

Annual Review of Neuroscience

Synaptic Plasticity Forms and Functions

Jeffrey C. Magee and Christine Grienberger

Department of Neuroscience and Howard Hughes Medical Institute, Baylor College of Medicine, Houston, Texas 77030, USA; email: jcmagee@bcm.edu

Annu. Rev. Neurosci. 2020. 43:95–117

First published as a Review in Advance on
February 19, 2020

The *Annual Review of Neuroscience* is online at
neuro.annualreviews.org

<https://doi.org/10.1146/annurev-neuro-090919-022842>

Copyright © 2020 by Annual Reviews.
All rights reserved

ANNUAL
REVIEWS **CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Keywords

learning, memory, synapses, dendrites, neuromodulation, eligibility traces

Abstract

Synaptic plasticity, the activity-dependent change in neuronal connection strength, has long been considered an important component of learning and memory. Computational and engineering work corroborate the power of learning through the directed adjustment of connection weights. Here we review the fundamental elements of four broadly categorized forms of synaptic plasticity and discuss their functional capabilities and limitations. Although standard, correlation-based, Hebbian synaptic plasticity has been the primary focus of neuroscientists for decades, it is inherently limited. Three-factor plasticity rules supplement Hebbian forms with neuromodulation and eligibility traces, while true supervised types go even further by adding objectives and instructive signals. Finally, a recently discovered hippocampal form of synaptic plasticity combines the above elements, while leaving behind the primary Hebbian requirement. We suggest that the effort to determine the neural basis of adaptive behavior could benefit from renewed experimental and theoretical investigation of more powerful directed types of synaptic plasticity.

Contents

INTRODUCTION	96
CORRELATIVE HEBBIAN SYNAPTIC PLASTICITY	97
NEUROMODULATION, ELIGIBILITY TRACES, AND THREE-FACTOR PLASTICITY RULES	99
SUPERVISED SYNAPTIC PLASTICITY	100
BEHAVIORAL TIMESCALE SYNAPTIC PLASTICITY	104
FUNCTIONS OF BEHAVIORAL TIMESCALE SYNAPTIC PLASTICITY	108
CLOSING	110

INTRODUCTION

A fundamental objective of neuroscience is to understand how brains produce the learned or adaptive behaviors that allow animals to thrive in complex environments. The prevailing core idea is that learning is associated with persistent experience-driven changes to animals' brains that subsequently aid them in the successful performance of important tasks such as the acquisition of necessities like food and shelter while avoiding the unpleasantness that accompanies injury or predation. Obviously, this behavior is complicated, requiring an accurate perception of the sensory elements of an environment, a capacity to make suitable decisions, and an ability to move appropriately within the environment. On top of this, sensory and motor systems are imperfect, and it remains impossible to perceive anything about the future. Thus, it is also important for brains to generate inferences and predictions about the world that can fill in perceptual gaps, correct for inappropriate motor commands, and guide decision-making. Because of this complexity, the process of acquiring a particular adaptive behavior involves modifications to many brain regions, with the actual site of the adjustments being found within the individual neurons that make up these regions. In the end, the adaptations should manifest themselves as alterations in the responsiveness of the neurons to particular input patterns. That is, the adaptive changes should alter the input-output transformations of neurons within a given region so that the population can best perform the specific computations that enable successful behaviors (e.g., improved perception, enhanced value to action mappings, efficient movement).

While it seems likely that there are multiple mechanisms involved in such an important function, we focus here on just one: experience-dependent alterations in the strength of synaptic connections between neurons. Long-lasting changes in the synaptic weights within particular networks have been hypothesized for many decades to mediate the task-relevant shaping of neuronal population activity. As a result, there has been an extraordinary amount of experimental, theoretical, and technical/engineering work that has resulted in tens of thousands of published manuscripts and numerous life-impacting devices. Not surprisingly, synaptic plasticity is nearly ubiquitous throughout brains, and a dizzying array of forms have been reported. Our goal here is not to survey these various forms but instead to broadly group them into several, perhaps overlapping, categories and coarsely present the current understanding of their different fundamental elements and functional capabilities. We then turn our attention to a plasticity that was recently discovered to underlie place fields in hippocampal area CA1, because it appears to uniquely combine features of the different plasticity forms, allowing it distinct abilities.

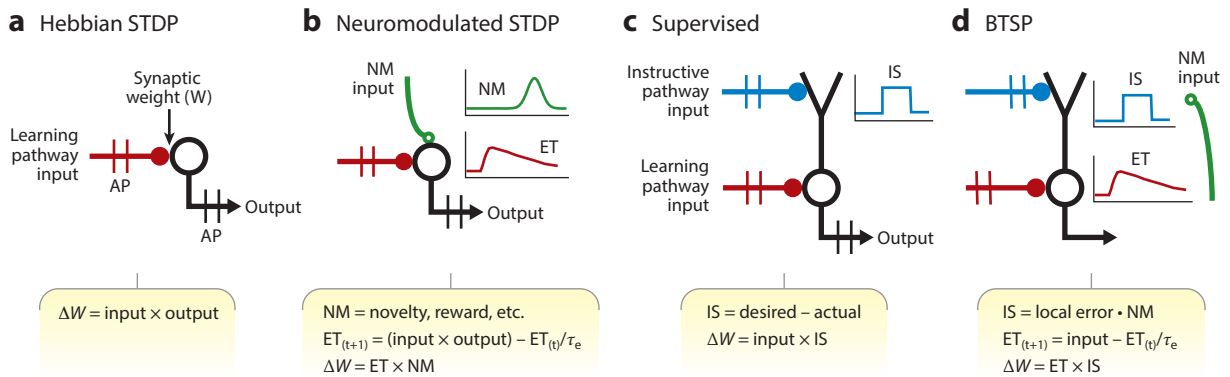


Figure 1

Common synaptic plasticity forms. (a) Standard correlative plasticity driven by repeated coactive input-output. (b) NM-gated correlative plasticity, where repeated coactive input-output generates ETs (τ_e is ET decay time constant) that are gated by a global signal (i.e., NM), potentially representing a behavioral variable. (c) Standard correlative plasticity, except that the postsynaptic term is replaced with an error term delivered by an instructive pathway that drives dendritic Ca^{2+} spikes (IS). (d) BTSP combines many of the above features without the requirement of correlated input-output (note lack of output action potential). In all panels, black circles and arrows labeled output are postsynaptic neurons. In panels c and d, the black Y-shaped element is an apical dendrite. Abbreviations: BTSP, behavioral timescale synaptic plasticity; ET, eligibility trace; IS, instructive signal; NM, neuromodulator; STDP, spike timing-dependent plasticity, ΔW , synaptic weight change.

CORRELATIVE HEBBIAN SYNAPTIC PLASTICITY

One of the most influential ideas about how learning-related changes might occur in brains was first articulated by Donald Hebb (1949) around seventy years ago. Hebb's famous postulate states that if cell A "repeatedly or persistently takes part" in firing cell B, then the strength of their connection should increase. In this statement, Hebb points out that learning-related changes should be found in the strength of the synaptic connections between individual neurons within a neuronal population or assembly and that these changes should, in his view, be based on causality and repetition. The underlying ideas here are that synapses responsible for correct responses (reward, escape, etc.) should be properly credited (this is known as credit assignment) and that this process should be gradual to avoid the errors associated with noise.

In the following decades, the pioneers of artificial neural networks (ANNs) interpreted Hebb's idea to be that weight changes among the units of single-layer networks should be based on coincidence, or the product of pre- and postsynaptic activity, thus transforming the causality into changes proportional to the coactivity or correlation of the input and output units (Rosenblatt 1959, Woodrow & Hoff 1960) (**Figure 1a**). The most straightforward of these network learning rules, the unsupervised forms, produces a rudimentary form of memory by linking a system's particular input patterns with the output patterns that are consistently associated with them. Here, inputs driving a given output will have their connection strengths enhanced, forming a simple association that allows even fragments of the associated input pattern to evoke the correct output activity (Andersen 1972, Kohonen 1972, Hopfield 1982). In addition, these unsupervised Hebbian learning rules autonomously find low-order statistical structure in the input stream, allowing them to mediate some types of fundamental feature selectivity and to generate self-organizing topographical representations of those features (von der Malsberg 1973, Grossberg 1976, Kohonen 1982, Oja 1982, Erwin & Miller 1998, Song & Abbott 2001, Brito & Gerstner 2016).

Actual synaptic plasticity was first discovered in the hippocampus in the early 1970s, where it was established that repeated, near-synchronous activation of both pre- and postsynaptic neurons

produced an increase in the strength of the synaptic input of only the stimulated connections (Bliss & Lomo 1973, Levy & Steward 1983). This phenomenon became known as long-term potentiation (LTP), and in the intervening decades, LTP and its counterpart, long-term depression (LTD), have been heavily studied in a variety of brain areas (see reviews by Malenka & Bear 2004, Caporale & Dan 2008, Kandel et al. 2014, Nicoll 2017). The primary mechanism thought to mediate the above input-output correlation revolves around the special properties of *N*-methyl-D-aspartate glutamate receptors (NMDARs), which require a coincidence of glutamate binding and membrane potential (V_m) depolarization for gating (Mayer et al. 1984). This aspect allows them to report, via an increase in postsynaptic intracellular Ca^{2+} levels, the temporal overlap of synaptic input and action potentials (APs) backpropagating from the axon to the synapse (Magee & Johnston 1997, Koester & Sakmann 1998, Luscher & Malenka 2012). Most synaptic plasticity forms investigated thus far fall under the standard autonomous correlative Hebbian category, and this includes practically all types of spike timing-dependent plasticity (STDP), the quintessential Hebbian protocol for experimentally inducing either LTP or LTD (Markram et al. 1997, Bi & Poo 1998). Indeed, STDP retains the fundamental autonomous correlation-based property while adding timing requirements that somewhat refine and extend its capabilities (Song & Abbott 2001, Young et al. 2007, Feldman 2012).

As would be expected from theory, most evidence of a role for autonomous correlative plasticity in actual network functional properties has come from studies examining the development of fundamental feature selectivity in sensory systems or in associative forms of learning where the conditioned stimulus (CS) (e.g., tone, odor, location) overlaps with the unconditioned stimulus (US) (generally aversive stimuli) (McKernan & Shinnick-Gallagher 1997, Rogan et al. 1997, Yao & Dan 2001, Li et al. 2008, Butts & Kanold 2010, Ryan et al. 2015, Tonegawa et al. 2015, Kim & Cho 2017). Several of these studies have used different forms of artificial activation (electrical or optogenetic) paired with sensory stimuli and examined the shifting of receptive fields or the formation of simple associations and changes in associative learning (Schuett et al. 2001, Yao & Dan 2001, Meliza & Dan 2006, Nabavi et al. 2014, Carrillo-Reid et al. 2016, El-Boustani et al. 2018). These effects are usually dependent on the repeated (many dozens) occurrence of tightly timed input-output pairings and thus implicate STDP-like synaptic plasticity phenomena. It is, of course, difficult to know for certain if these forms of plasticity were actually those that naturally occurred to produce the given feature selectivity. However, further support is offered by reports that the formation of certain types of basic associative memories is correlated with what appears to be LTP and can be blocked by manipulations that inhibit synaptic plasticity (Morris et al. 1986, McKernan & Shinnick-Gallagher 1997, Rogan et al. 1997, Nakazawa et al. 2004, Ryan et al. 2015, Kim Cho 2017). Thus, all things considered, there is appreciable evidence that synaptic plasticity mechanisms approximating those formulated by Hebb 70 years ago play a role in the development of the feature extraction capabilities of many brain networks and perhaps also in simple forms of associative memory.

