

Hebbian plasticity requires compensatory processes on multiple timescales

Friedemann Zenke^{1,*} and Wulfram Gerstner²

January 16, 2017

1) Dept. of Applied Physics, Stanford University, Stanford, CA 94305, USA
orcid.org/0000-0003-1883-644X

2) Brain Mind Institute, School of Life Sciences and School of Computer and Communication Sciences, Ecole Polytechnique Fédérale de Lausanne, CH-1015 Lausanne EPFL, Switzerland
orcid.org/0000-0002-4344-2189

This article is published as:

Zenke, F., and Gerstner, W. (2017). Hebbian plasticity requires compensatory processes on multiple timescales. *Phil. Trans. R. Soc. B* 372, 20160259. DOI: 10.1098/rstb.2016.0259
<http://rstb.royalsocietypublishing.org/content/372/1715/20160259>

Abstract

We review a body of theoretical and experimental research on Hebbian and homeostatic plasticity, starting from a puzzling observation: While homeostasis of synapses found in experiments is a slow compensatory process, most mathematical models of synaptic plasticity use rapid compensatory processes. Even worse, with the slow homeostatic plasticity reported in experiments, simulations of existing plasticity models cannot maintain network stability unless further control mechanisms are implemented. To solve this paradox, we suggest that in addition to slow forms of homeostatic plasticity there are rapid compensatory processes which stabilize synaptic plasticity on short timescales. These rapid processes may include heterosynaptic depression triggered by episodes of high postsynaptic firing rate. While slower forms of homeostatic plasticity are not sufficient to stabilize Hebbian plasticity, they are important for fine-tuning neural circuits. Taken together we suggest that learning and memory rely on an intricate interplay of diverse plasticity mechanisms on different timescales which jointly ensure stability and plasticity of neural circuits.

Keywords: Hebbian plasticity, homeostasis, rapid compensatory processes, heterosynaptic plasticity, synaptic scaling, metaplasticity

*Author for correspondence (fzenke@stanford.edu)

Introduction

Homeostasis refers to a family of compensatory processes at different spatial and temporal scales whose objective is to maintain the body, its organs, the brain, or even individual neurons in the brain in a dynamic regime where they function optimally. A well-known example is the homeostatic regulation of body temperature in mammals, maintained at about 37 degrees Celsius independent of weather condition and air temperature. In neuroscience, homeostasis or homeostatic plasticity often refers to a compensatory process that stabilizes neural firing rates. In a classic experiment, cultured neurons that normally fire at, say 5Hz, change their firing rate after a modulation of the chemical conditions in the culture, but eventually return to their target rate of 5Hz during the following 24 hours [1]. Thus, the experimentally best-studied form of homeostasis acts on a timescale of hours to days. This slow form of homeostatic plasticity manifests itself as the rescaling of the efficacy or weight of all afferent synapses onto a single neuron by a fixed fraction, for instance 0.78. This phenomenon is called “synaptic scaling” [1].

Mathematical models of neural networks often make use of compensatory processes similar to synaptic scaling to stabilize firing rates in the presence of Hebbian plasticity. Hebbian plasticity is a form of synaptic plasticity which is induced by and further amplifies correlations in neuronal activity. It has been observed in many brain areas and can be induced quickly on a timescale of seconds to minutes. Its effect, however, is often long-lasting. It can last hours, days and possibly a lifetime. Due to these properties Hebbian plasticity is widely assumed to be the neural basis of associative long-term memory [2–4]. Moreover, Hebbian learning is thought to be the basis of developmental changes such as receptive field development [5–9].

However, Hebbian plasticity alone leads to a positive feedback loop in which correlations of pre- and postsynaptic firing drive potentiation of synapses that increase postsynaptic rates and correlations further, which is unstable. To avoid pathological run-away dynamics of neural activity in mathematical models, it is necessary to add appropriate constraints to plasticity models [10, 11]. A typical example of such a constraint is the normalization or rescaling of the sum of afferent synaptic weights: When the weight of one synaptic connection increases, weights of other connections onto the same neuron are algorithmically decreased to keep the total input constant or close to the optimal target regime. At a first glance, this form of multiplicative normalization [10] seems virtually identical to homeostatic “synaptic scaling” introduced above. However, these two mechanisms are fundamentally distinct because they act on vastly different timescales. While normalization in models typically takes place on a timescale of seconds or less [10, 12–14], in biology the effects of synaptic scaling manifest themselves only after hours [15, 16]. A similar observation holds for homeostatic metaplasticity, which exists on timescales ranging from some tens of minutes to days (Fig. 1; [17, 18]). Moreover, the difference between experimental data and models cannot be explained by a simple rescaling of time in the models, because the problem persists for quantitative plasticity models which capture the time course of biological data.




Figure 1: **The timescales of synaptic scaling or metaplasticity are faster in models than reported in experiments.** Here we plot the timescale of either synaptic scaling or homeostatic metaplasticity as used in influential modeling studies (gray). For comparison we plot the typical readout time for experimental studies on synaptic scaling and metaplasticity (red). Publications suffixed with * describe network models as opposed to the other studies which relied on single neurons. Note that the model marked with † by Toyoizumi et al. [23] is an interesting case which has both RCPs and a slow form of homeostasis. Here we aligned it according to its homeostatic timescale.

However, this difference in timescales may challenge the popular view that in biology Hebbian plasticity is constrained through homeostatic plasticity [16, 19–22]. The algorithmic normalization of synaptic weights every second is not the same mechanism as the biological rescaling of synaptic weights over hours. Although, in the theoretical literature, a rapid stabilizing mechanism is typically called “homeostatic”, here we will refer to this class of control mechanisms as Rapid Compensatory Processes (RCPs). The term “homeostatic plasticity” is in the following reserved for slow negative feedback processes on the timescale of hours or days — a terminology that seems consistent with the available experimental literature [15, 16]. In this review we focus on this discrepancy of timescales and ask which biologically plausible processes could constrain Hebbian plasticity. Specifically, we will try to answer the following questions: Why do we need RCPs to stabilize Hebbian plasticity? How fast do these processes have to be — hours, minutes, seconds or less? Which mechanisms could fill this role in Hebbian learning? Moreover, what are the consequences of fast control mechanisms on memory formation and recall in network models? And finally, if RCPs are a requirement, what is the role of slower forms of negative feedback implemented by known forms of homeostatic plasticity?

Models of Synaptic Plasticity

Synaptic plasticity exists across different timescales. For instance, synaptic changes induced by a sequence of four presynaptic spikes in rapid sequence typically decay within a few hundred milliseconds [42–44] and are called short-term plasticity. The rapid decay implies that the changes are not

useful for long-term memory formation, but more likely involved in gain control [42].

Other forms of plasticity induced by classic induction protocols [45–47], can have long-term effects on the timescale of hours or more. Long-term plasticity is therefore potentially useful for memory formation [2]. We remind the reader that the induction of long-term plasticity can be as fast as seconds, but the induced changes persist for much longer. Depending on the direction of synaptic change, we speak of long-term potentiation (LTP) or long-term depression (LTD).

Under suitable conditions the changes induced by a protocol of LTP or LTD are further consolidated after about an hour [48–50]. These effects are often referred to as late-phase long-term plasticity. In the rest of the paper we focus on plasticity induction and the early phase of long-term plasticity and neglect consolidation and maintenance.

The diverse effects of long-term plasticity can be cast into a mathematical framework which describes the change of synaptic efficacy over time. Apart from a few notable exceptions [23, 51–54], the vast majority of models of long-term plasticity assume a one-dimensional synaptic state space which represents the synaptic efficacy or weight w_{ij} of a synapse from neuron j to neuron i [5, 8, 9, 33, 55–65]. The evolution w_{ij} is then characterized by the differential equation

$$\frac{dw_{ij}}{dt} = G \tag{1}$$

in which the function G , often called the “learning rule”, is a member of an infinite dimensional function space \mathcal{G} , the space of all possible learning rules. This space is strongly constrained if we only focus on plausible learning rules, which are the rules in agreement with experimental findings.

For example, classical stimulation protocols for LTP [45–47], LTD [66, 67], or spike-timing-dependent plasticity (STDP) [68–70], combine the activation of a presynaptic neuron, or a presynaptic pathway, with an activation, depolarization, or chemical manipulation of the postsynaptic neurons, to induce synaptic changes. In models this is typically formalized by stating that G only depends on quantities which are locally accessible to the synapse. It is customary to assume that the main locally accessible variables include: (i) the current synaptic state w_{ij} ; (ii) the activity pre_j of the presynaptic neuron; and (iii) the state post_i of the postsynaptic neuron [64, 71, 72]. Thus we can write $\frac{dw_{ij}}{dt} = G(w_{ij}, \text{post}_i, \text{pre}_j)$. Additionally, G could also depend on low dimensional information carried by chemical signals such as neuromodulators (see Frémaux and Gerstner [73] for a review).

Most published learning rules G can be written as the linear sum of different terms in which each term can be interpreted as a specific manifestation of plasticity. These terms act together to explain the measured outcome in plasticity experiments. Let us explain the most common ones

using the following example learning rule:

$$G(w_{ij}, \text{post}_i, \text{pre}_j) = \underbrace{g_0(w_{ij})}_{\text{e.g. decay term}} + \underbrace{g_1(w_{ij}, \text{pre}_j)}_{\text{“transmitter-induced”}} + \underbrace{g_2(w_{ij}, \text{post}_i)}_{\text{heterosynaptic}} + \underbrace{H(w_{ij}, \text{post}_i, \text{pre}_j)}_{\text{homosynaptic}} \quad (2)$$

We will discuss each of the terms, going from right to left. H : Co-stimulation of a presynaptic pathway (or presynaptic neuron) and a postsynaptic neuron, as used in many classic induction protocols of LTD or LTP, changes the activated synapses. Such input specific changes are called “homosynaptic” because they affect *the same* synapses that are stimulated. In the above equation the homosynaptic changes are characterized by the term H on the right, where H is short-hand for “Hebbian”. The homosynaptic changes are often further separable into individual contributions of LTD and LTP (e.g. Song et al. [8], Gerstner and Kistler [74]).

g_2 : If a stimulation protocol induces a change at *other* (unstimulated) synapses onto the same postsynaptic neuron, the effect is called “heterosynaptic” [66, 75–77]¹. In Eq. (2) heterosynaptic effects are described by the function g_2 which does not depend on pre_j , but only on the state of the postsynaptic neuron. An example of heterosynaptic plasticity is synaptic scaling [1, 19] which has been modelled using a heterosynaptic term g_2 with a linear weight dependence [31].

g_1 : If presynaptic activity alone is sufficient to induce plasticity — one could think of non-associative LTP at the parallel fiber to Purkinje cell synapses [79, 80] or spontaneous spine growth in the presence of glutamate [81] — this is captured by the function g_1 , which depends on the presynaptic activity pre_j , but not on post_i .

g_0 : Finally, a slow drift, a spontaneous growth or decay of the synaptic strength that does not depend on the input or the state of the postsynaptic neuron is captured by the function $g_0(w_{ij})$.

