paper describes a computational model which formalizes synaptic tagging and capture in terms of statistical learning mechanisms. According to this model, synaptic strength encodes a probabilistic inference about the dynamically changing association between pre- and post-synaptic firing rates. The rate of change is itself inferred, coupling together different synapses on the same neuron. When the inputs to one synapse change rapidly, the inferred rate of change increases, amplifying learning at other synapses.

Keywords: *Synaptic plasticity, synaptic tagging and capture, Kalman filter*

The individual synapse has long served as a microcosm for studying the cellular basis of memory formation. In particular, memory formation is commonly believed to rely on long-term potentiation (LTP), the sustained increase in synaptic strength induced by repeated high-frequency stimulation of a neuron (Martin et al., 2000). However, research over the last two decades has demonstrated that memory formation at one synapse can depend strongly on what happens at other synapses.

Correspondence: Samuel J. Gershman, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 77 Massachusetts Ave., Room 46-4053, Cambridge, MA, 02139. E-mail: sjgershm@mit.edu

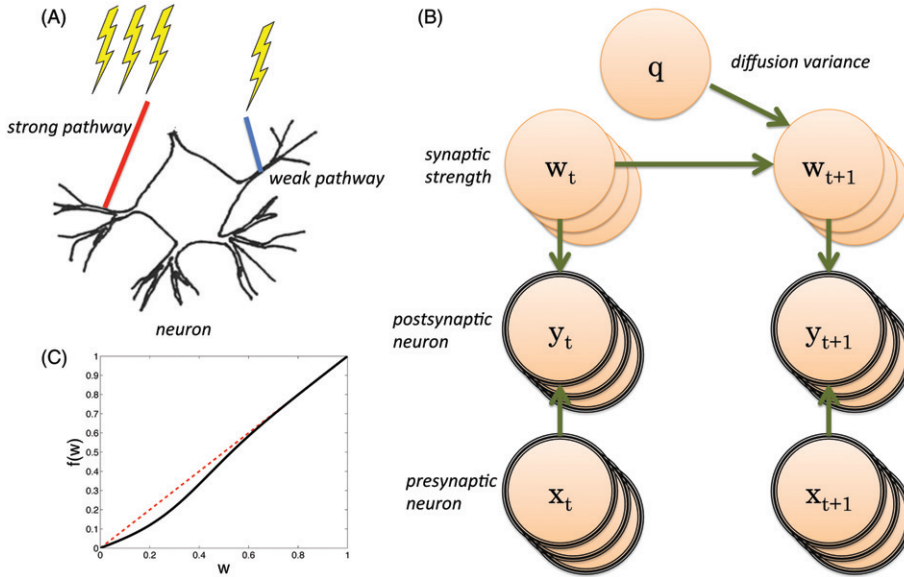


Figure 1. Model schematic. (A) Experimental setup: a neuron is tetanized at two different synaptic sites. The strong pathway (red) is subjected to repeated trains of high-frequency stimulation, while the weak pathway (blue) is subjected to a single train. (B) Graphical model showing two time slices of the generative process. The model consists of multiple synapses, all of which share the same diffusion variance q . Observed variables are indicated by a black border. (C) The transition function $f(w)$ of synaptic strength as a function of the current strength w . (A color version of this figure is available in the online edition of this article.)

Frey and Morris, (1997) showed that weak tetanic stimulation of a set of synapses, or strong stimulation in the presence of a protein-synthesis inhibitor, induces an early-phase of LTP (E-LTP) which decays over the course of hours, but can be transformed into sustained, late-phase LTP (L-LTP) if preceded by strong stimulation of an independent set of synapses in the same population of neurons (Figure 1A). Capture of plasticity-related proteins by a weakly stimulated synapse can also occur if strong stimulation of another synapse occurs *after* the weak stimulation (Frey and Morris, 1998). These discoveries have led to the synaptic tagging and capture hypothesis (Martin and Kosik, 2002; Redondo and Morris, 2011), according to which stimulation sets a synaptic tag that enables the synapse to capture plasticity-related proteins synthesized at another synapse.

Synaptic tagging and capture provides a mechanism by which learning at one synapse casts a penumbra over other synapses, facilitating learning. This penumbra is accompanied by a behavioral correlate: Learning one task is enhanced by learning an unrelated task shortly before or after, an observation replicated across several different paradigms and species (Moncada and Viola, 2007; Merhav and Rosenblum, 2008; Wang et al., 2010; Duncan et al., 2012). For example, learning in a variety of tasks (spatial object recognition, contextual fear conditioning, conditioned taste aversion) can be enhanced by allowing animals to explore a novel spatial environment before or after training (Ballarini et al., 2009).

This paper presents a computational theory of synaptic tagging and capture, formalizing its role in a statistical learning system. The central idea is that each synapse estimates a time-varying association between the firing rates of pre- and

post-synaptic neurons. The rate at which the association changes over time is itself unknown and must be inferred. When change is inferred to be faster, the learning rate is increased. Importantly, the rate of change is shared across multiple synapses at a single neuron, thereby coupling the synapses together within a penumbra of learning: A change in one association is evidence for change in all the associations (although not necessarily in the same direction). The availability of plasticity-related proteins, on this view, signals the inferred rate of change, with the diffusion of proteins between synapses propagating this inference across the probabilistic model.

Methods

Generative model

We begin by describing a probabilistic generative model of the random variables represented by a neural population, schematized in Figure 1B. For simplicity, we will consider two real-valued random variables (x and y), but the full theory deals with multiple pairs of random variables. The random variables are represented by the firing rates of two synaptically connected neurons, with x_t denoting the firing rate of the pre-synaptic neuron at time t , and y_t denoting the firing rate of the post-synaptic neuron. The relationship between x_t and y_t is governed by a time-varying association w_t . When $w_t > 0$, pre-synaptic firing tends to evoke post-synaptic firing, and when $w_t < 0$, pre-synaptic firing tends to suppress post-synaptic firing.

Two processes govern the evolution of w_t over time. The first is a decay towards 0 specified by a transition function $f(w)$, shown in Figure 1C. The transition function expresses a strength-dependent decay, such that stronger associations decay less than weak associations. In other words, this amounts to the assumption that strong relationships between variables tend to persist, while weak relationships tend to dissipate quickly. From a neurobiological perspective, this transition function induces a form of cellular consolidation (McGaugh, 2000): If a synaptic weight is sufficiently strong, it will persist almost indefinitely (i.e., it will be consolidated), whereas a weak synaptic weight will decay back to baseline.

The second process governing the evolution of w_t over time is a Gaussian diffusion process determined by the diffusion variance q . Larger values of q produce faster rates of change. Crucially, we assume that the diffusion variance is unknown, and that when there are multiple synapses on the same post-synaptic neuron, the diffusion variance is shared across all of them. This means that synapses are coupled not in terms of their specific associations, but rather in terms of the rate at which the associations change. As a consequence, observing a change in one synapse provides information that other synapses have changed as well (but possibly in a different direction).

Formally, the assumptions described above are implemented by the following stochastic dynamics:

$$w_t = f(w_{t-1}) + \epsilon_w \quad (1)$$

$$y_t = w_t x_t + \epsilon_y, \quad (2)$$

where $w_0 = 0$, $\varepsilon_w \sim \mathcal{N}(0, q)$, and $\varepsilon_y \sim \mathcal{N}(0, r)$. The parameter $q > 0$ determines the rate of change, and $r > 0$ determines the observation noise. We used $r = 0.1$ for all our simulations.