However, having said this, there are a few fundamental issues concerning the standard varieties of synaptic plasticity discussed above that should also be considered. Simple autonomous categories of plasticity, such as correlative Hebbian forms, all support a kind of unsupervised learning that is inherently limited in the classes of problems it can solve. To begin with, without some sort of regulation over which synapses can change and when, synaptic weight modifications will be frequent and storage capacity quickly saturated, resulting in a ruinous overwriting of old memories (Fusi 2002, Fusi & Abbott 2007). In addition, it is difficult for synaptic plasticity mechanisms operating on physiological timescales (tens of milliseconds) to handle the time course of behavior (at least seconds) (Sutton & Barto 1981). Furthermore, because these forms lack a target for neuronal activity to approach, there simply is no obvious way for a neuron to know if a given change

in a particular synaptic weight is useful, thus limiting these self-organizing mechanisms to the formation of fundamental types of representations that are not always sufficient for achieving a specific learning objective (Fremaux & Gerstner 2016, Marblestone et al. 2016). We discuss several proposed solutions to these problems below as well as present evidence for a type of synaptic plasticity, recently uncovered in the hippocampus, that may point toward additional routes taken by the mammalian nervous system.

NEUROMODULATION, ELIGIBILITY TRACES, AND THREE-FACTOR PLASTICITY RULES

A large and complex literature exists describing the relationship between neuromodulatory activity and certain behaviors. The dependence of reward-motivated behavior on dopamine (DA) neuron activity and DA receptor activation is a well-documented example (Schultz 2007, Steinberg et al. 2013, Pignatelli & Bonci 2015, Berke 2018). Because most modulatory regions project broadly to numerous cortical and subcortical brain areas, and because neuromodulator release is mediated through nonsynaptic volume transmission, the action of these transmitters is thought to be global; that is, they more or less equally effect most neurons within a large volume of brain tissue. In addition, neuromodulators have been shown to profoundly alter standard correlative types of synaptic plasticity in a concentration- and receptor subtype-dependent manner, and this has led to the concept that specific neuromodulators could gate Hebbian plasticity (Gu 2002, Nadim & Bucher 2014, Brzosko et al. 2017). Such global neuromodulatory gating potentially obviates the ruinous effects of uncontrolled correlative plasticity, and it could also provide a degree of supervision if the neuromodulators communicate some type of behaviorally related message (e.g., reward, novelty).

Of course, even here there remains a problem related to the fact that much of the sensory experience of an animal is separated in time by at least seconds from any associated consequences [this is known as the distal reward problem (Hull 1943)]. It was demonstrated many decades ago that correlative Hebbian learning rules alone are of limited use in explaining the predictive responses observed in many classical conditioning experiments, since the timescales of these rules do not allow them to directly adjust the weights of inputs preceding the US by relevant amounts of time (Sutton & Barto 1981). To correct for this deficiency, Sutton & Barto (1981) implemented non-stimulating eligibility traces (ETs): internal signals, local to the synapse and not related to electrical output activity, that decay away over several seconds, thereby marking a given synapse as eligible for modification (**Figure 1b**). The use of such traces, basically filtered versions of synaptic input, alleviates the distal reward problem by linking presynaptic activity with distant postsynaptic activations. In essence, ETs generate plasticity rules that are asymmetric in time, in that the time course over which they elevate synaptic weights stretches farther backward than forward, and such rules are known to produce predictive neuronal activity (Abbott & Blum 1996, Blum & Abbott 1996, Bittner et al. 2017). It was even shown that the asymmetric plasticity rule present in STDP can itself produce predictions (but of course these are limited by the tens-of-millisecond timescales of the plasticity rule) (Mehta et al. 2000; Rao & Sejnowski 2001, 2002).

Eventually, the inclusion of a number of other modifications to the standard Hebbian rule, such as replacing postsynaptic activity with a reinforcement signal (reward minus expected reward or temporal difference), allows many models to accurately capture details of classical conditioning experiments, including the linking of seconds-long CS-US intervals and the production of the useful predictive responses, even with short ETs (Sutton 1988, Sutton & Barto 1990, Montague et al. 1996, Ludvig et al. 2008). As such, these models have been connected to the midbrain DA reward system and reinforcement learning in general (Houk et al. 1995, Montague et al. 1996, Schultz 1998, Sutton & Barto 2018).

However, perhaps because the above discussed models require the use of specially prescribed, and potentially fictional, network input patterns, the concept of ETs has persisted. Currently, ETs can be found in three-factor plasticity rules based on original STDP, except that now the repeated coincident pairing of pre- and postsynaptic activity does not itself induce plasticity but instead generates ETs that are then gated by a later-occurring neuromodulator to drive plasticity (Izhikevich 2007, Fremaux & Gerstner 2016, Gerstner et al. 2018). Here, by using STDP rules for credit assignment, the resulting plasticity is restricted to only those inputs that are correlated with AP output. Unfortunately, the biological data supporting these ideas are still somewhat uncertain, as the effect sizes are small and variable in the mammalian circuits examined (usually striatum but also neocortex) and the reported effective STDP-neuromodulator intervals are highly inconsistent [coincident (Yagishita et al. 2014), approximately 1 s (Fisher et al. 2017), exactly 2 s (Shindou 2019), and also coincident in the neocortex (He et al. 2015)] (see also Cassenaer & Laurent 2012).

In the end, the concept that STDP-like correlations actually generate ETs that function as a transient memory of these coincidences, instead of directly producing synaptic weight changes, extends the functionality of correlative-type plasticity by overcoming the distal reward problem. Furthermore, the use of neuromodulatory signals, linked to behavioral variables such as reward and novelty, allows some type of regulation or supervision of the learning-related plasticity. Now, plasticity only proceeds as long as a stimulus remains novel or associated with a reward or punishment. However, the non-synaptic volume transmission mode of neuromodulator release impedes the ability of ascending modulatory systems to selectively adjust the activity of individual neurons in a fine-grained manner. This global nature of neuromodulation thus limits the ability of three-factor plasticity rules to function in a truly supervised manner (Richards & Lillicrap 2019). In addition, certain patterns of input alone seem capable of generating ETs in particular circuits, raising the idea that input-output correlations may not always be the optimal structure on which to base credit assignment for learning adaptive behaviors (Bittner et al. 2017). At any rate, the preponderance of evidence indicates that neuromodulators impact synaptic plasticity, and the interesting relationships between behavioral variables and neuromodulatory area activities leave little doubt that neuromodulation could play a role in the regulation of learning-related neuronal changes. The exact nature of this role, however, remains to be determined, and it may turn out that neuromodulation works more in a permissive-type capacity than as a direct instructor (see below).

SUPERVISED SYNAPTIC PLASTICITY

A variation on standard Hebbian plasticity produces a true supervised form of learning. Although Hebb did not consider teaching or target inputs, these rules are still based on the product of pre- and postsynaptic activity, except that here, the postsynaptic term is the difference between the actual output and a target or teacher activity pattern (**Figure 1c**). Supervised learning rules are part of error-correcting algorithms, where networks are trained to optimize their performance by incrementally reducing this error or difference. As such, these systems require the formulation of a goal or objective for learning (objective functions); the determination of how far current network activity is from that objective; and, finally, a method for minimizing this difference. The objective functions can be overt and externally driven, as in a teacher that explicitly indicates the exact target pattern, or something less specific, internally generated, and more dynamic, such as a target toward sparse output activity or the minimum information needed to encode sensory input (Olshausen & Field 1996, Rao & Ballard 1999, Marblestone et al. 2016). Either way, there is a goal and an incremental process to achieve that target.

Even the simplest single-layer networks using supervised learning procedures (**Figure 2a,b**) are quite powerful and are in widespread use in signal processing and adaptive control (Woodrow

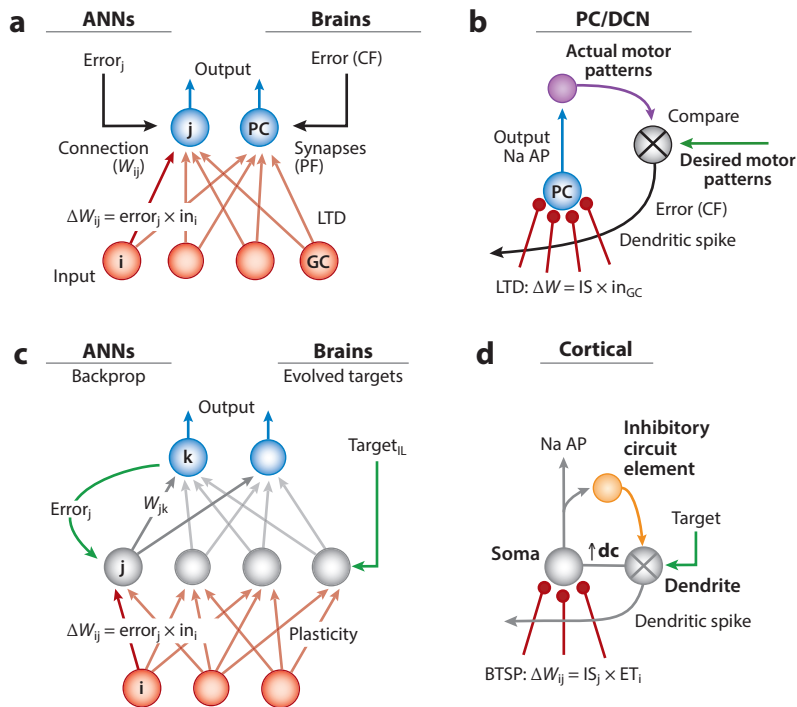


Figure 2

Supervised learning in ANNs and brain regions. (a) Single-layer network with input (red) and output (blue) units showing ANN elements in a single-layer net and potential neuronal equivalents in brains. (b) Single PC (blue) receiving GC input (red) and AP output. Here, error signal calculation or comparison (Compare), the difference between the desired (green) and actual (purple) motor patterns, is performed external (black circle) to the PC and delivered via CF (black arrow), which drives dendritic spiking and standard LTD. (c) Multilayer network where the calculation of intermediate-layer (gray) weight changes (ΔW_{ij}) uses an error signal based on properties of the subsequent layer (left side, green). Such information seems unlikely to be available for real neurons and, as such, specialized target signals may be used (right side, green). (d) Cortical circuits, which may use specialized target input that varies from neuron to neuron to produce dendritic V_m changes (dendrite) that drive dendritic spikes and LTP of feedforward input until somatic output (soma) matches the target-canceling local error signal through inhibitory circuit elements (yellow). Abbreviations: ANN, artificial neural network; AP, action potential; backprop, backpropagation of error algorithm; CF, climbing fiber; dc, dendritic compartmentalization; DCN, deep cerebellar nuclei; GC, granule cell; LTD, long-term depression; LTP, long-term potentiation; PC, Purkinje cell; PF, parallel fiber.