In our example all terms explicitly depend on w_{ij} . While this is not a strict requirement it is customary to limit the allowed range of w_{ij} to avoid infinite weight growth. Since big weights are associated with physically large synapses, while the total space in the brain is limited, a bound on synaptic weights is reasonable. Depending on the implementation details, the limit can be implemented either as a “hard bound” or “soft bound” (e.g. Gerstner and Kistler [74], Rubin et al. [82], Gütig et al. [83]).

Virtually all existing plasticity models can be written in a form similar to Equation (2). Differences between model formulations arise if (i) pre_j is interpreted as presynaptic firing rate, presynaptic spikes, or as chemical traces left by spikes (e.g. glutamate); (ii) post_i is interpreted as postsynaptic firing rate, postsynaptic spikes, chemical traces left by postsynaptic spikes, postsynaptic calcium, postsynaptic voltage, or combinations thereof; (iii) the weight-dependence is chosen identical or differently for each term. With this framework we can now state what we mean by compensatory

¹As an aside, we note that the term “heterosynaptic plasticity” is sometimes also used for synaptic changes that are visible at the connection from a presynaptic neuron j to a postsynaptic neuron i , but induced by the activation of a third, typically modulatory neuron [78]. However, in this article we do not consider this possibility.

processes and address the question why we need them to be fast.

Why do we need rapid compensatory processes to stabilize Hebbian plasticity?

Intuitively, synaptic plasticity that is useful for memory formation must be sensitive to the present activation pattern of the pre- and postsynaptic neuron. Following Hebb’s idea of learning and cell assembly formation, the synaptic changes should make the same activation pattern more likely to re-appear in the future, to allow contents from memory to be retrieved. However, the re-appearance of the same pattern will induce further synaptic plasticity. This forms an unstable positive feedback loop. Anybody who was sitting in the audience when the positive feedback loop between the speaker’s microphone and the loudspeaker resulted in an unpleasant shriek, knows what this means. In many cases an unstable system can be made stable by adding sensible control mechanisms [84] which are thus typically integrated in theoretically motivated plasticity models.

Let us now consider one such classic example of a learning rule. To that end we consider Oja’s rule [57]

$$\frac{dw_{ij}}{dt} = G = \eta (x_j y_i - w_{ij} y_i^2) \quad (3)$$

where η is a small constant called learning rate. Since Oja’s rule corresponds to a specific choice of G in Equations (1) and (2), let us highlight the relation. First, in Oja’s rule the presynaptic activity pre_j is characterized by the presynaptic rate x_j and the state of the postsynaptic neuron post_i by its firing rate y_i . Second, and with this in mind, we can now identify two terms on the right-hand side of Equation (3). Oja’s rule contains a Hebbian term $H = \eta x_j y_i$ which does not have any weight dependence as well as a heterosynaptic term $g_2 = -\eta w_{ij} y_i^2$ which comes with a negative sign and is linear in the weight. Following our convention from above (Eq. 2) we call the term heterosynaptic because it acts on all synapses, even those that do not receive presynaptic stimulation. For simplicity, and following the tradition [57] we combine Oja’s rule with a linear neuron model $y_i = \sum_j w_{ij} x_j$.

It is quite intuitive to see how stability arises in this model. As the synaptic weights w_{ij} grow, due to the Hebbian term, the firing rate y_i of the postsynaptic neuron increases and therefore the influence of the negative heterosynaptic term gets stronger. Because the heterosynaptic term has a superlinear dependence on y_i , it is guaranteed to “catch up” with the Hebbian term eventually. It can be shown that for a linear neuron model and sufficiently small η , Oja’s rule ensures that the weights converge such that $\vec{w}_i = (w_{i1}, \dots, w_{iN})$ aligns with the first principal component of the data x , while the squared sum of all afferent weights remains normalized [57].

We interpret the heterosynaptic term in Oja’s rule as RCP: First, it is rapid because it responds instantaneously to activity fluctuations in y_i . Second, it is compensatory because it ensures stability by effectively enforcing a constraint on the afferent synaptic weights [10, 57]. Biologically, such a

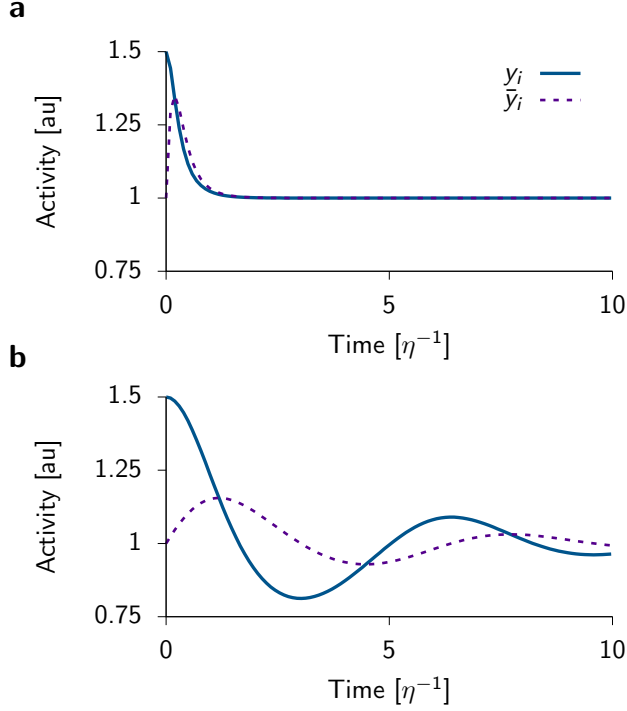


Figure 2: **Illustration of a rapid and slow compensatory processes in a variation of Oja's rule.** (a) Here we used a fast filter time constant $\tau = 0.1\eta^{-1}$ (cf. Eq. (4)) and plot the output firing rate y_i (solid) and the delayed estimate \bar{y}_i (dashed). (b) Same as in a), but with $\tau = 2\eta^{-1}$. Model: We simulated $\frac{dw}{dt} = xy - w\bar{y}^2$, $\tau \frac{d\bar{y}}{dt} = y - \bar{y}$ and $y = wx$ with $x \equiv 1$.

heterosynaptic effect could be obtained, for instance, when synapses have to compete for a shared resource [57, 85] or send chemical signals to each other [86].

One could now ask if we really need this compensatory processes to be rapid. Could we not simply replace the instantaneous firing rate y_i in the heterosynaptic term by a slower variable? The timescale of the slow variable could be related in a biological system to the time necessary to estimate the firing rate from, e.g. calcium concentration and translate these into metaplastic changes in the learning rule. To illustrate the general idea by a concrete example, we take Oja's rule, as in Equation (3), except that, in the heterosynaptic term, we replace y_i^2 by \bar{y}_i^2 , where \bar{y}_i is a low-pass filtered version of the postsynaptic rate

$$\tau_y \frac{d\bar{y}_i}{dt} = y_i - \bar{y}_i. \quad (4)$$

If we choose $\tau_y = 1\text{ms}$ (for a fixed η of, e.g. $\eta^{-1} = 10\text{ms}$), the processes $g_2 = -\eta w_{ij} \bar{y}_i^2$ would still be considered as rapid (Fig. 2a), but if we choose $\tau_y = 1\text{h}$, it would be considered as slow. When the compensatory processes is too slow, positive feedback induced by the Hebbian term is prone to take over and oscillations (Fig. 2b) or even run-away dynamics arise. This is why we generally want the compensatory processes to be *rapid*.

The same problematic has also been demonstrated nicely in the Bienenstock-Cooper-Munro (BCM) model [5]:

$$\tau \frac{dw_{ij}}{dt} = \eta x_j \phi(y_i, \bar{y}_i) \quad (5)$$

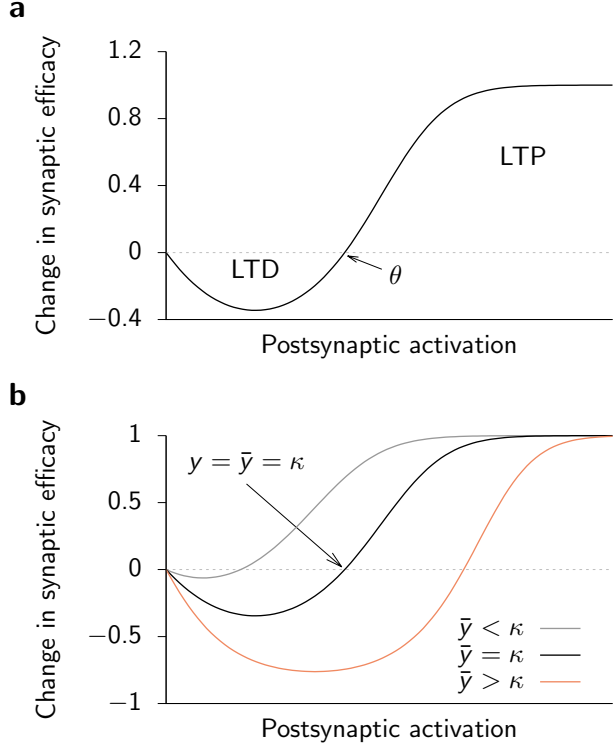


Figure 3: Most plasticity models can reproduce the notion of a plasticity threshold reported in experiments.