The transition function $f(w)$ determines the rate at which w_t decays back to 0 over time. As discussed above, we want a transition function that decays asymptotically to 0, but with slower decay for stronger associations. The following functional form satisfies these desiderata:

$$f(w) = \frac{w}{1 + \exp\{-\beta w^2\}}, \quad (3)$$

where $\beta \geq 0$ is a parameter that governs the nonlinear relationship between w and the decay rate; in particular, smaller absolute values of w decay more rapidly than larger values. This function is shown in Figure 1C. We used $\beta = 10$ for all our simulations unless otherwise mentioned. The precise form of the transition function is not important, only that it exhibits strength-dependent decay.¹

We approximate the distribution over the diffusion variance q with a discrete set of points, $\{q^1, \dots, q^K\}$. Specifically, the diffusion variance is drawn according to $P(q = q^k) \propto \exp\{ak^{-b}\}$ where a and b are constants and q^k ranges over K equally spaced values between 0.1 and 1 (note that the k superscript is an index, not a power). This distribution embodies the assumption that a slow rate of change is *a priori* more probable. For all our simulations, we used $K = 50$, $a = 8$ and $b = 5$, but the results are not sensitive to small variations in these values. Further explorations of the parameter space can be found in Appendix B.

It is important to note that the generative model described here is not a description of the learning dynamics (which we describe in the next section). Although w_t evolves over time independently of x_t and y_t , the synaptic *estimate* of w_t is *not* independent of x_t and y_t . This is because the estimate of w_t is formulated in terms of the conditional distribution $P(w_t | x_{1:t}, y_{1:t})$.

Synaptic plasticity as statistical inference

Given the generative model described in the previous section, the computational problem is to infer both the associative weight w and the diffusion variance q after observing data $(\mathbf{x}_{1:t}, \mathbf{y}_{1:t})$. The optimal statistical inference is stipulated by Bayes' rule:

$$P(w_t, q | \mathbf{x}_{1:t}, \mathbf{y}_{1:t}) \propto P(y_t | x_t, w_t) P(q) P(w_t | \mathbf{x}_{1:t-1}, \mathbf{y}_{1:t-1}, q), \quad (4)$$

where $\mathbf{x}_{1:t}$ and $\mathbf{y}_{1:t}$ are the pre- and post-synaptic (respectively) histories from time 1 to t . When there are multiple synapses, the likelihood becomes a product over the different synapses: $P(y_t | x_t, w_t) = \prod_d P(y_{td} | x_{td}, w_{td})$, where d indexes synapses. For ease of exposition, we present the posterior computations for a single synapse.

The posterior distribution admits a simple approximation scheme. We first describe the exact Bayesian update equations when q is known—i.e., the computation of $P(w_t | \mathbf{x}_{1:t}, \mathbf{y}_{1:t}, q)$. If $f(w)$ were linear, then Bayes' rule could be implemented by the Kalman filtering equations (Kalman, 1960). However, the nonlinearity in $f(w)$ forces us to make an approximation; we follow the standard practice in engineering and use the extended Kalman filter (Anderson and Moore, 1979), which linearizes around the previous estimate, yielding recursive

updates for the parameters (mean and variance) of a Gaussian posterior over the weights (see Appendix A for more details). We denote this posterior by $P(w_t | \mathbf{x}_{1:t}, \mathbf{y}_{1:t}, q) = \mathcal{N}(w_t; \hat{w}_t, v_t)$, with mean and variance defined as follows:

$$\text{Posterior mean: } \hat{w}_t = f(\hat{w}_{t-1}) + \alpha_t x_t (y_t - \hat{w}_{t-1} x_t) \quad (5)$$

$$\text{Posterior variance: } v_t = (1 - \alpha_t x_t^2) \sigma_t^2 \quad (6)$$

where $\sigma_t^2 = [f'(\hat{w}_{t-1})]^2 v_{t-1} + q$ is the predictive variance after weight decay $f(\hat{w}_{t-1})$ but before observing (x_t, y_t) :

$$P(w_t | \mathbf{x}_{1:t-1}, \mathbf{y}_{1:t-1}, q) = \mathcal{N}(w_t; f(\hat{w}_{t-1}), \sigma_t^2). \quad (7)$$

The derivative of the transition function with respect to the weight, $f'(w)$, is given in Appendix A. When σ_t^2 is high, there is greater uncertainty about the associative strength, and therefore observing new data (x_t, y_t) will have a greater impact on the weight posterior.

The learning rate α_t is given by $\alpha_t = \sigma_t^2 / \lambda_t$, where $\lambda_t = x_t^2 \sigma_t^2 + r$ is the variance of the posterior marginal likelihood:

$$\begin{aligned} P(y_t | \mathbf{x}_{1:t-1}, \mathbf{y}_{1:t-1}, q) &= \int_{w_t} P(y_t | x_t, w_t) P(w_t | \mathbf{x}_{1:t-1}, \mathbf{y}_{1:t-1}, q) dw_t \\ &= \mathcal{N}(y_t; \hat{w}_{t-1} x_t, \lambda_t). \end{aligned} \quad (8)$$

Intuitively, λ_t expresses the overall unpredictability (or noisiness) of the data; this unpredictability grows monotonically with the noise variance. The expression for the learning rate given above indicates that learning will be faster to the extent that the weight uncertainty is high (large σ_t^2) relative to the overall unpredictability (low λ_t).

The learning rule for synaptic strength can be viewed as a form of predictive Hebbian learning (Montague and Sejnowski, 1994) in which the pre-synaptic neuron is associated with the prediction error $y_t - \hat{w}_{t-1} x_t$. Intuitively, $\hat{w}_{t-1} x_t$ represents a prediction of the post-synaptic firing rate, and the synaptic strength is increased when the post-synaptic firing rate is greater than expected, or decreased when less than expected. Another way to look at this learning rule is as a Hebbian rule with a sliding plasticity threshold $\hat{w}_{t-1} x_t$, similar to the Bienenstock-Cooper-Munro (BCM) theory (Bienenstock et al., 1982). According to BCM theory, the plasticity threshold increases monotonically with the previous activity of the post-synaptic cell; according to the theory presented here, the plasticity threshold increases monotonically with the *predicted* post-synaptic activity.

This learning rule has a number of neurobiologically plausible properties: (1) As mentioned above, the transition function induces a form of cellular consolidation (McGaugh, 2000), whereby sufficiently strong synapses become resistant to decay; (2) synapses are metaplastic (Abraham, 2008), with a plasticity threshold that adjusts dynamically based on the state of the synapse; (3) the learning rate adjusts in response to environmental volatility, as observed behaviorally and neurally (Behrens et al., 2007; Roesch et al., 2010). In particular, the learning rate α_t will increase when the uncertainty due to weight diffusion (expressed by σ_t^2) is high relative to the overall noisiness of the data (expressed by λ_t).

Inferring the diffusion variance

With q unknown, the posterior marginal over w_t is a mixture of Gaussians:

$$\begin{aligned} P(w_t | \mathbf{x}_{1:t}, \mathbf{y}_{1:t}, q) &= \sum_{k=1}^K P(q = q^k | \mathbf{x}_{1:t}, \mathbf{y}_{1:t}) P(w_t | \mathbf{x}_{1:t}, \mathbf{y}_{1:t}, q = q^k) \\ &= \sum_{k=1}^K P(q = q^k | \mathbf{x}_{1:t}, \mathbf{y}_{1:t}) \mathcal{N}(w_t; \hat{w}_t^k, v_t^k). \end{aligned} \quad (9)$$

where we have indexed \hat{w} and v by k to indicate their implicit dependence on $q = q^k$. To obtain $P(q = q^k | \mathbf{x}_{1:t}, \mathbf{y}_{1:t})$, we apply the update equations 5 and 6 for each q^k in parallel, and use Bayes' rule to compute the posterior over q :

$$\begin{aligned} P(q = q^k | \mathbf{x}_{1:t}, \mathbf{y}_{1:t}) &\propto P(q = q^k) P(\mathbf{y}_{1:t} | \mathbf{x}_{1:t}, q) \\ &= P(q = q^k) \prod_{\tau=1}^t \mathcal{N}(y_\tau; \hat{w}_{\tau-1}^k x_\tau, \lambda_\tau^k), \end{aligned} \quad (10)$$

where the second line was obtained by plugging in Eq. 8.