et al. 1994). While the incorporation of additional intermediate or hidden layers allows for more complex internal representations and significantly more powerful functionality, the learning rules used are essentially the same (Woodrow & Hoff 1960, Rumelhart et al. 1986, LeCun et al. 2015) (Figure 2c–e). The main difference in multilayer or deep networks is that the method employed to achieve the objective is more involved, since it is also necessary to calculate how to change intermediate-layer connection weights such that those changes improve final network output. While there are relatively straightforward methods to accomplish this in ANNs (Rumelhart et al. 1986, LeCun et al. 2015), the commonly used back-propagation of error algorithm requires knowledge of downstream weights, the derivatives of output nonlinearities, and other aspects that have led authors to question its direct biological plausibility (Crick 1989, Bengio et al. 2016). However, this is an active research area, and recent work suggests that there may be more compatibilities

between the biology and theory than were originally appreciated (Bengio et al. 2016, Lillicrap et al. 2016, Roelfsema & Holtmann 2018, Whittington & Bogacz 2019).

There is, in fact, a rising appreciation that many neuronal circuits could and indeed may engage in certain forms of supervised learning, although the implementation may be quite distinct from the artificial types. The most familiar example is that of the mammalian cerebellum and associated nuclei (**Figure 2a,b**). These areas have been frequently modelled as single-layer neuronal networks where the adaptive elements (Purkinje neurons) receive two different inputs. The first is delivered by tens of thousands of feedforward synaptic inputs (parallel fibers) from granule cells whose weights are adjusted by a standard anti-Hebbian LTD to modulate axonal Na^+ AP output. The other input is a single very strong climbing fiber (CF) feedback that drives dendritic Ca^{2+} spikes, which in turn induce the above LTD (Kawato 1990, Ito 2008, Raymond & Medina 2018). Recent evidence suggests that ETs may also be present within Purkinje neurons to link the PF input with temporally delayed CF feedback (Suvrathan et al. 2016, Suvrathan 2019). Thus, in contrast to standard correlative learning forms, there are two separate voltage signals at work here: one that controls learning (dendritic Ca^{2+} spikes) and another that transmits the activation level of the adaptive element to downstream areas (axonal Na^+ APs). The activity of the CF input is thought to be related to a comparison, performed external to the cerebellum, of the desired and actual movement trajectories and therefore is considered to function as a type of error signal whose minimization is the objective of the system (Kawato 1990, Lang et al. 2016, Kostadinov et al. 2019). That each Purkinje neuron receives only a single CF input allows for a fine-grained synaptic credit assignment within the layer by allowing the error signal to drive plasticity within specific individual neurons (Richards & Lillicrap 2019). Thus, the single-layer cerebellum uses an external comparator to calculate and deliver an error signal that, by varying from neuron to neuron, can appropriately drive synaptic plasticity within different Purkinje neurons such that the error signal is minimized as the network output is optimized.

There are three new elements introduced here that are not found in the standard correlative plasticity forms considered above. The first is the concept of an objective output pattern to be achieved. The second is an additional excitatory input from another brain region that we call an instructive pathway (**Figure 1c**), since it carries information related to reaching the objective. Third is an additional voltage signal within neuronal dendrites (dendritic spikes) that are primarily involved in the induction of synaptic plasticity. We call this third element an instructive signal (**Figure 1c**), since these dendritic spikes are evoked by instructive pathway input and trigger learning-related synaptic plasticity. We can contrast this situation with that seen in correlative Hebbian forms, where the output APs responsible for inducing plasticity are driven by the same input that is being modified by learning-related weight changes (**Figure 1a,b** versus **Figure 1c,d**). The additional instructive elements play a critical role in the cerebellum's ability to improve performance by incrementally optimizing its output (Raymond & Medina 2018).

A variation on this theme that is potentially more relevant for other brain areas has been heavily studied in cerebellar-like structures in nonmammalian species. These structures appear similar to a multilayer network, which theoretically requires them to make or find a signal that can be used to appropriately adjust hidden-layer synaptic weights (Bell et al. 2008, Enikolopov et al. 2018, Muller et al. 2019) (**Figure 2c,d**). Thus, it is perhaps not surprising that instead of a standard error signal, the instructive pathway for the intermediate layer appears to bring a specialized target directly to the dendrites of these neurons, where an internal comparison is performed between the target and the feedforward synaptic input (Knudsen 1994, Enikolopov et al. 2018, Muller et al. 2019) (**Figure 2c,d**). Any deviation between the target and the feedforward input alters dendritic spiking, thereby generating an internal error signal that drives adaptive synaptic plasticity in each intermediate-layer neuron (Bell et al. 1997). Since the instructive pathway to the hidden layer

carries a target and not an error signal, the activity in the instructive pathway does not change in response to the generation of appropriate network output. What does change is the internal error signal, that is, the rate of dendritic spiking in the individual intermediate-layer neurons. This is an example where nature has discovered a target whose use in driving instructive signals in intermediate-layer neurons is guaranteed to reduce the output error when combined with the correct type of synaptic plasticity and circuit wiring.

The new elements introduced by the above circuit, a specialized target input driving local error signals within the dendrites of each individual intermediate-layer neuron, may also be applicable to learning in other circuits, such as those in mammalian cortical areas (e.g., neocortex and hippocampus). Cortical areas are part of complex, perhaps hierarchical, networks that appear to form progressively higher-order representations (Hubel & Wiesel 1962, 1968; Cadieu et al. 2014; Issa et al. 2018). This organization may also require them to find specialized signals that are useful in promoting proper credit assignment within what are the equivalent of intermediate-layer cortical synapses. What mechanisms are available to cortical areas that could allow them to operate within such a supervised learning theoretical framework (**Figure 2c,d**)?

The individual pyramidal neurons composing these networks receive both feedforward excitatory inputs from lower levels (on the order of tens of thousands of synapses) and feedback excitatory synapses (thousands) from a variety of brain regions, some of which could be considered farther up a hierarchy (Steward & Scoville 1976, Cauller et al. 1998, Megias et al. 2001, Petreanu et al. 2009, Fiser et al. 2016, Marques et al. 2018) (**Figure 3a, subpanel i**). The cable filtering and active membrane properties of cortical pyramidal neuron dendrites produce at least two separate integrative compartments, one around the soma (the perisomatic or proximal region) that is near the axonal output region of the neuron, and another in the apical tuft dendrite (or distal region) that is as far away from the output site of the neuron as possible (Magee 1998, Stuart & Spruston 1998, Harnett et al. 2013, Harnett et al. 2015) (**Figure 3a**). The proximal compartment mainly receives the feedforward input, while the feedback primarily innervates the distal compartments (**Figure 3a**). Each compartment also has its own feedforward and feedback inhibitory elements




Figure 3

Cortical microcircuit. (*a, i*) Pyramidal cell-based principal neuron with proximal and distal integrative compartments. Each compartment receives separate excitatory input (*red, blue*) and inhibitory (*ii*) feedforward (FF) and (*iii*) feedback (FB) inputs. These compartments are also under inhibitory control by (*iv*) disinhibitory elements as well as neuromodulatory inputs (*v*). Active signals are generated in (*b*) distal (Ca²⁺ plateau) and (*c*) proximal (Na⁺ action potentials) compartments.

(**Figure 3a**, subpanels *ii,iii*). The level of dendritic compartmentalization is such that the interaction between these distinct feedforward and feedback excitatory inputs is minimal. In addition, each integrative compartment generates its own active signal, where the proximal compartment initiates Na⁺ APs in the axon (Turner et al. 1991, Stuart & Sakmann 1994) (**Figure 3c**), while the distal compartment produces Ca²⁺ plateau potentials (Schiller et al. 1997, Larkum et al. 1999, Magee & Carruth 1999, Xu et al. 2012) (**Figure 3b**). These dendritic plateaus are mediated by voltage-gated Ca²⁺ channels and NMDAR currents whose activation requires strong depolarization delivered by the distal excitatory input (Tsay et al. 2007, Takahashi & Magee 2009, Xu et al. 2012, Williams & Fletcher 2019). Plateau probability and duration (from tens to hundreds of milliseconds) are subject to inhibitory regulation (**Figure 3a**, subpanel *iv*) and neuromodulatory control (**Figure 3a**, subpanel *v*), with noradrenaline (NA) and acetylcholine (ACh) producing decreases in threshold and increases in duration, respectively (Palmer et al. 2012, Royer et al. 2012, Harnett et al. 2013, Milstein et al. 2015, Labarrera et al. 2018, Williams & Fletcher 2019, Williams & Holtmaat 2019). As in the cerebellar structures, cortical dendritic plateau potentials are very effective drivers of synaptic plasticity (Golding et al. 2002; Kampa et al. 2006; Sjöström & Häusser 2006; Takahashi & Magee 2009; Gambino et al. 2014; Bittner et al. 2015, 2017).

Several computational studies have suggested that the cortical circuit architecture and physiology discussed above, where feedback inputs carrying information related to the output of network innervate the distal dendritic regions in relative isolation from the feedforward input, allow error-like signals to be locally calculated within the dendrites of individual hidden-layer neurons (Körding & König 2001, Urbanczik & Senn 2014, Guerguiev et al. 2017, Sacramento et al. 2018). Since these local error signals, perhaps mediated by dendritic plateau potentials, reflect the difference between the desired and actual network output, intermediate-layer neurons are able to adjust their feedforward synaptic weights as necessary to enhance the final output of the network. In addition, regulation of dendritic voltage signals by inhibitory and disinhibitory microcircuit elements as well as neuromodulatory input could provide these components a permissive control over learning, or even a direct role in credit assignment (Roelfsema & Holtmaat 2018, Sacramento et al. 2018). Thus, the physiology of cortical-type pyramidal neurons and the circuits in which they are embedded may allow them to implement the neuron-by-neuron attribution of credit at intermediate layers that could also be regulated by attention or other factors (Richards & Lillicrap 2019). Of course, in the end it is completely unknown if cortical circuits are functioning within a supervised learning-type framework; however, it has come to light recently that they do appear to possess several mechanisms that could allow them to perform some of the required computations. The above observations justify continued research along these lines.