(a) The change in synaptic efficacy in many plasticity models is a function of variables related to postsynaptic activation. The parameter θ is fixed and marks the transition point (“threshold”) between LTP and LTD. (b) Schematic illustration of the action of the homeostatic moving threshold $\theta(t)$ in the BCM model [5]. When the average \bar{y} is larger than the target value κ , $\theta(t)$ shifts to higher values. Likewise, $\theta(t)$ shifts to lower values when \bar{y} is too low. For $y = \bar{y} = \kappa$ changes in synaptic efficacy are naught.

where ϕ is a nonlinear function with a characteristic shape characterized by a threshold θ between LTP and LTD (Fig. 3a), consistent with some induction protocols [70, 87]. The threshold θ depends on the moving average \bar{y}_i over past neuronal activity (Fig. 3b) where \bar{y}_i is defined in Equation (4). This is the reason why the model is said to have a “sliding” threshold.

To ensure stability, the BCM model requires two independent assumptions. First, the sliding threshold has to be a superlinear function of \bar{y}_i [5]. A standard choice is [88]

$$\theta(t) = \frac{\bar{y}_i^2}{\kappa} \quad (6)$$

where κ is the “target rate” to which the moving average of the postsynaptic firing rate should converge. Second, τ_y cannot be “too large” compared to τ , because otherwise oscillations or runaway activity occur [23, 32, 88]. In fact, the ratio $\frac{\tau_y}{\tau}$ determines the stability of the model.

Oscillations and instabilities are generic to many nonlinear systems and not limited to the above models. Control theory enables theoreticians to identify parameter ranges that lead to stable behavior and avoid instabilities [31, 84, 89]. The control theoretic analysis of several plasticity models relying on moving averages of the postsynaptic firing rate shows that the response timescale of the compensatory processes is constrained from above [23, 32, 88, 90]. In other words, the response time of the firing rate control has to be “relatively fast” compared to Hebbian plasticity. But how fast is fast enough? Is it seconds, hours or days?

How fast do compensatory processes have to be?

Because time can be rescaled arbitrarily in the above model, a quantitative answer to the question can only be given for a specific combination of neuronal, network and plasticity model parameters once units of time are calibrated with biological data. In other words, we need to put a numerical value on τ to set the timescale of τ_y . To fix a timescale² one can thus use any quantitative plasticity model which has been fitted to experimental data in combination with plausible spiking neuron models embedded into a spiking neural network with a biologically inspired activity state.

Such an analysis was done in Zenke et al. [32] using the plasticity model by Pfister and Gerstner [33], combined with negative feedback via either a sliding threshold or synaptic scaling. The critical timescale τ_{crit} was determined as the value τ_y (Eq. (4)) above which a recurrent neural network, initially tuned to a low activity asynchronous state [91, 92], generates run-away activity. Using simulations and mean field theory τ_{crit} was found to be on the order of seconds to minutes. Thus, the negative feedback needs to be too rapid to be linked to the known experiments of homeostatic synaptic plasticity reviewed in Figure 1.

Several remarks are in order. First, although signatures of the stereotypical activity dependence of the BCM model (Fig. 3a) are also present in STDP data and captured by many modern plasticity models [33, 61–63, 65, 93], the existence of a sliding threshold mechanisms is still a topic of ongoing debate. However, we have shown analytically, and confirmed in simulations, that the instability that arises through slow feedback in the BCM model is virtually identical to the situation in which the sliding threshold in Eq. (5) is replaced by a fixed threshold and instead synaptic scaling is added to the model [32]. Additionally, the analysis suggests that similar temporal requirements hold for an entire family of plasticity models with an explicit rate dependence (see Yger and Gilson [90] for a review). Note, however, that additional instabilities can arise in the case of synaptic scaling [89].

Second, the critical timescale τ_{crit} depends not only on the plasticity model, but also on multiple parameters of the neuron and network model. Moreover, the results showed a strong dependence on background firing rate which was comparatively high in the Zenke et al. [32] model ($\sim 3\text{Hz}$). On the other hand, robust stability is only possible if the actual value of τ_y is chosen much smaller than τ_{crit} . The precise value of the critical timescale has therefore to be taken with care: We believe that any published numerical value for τ_{crit} may be off by a factor of 5 or 10 (because of uncertainty in choices of neuronal and network parameters), but it is unlikely to be off by a factor 100. In any case, despite the remaining uncertainty, these numerical results suggest that most experimental forms of homeostatic plasticity are too slow to stabilize Hebbian plasticity as captured by current models and that RCPs are required to maintain stability.

A recent voltage based plasticity model by Jedlicka et al. [29] corroborates the above findings.

²Strictly speaking, the notion of a “timescale” is meaningless for a nonlinear differential equation like Eq. (5), it is only defined for linear dynamical systems. The notion can be rescued, however, in the vicinity of a fixed point around which the system can be linearized. We refer the interested reader to the Supplementary Material where we provide additional information on this important, but technical issue.

By fitting their model with a rapid sliding threshold to *in vivo* data from dentate gyrus, the authors find $\tau_y \approx 12\text{s}$ which allows them to speculate that the sliding threshold could be linked to CaMKII inactivation.

Interestingly, Toyozumi et al. [23] arrive at qualitatively similar conclusions by carefully analysing the different phases of synaptic dynamics following monocular deprivation (MD) [39]. Specifically, they find that a fast sliding threshold guarantees stability, but provides a poor fit to experimental data, whereas a slow sliding threshold compromises stability altogether. Consequently, they suggest a model in which LTP and LTD saturate quickly to attain steady states. Additionally, a slow form of homeostatic plasticity is required to capture the data (cf. Fig. 1), but is no longer required to provide stability. In their model LTP and LTD saturate due to soft weight bounds. However, the model does not crucially depend on this point and would presumably also work with other RCPs.

Finally, these findings are in good empirical agreement with many existing simulation studies of plastic network models (Fig. 1) — in each of these, a rapid homeostatic control mechanism on a timescale of seconds to minutes was implemented to maintain stability [14, 25, 27, 28, 33, 63, 94].

We can summarize our insights as follows. The fact that Hebbian plasticity has to be appropriately constrained through stabilizing mechanisms to avoid run-away activity is well known. Classic models such as Oja’s rule or the BCM model, for example, explicitly include appropriate mechanisms based on the postsynaptic firing rate as an indicator and driver of stabilizing processes. However, the fact that these processes have to be rapid in absolute terms, only becomes apparent when units of time are fixed to a biologically meaningful timescale. Moreover, RCPs need to be even more rapid in large recurrent network models, because a large number of plastic synapses tends to amplify instabilities unless the learning rate is scaled with the inverse number of synapses. Accumulating evidence suggests that biological forms of LTP and LTD have to be accompanied by RCPs which operate on timescales of seconds to minutes and are thus orders of magnitude faster than most known forms of homeostatic plasticity (cf. Fig. 1). This answers the questions as to why RCPs are needed and gives us first upper limits on the intrinsic timescale of RCPs to successfully stabilize LTP and LTD. However, do we want RCPs to be a rapid version of homeostatic plasticity with a single set point? We will now discuss this question in some more detail, before we turn to potential mechanisms.

Functional consequences of enforcing constraints on short timescales

Homeostatic mechanisms are typically interpreted as negative feedback processes [22], which rely on an error signal to maintain a given control variable of a dynamical system at designated target values, or set points. Many control systems have a single set point. For instance, an autopilot tries to maintain a single given course at any given time. Similarly, most models of homeostatic plasticity have a single target value, such as the average postsynaptic firing rate (see κ in the BCM model Eqs. (5) and (6)). Suppose we are dealing with a fast negative feedback processes, what are the

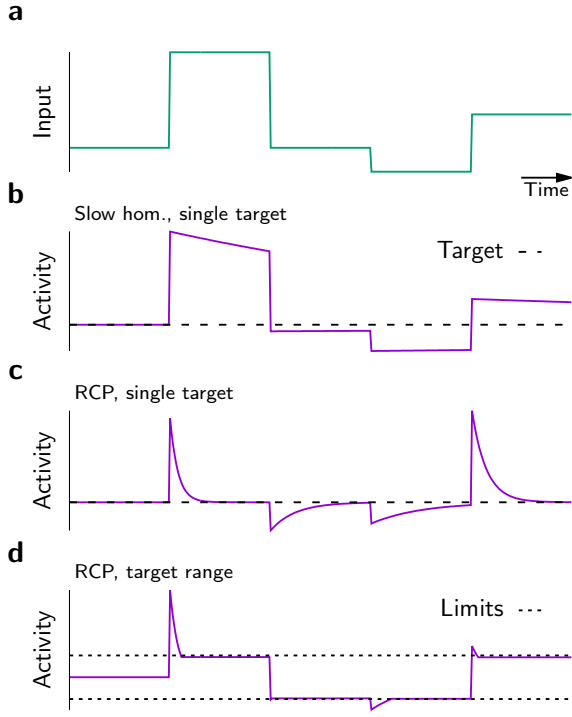


Figure 4: **The effect of slow homeostatic plasticity and rapid compensatory processes on the neuronal code.** (a) Fluctuating external world stimulus over time. (b) Neural activity “ $A = w \times \text{input}$ ” over time. A slow negative feedback mechanism (homeostasis) adjusts w to push A towards a single target (dashed line) which is rarely reached; negative feedback is modelled as $\tau_{\text{slow}} \frac{dw}{dt} = (\text{Target} - A)$. With a slow homeostatic mechanism, a neuron can track input relatively accurately. (c) Same as b, but for a rapid compensatory process ($\tau_{\text{slow}} = 50\tau_{\text{fast}}$) which drives the activity quickly to the target value. If the timescale of feedback is comparable to that of the stimulus, negative feedback interferes with the neuron’s ability to track the stimulus. (d) RCPs enforcing an allowed range (limits indicated by dotted lines). Even though the neuronal activity does not capture all the diversity of the input, it does capture some of it. Here we modeled the RCPs as the following nonlinear extension of the above model: $\tau_{\text{fast}} \frac{dw}{dt} = f(\text{Target} - A)$ with $f(x) = x$ for $\|x\| > \text{Limit}$ and $f(x) = 0$ otherwise.

functional consequences for plasticity and circuit dynamics? To do so, we focus on commonly found forms of firing rate homeostasis (FRH) with a *single* firing rate set point [15, 16, 95].