We can think of this model as positing for each synapse a set of K ‘‘micro-weights’’ that vary in their learning rates (α_t monotonically increases with q^k). This is similar to the idea that neurons have a ‘‘reservoir’’ of time constants (La Camera et al., 2006; Bernacchia et al., 2011; Shankar and Howard, 2012); in our theory, each neuron undergoes plasticity over a range of temporal scales, and the contribution of scale k depends on the posterior probability $P(q = q^k | \mathbf{x}_{1:t}, \mathbf{y}_{1:t})$. When one synapse changes its micro-weights, the posterior over q (which is shared across synapses) shifts its mass onto higher values. In this way, information about the diffusion variance is propagated across synapses.

There are two important aspects of this architecture that are worth highlighting. First, the change in q does not depend on the direction of the weight change: Both increases and decreases shift the posterior towards higher values of q . Second, shifts in the posterior over q affect not only *subsequent* learning but also *previous* learning, such that weak synapses can be retroactively strengthened. Both of these aspects have interesting empirical implications, as we explore in the next section.

We propose that the availability of plasticity-related proteins at time t is proportional to the mean of q under the posterior $P(q | \mathbf{x}_{1:t}, \mathbf{y}_{1:t})$. The logic of this proposal is that L-LTP (which depends on plasticity-related proteins) should occur only when change is likely (i.e., q is high). If change is unlikely, a perturbation of the firing rates should be treated as noise and no learning should occur. Since the diffusion variance is treated as a global variable, different synapses can influence each other through their affect on the posterior over q . Protein capture, on this view, is the mechanism by which synapses share information about environmental change, modulating the learning rate at other synapses.

Although any change in the inferred diffusion variance should in principle influence other synapses, there are a number of situations in which this will have no discernible effect. First, there must be *some* stimulation of a synapse in order for that synapse to be affected by changes in the inferred diffusion variance. These changes modulate the learning rate, but if nothing has been learned at all then such modulation will have no effect. Second, strong stimulation of a

synapse will increase the inferred diffusion variance, but this change will have little impact on the inferred synaptic weight, since a sufficiently strong synapse (e.g., after a standard LTP protocol) will achieve L-LTP regardless of the learning rate.

Results

Illustrations

We begin by illustrating the behavior of the model on random draws from the generative process described above (Figure 2). Two different settings of the diffusion variance q were explored: fast diffusion ($q=0.4$) and slow diffusion ($q=0.05$). The presynaptic firing rates x_t were drawn independently from a zero-mean Gaussian with a variance of 3. These random draws provide a sense of what kind of data would be generated under the statistical assumptions the model makes about the environment.

Figure 2 also shows that the model is able to track the weights over time and that posterior mean of q converges to the true diffusion variance (although convergence is slower when the environment changes more quickly). It is important to note that the model converges to the correct answer even when the true diffusion variance is much higher than the mean of the prior $P(q)$; thus, the prior provides a transient bias that dissipates as more data are observed.

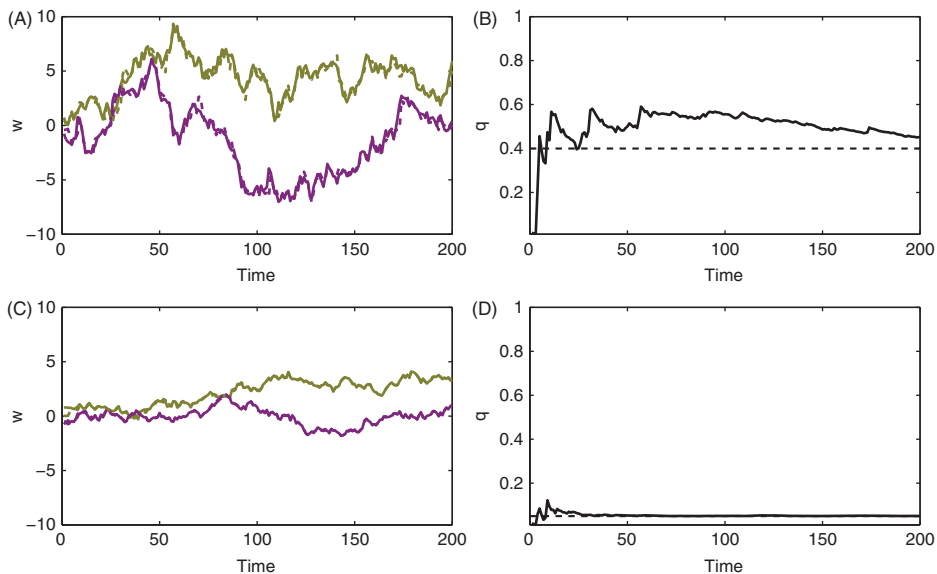


Figure 2. Illustration of the model. The top row (A, B) shows simulations with $q=0.4$ and the bottom row (C, D) shows simulations with $q=0.05$. The left panels show two random weight trajectories drawn from the generative process (solid line) along with the posterior mean estimator (dashed line). The right panels show the posterior mean diffusion variance; ground truth is indicated by a dashed horizontal line. (A color version of this figure is available in the online edition of this article.)

One advantage of our model, relative to a model which assumes fixed q , is that it can flexibly adapt to the appropriate timescale of neuronal events. This adaptation is facilitated by sharing of information between neurons in the form of plasticity-related proteins; as described above, the model postulates that the level of plasticity-related proteins represents an inference about the diffusion variance. We can quantify the advantage of this sharing by formulating a learning task and then measuring how performance changes as a function of population size. We drew random data (associative strength and neuronal firing rates) from the generative model described above for populations varying in size from 2 neurons to 20 neurons and with different settings of the diffusion variance. We then applied the synaptic update equations to estimate the associative strength and diffusion variance. For each random draw we computed the “prediction error” between the actual post-synaptic firing rate y_t and the predicted post-synaptic firing rate $x_t \hat{w}_{t-1}$ using the associative strength estimate. Performance was measured by the mean squared prediction error (lower squared error indicates better performance). We compared our model to a version with the same parameters, except the diffusion variance was fixed to 0.2 rather than inferred.

The results, shown in Figure 3, reveal that the fixed variance model always performs better than the inferred variance model when the actual variance is in the vicinity of the fixed variance (Figure 3A). However, this vicinity is relatively small, and importantly it shrinks as the population size grows (Figure 3B). This shrinking is the result of sharing information about the diffusion variance across neurons, allowing the population to efficiently home in on the true diffusion variance. Overall, the advantage of the inferred variance model increases with the number of neurons (Figure 3C).

Simulations of synaptic tagging and capture

In the simulations reported here, the post-synaptic neuron’s firing rate is clamped at $y_t = 1$, while the pre-synaptic neuron’s firing rate varies between -1 (below baseline) and $+1$ (above baseline). Strong stimulation is modeled as 3 pulses of $x_t = 1$, and weak stimulation is modeled as a single pulse.