BEHAVIORAL TIMESCALE SYNAPTIC PLASTICITY

A form of synaptic plasticity that appears to uniquely combine several features of the above learning forms such as instructive pathways, instructive signals, seconds-long ETs, and neuromodulation has been found to underlie place fields in the CA1 hippocampal region (**Figure 1d**). In addition, this plasticity is rapidly induced (within as few as a single trial), allowing for single-shot forms of learning and memory formation. The seconds-long time course suggests the name: behavioral timescale synaptic plasticity (BTSP) (Bittner et al. 2015, 2017; Grienberger et al. 2017; Zhao et al. 2019). Remarkably, at least for cortical areas, BTSP appears to break the primary Hebbian requirement, because the weight changes driven by it are not proportional to input-output correlations and there is a solid biophysical basis for this non-Hebbian property. We devote much of the remainder of this review to a detailed discussion of the basic properties of BTSP as well as its functional implications.

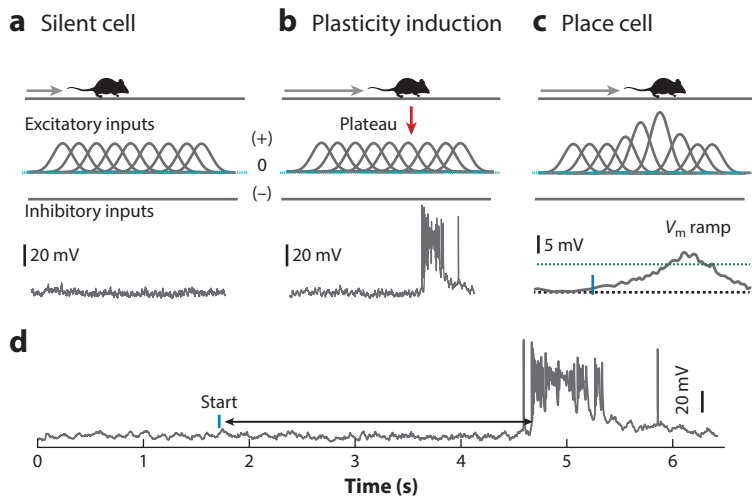


Figure 4

CA1 place field formation via behavioral timescale synaptic plasticity. (a) Balance between unpotentiated, tuned CA3 excitatory inputs and a uniform, relatively constant level of inhibition. (b) Initiation of plateau potential drives plasticity. (c) Plasticity increases the weight of a set of tuned inputs that in turn produce membrane potential (V_m) ramp. Green line is action potential (AP) threshold. (d) Arrow indicates how far back in time inputs are potentiated before plateau. The start position of V_m ramp is indicated by blue lines in panels c and d. Note the lack of APs (or even somatic depolarization) anywhere other than those driven by the plateau itself.

Recent data suggest that CA1 place fields are formed by BTSP according to the following scheme (**Figure 4**). It appears that all CA1 pyramidal neurons receive a constant barrage of tuned excitatory inputs, perhaps from upstream place cells in area CA3 (Bittner et al. 2015, Grienberger et al. 2017, Davoudi & Foster 2019, Zhao et al. 2019). In addition, the local population of inhibitory interneurons in CA1 produces a sustained level of inhibition that more or less exactly balances this excitation, leaving the soma of these neurons fluctuating around their resting V_m (Grienberger et al. 2017). At this point, these neurons are so-called silent cells that do not exhibit any selective firing in a given environment (Harvey et al. 2009, Epsztein et al. 2011) (**Figure 4a**). At some location in the environment, a Ca^{2+} plateau potential is initiated in the distal apical dendrites of a silent CA1 cell, and the large, long-lasting depolarization associated with the plateau effectively spreads throughout the neuron (Grienberger et al. 2014), inducing a potentiation of excitatory synaptic inputs that arrive around the plateau (Golding et al. 2002; Bittner et al. 2015, 2017) (**Figure 4b**). On the next trial, the enhanced weights of the now potentiated inputs drive the V_m out of excitation-inhibition balance to produce a slow ramp of depolarization that eventually crosses the threshold when the animal is near the location of the initial plateau potential, thus producing location-specific or place field firing (Harvey et al. 2009, Epsztein et al. 2011, Bittner et al. 2015, Diamantaki et al. 2018) (**Figure 4c**). Because of this arrangement, every CA1 pyramidal neuron is capable of rapidly expressing any place field relevant to a given environment with the exact location of the field being determined, perhaps solely, by the location of the plateau potential. These properties potentially provide a mechanism whereby the hippocampal place field representation in CA1 can be shaped by the behavioral experience of the animal (Bittner et al. 2015, 2017; Grienberger et al. 2017; Fischler et al. 2019; Zhao et al. 2019).

What is the evidence that synaptic plasticity is responsible for CA1 place field firing? First, plateau-driven synaptic plasticity is found in CA1, and it shows the same large magnitude

(threefold peak potentiation), extended time course (asymmetric, seconds-long timescale), and pharmacology (sensitive to NMDAR and Ca^{2+} channel antagonists) as plateau-driven CA1 place field formation (Bittner et al. 2015, 2017). Second, V_m fluctuation analyses suggest that the depolarization responsible for driving place field firing is the result of an increased input amplitude as opposed to input frequency. Furthermore, the fact that place field induction is so rapid (within a single trial, <10 s) is inconsistent with the amount of rewiring that would be required to double or triple the number (or frequency) of similarly tuned CA3 inputs received by the CA1 neuron (Bittner et al. 2015). Finally, data indicate that the level of inhibition received by a given CA1 pyramidal neuron does not change following place field induction, strongly suggesting that location-dependent changes in inhibitory input are not involved (Grienberger et al. 2017).

The asymmetric, seconds-long time course of BTSP is interesting and unique compared to most plasticity forms discussed above. Modelling results demonstrate that this time course could arise through the interaction of two biochemical signals, one a slowly decaying ET generated by active input from presynaptic place cells and another gating-type signal evoked by the plateau potential that decays somewhat faster (~ 1 s) (Bittner et al. 2017). The overlap of these two signals initiates synaptic plasticity expression mechanisms that are perhaps the same as those found for standard LTP (Luscher & Malenka 2012, Kandel et al. 2014).

While it seems fairly straightforward for a plateau potential actively propagating throughout the neuron to produce a global biochemical signal (Larkum et al. 1999, Takahashi & Magee 2009, Grienberger et al. 2014, Li et al. 2016), it is less obvious how long-lasting ETs are generated, since there are essentially no APs preceding the plateau potentials that induce CA1 place fields (Bittner et al. 2017) (**Figure 4d**). In this case, ETs are produced by subthreshold synaptic input alone, and conventionally this would not only be biophysically problematic but also raise certain conceptual issues concerning the appropriate limitation of synapse eligibility (i.e., to only properly tuned inputs). As far as the mechanisms go, several studies have found that pyramidal neuron spine neck resistance is relatively high (approaching $1\text{ G}\Omega$) and, as a result, average-sized unpotentiated excitatory postsynaptic potentials reach ~ 25 mV in spines (Bloodgood et al. 2009, Gullledge et al. 2012, Harnett et al. 2012, Jayant et al. 2017, Beaulieu-Laroche & Harnett 2018). This level of depolarization is enough to begin gating NMDARs, and the addition of a small amount of depolarization from neighboring active synapses nonlinearly increases this activation (Mayer et al. 1984, Harnett et al. 2012) (**Figure 5a**). Thus, repetitive stimulation of individual or small groups of synapses alone, without any correlated AP output, is adequate to generate an NMDAR-mediated Ca^{2+} entry that, in turn, could activate a biochemical signaling molecule functioning as an ET (e.g., CaMKII) (Lee et al. 2009, Yagishita et al. 2014) (**Figure 5b**). In addition, plasticity of spine neck resistance (Bloodgood & Sabatini 2005, Grunditz et al. 2008) can protect synapses from saturation following potentiation without adversely affecting the current delivered (Johnston & Wu 1995), and an inverse relationship between spine resistance and synaptic conductance has been reported in cortical pyramidal neurons (Beaulieu-Laroche & Harnett 2018) (**Figure 5c**). These same mechanisms also allow neurons to deal with the above conceptual issues, since neurons possessing high-resistance spines could implement credit assignment rules based on the spatiotemporal structure of the synaptic input, where only those inputs showing a sufficient amount of repeated activation or spatial clustering will drive ETs (Harvey & Svoboda 2007, Yu et al. 2009, Makino & Malinow 2011, Harnett et al. 2012, Druckmann et al. 2014; Kerlin et al. 2018). Although a small amount of theory work has explored the use of ETs driven by input alone in reinforcement learning (so-called noncontingent ETs) (Sutton & Barto 2018), additional thought on the unique credit assignment capabilities provided by noncontingent ETs is justified. For instance, action-related areas (e.g., striatum) might use a credit assignment rule based on causality

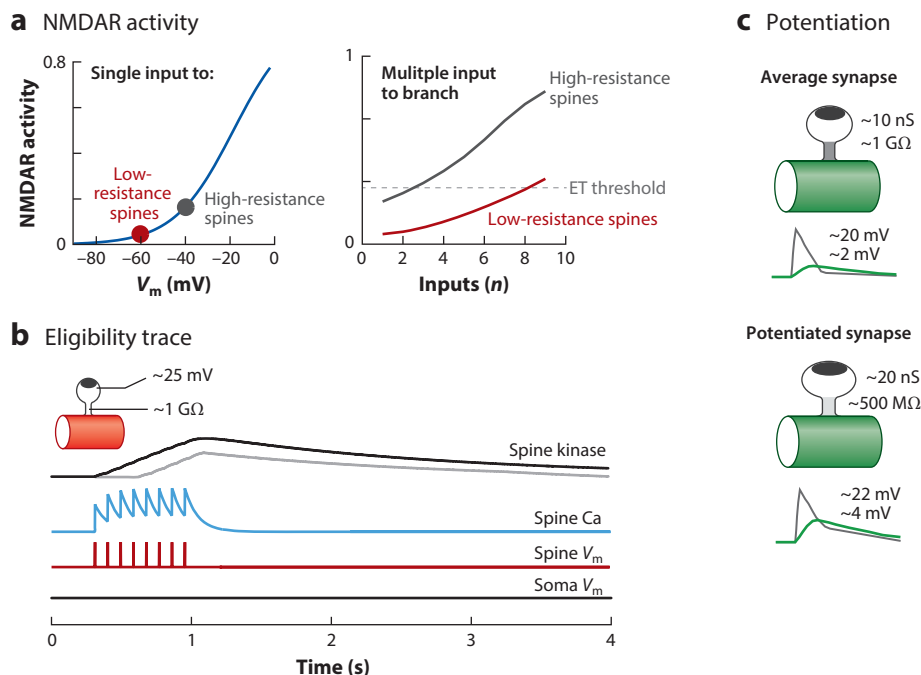


Figure 5

High-resistance spines eliminate the input-output pairing requirement. (a) NMDA receptor (NMDAR) activation by a single input is increased by the excitatory postsynaptic potential (EPSP) amplitude boosting of high-resistance spines. Enhanced input is cooperatively reflected in increased nonlinear interaction between nearby inputs. High-resistance spines lower the number of neighboring inputs required to activate the eligibility trace (ET). (b) A few coactive nearby inputs are sufficient to evoke ETs without the need of backpropagating action potentials. The gray trace represents spine kinase with lower sensitivity to spine Ca^{2+} , requiring a repetitive synaptic input (five times) before ET activation. (c) Decreasing spine neck resistance after potentiation can counter the impact of an increased conductance on local spine EPSP amplitude without adversely affecting the elevated current delivered to the rest of the neuron. Spine membrane potential (V_m) is shown in black, and dendrite V_m is shown in green.