Neurons encode information in changes of their electrical activity levels. For instance, subsets of simple cells in the visual system fire spikes in response to specific edge-like features in the visual field [96]; cells in higher brain areas respond with high specificity to complex concepts and remain quiescent when the concept they are coding for is not brought to mind [97–99]; and finally certain neurons respond selectively with elevated firing rates over extended periods during working memory tasks [100–102]. The ability of neurons to selectively indicate through periods of strong activity the presence of specific features in the input or specific concepts in working memory is an important condition for computation.

Is the notion of a single set point compatible with the task of neurons to selectively respond to stimulation? If negative feedback control of firing rates is slow (e.g. synaptic homeostasis), neuronal firing can deviate substantially from the mean firing rates during short times and thus encode information (Fig. 4a,b). However, we have strong reasons to believe that slow homeostatic control mechanism cannot stabilize the ravaging effects of Hebbian plasticity. So what can we say about a putative RCP? If it were to act like FRH, but on a short timescale (e.g. seconds to minutes), neural codes based on neuronal activity become problematic because synaptic plasticity starts to suppress activity fluctuations which could be carrying important information (Fig. 4c). For example, if the RCP has a timescale of two seconds, rapid stimuli that change on a timescale of 0.5 seconds would be transmitted as a rate signal of the postsynaptic neuron while stimuli sustained for more than 5 seconds would be suppressed by compensatory synaptic changes. Even more alarmingly, certain forms of homosynaptic plasticity, like the BCM [5] or the triplet STDP [33] model endowed with a rapid sliding threshold, not only suppress high activity periods, but also “unlearn” previously acquired selectivity and erase memories (Figs. 4a,c and 5a–d). Therefore rapid compensatory processes which enforce a single set point are hardly desirable from a functional point of view. Thus the requirement of fast negative feedback control over Hebbian plasticity with a single set point poses a problem in itself.

It is important to appreciate that this problem arises from the combination of a *single* target with the requirement to implement negative feedback on a short timescale. Fortunately, most forms of homeostatic plasticity are slow (cf. Fig. 1). Thus, known homeostatic mechanisms do not interfere with neuronal coding. For RCPs not to interfere with coding either, it thus seems important that they do not enforce a single set point constraint on postsynaptic activity. Nevertheless, these RCPs have to prevent run-away activity.

There could be at least one simple solution to this conundrum [23, 103]. Suppose there are two, or more, set points enforced by two or more RCPs. For instance, one RCP could activate above a certain activity threshold and ensure that neuronal activity does not exceed this threshold. Similarly, a second mechanisms could activate below another lower activity threshold. The combined action

of the two mechanisms enforces neural activity to stay within an allowed range, but still permits substantial firing rate fluctuations inside that range (Fig. 4d).

When such a pair of RCP is combined with a form of Hebbian plasticity which has its plasticity threshold within the limits of the allowed activity regime, the neural activity of the compound system naturally becomes multistable for prolonged stimulation with the same stimulus. Within the allowed range no RCP is active, but Hebbian plasticity is intrinsically unstable. Thus, for a stimulus sustained longer than the timescale of the RCP and Hebbian learning, any value of the postsynaptic rate within the allowed region will lead to either LTP or LTD until the system reaches the limits at which either RCP rapidly intervenes by undoing any excess LTP or LTD from there on. The compound system exhibits therefore two, potentially stimulus dependent, stable equilibrium points, one at low and one at high activity.

Let us apply these considerations to two different systems, viz., a sensory system and a memory system. To be concrete, we assume that in either system the timescales of both LTP induction and RCP are two seconds. In the sensory system, each neuron will respond in a graded manner to short stimuli (say, with a duration of half a second) because synapses hardly change during a single stimulus duration. However, the repeated stimulation with different stimuli will cause long-lasting weight changes. The location of the high-activity fixed point depends on the stimulus ensemble used during stimulation. Moreover, if we drive the neuron with a single sustained stimulus, the high-activity fixed point adjusts on the timescale of a few seconds and reflects the value of the input.

The case of a memory system was considered in [103]. Suppose the high-activity fixed point corresponds to a memory retrieval state, while the low-activity equilibrium is associated with a quiescent memory which is not being recalled. Because both states are stable, it is irrelevant whether the memory is recalled every other minute or once a year. Importantly, this is different from models with rapid FRH, which might require neuronal activity to regularly turn on and off to satisfy the constraint. An example for this is the network model by Litwin-Kumar and Doiron [14] in which inhibitory synaptic plasticity (ISP) acts as rapid FRH with a single activity fixed point.

We can therefore answer the third of the questions raised in the introduction: *The functional consequences of rapid control mechanism with a single set point are that neurons lose the flexibility that is necessary for coding.* The consequences are therefore undesirable and the proposed solution is to design RCPs that allow for several set points or a target range of permissible activity. In the next sections we will first discuss common ways to constrain unbounded weight growth and explain why they are insufficient to provide stability, before we turn to potential candidate plasticity mechanisms which could act as RCP. Finally, we show an example of a spiking network model based on these principles which forms and recalls memories encoded in cell assemblies.

Potential mechanisms

To stabilize Hebbian plasticity, any RCP at the synaptic, dendritic, neuronal or network level can be considered. Due to temporal and spatial constraints of the biological substrate it seems most likely, however, that the fastest mechanisms are implemented as close to the synapse as possible.

At the synaptic level, excessive growth has traditionally been limited by soft or hard bounds on individual synaptic weights or other choices of explicit weight dependence of the Hebbian and heterosynaptic terms in Eq. (2) [8, 10, 31, 65, 82, 83, 104, 104–107]. For example, to avoid bimodal weight distributions, which sometimes arise in competitive models, but are not observed in biology, a range of multiplicative plasticity models [31, 82, 83, 104, 105, 108], with an appropriate choice of the weight-dependence of H in Eq. (2), have been devised. However, bounds on individual synaptic weights only impose an implicit constraint on the postsynaptic activity. To see this, consider a permissible range of individual synaptic strength of, say, $\pm 50\%$ around the initial efficacy, which seems not uncommon for plasticity induction protocols. However, by setting this range we do not automatically exclude the situation in which *all* synapses increase their efficacy by 50% which would in all likelihood correspond to pathological activity levels.

To avoid such run-away activity, plasticity has to ensure that not all synapses are potentiated or depressed equally. Rather there should be some form of competition which ensures that when one set of synapses is potentiated other synapses are depressed by a certain amount. While some degree of competition can be seen in STDP models, in which presynaptic spikes compete in time to elicit a postsynaptic spike [8, 58, 109, 110], this competition is generally weak [31] and without additional constraints, activity levels still succumb to run-away effects with detrimental consequences in recurrent neural networks [14, 32, 103]. Robust competition, for instance, through a BCM-like threshold [5] or explicit constraints on the sum of weights [10], are therefore of paramount importance for plasticity models.

In summary, there exist multiple mechanisms to limit growth of individual synaptic weights. However, to achieve robust synaptic competition and stability of output firing rates, more explicit activity constraints are required, as exemplified in the BCM model, or through explicit heterosynaptic interactions, similar to Oja’s rule (cf. Eq. (3)). We have already argued that these constraints need to be enforced rapidly. We now ask what possible mechanisms at the neuronal or network level could achieve that.

At the network level RCPs might be implemented by inhibition and ISP which could influence plasticity at excitatory synapses either directly or indirectly. Some theoretical forms ISP are known to implement a rapid form of FRH for individual neurons [14, 111]. With accumulating experimental evidence for ISP [95, 112, 113] it therefore seems likely that synaptic inhibition influences plasticity of excitatory synapses at least indirectly through changes in activity. However, in experiments the timescale of FRH mediated through ISP appears to be relatively slow [95] and it remains to be seen whether biological forms of ISP can act as RCP or whether they have a rather homeostatic role.

However, in some cases, inhibition without ISP can have a stabilizing effect. Lim et al. [114] have recently demonstrated that this can indeed lead to stability of certain forms of Hebbian plasticity. Moreover, inhibition can also directly affect and regulate excitatory plasticity [115]. Particularly interesting in this context are results by Delgado et al. [116] who observed total-conductance-dependent changes of the STDP curve depending on excitatory and inhibitory background input. Their results suggests that increased, but balanced, excitatory and inhibitory input biases the STDP window towards LTD and can thus act as a RCP. Delgado et al. [116] demonstrated this in a single-neuron feed-forward model, but it is not yet known whether these results generalize to larger networks.

At the neuronal level, algorithmic normalization of afferent synaptic weights is a commonly used mechanism to stabilize Hebbian plasticity in network models while simultaneously allowing structure formation [9, 10, 13, 14]. While such rapid and precise scaling at the neuronal level has been criticized as biologically implausible [5], an “approximate” scaling could potentially be achieved through heterosynaptic plasticity at the dendritic level [117].

Heterosynaptic plasticity has moved back in the focus recently, because of its potential role as RCP [20, 21, 86, 103, 118]. Importantly, some forms of heterosynaptic plasticity are fast, and provide primarily negative feedback in response to high postsynaptic activity levels [20, 119] or in the presence of strong LTP on a dendritic segment [86]. This is reminiscent of Oja’s rule (Eq. (3)) and seems well suited to counteract run-away LTP. In contrast to Oja’s rule, these heterosynaptic changes are induced by bursts of postsynaptic activity which implies that the quadratic term y_i^2 in Eq. (3) should be replaced by a term that is triggered either by firing rates y_i above some threshold [118] or by a higher power such as y_i^4 [103].

In models which also show run-away LTD at low activities (e.g. [33, 63]), an additional RCP is needed which either saturates or counteracts LTD. Possible forms of plasticity include, but are not limited to, transmitter-induced plasticity, homeostatic scaling-up or spontaneous spine formation.

In the following section we review a plasticity model which combines Hebbian plasticity with two RCPs that enable more than a single set point of neuronal activity. We also discuss, in the context of the model, the potential role of additional slow homeostatic mechanisms.