Simulations of the classic synaptic tagging paradigm are shown in Figure 4 (left), along with the posterior mean diffusion variance (right). Weak stimulation produces a transient potentiation of the synapse (E-LTP) which decays back to baseline (Figure 4A). The transience of potentiation occurs due to the non-linear nature of the decay function: Weak weights decay much faster than strong weights. Because weak stimulation produces a weak synaptic strength, the synapse decays rapidly. Another consequence of weak stimulation is that the inferred diffusion variance is low (Figure 4B). This occurs because a weakly potentiated synapse has highest likelihood under the hypothesis that the environment is changing slowly. The low value of the inferred diffusion variance mirrors the low level of plasticity-related proteins observed cellularly in this paradigm.

Strong stimulation at another synapse can rescue E-LTP if applied shortly after (Figure 4C,D) or before the weak stimulation (Figure 4E,F). Strong stimulation provides evidence that the environment is changing more quickly. As a consequence, the contribution of the “fast” weights (those associated with higher

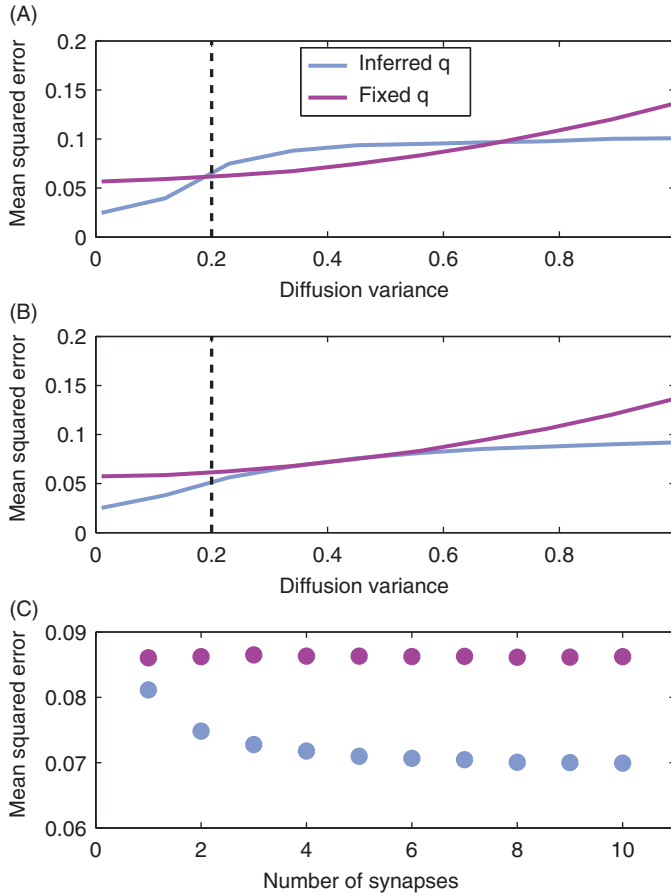


Figure 3. Performance of the model. Prediction error as a function of true diffusion variance with 2 synapses (A) and 20 synapses (B), shown for the full model in which the diffusion variance is inferred, as well as with the diffusion variance fixed to 0.2 (indicated by the dashed vertical line). Results are averaged across 1000 random draws from the generative model. (C) Average error across all values of the true diffusion variance, shown as a function of the number of synapses. (A color version of this figure is available in the online edition of this article.)

values of q) is amplified by the posterior $P(q|\mathbf{x}_{1:n}, \mathbf{y}_{1:n})$. Also notice that the inferred diffusion variance only increases substantially once both pathways are stimulated. This occurs because the data must overcome a prior that is biased towards a low diffusion variance.

To demonstrate the importance of inferring q , we ran the same simulations with q fixed to 0.2 (Figure 5). This resulted in all stimulation protocols (both weak and strong) producing L-LTP, contrary to the empirical observations described above.

Following the seminal work of (Frey and Morris, 1997), and parallel investigations in the invertebrate sea slug *Aplysia* (Martin et al., 1997), subsequent studies have revealed several other properties of synaptic tagging and capture. We address each in turn.

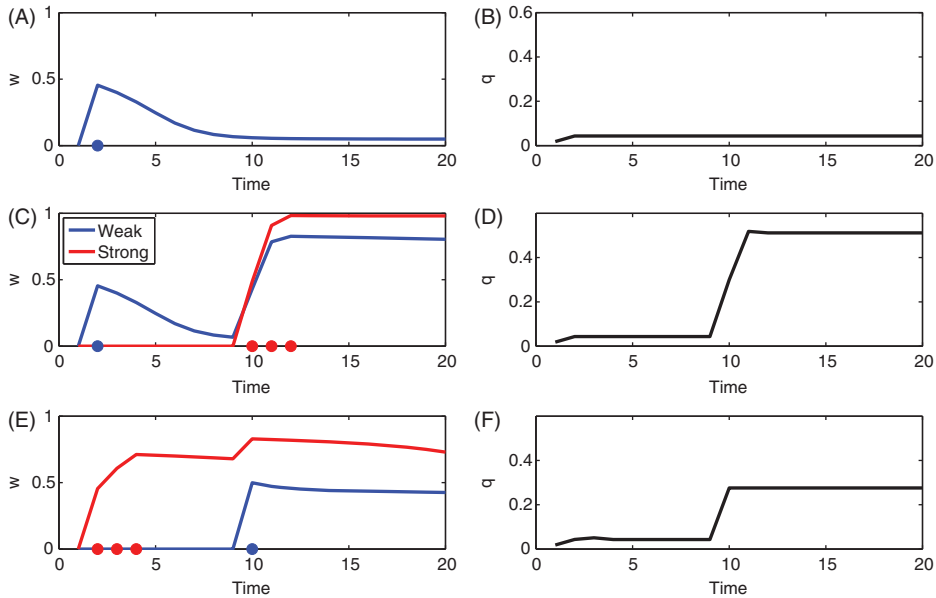


Figure 4. Tagging and capture. Each plot on the left shows the evolution of the posterior mean estimator of synaptic strength (Eq. 5) over time as a function of transient high-frequency stimulation, indicated by filled circles along the abscissa. “Weak” stimulation refers to a single stimulus train (i.e., a single circle), whereas “strong” stimulation refers to three trains. Each plot on the right shows the mean of the posterior over diffusion variance, q . (A,B) Weak stimulation results in E-LTP, a transient increase in synaptic strength that then decays back to baseline. E-LTP can be transformed into L-LTP if followed (C,D) or preceded (E,F) by strong stimulation at another synapse. (A color version of this figure is available in the online edition of this article.)

Cross-capture. Repeated low-frequency stimulation induces a sustained weakening of synaptic strength known as long-term depression (LTD). Synaptic tagging and capture can also be observed in the LTD protocol, although the molecular tags for LTP and LTD may be different (Sajikumar et al., 2007). Capture, on the other hand, does not distinguish between potentiation and depression: E-LTD can be transformed into L-LTD by LTP at another synapse, and vice versa (Sajikumar and Frey, 2004b). As shown in Figure 6, weak low-frequency stimulation (modeled as $x_t = -1$) induces E-LTD, which can then be rescued by strong high-frequency stimulation at another synapse. This reflects the fact that the inferred diffusion variance increases for both above-baseline and below-baseline perturbations. In other words, both LTD and LTP provide evidence that the environment is changing, leading to a higher value of q and hence amplification of weakly potentiated synapses.