(i.e., pre/postactivity correlation), while areas more involved in representing external stimuli or in episodic memory formation and one-shot learning (e.g., cortical regions) could take advantage of ETs driven by input alone, as the causality requirement or contingency is less relevant (Sutton & Barto 2018).

Another consequential element of BTSP, reminiscent of the learning-related plasticity found in cerebellar structures, is the ability of dendritic plateau potentials to act as an instructive signal to induce synaptic plasticity in a feedforward pathway. This raises the obvious question of what is driving dendritic Ca^{2+} plateaus in CA1? As discussed above for cortical circuits, the current biophysical understanding suggests that a combination of membrane excitability enhancing neuromodulation (ACh, NA), reduced levels of dendritic inhibition, and strong excitatory input from the thousands of entorhinal cortex feedback projections (EC3) targeting the distal integrative compartment are involved in regulating plateau initiation (Witter et al. 2000, Suh et al. 2011, Lovett-Barron et al. 2012, Bloss et al. 2018, Labarrera et al. 2018, Turi et al. 2019, Williams & Fletcher 2019, Kauffman et al. 2020). Given this, it appears that a complex circuit-level computation is involved in generating what could be a fine-grained instructive signal permissively

controlled by behavioral variables delivered by neuromodulatory input to the circuit (Turi et al. 2019, Kauffman et al. 2020). Determining what information and behavioral variables are represented by the various elements involved could shed light on how proper credit assignment is achieved in CA1. Future experiments should also explore the possibility that, in analogy to the cerebellum-like structures, EC3 provides CA1 with a specialized target that generates local error signals mediated by plateau potentials in the apical dendrites of pyramidal neurons (Figure 2c).

In sum, the synaptic plasticity responsible for CA1 place fields appears to combine several signature elements of the forms discussed above but with substantial alterations. When compared with standard Hebbian plasticity types, BTSP is not driven by input-output correlations; it changes the weights of inputs over a period of time that covers seconds instead of milliseconds and only requires a single instance for induction, as opposed to the many dozens of repetitions required in STDP protocols. The extended time course suggests an ET-like filtering of input that, in contrast to the three-factor rules thought to be in action in the striatum, requires certain patterns of input alone. Furthermore, although it is still incompletely understood, neuromodulation does not appear to be functioning as a third factor in BTSP but instead, perhaps, plays more of a permissive role in regulating when the plasticity can occur by altering the probability of plateau potential initiation or the generation of ETs. In addition, an instructive-like signal reminiscent of that observed in cerebellar areas appears to control BTSP through an interaction with ETs, except that in CA1 this signal is so strong that it requires only a single trial for induction. The signal itself is also somewhat different from that found in Purkinje neurons in that it potentially reflects a calculation of error within the individual pyramidal neurons themselves, as opposed to the externally generated error signal that is subsequently brought to Purkinje neurons by CFs. Thus, CA1 appears to possess a directed plasticity that uses ETs produced by input alone and a very strong local instructive signal generated in the dendrites of individual neurons by a second input carrying information related to overall hippocampal output.

FUNCTIONS OF BEHAVIORAL TIMESCALE SYNAPTIC PLASTICITY

The unique features of this synaptic plasticity produce several valuable functions within area CA1. First, the capability to form accurate predictions of near-future events and states is highly beneficial to an agent operating in a complex environment. Asymmetric plasticity rules automatically produce predictions of the future, and indeed, BTSP produces place fields that have a peak and center of mass shifted some 10–15 cm before the location in the environment where the plateau was initiated (Abbott & Blum 1996, Bittner et al. 2017) (Figure 6). That is, CA1 place fields are encoding something about a location approximately 1 s ahead of the mouse. If this interpretation is correct, the output of hippocampus is not actually a cognitive map of the animal's current location; instead, the hippocampus has computed and is transmitting some quantity concerning a property of the near future, perhaps the value of a future state (Stachenfeld et al. 2017, Wayne et al. 2018).

The timescales that the rule is operating on (seconds versus milliseconds) as well as the asymmetry in the kernel primarily determine how much of a prediction, if any, is produced (Abbott & Blum 1996) (Figure 6). This property perhaps provides an explanation of the right (acausal) side of the BTSP plasticity curve. Since the shape of the BTSP time course determines the amount of prediction produced by the rule, it appears that CA1 arrived at a rule to produce exactly the observed level of prediction. Thus, different circuits could be tuned to different levels of prediction depending upon their functional requirements, and all this necessitates is a judicious use of either

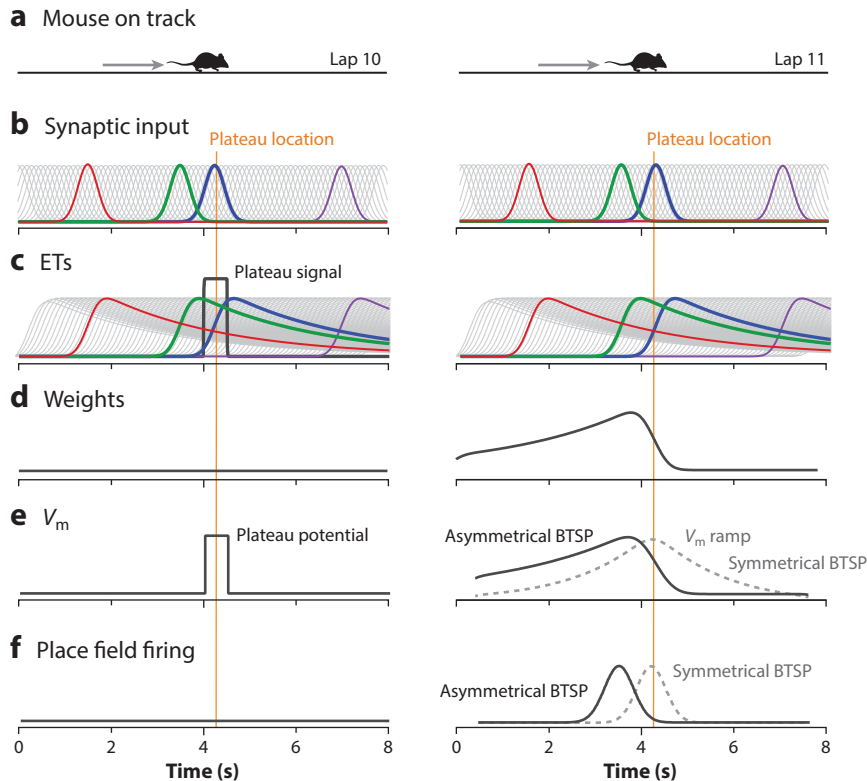


Figure 6

Eligibility trace (ET) and gating signal kinetics determine the level of prediction. On the left side of each panel is a silent cell with the first occurrence of a plasticity-inducing plateau near the center of the track (lap 10). (a) Mouse on track with arrow indicating direction of travel. (b) Tuned synaptic inputs arriving as mouse enters presynaptic neurons' place field. Colored traces in panels *b* and *c* highlight four particular inputs, and all others are gray. (c) ETs (convolution of synaptic inputs in panel *b* and exponential with $\tau_{ET} = 3$ s) filter the input, such that synapses active (green in panel *b*) hundreds of milliseconds before the plateau potential produce ETs (green in panel *c*) that overlap more with the plateau gating signal (black; convolution of square step in panel *e* and exponential with $\tau_P = 0.1$ s) than inputs arriving at exactly the same time (blue in panels *b,c*). For lap 10, synaptic weights are unpotentiated in panel *d*, membrane potential (V_m) trace shows only plateau in panel *e*, and there is no place field-related action potential firing in panel *f*. On the right side of each panel, the next lap (lap 11) is shown. The difference in the overlap of ETs and gating signals increases the weights of synapses arriving before the plateau more than those arriving coincidentally (panel *d*). Thus, information arriving current with the plateau is stored in the weights of synaptic inputs that arrived seconds in the past. The weight profile heavily shifts V_m ramp (dark gray line in panel *e*), causing the resulting place cell to fire well before the location of the initiating plateau (orange line in panels *b-f*), thus reporting on an aspect of future locations. Compare predictive V_m ramp and place field (dark gray line) with those produced if the two signals have the same time constants (gray dashed lines in panels *e,f*; lines were generated by running the simulation again with $\tau_{ET} = 1.5$ s, $\tau_P = 1.5$ s). Here, a symmetric V_m ramp produces a place field that reports on the present location. The duration of the left side of the BTSP time course is set by τ_{ET} , while the right side is set by τ_P . When τ_{ET} is greater than τ_P , the V_m ramp shifts backwards; if they are equal there is no shift, and finally, if τ_P is greater than τ_{ET} , the resulting V_m ramp will actually shift to the right.

different signaling molecules or posttranslational modifications to produce the desired activation time courses (Suvrathan et al. 2016, Raymond & Medina 2018, Suvrathan 2019).

A second function provided by plateau-driven plasticity is the ability to tailor CA1 population activity to the environmental experience of the animal such that behavior is potentially improved. Of course, this appears to be similar to the error-correcting abilities of a supervised form of plasticity. Recent data suggest that CA1 may be involved in producing a map related to the value of a given environment, where salient regions are more densely represented by particular neurons (Hollup et al. 2001, Dupret et al. 2010, Zaremba et al. 2017, Turi et al. 2019). Putting this in the context of the known properties of plateau-driven plasticity in CA1, it is possible that the thousands of excitatory EC3 inputs received by each CA1 pyramidal neuron are able to provide a unique instructive input to every cell, which drives learning with the objective of producing hippocampal output that reconstructs network input to the hippocampus (Gluck & Myers 1993, Gluck & Granger 1993, Wayne et al. 2018). If this is correct, the EC3 should also show enhanced activity around particularly salient locations, as has recently been reported (Boccarda et al. 2019, Butler et al. 2019). In the end, it remains unknown if CA1 is operating within a supervised learning framework, but it does appear that a plasticity form combining elements of three-factor and supervised learning rules allows CA1 to generate unique internal representations of the world that are rapidly shaped by the exploratory experience of the animal to produce predictions of the near future. These fundamental properties of BTSP could potentially mediate the ability of the hippocampus to very rapidly (one-shot) learn and store episodes of events.