Learning and recall in a recurrent spiking network model

We now discuss a learning rule which combines a plausible model of Hebbian plasticity with two additional RCPs [103]. For sensible combinations this compound model does not suffer from the run-away effects of purely Hebbian plasticity and exhibits intrinsic multistability instead (cf. Fig. 4d).

The basic logic of multistable plasticity can be summarized as follows. At high activity levels a rapid form of heterosynaptic plasticity limits run-away LTP and creates synaptic competition. Similarly, at low activity levels an unspecific form of plasticity which only depends on presynaptic activity prevents run-away LTD. The well-orchestrated interplay between these adversarial plasticity

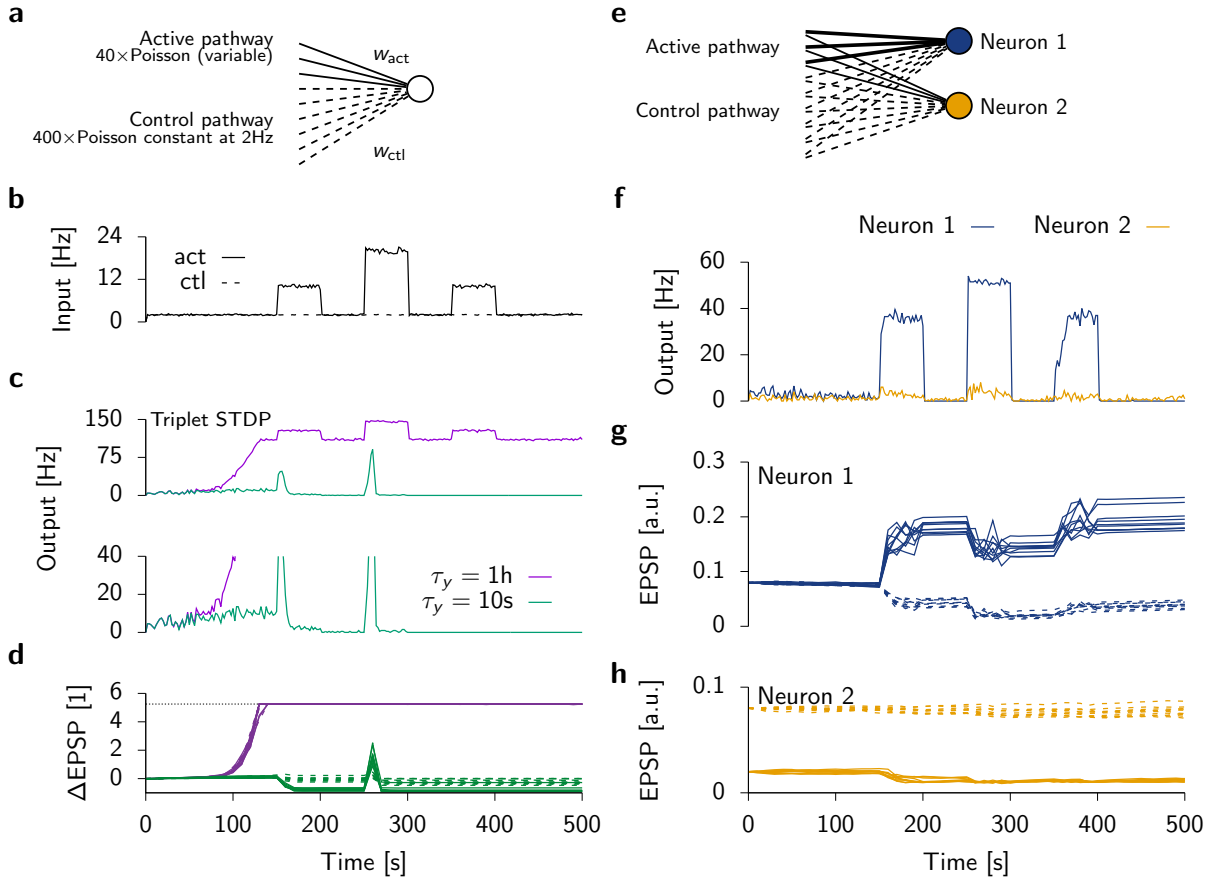


Figure 5: Firing rate stabilization, synaptic weights, and feature selectivity through RCPs. (a) Schematic figure of single neuron with two distinct input pathways. The “active” pathway consists of 40 Poisson neurons switching their rates synchronously between 2, 10 and 20Hz. The control pathway consists of 400 neurons all firing constantly at 2Hz. Synaptic plasticity is modeled with triplet STDP [33] with a BCM-like sliding threshold as defined in Eqs. (4–6) All weights are initialized at the same value and can freely move between hard bounds at zero and ≈ 6 times the initial value. (b) Population firing rates of the input populations averaged over 2s bins. Firing rates in the active pathway (solid line) are switched three times from 2Hz to a higher rate and back (10Hz, 20Hz and 10Hz for 50 seconds each time), whereas firing rates in the control pathway are constant at 2Hz. (c) Output firing rates of a single postsynaptic neuron. Purple: Slow sliding threshold, with time constant $\tau_y = 1h$; Green: Fast sliding threshold, $\tau_y = 10s$; (see Zenke et al. [32], $\kappa=3Hz$ in Eq. (6)). Top and bottom show the same firing rate plot for different y-axis ranges. (d) Relative weight changes for 10 randomly chosen weights from each pathway for the slow (purple) and the fast (green) sliding threshold. Solid lines correspond to active pathway weights and dashed lines to the control pathway. Note that for the fast sliding threshold the active pathway develops weaker synapses than the control pathway. For the slow sliding threshold all weights saturate. (e) Simplified sketch of the same setup as in a, but with two postsynaptic neurons. The only difference between the two neurons is the choice of initial conditions of the synaptic weights. For Neuron 2 the active pathway weights are initialized at a lower value than for Neuron 1. All synaptic weights exhibit triplet STDP combined with heterosynaptic plasticity (see Zenke et al. [103]). (f) Output firing rates of the two neurons over time. Neuron 1 (blue) responds selectively (with rates $>30Hz$) to the elevated inputs in the active pathway (cf. b). Neuron 2 (orange) continues to fire with low firing rates. (g) Evolution of weights over time for Neuron 1. Active pathway weights are plotted as solid lines and control pathway weights are dashed. For Neuron 1, the synapses in the active pathway undergo LTP during the first strong stimulation of the active pathway. However, weights quickly saturate. Synapses in the control pathway exhibit heterosynaptic depression. (h) Same as g, but for Neuron 2. The weights in the active pathway are slightly depressed during initial stimulation.

mechanisms dynamically creates multistability of neuronal activity and prevents pathological run-away effects.

Our approach is quite general and many Hebbian plasticity models can be stabilized through the addition of two non-Hebbian forms of plasticity. For illustration purposes, we will now focus on the triplet STDP model for which biologically plausible sets of model parameters exist [33]. To prevent run-away LTP we require a mechanism which balances out potentiation at high activity levels. To that end, we use a form of weight-dependent, multiplicative heterosynaptic depression [21, 103, 118]. Our choice of a purely heterosynaptic RCP is one possibility, but other homosynaptic forms of plasticity could achieve similar results. For instance, “heterosynaptic” LTD [29, 120] which requires low presynaptic activity for depression in the unstimulated pathway, is one possibility. In short, as long as LTP in a strongly stimulated pathway is accompanied by either “no change” or synaptic depression of synapses with low levels of presynaptic activity, run-away LTP can be avoided. To similarly prevent run-away LTD in our model we introduced a hypothetical form of transmitter-induced plasticity. Together the three plasticity mechanisms — Hebbian plasticity and two RCPs — work in symphony to generate stable levels of neuronal activity.

Let us consider the weight w_{ij} from a presynaptic neuron j to a postsynaptic neuron i . Although the full model is an STDP model, we now express its core ideas in terms of a rate model of the pre and postsynaptic rates x_j and y_i :

$$\frac{dw_{ij}}{dt} = \underbrace{\delta \cdot x_j}_{\text{Transmitter-induced}} + \underbrace{+\eta \cdot x_j y_i (y_i - \theta)}_{\text{Triplet model}} - \underbrace{\beta \cdot (w_{ij} - \tilde{w}_{ij}) y_i^4}_{\text{Heterosynaptic}} . \quad (7)$$

Here, δ and β are strength parameters for the two non-Hebbian components of the plasticity model; η is the strength parameter (learning rate) for Hebbian plasticity; and \tilde{w}_j serves as a reference weight that can be related to consolidation dynamics [54, 103]. Note, because the negative feedback plasticity mechanisms are “rapid”, there is no low-pass filtered variable \bar{y} in this expression. Runaway effects which normally would occur in the triplet STDP model without RCPs (Fig. 5a–d) are avoided with the additional plasticity mechanisms. Due to its rapid action and the high power, the heterosynaptic term in Eq. (7) acts as a burst detector [118] which dominates at high activity levels and prevents LTP run-away dynamics (Fig. 5e,f).

For sensible choices of δ and β , neuronal firing rates remain in intermediate regimes (Fig. 5f) and synaptic weights in the model converge towards stable weights \hat{w} whose values are dependent on the activation history and enable the formation of long-term memories (Fig. 5g). Importantly, the model preserves the plasticity threshold between LTD and LTP of the original triplet STDP model. The triplet STDP model together with the non-Hebbian plasticity mechanisms, dynamically creates one unstable and two stable equilibrium points. The activity level of the higher stable fixed point depends on the stimulus. In particular, it is sensitive to the number, firing rate and temporal structure of the active synaptic inputs and a stronger stimulus will typically result in a higher




Figure 6: **Reproducible responses to input stimuli with varying strength.** (a) Population firing rate of active input pathway for the continuation of the simulation shown in Figure 5b,e-h. The neuron is stimulated 16 times with 4 interleaved steps of increasing input firing rate in the active pathway. (b) Output firing rate of Neuron 1 (200ms bins; cf. Fig. 5f). After the first set of stimuli, responses to the later sets remain graded and are overall reproducible. (c) Data points from a and b (input and output) plotted against each other.

steady state response. For any given stimulus, synaptic weights converge rapidly towards one of two possible stable equilibrium states (Fig. 5f-h). First, there is a “selective” equilibrium state associated with high postsynaptic activity. In this state some weights are strong while other weights onto the same postsynaptic neuron remain weak. Thus the neuron becomes selective to features in its input (Neuron 1 in Fig. 5f and g). Second, there is a “non-selective” equilibrium at low-activity (Fig. 5f and h). Which fixed point a neuron converges to depends on its initial conditions, lateral interactions and the details of the activation pattern (Fig. 5g,h). Once weights have converged to one of the respective stable states, they keep fluctuating, but do not change on average. Since the RCPs do not impose a single set-point, activity patterns are not unlearned when a certain input is kept active (inactive) for extended periods of times (compare Fig. 5d with 5g,h).