Tag resetting. Low-frequency stimulation shortly following E-LTP induction at the same synapse resets the tag, preventing subsequent capture (Sajikumar and Frey, 2004a). Figure 7A shows simulations of this paradigm. Despite the inferred diffusion variance increasing (Figure 7B), the low-frequency stimulation causes the

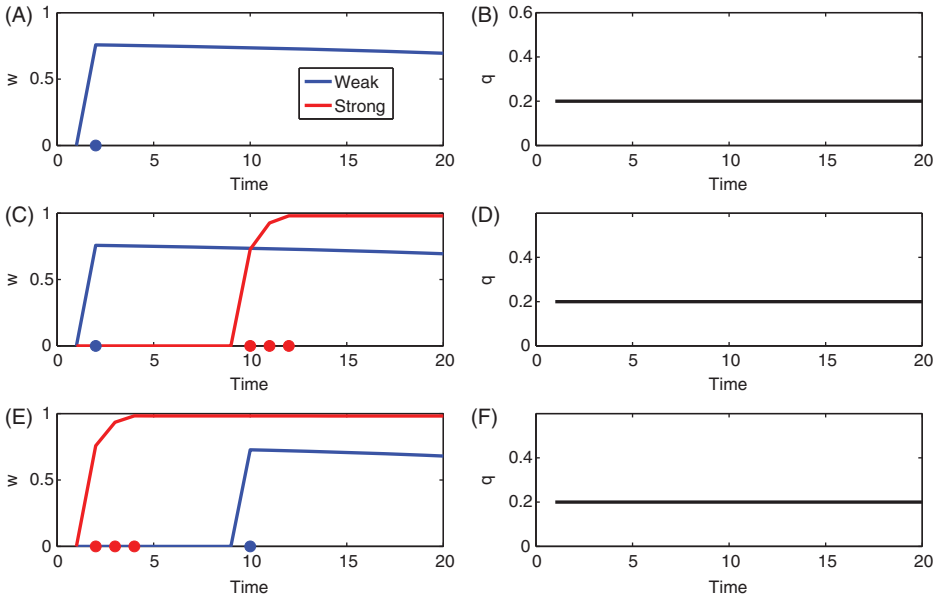


Figure 5. Tagging and capture with fixed diffusion variance. Plots are in the same format as Figure 2, with the left plots showing the posterior mean estimator of synaptic strength (Eq. 5) and the right plots showing the posterior mean diffusion variance. Here the diffusion variance q is fixed to 0.2 rather than inferred by the model. The result is that all stimulation protocols produce L-LTP, contrary to empirical observations. (A color version of this figure is available in the online edition of this article.)

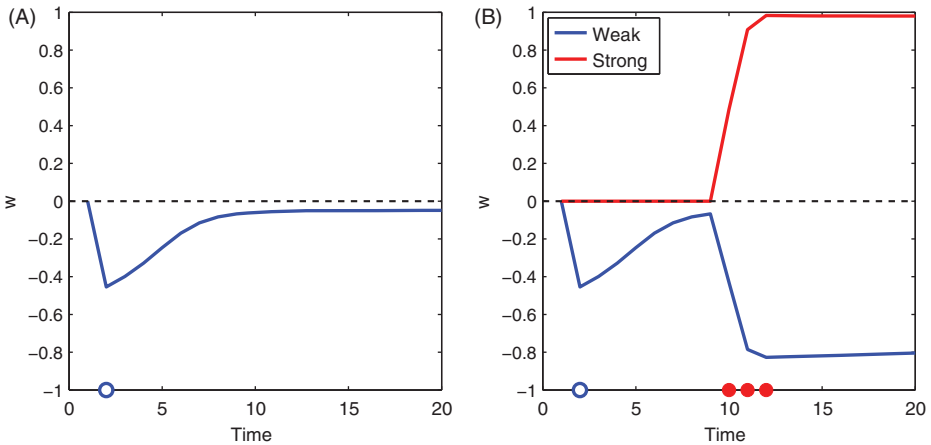


Figure 6. Cross-capture. Each plot shows the posterior mean estimator of synaptic strength. (A) Weak low-frequency stimulation (unfilled circles) induces E-LTD that decays back to baseline (dashed line). (B) E-LTD can be transformed to L-LTD if followed by strong high-frequency stimulation (filled circles) of another synapse. (A color version of this figure is available in the online edition of this article.)

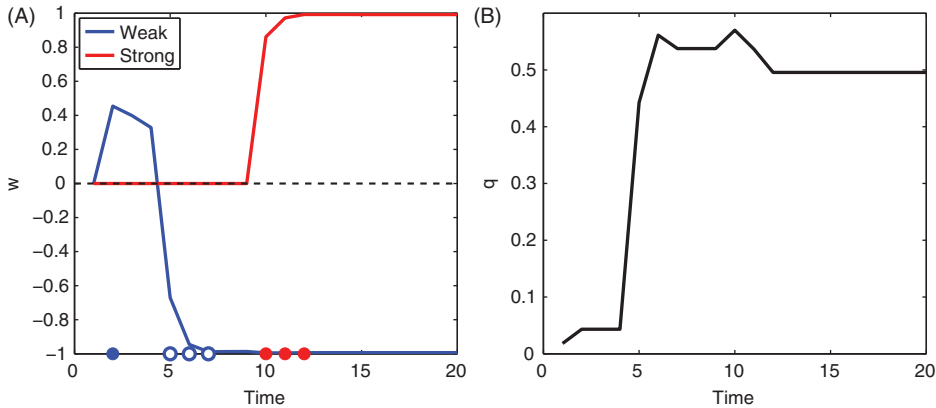


Figure 7. Tag resetting. (A) If low-frequency stimulation (unfilled circles) is applied following induction of E-LTP, the posterior mean synaptic strength falls below baseline. (B) Evolution of the posterior mean diffusion variance. (A color version of this figure is available in the online edition of this article.)

synaptic strength to drop below baseline, and increasing the diffusion variance following strong stimulation pushes the synapse even further below baseline.

Protein-synthesis inhibitors. We next examined the effects of protein-synthesis inhibition, which we modeled by resetting the posterior over q back to the prior (thereby inducing a belief that change is unlikely). This fits with our characterization of plasticity-related protein levels as signaling the inferred diffusion variance. When applied to the weakly stimulated pathway just before stimulation, protein synthesis inhibition has no effect (Figure 8A), consistent with experimental findings (Frey and Morris, 1997). This happens because at this time point the inferred diffusion variance is already low; protein-synthesis inhibition (since it is transient) does not affect the increase in diffusion variance following strong stimulation, which comes later. However, when applied prior to strong stimulation, protein-synthesis inhibition prevents the transformation of E-LTP to L-LTP (Figure 8B), due to the fact that the inferred diffusion variance is reduced by the resetting of the posterior, consistent with experimental findings (Frey and Morris, 1997).

Memory maintenance. A number of experiments have suggested that the maintenance of LTP depends on the activity of an autonomously active isoform of protein kinase C, known as PKM ζ (Sacktor, 2010). One computational interpretation of this molecule is that it maintains LTP by regulating the decay rate, controlled by the parameter β (see Eq. 3). We modeled PKM ζ inhibitors by transiently decreasing β from 10 to 0.1, which has the effect of making memories decay faster. (Sajikumar et al., 2005) found that while E-LTP was preserved under PKM ζ inhibition following weak stimulation, L-LTP was abolished despite prior strong stimulation at another synapse, a finding reproduced in our simulations (Figure 9). This occurs because low values of β cause even strong synapses to decay back to zero.