CLOSING

The successful learning of adaptive behaviors requires experience-driven alterations in the complex representations and computations that brains produce. These changes enhance the perception of relevant environmental objects and the generation via inference, prediction, and valuation of features that are not directly perceivable. Ultimately, this activity supports the appropriate decisions, action selection, and implementations that promote survival. Producing this range of brain functions is challenging for any intelligent system and may require types of neural plasticity that go beyond autonomous synaptic plasticity based exclusively on pre-post activity correlations. Instead, what seems to be needed for learning through experience are plasticity forms that enable a neuronal population to refine its activity toward a target or goal pattern that is, in the end, useful to the animal in a specific behavior. The recent discovery that such a directed type of plasticity (BTSP) is involved in shaping hippocampal place cell representations suggests that comparable forms may also be prevalent in cortical areas as well as in cerebellar structures.

Thus, future experiments should attempt to determine how widespread directed forms of plasticity are within the mammalian brain. Two of the more important concepts here are ETs and instructive signals, since these elements determine the impact that a plasticity will have on population activity. Along these lines, a comparison of the types of neuronal activity required to evoke ETs in various brain regions (e.g., action-related striatum versus representation-related cortical regions) could shed light on the basic logic used to drive learning-related changes (rules using pre-postcorrelation versus presynaptic activity patterns alone) (Sutton & Barto 2018). In addition, the kinetics of the ETs and perhaps their biochemical identity could vary from one region to another, depending on the computations performed (internal representations of the present versus predictions of the future) (Suvrathan et al. 2016, Bittner et al. 2017, Suvrathan 2019). As for instructive signals, determining what information feedback pathways are carrying and how their activity changes during learning will be critical to revealing the role of these pathways in adaptive behavior, particularly if these signals are involved in the optimization of something similar to an

objective function. Additionally, a thorough comparison of the basic physiological mechanisms of instructive signals and how they are modulated and regulated by microcircuit elements in a variety of regions will lead to an improved understanding of how behavioral and cognitive variables (e.g., attention) control learning in different areas.

In the past, the experimental examination of synaptic plasticity was heavily biased toward in vitro preparations, and an inability to accurately recreate the myriad conditions found in the brains of behaving animals may have led to the excessive variability present within the results. While in vitro work still has its place, the inclination in the future should be more toward in vivo experiments capable of determining the plasticity rules responsible for regional feature selectivity and its alteration during actual learning. Such experiments will surely rely on an ability to measure and manipulate intracellular V_m . While certain experimental tools are currently available, additional devices are needed for these experiments to be most fruitful (see review by Humeau & Choquet 2019). Furthermore, new and improved biochemical sensors and rapidly activatable antagonists for a variety of kinases and phosphatases will be beneficial in the hunt for ETs. Finally, improved labeling and recording methods to examine how the activity of feedback pathways to individual neurons change during learning would be extraordinarily useful. In the end, the endeavor to find the biological basis of adaptive behavior has had a good start. What is needed now is a refocusing of neuroscientists' attention toward more powerful and directed forms of learning-related plasticity.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

The authors thank L. Abbott, A. Milstein, S. Romani, N. Sawtell, and M. Shadlen for discussions and comments on the manuscript. This work was supported by the Howard Hughes Medical Institute.

LITERATURE CITED

- Abbott LF, Blum KI. 1996. Functional significance of long-term potentiation for sequence learning and prediction. *Cereb. Cortex* 6:406–16
- Andersen JA. 1972. A simple neural network generating an interactive memory. *Math. Biosci.* 14:197–220
- Beaulieu-Laroche L, Harnett MT. 2018. Dendritic spines prevent synaptic voltage clamp. *Neuron* 97:75–82
- Bell CB, Han V, Sawtell NB. 2008. Cerebellum-like structures and their implications for cerebellar function. *Annu. Rev. Neurosci.* 31:1–24
- Bell CB, Han V, Sugawara Y, Grant K. 1997. Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* 387:278–81
- Bengio Y, Lee D, Bornschein J, Mesnard T, Lin Z. 2016. Towards biologically plausible deep learning. arXiv:1502.04156 [cs.LG]
- Berke JD. 2018. What does dopamine mean? *Nat. Neurosci.* 21:787–93
- Bi GQ, Poo MM. 1998. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18:10464–72
- Bittner KC, Grienberger C, Vaidya SP, Milstein AD, Sacklin JJ, et al. 2015. Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. *Nat. Neurosci.* 18:1133–42
- Bittner KC, Milstein AD, Grienberger C, Romani S, Magee JC. 2017. Behavioral timescale synaptic plasticity underlies CA1 place fields. *Science* 357:1033–36

- Bliss TV, Lomo T. 1973. Long-lasting potentiation of synaptic transmission in the dentate area of the anesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232:331–56
- Bloodgood BL, Sabatini BL. 2005. Neuronal activity regulates diffusion across the neck of dendritic spines. *Science* 310:866–69
- Bloodgood BL, Giessel AJ, Sabatini BL. 2009. Biphasic synaptic Ca influx arising from compartmentalized electrical signals in dendritic spines. *PLOS Biol.* 7(9):e1000190
- Bloss EB, Cembrowski MS, Karsh B, Colonell J, Fetter RD, Spruston N. 2018. Single excitatory axons form clustered synapses onto CA1 pyramidal cell dendrites. *Nat. Neurosci.* 21:353–63
- Blum K, Abbott L. 1996. A model of spatial map formation in the hippocampus of the rat. *Neural Comput.* 8:85–93
- Boccaro C, Nardin M, Stella F, O'Neill J, Csicsvari J. 2019. The entorhinal cognitive map is attracted to goals. *Science* 363:1443–47
- Brito CSN, Gerstner W. 2016. Nonlinear Hebbian learning as a unifying principle in receptive field formation. *PLOS Comput. Biol.* 12(9):e1005070
- Brzosko Z, Zannone S, Schultz W, Clopath C, Paulsen O. 2017. Sequential neuromodulation of Hebbian plasticity offers mechanism for effective reward-based navigation. *eLife* 6:e27756
- Butler WN, Harcastle K, Giocomo LM. 2019. Remembered reward locations restructure entorhinal spatial maps. *Science* 363:1447–52
- Butts DA, Kanold PO. 2010. The applicability of spike time dependent plasticity to development. *Front. Synaptic Neurosci.* 2:30
- Cadiou CF, Hong H, Yamins DLK, Pinto N, Ardila D, et al. 2014. Deep neural networks rival the representation of primate IT cortex for core visual object recognition. *PLOS Comput. Biol.* 10(12):e1003963
- Caulley LJ, Clancy B, Connors BW. 1998. Backward cortical projections to primary somatosensory cortex in rats extend long horizontal axons in layer I. *J. Comp. Neurol.* 390:297–310
- Caporale N, Dan Y. 2008. Spike timing-dependent plasticity: a Hebbian learning rule. *Annu. Rev. Neurosci.* 31:25–46
- Carrillo-Reid L, Yang W, Bando Y, Peterka DS, Yuste R. 2016. Imprinting and recalling cortical ensembles. *Science* 353:691–94
- Cassenaer S, Laurent G. 2012. Conditional modulation of spike-timing-dependent plasticity for olfactory learning. *Nature* 482:47–53
- Crick F. 1989. The recent excitement about neural networks. *Nature* 337:129–32
- Davoudi H, Foster D. 2019. Acute silencing of hippocampal CA3 reveals a dominant role in place field responses. *Nat. Neurosci.* 22:337–42
- Diamantaki M, Coletta S, Nasr K, Zeraati R, Laturnus S, et al. 2018. Manipulating hippocampal place cell activity by single-cell stimulation in freely moving mice. *Cell Rep.* 23:32–38
- Druckmann S, Feng L, Lee B, Yook C, Zhao T, et al. 2014. Structured synaptic connectivity between hippocampal regions. *Neuron* 81:629–40
- Dupret D, O'Neill J, Pleydell-Bouverie B, Csicsvari J. 2010. The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat. Neurosci.* 13:995–1002
- El-Boustani S, Ip JPK, Brenton-Provencher V, Knott GW, Okuno H, et al. 2018. Locally coordinated synaptic plasticity of visual cortex neurons in vivo. *Science* 360:1349–54
- Enikolopov A, Abbot LF, Sawtell NB. 2018. Internally generated predictions enhance neural and behavioral detection of sensory stimuli in an electric fish. *Neuron* 99:135–46
- Epsztein J, Brecht M, Lee AK. 2011. Intracellular determinants of hippocampal CA1 place and silent cell activity in a novel environment. *Neuron* 70:109–20
- Erwin E, Miller KD. 1998. Correlation-based development of ocularly matched orientation and ocular dominance maps: determination of required input activities. *J. Neurosci.* 18:9870–95
- Feldman DE. 2012. The spike-timing dependence of plasticity. *Neuron* 75:556–71
- Fischler WM, Joshi NR, Devi-Chou V, Kitch L, Schnitzer MJ, Abbott LF, Axel R. 2019. Olfactory landmarks and path integration converge to form a cognitive spatial map. bioRxiv 752360. <https://doi.org/10.1101/752360>
- Fiser A, Mahringer D, Oyibo HK, Petersen AV, Leinweber M, Keller GB. 2016. Experience-dependent spatial expectations in mouse visual cortex. *Nat. Neurosci.* 19:1658–64