As a result, Neuron 2 shown in Fig. 5h shows “no learning” which seems undesirable at first. However, it is in fact useful for such a stable equilibrium to exist when learning is considered as a network phenomenon. Other neurons (Neuron 1 in our case) already are selective and code for a given stimulus. Analogously, Neuron 2 might in fact code for a different stimulus which is not active at the moment, in which case we would like to perturb it as little as possible while other neurons “learn” (Fig. 5h). Similar dynamics can be achieved in learning models with strong lateral inhibition which completely suppresses neuronal activity and thus also associative plasticity. In the present scenario, however, this is not the case. Neuron 2 is still firing with some low background activity throughout (Fig. 5f).

There are several aspects worth noting about the model. First, heterosynaptic plasticity does not only stabilize Hebbian plasticity in the active pathway, it also introduces synaptic competition between the active and the control pathway (Fig. 5g). In contrast to BCM-like models in which heterosynaptic depression of the inactive pathway depends on intermediate periods of background

activity in between stimuli [29, 88], here the heterosynaptic depression happens simultaneously to LTP induction (Fig. 5g). Second, although the learning rule effectively implements a rapid redistribution of synaptic weights reminiscent of synaptic scaling, it is still a fully local learning rule which only depends on information which is available at the synapse (cf. Eq. (7)). Third, although the learning rule effectively implements bistable dynamics for each stimulus, the “selective” equilibrium level remains stimulus dependent, which allows the neuron to respond in a graded and reproducible way to input stimuli of varying intensity (Fig. 6). Fourth, in general the reference weight \tilde{w} is not fixed, but follows its own temporal dynamics on a slower timescale ($\approx 20\text{min}$ and more). Such slow complex synaptic dynamics are essential to capture experiments on synaptic consolidation [50, 54, 121], but could similarly be used to model slow forms of homeostatic plasticity [23].

Finally, the stability properties of the learning rule in Eq. (7) are not limited to simple feed-forward circuits, but generalize to more realistic scenarios. Specifically, the combination of heterosynaptic and Hebbian plasticity enables stable on-line learning and recall of cell assemblies in large spiking neural networks (Fig. 7a,b; [103]). Stationary firing rates in the model depend on the connectivity pattern and the spiking statistics of active inputs. In a recurrent network, however, output spike trains pose as inputs to other neurons. As a nontrivial consequence, stationary solutions of the network state exhibit firing rate distributions which are uni-modal and long-tailed (Fig. 7c,d). Individual neuronal firing rates only stabilize under stationary conditions. If the rates are non stationary, for instance due to the inclusion of additional adaptation processes in the neuron model, rates in the model drift on behavioral timescales (see Zenke et al. [103] for details).

In summary, orchestrating Hebbian plasticity and RCPs on comparable timescales dynamically generates multistability. This reconciles the experimentally observed fast induction of synaptic plasticity with stable synaptic dynamics and stability of learning and memory at the single neuron level as well as in large networks. However, there are a few caveats with this approach which we will discuss in the following.

Problems of RCPs at the single neuron level

Consider a population of neurons with plastic synapses which follow intrinsically stable plasticity dynamics such as the ones described in the last section. To encode and process information efficiently, neuronal populations need to create internal representations of the external world. Doing this efficiently requires the response to be sparse across the population. In other words, only a subset of neurons should respond for each stimulus. Moreover, different stimuli should evoke responses from different subsets of neurons within the population to avoid that all stimuli look “the same” to the neural circuit. Finally, individual neurons should respond sparsely over time. Imagine a neuron which is active for all possible stimuli. It would be as uninformative as one which never responds to any of the inputs. Therefore, to represent and process information in neural populations efficiently,

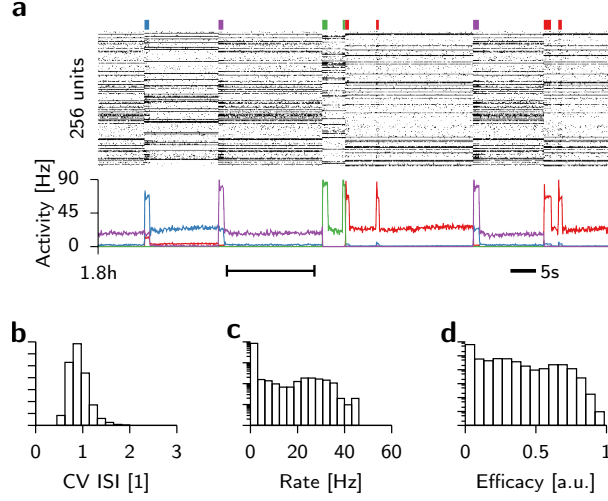


Figure 7: **Stable activity in a recurrent neural model with ongoing plasticity.** (a) Memory recall in associative cell assemblies through selective delay activity in a network which previously has learned to distinguish between four repeating input patterns [103]. The colored bars at the top indicate time and duration of external stimulation with one out of four stimuli. The color indicates stimulus identity. The spike raster in the middle shows spiking activity of 256 randomly chosen excitatory cells from the network. The graph at the bottom shows the firing rate of the four subpopulations defined by the cell assemblies in the network. The multistable plasticity model of Eq. (7) is active throughout the simulation. (b) Histogram of the coefficient of variation of the inter-spike-interval of excitatory cells in the network during the interval indicated by the black range bar in a. (c) As in b, but for mean neuronal firing rates over the same interval. (d) Distribution of synaptic efficacies of plastic recurrent synapses at the end of the network simulation. Figure adapted from Zenke [122] and Zenke et al. [103].

different neurons in the population have to develop selectivity to different features.

Multistable plasticity at the neuronal level as described above does not prevent neurons from responding weakly to all stimuli (see, for example, Neuron 2 in Fig. 5f). This is a direct consequence of the fact that the model presented here does not have a sliding threshold like the BCM model. Moreover, with more similar initial conditions and in the absence of lateral inhibition both Neuron 1 and 2 could have developed selectivity to the *same* input. Thus, in a large network in which all synapses are changed by the intrinsically stable plasticity rule introduced above, all neurons could end up responding to the same feature. How can such an undesired outcome be avoided?

To successfully implement network functions like the ones shown in our example (Fig. 7), several network parameters and properties of the learning rules themselves need to be fine-tuned and maintained in sensible parameters regimes. For instance, successful learning as demonstrated in Zenke et al. [103] depends on sensible choices of the initial synaptic weights and connectivity values. To achieve the necessary degree of tuning and maintenance, biological networks presumably rely on additional forms of plasticity which drive the network towards a dynamical state which is conducive for learning. However, due to the intrinsic stability of the learning rule, these additional mechanisms, for instance a BCM-like sliding threshold, can now safely operate on much longer timescales. This suggests that homeostatic plasticity and metaplasticity could fulfill this fine-tuning and maintenance role.

What is the role and scope of slower homeostatic plasticity mechanisms?

Diverse homeostatic mechanisms exist in the brain at different temporal and spatial scales [123–127]. We have argued that RCPs are important for stability, but what advantages do slow homeostatic mechanisms have and what is their computational role?

An advantage of slow homeostatic processes is that they can integrate activity over long timescales to achieve precise regulation of neural target set points [19, 31, 89]. Longer integration times also allow to integrate signals from other parts of a neural network which take time to be transmitted as diffusive factors [124, 128]. Slower homeostasis thus seems well suited for control problems which either require fine-tuning or a spatially distributed homeostatic regulation of functions at the network level.

There are at least two important dynamical network properties which are not directly controllable by neuronal-level RCPs (Eq. (7)). First, temporal sparseness at the neuronal level is not automatically guaranteed. A neuron that never responds to any stimulus will never learn to do so under multistable plasticity, if the LTP threshold is too high. Similarly, a neuron that always responds is uninformative, but will not change its behavior if the LTD threshold is too low. Second, spatial sparseness at the network level, in the sense that for any stimulus a subset of neurons

responds, is also not automatically guaranteed. Lateral inhibition is a suitable candidate to decorrelate responses of different neurons in a network, but, as excitatory synapses change during learning, the strength of lateral inhibition needs to be co-regulated.

The problem of *temporal sparseness* can be solved by any mechanism which ensures that a neuron which has been completely silent for very long, eventually “gets a chance” to reach an activity level above the LTP threshold. This can be achieved by either lowering the threshold as in the BCM theory [5, 88, 103] or by slowly increasing the gain of either the neuron itself or the excitatory synapses through other forms of slow homeostatic plasticity [23, 129, 130]. Finally, similar homeostatic effects could be achieved by dis-inhibition through the action of neuron specific ISP [112] or by decreasing the response of inhibitory neurons [40, 131]. Conversely, a neuron that is uninformative because it is always active could decrease its response to some stimuli by the opposite action of one or several of the homeostatic mechanisms mentioned above, such as increased inhibition, reduced excitation, or reduced excitability.