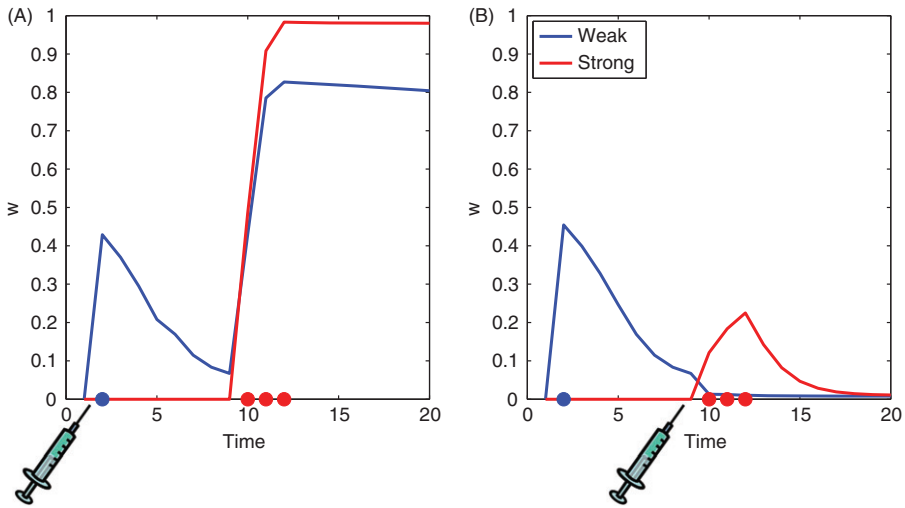


Figure 8. Effects of protein-synthesis inhibitors. Each plot shows the posterior mean estimator of synaptic strength. Protein synthesis inhibition was modeled by resetting the posterior to the prior. (A) The transformation of E-LTP to L-LTP is insensitive to protein-synthesis inhibition at the time of weak stimulation. (B) Protein-synthesis inhibition at the time of strong stimulation prevents the transformation of E-LTP to L-LTP. (A color version of this figure is available in the online edition of this article.)

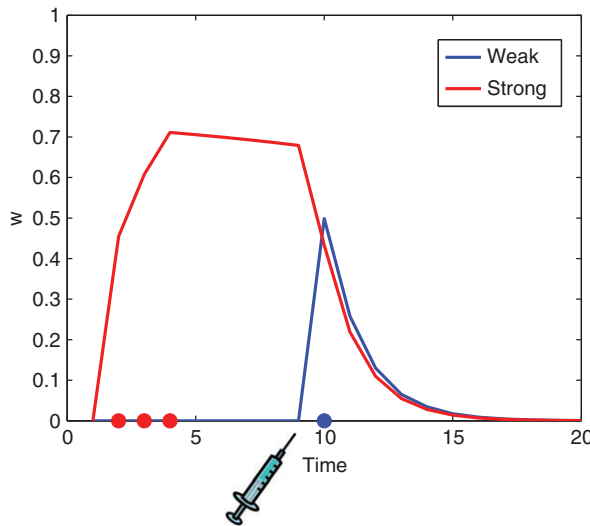


Figure 9. Memory maintenance. Application of a PKM ζ inhibitor (modeled as decreasing β) following weak stimulation prevents the transformation of E-LTP to L-LTP. (A color version of this figure is available in the online edition of this article.)

New predictions

The model presented in this paper makes a number of new, untested predictions. First, the availability of plasticity-related proteins should be correlated not with novelty *per se* but with inferred environmental volatility—i.e., the rate of change.

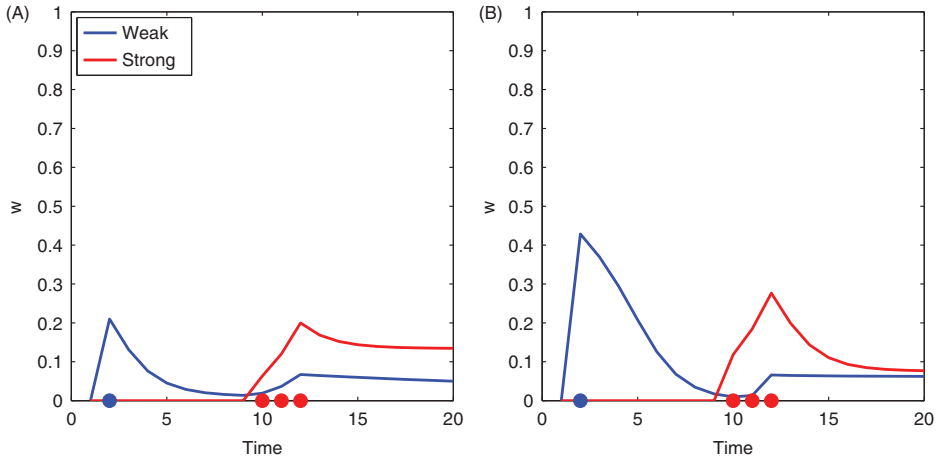


Figure 10. New predictions. Synaptic tagging and capture is reduced in a noisy (A) or slowly changing (B) environment. (A color version of this figure is available in the online edition of this article.)

If the environment is very noisy but slowly changing, then protein availability should fluctuate around zero and learning should remain low. Figure 10A shows a simulation of this hypothetical experiment, using the weak-before-strong protocol. We set $r=0.3$ (three times the level used in the simulations reported above) to simulate a noisy environment. The simulation shows that LTP is reduced following both strong and weak stimulation, and no rescuing of E-LTP is observed following strong stimulation.

A similar outcome should occur if the noise level is low but the prior over q is strongly biased towards slow rates of change. This predicts that pre-exposing animals to an unchanging context should reduce subsequent sensitivity to synaptic tagging and capture protocols. We simulated this by setting the parameters of the prior over q to more strongly favor small values of q (and hence slower rates of change); specifically, we set $a=15$ and $b=5$. This produced weakened LTP and no rescuing of early LTP following strong stimulation (Figure 10B).

Discussion

The model developed in this paper rationalizes the penumbra of learning as a consequence of optimal statistical learning. In the process, we have shed new light on the putative cellular underpinning of the penumbra, synaptic tagging and capture. Simulations demonstrated that the model reproduces main properties of synaptic tagging and capture, including tag resetting, cross-capture and the effects of protein-synthesis inhibitors.

Several other computational theories have been developed to account for synaptic tagging and capture (Clopath et al., 2008; Barrett et al., 2009; Paper et al., 2011). The existing theories share in common the idea that a synapse can be in one of several abstract states, representing different phases of LTP (Smolen et al., 2012). Transitions between these states determine the dynamics

of synaptic tagging and capture. Computational investigations suggest that tagging and capture functions to prolong memory lifetimes, and may also aid in learning associations between experiences and distal rewards (Päpper et al., 2011). In contrast to these essentially descriptive models, our point of departure is a normative characterization of the learning problem: Starting from a probabilistic description of the environment, we derive the optimal synaptic learning rules (Fiser et al., 2010).

The idea that learning rates are modulated by the inferred rate of change is central to many theories of learning (Pearce and Hall, 1980; Dayan et al., 2000; Courville et al., 2006; Behrens et al., 2007; Kording et al., 2007; Wilson et al., 2010) and is supported by abundant behavioral and neural data (Roesch et al., 2012). These theories typically consider only a single association. However, synaptic tagging and capture indicates that learning rates are also modulated by the inferred *global* rate of change (i.e., across associations). According to the model developed here, the functional role of synaptic tagging and capture is to compute the inferred global rate of change (operationalized here as the diffusion variance) and propagate this signal across synapses.

While our simulations focused on phenomena studied with *in vitro* slice preparations, we believe that the penumbra of learning serves a more general function. Whenever there is uncertainty about the global rate of change, one should see the effect of change in one task on the learning rate in another task. For example, (Nassar et al., 2012), using a number prediction task, observed that a surprising (but task-irrelevant) change in an auditory cue accelerated learning. There is some evidence that a surprising word in a sequentially-presented list can improve verbal memory for words in neighboring serial positions (Wallace, 1965). Thus, the penumbra of learning appears to be a general phenomenon.

There are a number of limitations of the theory as it currently stands.