- Fisher SD, Robertson PB, Blak MJ, Redgrave P, Sagar MA, et al. 2017. Reinforcement determines the timing dependence of corticostriatal synaptic plasticity in vivo. *Nat. Comm.* 8:334
- Fremaux N, Gerstner W. 2016. Neuromodulated spike-timing-dependent plasticity and theory of three-factor learning rules. *Front. Neural Circuits* 9:85
- Fusi S. 2002. Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates. *Biol. Cybern.* 87:459–70
- Fusi S, Abbott LF. 2007. Limits on the memory storage capacity of bounded synapses. *Nat. Neurosci.* 10:485–93
- Gambino F, Grimaldo F, Lopez-Inesta E, Mehmani B, Squazzoni F. 2014. Sensory-evoked LTP driven by dendritic plateau potentials in vivo. *Nature* 515:116–19
- Gerstner W, Lehmann M, Liakoni V, Corneil D, Brea J. 2018. Eligibility traces and plasticity on behavioral time scales: experimental support of neo Hebbian three-factor learning rules. *Front. Neural Circuits* 12:53
- Gluck MA, Granger R. 1993. Computational models of the neural bases of learning and memory. *Annu. Rev. Neurosci.* 16:667–706
- Gluck MA, Myers C. 1993. Hippocampal mediations of stimulus representation: a computational theory. *Hippocampus* 3:491–516
- Golding NL, Staff NP, Spruston N. 2002. Dendritic spikes as a mechanism for cooperative long-term potentiation. *Nature* 418:326–31
- Grienberger C, Chen X, Konnerth A. 2014. NMDA receptor-dependent multidendrite Ca^{2+} spikes required for hippocampal burst firing in vivo. *Neuron* 81:1274–81
- Grienberger C, Milstein AD, Bittner KC, Romani S, Magee JC. 2017. Inhibitory suppression of heterogeneously tuned excitation enhances spatial coding in CA1 place cells. *Nat. Neurosci.* 20:417–26
- Grossberg S. 1976. On the development of feature detectors in the visual cortex with applications to learning and reaction-diffusion systems. *Biol. Cybern.* 21:145–59
- Grunditz A, Holbro N, Tian L, Zuo Y, Oertner TG. 2008. Spine neck plasticity controls postsynaptic calcium signals through electrical compartmentalization. *J. Neurosci.* 28:13457–66
- Gu Q. 2002. Neuromodulatory transmitter systems in the cortex and their role in cortical plasticity. *Neuroscience* 111:815–35
- Guerguiev J, Lillicrap TP, Richards BA. 2017. Towards deep learning with segregated dendrites. *eLife* 6:e22901
- Gulledge AT, Carnevale NT, Stuart GJ. 2012. Electrical advantages of dendritic spines. *PLOS ONE* 7:e36007
- Harnett MT, Magee JC, Williams S. 2015. Distribution and function of HCN channels in the apical dendritic tuft of neocortical pyramidal neurons. *J. Neurosci.* 29:516–29
- Harnett MT, Makara JK, Spruston N, Kath WL, Magee JC. 2012. Synaptic amplification by dendritic spines enhances input cooperativity. *Nature* 491:599–602
- Harnett MT, Xu NL, Magee JC, Williams S. 2013. Potassium channels control the interaction between active dendritic compartments in layer 5 cortical pyramidal neurons. *Neuron* 79:516–29
- Harvey CD, Collman F, Dombeck DA, Tank DW. 2009. Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature* 461:941–46
- Harvey CD, Svoboda K. 2007. Locally dynamic synaptic learning rules in pyramidal neuron dendrites. *Nature* 450:1195–200
- He K, Huertas M, Hong SZ, Tie X, Hell JW, et al. 2015. Distinct eligibility traces for LTP and LTD in cortical synapses. *Neuron* 88:528–38
- Hebb D. 1949. *The Organization of Behavior*. New York: Wiley
- Hollup SA, Molden S, Donnett JG, Moser M-B, Moser EI. 2001. Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *J. Neurosci.* 21:1635–44
- Hopfield JJ. 1982. Neural networks and physical systems with emergent collective computational abilities. *PNAS* 79:2554–58
- Houk JC, Adams JL, Barto AG. 1995. A model of how the basal ganglia generate and use neural signals that predict reinforcement. In *Models of Information Processing in the Basal Ganglia*, ed. JC Houk, JL Davis, DG Beiser, pp. 249–70. Cambridge, MA: MIT Press
- Hubel DH, Wiesel TN. 1962. Receptive fields, binocular interactions and functional architecture in the cat's visual cortex. *J. Physiol.* 160:106–54
- Hubel DH, Wiesel TN. 1968. Receptive fields and functional architecture of monkey striate cortex. *J. Physiol.* 195:215–43

- Hull CL. 1943. *Principles of Behavior: An Introduction to Behavior Theory*. Oxford, UK: Appleton-Century
- Humeau Y, Choquet D. 2019. The next generation of approaches to investigate the link between synaptic plasticity and learning. *Nat. Neurosci.* 22:1536–43
- Issa EB, Cadieu CF, DiCarlo JJ. 2018. Neural dynamics at successive stages of the ventral visual stream are consistent with hierarchical error signals. *eLife* 7:e42870
- Ito M. 2008. Control of mental activities by internal models in the cerebellum. *Nature* 9:304–13
- Izhikevich EM. 2007. Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cereb. Cortex* 17:2443–52
- Jayant K, Hirtz JJ, Plante IJ, Tsai DM, De Boer WD, Semonche A, et al. 2017. Targeted intracellular voltage recordings from dendritic spines using quantum-dot-coated nanopipettes. *Nat. Nanotechnol.* 12:335–42
- Johnston D, Wu SM-S. 1995. *Foundations of Cellular Neurophysiology*. Cambridge, MA: MIT Press
- Kampa B, Letzkus J, Stuart G. 2006. Requirement of dendritic calcium spikes for induction of spike-timing-dependent synaptic plasticity. *J. Physiol.* 574(1):283–90
- Kandel ER, Dudai Y, Mayford MR. 2014. The molecular and systems biology of memory. *Cell* 157:163–86
- Kauffman AM, Geiller T, Losonczy A. 2020. A role for the locus coeruleus in hippocampal CA1 place cell reorganization during spatial reward learning. *Neuron*. In press. <https://doi.org/10.1016/j.neuron.2019.12.029>
- Kawato M. 1990. Internal models for motor control and trajectory planning. *Curr. Opin. Neurobiol.* 9:718–27
- Kerlin A, Mohar B, Flickinger D, MacLennan B, Davis C, et al. 2018. Functional clustering of dendritic activity during decision-making. *bioRxiv* 440396. <https://doi.org/10.1101/440396>
- Kim WB, Cho J. 2017. Encoding of discriminative fear memory by input-specific LTP in the amygdala. *Neuron* 95:1129–46
- Knudsen EI. 1994. Supervised learning in the brain. *J. Neurosci.* 14:1985–97
- Koester HJ, Sakmann B. 1998. Calcium dynamics in single spines during coincident pre- and postsynaptic activity depend on relative timing of back-propagating action potentials and subthreshold excitatory postsynaptic potentials. *PNAS* 95:9596–601
- Kohonen T. 1972. Correlation matrix memories. *IEEE Trans. Comput.* 21:353–59
- Kohonen T. 1982. Self-organized formation of topologically correct feature maps. *Biol. Cybern.* 43:59–69
- Körding KP, König P. 2001. Supervised and unsupervised learning with two sites of synaptic integration. *J. Comput. Neurosci.* 11:207–15
- Kostadinov D, Beau M, Pozo MB, Häusser M. 2019. Predictive and reactive reward signals conveyed by climbing fiber inputs to cerebellar Purkinje cells. *Nat. Neurosci.* 22:950–62
- Labarrera C, Deitcher Y, Dudai A, Weiner B, Amichai AK, et al. 2018. Adrenergic modulation regulates the dendritic excitability of layer 5 pyramidal neurons in vivo. *Cell Rep.* 23:1034–44
- Lang EJ, Apps R, Bengtsson F, Cerminara NL, De Zeeuw CI, et al. 2016. The roles of the olivocerebellar pathway in motor learning and motor control. A consensus paper. *Cerebellum* 16:230–52
- Larkum ME, Zhu JJ, Sakmann B. 1999. A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* 389:338–41
- LeCun Y, Bengio Y, Hinton G. 2015. Deep learning. *Nature* 521:436–44
- Lee SR, Escobedo-Lozoya Y, Szatmari EM, Yasuda R. 2009. Activation of CaMKII in single dendritic spines during long-term potentiation. *Nature* 458:299–304
- Levy WB, Steward O. 1983. Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience* 8:791–97
- Li B, Tadross MR, Tsien RW. 2016. Sequential ionic and conformational signaling by calcium channels drives neuronal gene expression. *Science* 351:863–67
- Li Y, Van Hooser SD, Mazurek M, White LE, Fitzpatrick D. 2008. Experience with moving visual stimuli drives the early development of cortical direction selectivity. *Nature* 456:952–56
- Lillicrap TP, Crownden D, Tweed DB, Akerman CJ. 2016. Random synaptic feedback weights support error backpropagation for deep learning. *Nat. Comm.* 7:13276
- Lovett-Barron M, Turi GF, Kaifosh P, Lee PH, Bolze F, et al. 2012. Regulation of neuronal input transformations by tunable dendritic inhibition. *Nat. Neurosci.* 15(3):423–30
- Ludvig EA, Sutton RS, Kehoe EJ. 2008. Stimulus representation and the timing of reward-prediction errors in models of the dopamine system. *Neural Comput.* 20:3034–54

- Luscher C, Malenka RC. 2012. NMDA receptor-dependent long-term potentiation and long-term depression (LTP/LTD). *Cold Spring Harb. Perspect. Biol.* 4(6):a005710
- Magee JC. 1998. Dendritic hyperpolarization-activated currents modify the integrative properties of hippocampal CA1 pyramidal neurons. *J. Neurosci.* 18:7613–24
- Magee JC, Carruth M. 1999. Dendritic voltage-gated ion channels regulate the action potential firing mode of hippocampal CA1 pyramidal neurons. *J. Neurophysiol.* 82:1895–901
- Magee JC, Johnston D. 1997. A synaptically-controlled, associative signal for Hebbian plasticity in hippocampal neurons. *Science* 275:209–13
- Makino H, Malinow R. 2011. Compartmentalized versus global synaptic plasticity on dendrites controlled by experience. *Neuron* 72:1001–11
- Malenka RC, Bear M. 2004. LTP and LTD: an embarrassment of riches. *Neuron* 30:5–21
- Marblestone AH, Wayne G, Kording K. 2016. Toward an integration of deep learning and neuroscience. *Front. Comput. Neurosci.* 10:94
- Markram H, Lubke J, Frotscher M, Sakmann B. 1997. Regulation of synaptic efficacy by coincidence of postsynaptic AP and EPSPs. *Science* 275:213–15
- Marques T, Nguyen J, Fioreze G, Petreanu L. 2018. The functional organization of cortical feedback inputs to primary visual cortex. *Nat. Neurosci.* 21:757–64
- Mayer ML, Westbrook GL, Guthrie PB. 1984. Voltage-dependent block by Mg^{2+} of NMDA responses in spinal cord neurones. *Nature* 309:261–63
- McKernan MG, Shinnick-Gallagher P. 1997. Fear conditioning induces a lasting potentiation of synaptic currents in vivo. *Nature* 390:607–11
- Megias M, Emri Z, Freund TF, Gulyas AI. 2001. Total number and distribution of inhibitory and excitatory synapses on hippocampal CA1 pyramidal cells. *Neuroscience* 102:527–40
- Mehta MR, Quirk MC, Wilson MA. 2000. Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* 25:707–15
- Meliza CD, Dan Y. 2006. Receptive-field modification in rat visual cortex induced by paired visual stimulation and single-cell spiking. *Neuron* 49:183–89
- Milstein AD, Bloss EB, Apostolides PF, Vaidya SP, Dilly GA, et al. 2015. Inhibitory gating of input comparison in the CA1 microcircuit. *Neuron* 87:1274–89
- Montague PR, Dayan P, Sejnowski TJ. 1996. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J. Neurosci.* 16(5):1936–47
- Morris RG, Anderson E, Lynch GS, Baudry M. 1986. Selective impairment of learning and blockade of long-term potentiation by an *N*-methyl-D-aspartate receptor antagonist, AP5. *Nature* 319:774–76
- Muller SZ, Zadina A, Abbott LF, Sawtell NB. 2019. Continual learning in a multi-layer network of an electric fish. *Cell* 179(6):1382–92.e10
- Nabavi S, Fox R, Proulx CD, Lin JY, Tsien RY, Malinow R. 2014. Engineering a memory with LTD and LTP. *Nature* 511:348–52
- Nadim F, Bucher D. 2014. Neuromodulation of neurons and synapses. *Curr. Opin. Neurobiol.* 29:48–56
- Nakazawa K, McHugh TJ, Wilson M, Tonegawa S. 2004. NMDA receptors, place cells and hippocampal spatial memory. *Cell* 147:509–24
- Nicoll RA. 2017. A brief history of long-term potentiation. *Neuron* 93:281–90
- Oja E. 1982. Simplified neuron model as a principal component analyzer. *J. Math. Biol.* 15:267–73
- Olshausen BA, Field DJ. 1996. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381:607–9
- Palmer L, Murayama M, Larkum M. 2012. Inhibitory regulation of dendritic activity in vivo. *Front. Neural Circuits* 6:26–32
- Petreanu L, Mao T, Sternson SM, Svoboda K. 2009. The subcellular organization of neocortical excitatory connections. *Nature* 457:1142–45
- Pignatelli M, Bonci A. 2015. Role of dopamine neurons in reward and aversion: a synaptic plasticity perspective. *Neuron* 86:1145–57
- Rao RPN, Ballard DH. 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2:79–87