While it is conceivable that mechanisms addressing the issue of temporal sparseness could act locally at the neuronal level, it is clear that enforcing *spatially sparse activity* at the population level can only be achieved in a non-local manner. A common approach to guarantee spatial sparseness in models, is to include lateral inhibition, as done in subspace learning algorithms [132, 133], sparse coding paradigms [134, 135] or models of associative memory [14, 103, 136–140]. However, achieving appropriate levels of inhibition can be difficult, especially if excitatory synaptic weights are not static, but change over time and on a per neuron basis [111]. To solve this task in biological networks, ISP would be a natural candidate. However, most existing models of ISP are purely local and tune inhibition on a per neuron level [111, 112, 141]. More specifically, ISP acts as a neuronal RCP which rapidly drives firing rates to a single stable set point (cf. Fig. 4; Vogels et al. [111]). To achieve a certain level of spatial sparseness through any form of homeostatic plasticity, requires a signal with a wider scope which encodes network activity [124, 128]. Using such a signal it is then possible to modulate plasticity [73]. For example, in Zenke et al. [103] ISP is modulated by a low-pass filtered signal which encodes network activity. As a result, the combination of intrinsically multistable plasticity at excitatory synapses and ISP, ensures that recurrent inhibition is tuned to a level where only one cell assembly can be active at any given time. Importantly, this homeostatic mechanism does not have to be permanently active. For instance, once the inhibitory feedback within the model is tuned to the “sweet spot” at which the network can operate, ISP can be turned off safely without impairing stability. Similarly, it seems likely that some forms of homeostatic plasticity could be dormant for most of the time and spring into action only during the initial phase of development [142] or when an extreme external manipulation changes the network dynamics [1].

We are thus able to answer the final question from the introduction as follows: *Slow homeostatic mechanisms tune parameters of plasticity rules and neurons to enable efficient use of the available resources in networks.* For example, for the sake of efficiency, no neuron should be *never* active;




Figure 8: **Hypothetical evolution of synaptic state variable as a function of time during a hypothetical LTP induction experiment.** The purple, linear time course is a common assumption underlying many plasticity models. However, since the synaptic state during plasticity induction is typically not observable, many time courses and thus plasticity modes are compatible with the data. Note, that most plasticity models will interpret the “synaptic state” as the synaptic efficacy, whereas in experiments the synaptic efficacy may follow after a short delay.

no neuron should be *always* active; the number of neurons that respond to the exact same set of stimuli should stay limited.

Discussion

Taken together with the results from the previous sections, these insights suggest two distinct roles for negative feedback on different timescales. First, RCPs on short timescales stabilize Hebbian plasticity and make synapses onto the same neuron compete with each other. Heterosynaptic plasticity is likely to play a major role for these functionalities. Second, homeostatic mechanisms on slower timescales achieve fine-tuning of multiple network parameters. A slow shift of the threshold between LTD and LTP, the slow rescaling of all synaptic weights, or a slow regulation of neuronal parameters are good candidates for these functionalities. Some of these slow mechanisms could be important only in setting up the network initially or after a strong external perturbation to the circuit. This view, however, raises an important question: Why do many modern plasticity models not include built-in RCPs, whereas classic models do?

Why are RCPs missing in many STDP models?

Modern plasticity models try to capture a diversity of experimental data from rate-dependent [45], voltage-dependent [87] and spike-timing-dependent [68–70] plasticity experiments. One salient feature captured by most models [33, 61, 63, 65, 93] is the notion of a plasticity threshold which correlates with postsynaptic voltage, calcium concentration, postsynaptic firing rate, or other neuronal variables related to postsynaptic activation (Fig. 3a). Interestingly, most existing STDP models, although often explicitly fitted to data, are purely Hebbian and do not include the notion of a RCPs. If such a rapid mechanisms exist — which is what we argue here — then how can it be

that existing plasticity models without them can quantitatively capture the data from experiments?

There are presumably three main reasons for this. First, STDP experiments typically manipulate a single pathway, either by stimulating a presynaptic neuron or a bundle of presynaptic axons. Sometimes a designated control pathway (i.e. a second presynaptic neuron) is missing, or, if it is not missing, the effect size in the control pathway is considered as weak. However, from a theoretical perspective, we expect that heterosynaptic effects caused by stimulation of one presynaptic pathway are weak when measured at only one “control” synapse; a weak change at individual synapses could still have a strong accumulated effect over thousands of synapses. Therefore even weak heterosynaptic plasticity could act as a strong RCP [21, 103, 118].

Second, in an STDP experiment with 60 repetitions of pre-post-pairs, the total activation of the postsynaptic neuron is still in a reasonable regime. Therefore it is unclear whether the “burst-detector” for heterosynaptic plasticity would be triggered [21, 103, 118].

Third, experiments typically rely on repeated pre- and postsynaptic activation. Moreover, during the induction protocol, synaptic efficacy changes are usually not observable. Plasticity models are thus fitted to pairs of initial and final synaptic strength. However, the unobserved intermediate synaptic dynamics during a standard LTP induction protocol could be qualitatively very different (Fig. 8), but are obscured in experiments by measurement artifacts as well as short-term plasticity riding on top of the induced Hebbian LTP. These differences in the dynamics contain the answers to questions such as: Is the final synaptic strength stable or would it increase further with additional pairings? Is there a threshold of number of pairings that needs to be reached for an all or nothing effect?

Because the detailed internal dynamics of synapses during induction are not known, different plasticity models make different assumptions about the saturation of weights. Due to the limited amount of experimental data, it is possible to construct a diversity of different models which are all consistent with the data. For instance, the Zenke et al. [103] model discussed in this paper is based on the triplet STDP model, and therefore consistent with existing STDP data, but it includes additional non-Hebbian RCPs. Although the presence of these added processes is important for network stability, their overall contribution to simulated STDP protocols is negligible. So, how can one verify or falsify the existence of RCPs experimentally?

How can we further constrain plasticity models by experiments?

There are multiple ways in which synaptic plasticity models could be constrained better through additional data. In the past, a large body of research has focused on homosynaptic associative plasticity, also called Hebbian plasticity, using pairing experiments with various protocols such as STDP. Here, we argue that heterosynaptic plasticity as well as transmitter-induced plasticity or similar stabilizing plasticity mechanisms are as important as Hebbian plasticity due to their crucial role for network stability.

Heterosynaptic plasticity and heterosynaptic effects mediated through metaplasticity [29, 120, 143] are promising candidates to stabilize Hebbian plasticity models against run-away LTP [21, 32, 103, 118]. While heterosynaptic plasticity has been observed in various experiments [20, 117], a conclusive picture and data-driven models are still scarce. Is it possible to measure the timescale, frequency dependence, and weight-dependence of neuron-wide heterosynaptic depression by manipulating the stimulation of the postsynaptic neuron? Does “pure” heterosynaptic plasticity exist in the absence of presynaptic activity or is a slight activation of the presynaptic pathway always necessary to induce changes [120]? Another important question for the interpretation of heterosynaptic plasticity is whether it causes mostly synaptic depression similar to LTD or if it rather prevents or even resets early LTP through depotentiation at the unstimulated pathway [144]. Finally, the role of *heterosynaptic metaplasticity* [143] remains largely elusive.

Transmitter-induced plasticity is important in models and might be present in many experiments, even though it has not been reported as such. Here, we have interpreted transmitter-induced plasticity as a potentially weak form of long-term potentiation that is caused by presynaptic firing in the absence of postsynaptic activity. Why is this form of plasticity important? Suppose you have a network of neurons firing at low activity, so that any given neuron can be considered a weakly active postsynaptic neuron. Since low activity typically induces LTD, many plastic network simulations have the tendency to fall silent. To compensate for this theorists have either introduced lower bounds on synaptic weights or added weak LTP triggered by presynaptic activity [103, 114]. How realistic are these assumptions?

Direct experimental evidence for such terms would, for instance, be the growth of synaptic efficacy during low activity “pre only” stimulation. Such a term would manifest as a systematic positive drift of baseline in an experiment and could thus be easily interpreted as an unwanted instability [145, 146]. From a theoretical standpoint, the importance of such a term — even if only weak — makes it an interesting target for future studies.

Finally, transmitter-induced plasticity could be replaced by a growth term without explicit presynaptic dependence. A plausible candidate for such a mechanism would for instance be spontaneous spine-growth in the vicinity of a presynaptic axon. However, whether or not these rates would be on the correct timescale to compensate LTD effects requires further theoretical investigation.

Consolidation of synapses is summarized in the present model by a reference weight \tilde{w} [54, 103]. Simulations predict that synaptic consolidation renders synapses inert against heterosynaptic plasticity. Intuitively, the measured synaptic weights become “sticky” and are always attracted back to their momentary stable state, i.e. weak or strong. This prediction requires future experimental clarification.

The path towards saturation of synaptic weights during a pairing experiment (Fig. 8) is vital to building better plasticity models. Virtually any information which helps theorists to constrain how the synaptic weight increases would be helpful. Importantly, this also includes any information

about conditions (or experimental protocols) which do not induce plasticity, despite the fact that either the presynaptic or the postsynaptic neuron or both have been activated.

Conclusion

One of the most striking differences between plasticity models and experimental data concerns their timescales. Hebbian plasticity can be induced within seconds to minutes [45, 68, 69, 87]. In simulated network models, a similarly fast form of Hebbian plasticity leads to run-away activity within seconds, unless Hebbian plasticity is complemented with RCPs. Here, “rapid” means that these changes need to take effect after seconds or at most a few minutes [32]. This, however, is much faster than homeostatic plasticity observed in experiments. One of the most extensively studied forms of homeostasis in experiments is synaptic scaling [1] which proportionally scales synapses up or down when the network activity is too low or too high respectively. However, even the fastest known forms of scaling take hours to days to cause measurable changes to synaptic weights (Fig. 1; [15, 35, 36]).

This apparent difference of timescales between RCPs required for stability in models and experimental results is a challenge for current theories [23, 32, 118, 147]. To reconcile plasticity models and stability in networks of simulated neurons, we need to reconsider models of Hebbian plasticity and how they are fitted to data.

In most plasticity induction experiments neither the time course of the manipulated synaptic state nor the precise changes of other synapses are observable during stimulation. Quantitative models of synaptic plasticity thus make minimal assumptions about these unobserved temporal dynamics and generally ignore heterosynaptic effects entirely. In other words, missing experimental data makes it possible to build different models which all capture the existing experimental data, but make different assumptions about the unobserved dynamics. Importantly, some of these models become intrinsically stable [10, 57, 118] or even multistable [23, 103]. In most situations these models can be interpreted as compound models consisting of Hebbian plasticity and forms of RCPs which only rely on quantities that are locally known to the synapse, i.e. the pre- postsynaptic activity as well as its own synaptic weight. Although such local forms of plasticity can solve the problem of stability at a neuronal level, in practice, most network models require additional fine-tuning of parameters to achieve plausible activity levels across a network of neurons. This role can be fulfilled by slow homeostatic mechanisms which act on timescales of hours or days, consistent with experimental data on homeostatic plasticity.