- (1) It seems implausible that all synapses are parameterized by the same diffusion variance. A more sophisticated approach would be a hierarchical model in which each synapse is parameterized by a unique diffusion variance drawn from a higher-level distribution; in this way, the synapses can still share statistical strength, but can also diverge from one another.
- (2) There is evidence suggesting that low-frequency stimulation of one synapse can reset the tag on another synapse (Young and Nguyen, 2005); this finding poses a challenge to the theory, which only accounts for homosynaptic tag resetting.
- (3) We have not explicitly modeled the critical time window in which tagging and capture are effective (Frey and Morris, 1997); we have instead implicitly assumed in our simulations that events occurred within this time window. In fact, the results are relatively insensitive to the interval between the two stimulations. A more complete theory should explicitly incorporate the time window into the generative model (e.g., by having the diffusion variance decay back to zero).
- (4) (Fonseca et al., 2004) have shown that when the level of plasticity-related proteins in a cell is low, synapses compete with one another such that L-LTP at one synapse comes at the expense of L-LTP at another synapse. This finding does not seem to fit with our purely statistical characterization of synaptic tagging and capture.

- (5) The theory assumes that synaptic weights decay to zero, but empirically decay asymptotes at pre-stimulation levels.² One way to accommodate this observation is to add a neuron-specific baseline constant to the weight dynamics, such that each weight will eventually decay back to its baseline. This baseline could itself be learned through statistical inference by defining the input to be $x'_t = [x_t, 1]$ and $w'_t = [w_t, w_0]$, where w_0 is the baseline constant, and the output to be $y_t = w_t x_t + w_0$ (i.e., the inner product of w'_t and x'_t). The Kalman filtering equations can be applied to the augmented weight vector, thereby adapting the baseline over time.

Many molecules participate in synaptic tagging and capture; the theory described in this paper has not yet assigned computational roles to most of these. Of special relevance is the role of dopamine, a neuromodulator that is known to be necessary for producing the penumbra of learning in behavioral tagging experiments (Moncada and Viola, 2007; Wang et al., 2010). Activation of D1/D5 dopamine receptors in hippocampal neurons stimulates local protein synthesis (Smith et al., 2005), suggesting that dopamine might underlie diffusion variance estimation in our theory. This view resonates with recent ideas about dopamine function, particularly its role as an alerting signal for salient or novel sensory events (Bromberg-Martin et al., 2010), and may point towards a deeper computational synthesis of dopamine's multifaceted role in learning.

Acknowledgments

This work was supported by an MIT Intelligence Initiative postdoctoral fellowship. I am grateful to Roger Redondo for helpful discussions.

Declaration of interest: The authors report no conflicts of interest.

Notes

- [1] The weight dynamics in their current form violate Dale's Law (Eccles, 1964), which in its abstract form states that a synapse can only be either excitatory or inhibitory. If we wish to make the dynamics more biologically plausible, we could introduce a rectification that prevents the weight from changing sign. This modification is not explored here.
- [2] We thank an anonymous reviewer for pointing this out.

References

- Abraham WC. 2008. Metaplasticity: tuning synapses and networks for plasticity. *Nature Reviews Neuroscience* 9:387–399.
- Anderson BD, Moore JB. 1979. *Optimal Filtering*. Prentice-Hall.

- Ballarini F, Moncada D, Martinez MC, Alen N, Viola H. 2009. Behavioral tagging is a general mechanism of long-term memory formation. *Proceedings of the National Academy of Sciences* 106:14599–14604.
- Barrett AB, Billings GO, Morris RG, Van Rossum MC. 2009. State based model of long-term potentiation and synaptic tagging and capture. *PLoS Computational Biology* 5:e1000259.
- Behrens TE, Woolrich MW, Walton ME, Rushworth MF. 2007. Learning the value of information in an uncertain world. *Nature Neuroscience* 10:1214–1221.
- Bernacchia A, Seo H, Lee D, Wang X.-J. 2011. A reservoir of time constants for memory traces in cortical neurons. *Nature Neuroscience* 14:366–372.
- Bienenstock EL, Cooper LN, Munro PW. 1982. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *The Journal of Neuroscience* 2:32–48.
- Bromberg-Martin ES, Matsumoto M, Hikosaka O. 2010. Dopamine in motivational control: Rewarding, aversive, and alerting. *Neuron* 68:815–834.
- Clopath C, Ziegler L, Vasilaki E, Büsing L, Gerstner W. 2008. Tag-trigger-consolidation: a model of early and late long-term-potentiation and depression. *PLoS Computational Biology* 4:e1000248.
- Courville AC, Daw ND, Touretzky DS. 2006. Bayesian theories of conditioning in a changing world. *Trends in Cognitive Sciences* 10:294–300.
- Dayan P, Kakade S, Montague PR. 2000. Learning and selective attention. *Nature Neuroscience* 3:1218–1223.
- Duncan K, Sadanand A, Davachi L. 2012. Memory's penumbra: Episodic memory decisions induce lingering mnemonic biases. *Science* 337:485–487.
- Eccles JC. 1964. *The physiology of synapses*. New York: Academic Press.
- Fiser J, Berkes P, Orbán G, Lengyel M. 2010. Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences* 14:119–130.
- Fonseca R, Nägerl UV, Morris RG, Bonhoeffer T. 2004. Competing for memory: hippocampal LTP under regimes of reduced protein synthesis. *Neuron* 44:1011–1020.
- Frey U, Morris R. 1998. Weak before strong: dissociating synaptic tagging and plasticity-factor accounts of late-LTP. *Neuropharmacology* 37(4):545–552.
- Frey U, Morris RG. 1997. Synaptic tagging and long-term potentiation. *Nature* 385(6616):533–536.
- Kalman RE. 1960. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* 82:35–45.
- Kording KP, Tenenbaum JB, Shadmehr R. 2007. The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience* 10:779–786.
- La Camera G, Rauch A, Thurbon D, Lüscher H.-R, Senn W, Fusi S. 2006. Multiple time scales of temporal response in pyramidal and fast spiking cortical neurons. *Journal of Neurophysiology* 96:3448–3464.
- Martin KC, Casadio A, Zhu H, Rose Y, Chen M, Bailey CH, Kandel ER, et al. 1997. Synapse-specific, long-term facilitation of aplysia sensory to motor synapses: a function for local protein synthesis in memory storage. *Cell* 91:927–938.
- Martin KC, Kosik KS. 2002. Synaptic tagging: who's it? *Nature Reviews Neuroscience* 3:813–820.
- Martin S, Grimwood P, Morris R. 2000. Synaptic plasticity and memory: an evaluation of the hypothesis. *Annual Review of Neuroscience* 23:649–711.
- McGaugh JL. 2000. Memory—a century of consolidation. *Science* 287:248–251.
- Merhav M, Rosenblum K. 2008. Facilitation of taste memory acquisition by experiencing previous novel taste is protein-synthesis dependent. *Learning & Memory* 15:501–507.
- Moncada D, Viola H. 2007. Induction of long-term memory by exposure to novelty requires protein synthesis: evidence for a behavioral tagging. *The Journal of Neuroscience* 27:7476–7481.
- Montague PR, Sejnowski TJ. 1994. The predictive brain: temporal coincidence and temporal order in synaptic learning mechanisms. *Learning & Memory* 1:1–33.
- Nassar MR, Rumsey KM, Wilson RC, Parikh K, Heasley B, Gold JI. 2012. Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience* 15:1040–1046.
- Päpper M, Kempster R, Leibold C. 2011. Synaptic tagging, evaluation of memories the distal reward problem. *Learning & Memory* 18:58–70.
- Pearce JM, Hall G. 1980. A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review* 87:532–552.
- Redondo RL, Morris RG. 2011. Making memories last: the synaptic tagging and capture hypothesis. *Nature Reviews Neuroscience* 12:17–30.