- Rao RPN, Sejnowski TJ. 2001. Spike-timing-dependent Hebbian plasticity as temporal difference learning. *Neural Comput.* 13:2221–37
- Rao RPN, Sejnowski TJ. 2002. Predictive coding, cortical feedback, and spike-timing dependent plasticity. In *Probabilistic Models of the Brain: Perception and Neural Function*, ed. RPN Rao, BA Olshausen, MS Lewicki, pp. 297–315. Cambridge, MA: MIT Press
- Raymond JL, Medina JF. 2018. Computational principles of supervised learning in the cerebellum. *Annu. Rev. Neurosci.* 41:233–53
- Richards BA, Lillicrap TP. 2019. Dendritic solutions to the credit assignment problem. *Curr. Opin. Neurobiol.* 54:28–36
- Roelfsema PR, Holtmaat A. 2018. Control of synaptic plasticity in deep cortical networks. *Nat. Rev. Neurosci.* 19:166–80
- Rogan MT, Stäubli UV, LeDoux JE. 1997. Fear conditioning induces associative long-term potentiation in the amygdala. *Nature* 390:604–7
- Rosenblatt F. 1959. Two theorems of statistical separability in the perceptron. In *Mechanisation of Thought Processes: Proceeding of a Symposium Held at the National Physical Laboratory on 24th, 25th, 26th and 27th November 1958*, Vol. 1, pp. 421–56. London: HM Station. Off.
- Royer S, Zemelman BV, Losonczy A, Kim J, Chance F, et al. 2012. Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition. *Nat. Neurosci.* 15:769–75
- Rumelhart DE, Hinton GE, Williams RJ. 1986. Learning internal representations by error propagation. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, ed. DE Rumelhart, JL McClelland, pp. 318–62 Cambridge, MA: MIT Press
- Ryan TJ, Roy DS, Pignatelli M, Arons A, Tonegawa S. 2015. Memory engram cells retain memory under retrograde amnesia. *Science* 348:1007–13
- Sacramento J, Costa RP, Bengio Y, Senn W. 2018. Dendritic cortical microcircuits approximate the back-propagation algorithm. In *Neural Information Processing Systems 31 (NIPS2018)*, ed. S Bengio, H Wallach, H Larochelle, K Grauman, N Cesa-Bianchi, R Garnett. San Diego, CA: NeurIPS
- Schiller J, Schiller Y, Stuart G, Sakmann B. 1997. Calcium action potentials restricted to distal apical dendrites of rat neocortical pyramidal neurons. *J. Physiol.* 505:605–16
- Schuetz S, Bonhoeffer T, Hubener M. 2001. Pairing-induced changes of orientation maps in cat visual cortex. *Neuron* 32:325–37
- Schultz W. 1998. Predictive reward signals of dopamine neurons. *J. Neurophysiol.* 80:1–27
- Schultz W. 2007. Multiple dopamine functions at different time courses. *Annu. Rev. Neurosci.* 30:259–88
- Shindou T, Shindou M, Watanabe S, Wickens J. 2019. A silent eligibility trace enables dopamine-dependent synaptic plasticity for reinforcement learning in the mouse striatum. *Eur. J. Neurosci.* 49:726–36
- Sjöström PJ, Häusser M. 2006. A cooperative switch determines the sign of synaptic plasticity in distal dendrites of neocortical pyramidal neurons. *Neuron* 51:227–38
- Song S, Abbott LF. 2001. Cortical development and remapping through spike-timing-dependent plasticity. *Neuron* 32:339–50
- Stachenfeld KL, Botvinick MM, Gersman SJ. 2017. The hippocampus as a predictive map. *Nat. Neurosci.* 20:1643–53
- Steinberg E, Keiflin R, Boivin J, Witten I, Deisseroth K, Janak P. 2013. A causal link between prediction errors, dopamine neurons and learning. *Nat. Neurosci.* 16:966–73
- Steward O, Scoville SA. 1976. Cells of origin of entorhinal cortical afferents to the hippocampus and fascia dentata of the rat. *J. Comp. Neurol.* 169:347–70
- Stuart G, Sakmann B. 1994. Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature* 367:69–72
- Stuart G, Spruston N. 1998. Determinants of voltage attenuation in neocortical pyramidal neuron dendrites. *J. Neurosci.* 18:3501–10
- Suh J, Rivest AJ, Nakashiba T, Tominaga T, Tonegawa S. 2011. Entorhinal cortex layer III input to the hippocampus is crucial for temporal association memory. *Science* 334:1415–20
- Sutton RS. 1988. Learning to predict by the methods of temporal differences. *Mach. Learn.* 3:9–44

- Sutton RS, Barto AG. 1981. Toward a modern theory of adaptive networks: expectation and prediction. *Am. Psych. Assoc.* 88:135–70
- Sutton RS, Barto AG. 1990. Time-derivative models of Pavlovian reinforcement. In *Learning and Computational Neuroscience: Foundations of Adaptive Networks*, ed. M Gabriel, J Moore, pp. 497–537. Cambridge, MA: MIT Press
- Sutton RS, Barto AG. 2018. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press
- Suvrathan A. 2019. Beyond STDP—towards diverse and functionally relevant plasticity rules. *Curr. Opin. Neurobiol.* 54:12–19
- Suvrathan A, Payne HL, Raymond J. 2016. Timing rules for synaptic plasticity matched to behavioral function. *Neuron* 92:959–67
- Takahashi H, Magee JC. 2009. Pathway interactions and synaptic plasticity in the dendritic tuft regions of CA1 pyramidal neurons. *Neuron* 62:102–11
- Tonegawa S, Liu X, Ramirez S, Redondo R. 2015. Memory engram cells have come of age. *Neuron* 87:918–31
- Tsay D, Dudman JT, Siegelbaum SA. 2007. HCN1 channels constrain synaptically evoked Ca^{2+} spikes in distal dendrites of CA1 pyramidal neurons. *Neuron* 56:1076–89
- Turi G, Li W, Chavlis S, Bozelos P, Pandi I, et al. 2019. Vasoactive intestinal polypeptide-expressing interneurons in the hippocampus support goal-oriented spatial learning. *Neuron* 101:1150–65
- Turner RW, Meyers DE, Richardson TL, Barker JL. 1991. The site of action potential discharge over the somatodendritic axis of rat hippocampal CA1 pyramidal neurons. *J. Neurosci.* 11:2270–80
- Urbanczik R, Senn W. 2014. Learning by the dendritic prediction of somatic spiking. *Neuron* 81:521–28
- von der Malsberg C. 1973. Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik* 14:85–100
- Wayne G, Hung C, Amos D, Mirza M, Ahuja A, et al. 2018. Unsupervised predictive memory in a goal-directed agent. arXiv:1803.10760 [cs.LG]
- Whittington JCR, Bogacz R. 2019. Theories of error back-propagation in the brain. *Trends Cogn. Sci.* 23(3):235–50
- Williams LE, Holtmaat A. 2019. Higher-order thalamocortical inputs gate synaptic long-term potentiation via disinhibition. *Neuron* 101:91–102
- Williams S, Fletcher LN. 2019. Dendritic substrate for the cholinergic control of neocortical output. *Neuron* 101:486–99
- Witter MP, Naber PA, van Haeften T, Machielsen WC, Rombouts SA, et al. 2000. Cortico-hippocampal communication by way of parallel parahippocampal-subicular pathways. *Hippocampus* 10:398–410
- Woodrow B, Hoff ME Jr. 1960. Adaptive switching circuits. *IRE WESCON Conv. Rec.* 4:96–104
- Woodrow B, Rumelhart D, Lehr MA. 1994. Neural networks: applications in industry, business and science. *Commun. ACM* 37:93–105
- Xu NL, Harnett MT, Williams SR, Huber D, O’Conner DH, et al. 2012. Nonlinear dendritic integration of sensory and motor input produces an object localization signal. *Nature* 492:247–51
- Yagishita S, Hayashi-Takagi A, Ellis-Davies GCR, Urakubo H, Ishii S, Kasai H. 2014. A critical time window for dopamine actions on the structural plasticity of dendritic spines. *Science* 345:1616–17
- Yao H, Dan Y. 2001. Stimulus timing-dependent plasticity in cortical processing of orientation. *Neuron* 32:315–23
- Young JM, Waleszczyk WJ, Wang C, Calford MB, Dreher B, Obermayer K. 2007. Cortical reorganization consistent with spike timing- but not correlation-dependent plasticity. *Nat. Neurosci.* 10:887–95
- Yu YC, Bultje RS, Wang X, Shi SH. 2009. Specific synapses develop preferentially among sister excitatory neurons in the neocortex. *Nature* 458:501–4
- Zaremba J, Diamantopoulou A, Danielson N, Groszmark AD, Kaifosh P, et al. 2017. Impaired hippocampal place cell dynamics in a mouse model of the 22q11.2 deletion. *Nat. Neurosci.* 20:1612–23
- Zhao X, Wang Y, Spruston N, Magee JC. 2019. Synaptic mechanisms of context-dependent sensory responses in the hippocampus. bioRxiv 624262. <https://doi.org/10.1101/624262>