- Roesch MR, Calu DJ, Esber GR, Schoenbaum G. 2010. Neural correlates of variations in event processing during learning in basolateral amygdala. *The Journal of Neuroscience* 30:2464–2471.
- Roesch MR, Esber GR, Li J, Daw ND, Schoenbaum G. 2012. Surprise! neural correlates of Pearce–Hall and Rescorla–Wagner coexist within the brain. *European Journal of Neuroscience* 35:1190–1200.
- Sacktor TC. 2010. How does PKM ζ maintain long-term memory?. *Nature Reviews Neuroscience* 12:9–15.
- Sajikumar S, Frey J. 2004a. Resetting of ‘synaptic tags’ is time- and activity-dependent in rat hippocampal *ca in vitro*. *Neuroscience* 129:503–507.
- Sajikumar S, Frey JU. 2004b. Late-associativity, synaptic tagging, and the role of dopamine during *ltp* and *ltd*. *Neurobiology of Learning and Memory* 82:12–25.
- Sajikumar S, Navakkode S, Frey JU. 2007. Identification of compartment- and process-specific molecules required for “synaptic tagging” during long-term potentiation and long-term depression in hippocampal *ca1*. *The Journal of Neuroscience* 27:5068–5080.
- Sajikumar S, Navakkode S, Sacktor TC, Frey JU. 2005. Synaptic tagging and cross-tagging: the role of protein kinase $m\zeta$ in maintaining long-term potentiation but not long-term depression. *The Journal of Neuroscience* 25(24):5750–5756.
- Shankar KH, Howard MW. 2012. A scale-invariant internal representation of time. *Neural Computation* 24:134–193.
- Smith WB, Starck SR, Roberts RW, Schuman EM. 2005. Dopaminergic stimulation of local protein synthesis enhances surface expression of *glur1* and synaptic transmission in hippocampal neurons. *Neuron* 45:765–779.
- Smolen P, Baxter DA, Byrne JH. 2012. Molecular constraints on synaptic tagging and maintenance of long-term potentiation: a predictive model. *PLoS Computational Biology* 8:e1002620.
- Wallace WP. 1965. Review of the historical, empirical, and theoretical status of the von Restorff phenomenon. *Psychological Bulletin* 63:410–424.
- Wang S-H, Redondo RL, Morris RG. 2010. Relevance of synaptic tagging and capture to the persistence of long-term potentiation and everyday spatial memory. *Proceedings of the National Academy of Sciences* 107:19537–19542.
- Wilson RC, Nassar MR, Gold JJ. 2010. Bayesian online learning of the hazard rate in change-point problems. *Neural Computation* 22:2452–2476.
- Young JZ, Nguyen PV. 2005. Homosynaptic and heterosynaptic inhibition of synaptic tagging and capture of long-term potentiation by previous synaptic activity. *The Journal of Neuroscience* 25:7221–7231.

Appendix A: The extended Kalman filter

When conditioned on the diffusion variance q , the posterior distribution over the synaptic weight w_t is given by:

$$\begin{aligned} P(w_t | \mathbf{x}_{1:t}, \mathbf{y}_{1:t}, q) &\propto P(y_t | w_t, x_t) P(w_t | \mathbf{x}_{1:t-1}, \mathbf{y}_{1:t-1}, q) \\ &= \mathcal{N}(y_t; w_t x_t, r) \mathcal{N}(w_t; f(\hat{w}_{t-1}), q). \end{aligned} \quad (11)$$

This expression follows directly from the definition of the generative model (equations 1 and 2). The extended Kalman filter approximates the transition function with a first-order Taylor series expansion around the previous weight estimate, \hat{w}_{t-1} :

$$f(w_{t-1}) \approx f(\hat{w}_{t-1}) + f'(\hat{w}_{t-1})(w - \hat{w}_{t-1}), \quad (12)$$

where $f'(\hat{w}_{t-1})$ is the first derivative of $f(w_{t-1})$ with respect to \hat{w}_{t-1} :

$$f'(\hat{w}_{t-1}) = f(\hat{w}_{t-1}) [2\beta \hat{w}_{t-1} (1 - f(\hat{w}_{t-1})) + 1]. \quad (13)$$

The Taylor series expansion approximates the transition function as a linear function of the previous weight estimate. This yields a linear-Gaussian dynamical

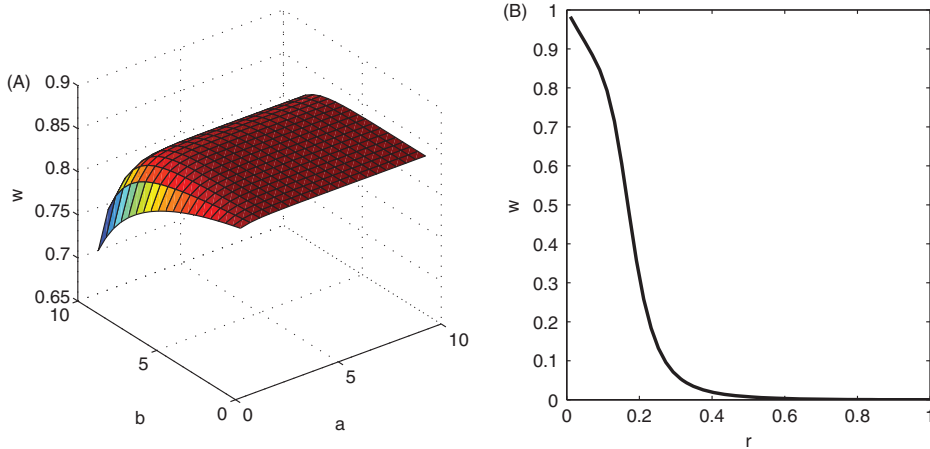


Figure 11. Parameter explorations. Effects of changing the parameter settings on the weakly-stimulated synapse in a weak-before-strong protocol. Each plot shows the posterior mean weight of the weakly-stimulated synapse after strong stimulation of another synapse. (A) Dependence of tagging and capture on the prior parameters a and b . (B) Dependence of tagging and capture on the observation noise variance r . (A color version of this figure is available in the online edition of this article.)

system:

$$\begin{aligned} w_t &= f(\hat{w}_{t-1}) + f'(\hat{w}_{t-1})(w_{t-1} - \hat{w}_{t-1}) + \epsilon_w \\ y_t &= w_t x_t + \epsilon_y. \end{aligned} \quad (14)$$

We can then apply the standard Kalman filtering equations (see for example Anderson and Moore, 1979) to this linearized system. The result is the posterior approximation given by equations 5 and 6.

Appendix B: Exploring the parameter space

This appendix considers the effects of parameter variations on synaptic tagging and capture. For illustration we use the weak-before-strong protocol described in the main text. We first varied the parameters of the prior, a and b , and measured the posterior mean weight of the weakly-stimulated synapse after strong stimulation of another synapse. Recall that a large value of w is indicative of successful synaptic tagging and capture (i.e., L-LTP occurs at the weak synapse). Figure 11A shows that there is a fairly large plateau of parameter values for which L-LTP occurs.

We next investigated the effect of the noise variance r . When this parameter is close to 0, synapses exert a stronger influence upon each other, because posterior uncertainty is lower. Accordingly, making r smaller results in stronger synaptic tagging and capture (Figure 11B). We can infer from this simulation that firing rate noise must be relatively small in order to observe synaptic tagging and capture.

Notice of correction:

Redundant text has been removed from the opening paragraph of this article.