In summary, several theoretical arguments suggest that Hebbian plasticity is intrinsically stabilized on short timescales by RCPs, likely to be implemented as heterosynaptic plasticity, or network wide negative feedback mechanisms. Slow forms of homeostatic plasticity, on the other hand, set the stage for stable learning. This hypothesis will now have to stand the test of time. It will thus be an important challenge for the coming years to go beyond homosynaptic Hebbian plasticity and to

gain a more complete understanding of its interactions with a diversity of compensatory processes across timescales.

Funding statement

WG was supported for this work by the European Research Council under grant agreement number 268689 (MultiRules) and by the European Community’s Seventh Framework Program under grant no. 604102 (Human Brain Project). FZ was supported by the SNSF (Swiss National Science Foundation).

References

- [1] Gina G. Turrigiano, Kenneth R. Leslie, Niraj S. Desai, Lana C. Rutherford, and Sacha B. Nelson. Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, 391(6670):892–896, February 1998. doi: 10.1038/36103.
- [2] Donald O. Hebb. *The Organization of Behavior: A Neuropsychological Theory*. Wiley & Sons New York, 1949.
- [3] R. G. M. Morris, E. Anderson, G. S. Lynch, and M. Baudry. Selective impairment of learning and blockade of long-term potentiation by an N-methyl-D-aspartate receptor antagonist, AP5. *Nature*, 319(6056):774–776, February 1986. doi: 10.1038/319774a0.
- [4] S. J. Martin, P. D. Grimwood, and R. G. M. Morris. Synaptic Plasticity and Memory: An Evaluation of the Hypothesis. *Annu Rev Neurosci*, 23(1):649–711, 2000. doi: 10.1146/annurev.neuro.23.1.649.
- [5] EL Bienenstock, LN Cooper, and PW Munro. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci*, 2(1):32–48, January 1982.
- [6] R. Linsker. Self-organization in a perceptual network. *Computer*, 21:105–117, 1988.
- [7] Brian S. Blais, N. Intrator, H. Shouval, and Leon N. Cooper. Receptive Field Formation in Natural Scene Environments: Comparison of Single-Cell Learning Rules. *Neural Computation*, 10(7):1797–1813, October 1998. doi: 10.1162/089976698300017142.
- [8] Sen Song, Kenneth D. Miller, and L. F. Abbott. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci*, 3(9):919–926, September 2000. doi: 10.1038/78829.

- [9] Carlos S. N. Brito and Wulfram Gerstner. Nonlinear Hebbian Learning as a Unifying Principle in Receptive Field Formation. *PLOS Comput Biol*, 12(9):e1005070, September 2016. doi: 10.1371/journal.pcbi.1005070.
- [10] Kenneth D Miller and David JC MacKay. The role of constraints in Hebbian learning. *Neural Comput*, 6(1):100–126, 1994.
- [11] L. F. Abbott and Sacha B. Nelson. Synaptic plasticity: taming the beast. *Nat Neurosci*, 3: 1178–1183, 2000. doi: 10.1038/81453.
- [12] Alfonso Renart, Pengcheng Song, and Xiao-Jing Wang. Robust Spatial Working Memory through Homeostatic Synaptic Scaling in Heterogeneous Cortical Networks. *Neuron*, 38(3): 473–485, May 2003. doi: 10.1016/S0896-6273(03)00255-1.
- [13] Ila R Fiete, Walter Senn, Claude Z H Wang, and Richard H R Hahnloser. Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron*, 65(4):563–576, February 2010. doi: 10.1016/j.neuron.2010.02.003.
- [14] Ashok Litwin-Kumar and Brent Doiron. Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nat Commun*, 5, November 2014. doi: 10.1038/ncomms6319.
- [15] Gina G. Turrigiano. The Self-Tuning Neuron: Synaptic Scaling of Excitatory Synapses. *Cell*, 135(3):422–435, October 2008. doi: 10.1016/j.cell.2008.10.008.
- [16] Alanna J. Watt and Niraj S. Desai. Homeostatic plasticity and STDP: keeping a neuron’s cool in a fluctuating world. *Front Synaptic Neurosci*, 2:5, 2010. doi: 10.3389/fnsyn.2010.00005.
- [17] Wickliffe C. Abraham. Metaplasticity: Tuning synapses and networks for plasticity. *Nat Rev Neurosci*, 9(5):387–387, May 2008. doi: 10.1038/nrn2356.
- [18] Ada X. Yee, Yu-Tien Hsu, and Lu Chen. A metaplastic view of the interaction between homeostatic and Hebbian plasticity. *Phil Trans R Soc B*, (accepted), 2016.
- [19] Gina G Turrigiano and Sacha B Nelson. Hebb and homeostasis in neuronal plasticity. *Curr Opin Neurobiol*, 10(3):358–364, June 2000. doi: 10.1016/S0959-4388(00)00091-X.
- [20] Marina Chistiakova, Nicholas M. Bannon, Maxim Bazhenov, and Maxim Volgushev. Heterosynaptic Plasticity Multiple Mechanisms and Multiple Roles. *Neuroscientist*, 20(5):483–498, October 2014. doi: 10.1177/1073858414529829.
- [21] Marina Chistiakova, Nicholas M. Bannon, Jen-Yung Chen, Maxim Bazhenov, and Maxim Volgushev. Homeostatic role of heterosynaptic plasticity: models and experiments. *Front Comput Neurosci*, 9:89, 2015. doi: 10.3389/fncom.2015.00089.

- [22] Gina G. Turrigiano. The Dialectic of Hebb and Homeostasis. *Phil Trans R Soc B*, (accepted), 2016.
- [23] Taro Toyozumi, Megumi Kaneko, Michael P. Stryker, and Kenneth D. Miller. Modeling the Dynamic Interaction of Hebbian and Homeostatic Plasticity. *Neuron*, 84(2):497–510, October 2014. doi: 10.1016/j.neuron.2014.09.036.
- [24] Christoph von der Malsburg. Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik*, 14(2):85–100, December 1973.
- [25] Andreea Lazar, Gordon Pipa, and Jochen Triesch. SORN: A Self-Organizing Recurrent Neural Network. *Front Comput Neurosci*, 3, October 2009. doi: 10.3389/neuro.10.023.2009.
- [26] Claudia Clopath, Lars BÜsing, Eleni Vasilaki, and Wulfram Gerstner. Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nat Neurosci*, 13(3):344–352, March 2010. doi: 10.1038/nn.2479.
- [27] Sami El Boustani, Pierre Yger, Yves Frégnac, and Alain Destexhe. Stable Learning in Stochastic Network States. *J Neurosci*, 32(1):194 –214, January 2012. doi: 10.1523/JNEUROSCI.2496-11.2012.
- [28] Julijana Gjorgjieva, Claudia Clopath, Juliette Audet, and Jean-Pascal Pfister. A triplet spike-timing-dependent plasticity model generalizes the Bienenstock–Cooper–Munro rule to higher-order spatiotemporal correlations. *Proc Natl Acad Sci U S A*, 108(48):19383 –19388, November 2011. doi: 10.1073/pnas.1105933108.
- [29] Peter Jedlicka, Lubica Benuskova, and Wickliffe C. Abraham. A Voltage-Based STDP Rule Combined with Fast BCM-Like Metaplasticity Accounts for LTP and Concurrent “Heterosynaptic” LTD in the Dentate Gyrus In Vivo. *PLoS Comput Biol*, 11(11):e1004588, November 2015. doi: 10.1371/journal.pcbi.1004588.
- [30] Lubica Benuskova and Wickliffe Abraham. STDP rule endowed with the BCM sliding threshold accounts for hippocampal heterosynaptic plasticity. *J Comput Neurosci*, 22(2):129–133, 2007. doi: 10.1007/s10827-006-0002-x.
- [31] M. C. W. van Rossum, G. Q. Bi, and G. G. Turrigiano. Stable Hebbian Learning from Spike Timing-Dependent Plasticity. *J Neurosci*, 20(23):8812–8821, December 2000.
- [32] Friedemann Zenke, Guillaume Hennequin, and Wulfram Gerstner. Synaptic Plasticity in Neural Networks Needs Homeostasis with a Fast Rate Detector. *PLoS Comput Biol*, 9(11):e1003330, November 2013. doi: 10.1371/journal.pcbi.1003330.

- [33] Jean-Pascal Pfister and Wulfram Gerstner. Triplets of Spikes in a Model of Spike Timing-Dependent Plasticity. *J Neurosci*, 26(38):9673–9682, September 2006. doi: 10.1523/JNEUROSCI.1425-06.2006.
- [34] Y Y Huang, A Colino, D K Selig, and R C Malenka. The influence of prior synaptic activity on the induction of long-term potentiation. *Science*, 255(5045):730–733, February 1992.
- [35] Jason Aoto, Christine I. Nam, Michael M. Poon, Pamela Ting, and Lu Chen. Synaptic Signaling by All-Trans Retinoic Acid in Homeostatic Synaptic Plasticity. *Neuron*, 60(2):308–320, October 2008. doi: 10.1016/j.neuron.2008.08.012.
- [36] Keiji Ibata, Qian Sun, and Gina G. Turrigiano. Rapid Synaptic Scaling Induced by Changes in Postsynaptic Firing. *Neuron*, 57(6):819–826, March 2008. doi: 10.1016/j.neuron.2008.02.031.
- [37] Brian R. Christie and Wickliffe C. Abraham. Priming of associative long-term depression in the dentate gyrus by theta frequency synaptic activity. *Neuron*, 9(1):79–84, July 1992. doi: 10.1016/0896-6273(92)90222-Y.
- [38] Stuart David Greenhill, Adam Ranson, and Kevin Fox. Hebbian and Homeostatic Plasticity Mechanisms in Regular Spiking and Intrinsic Bursting Cells of Cortical Layer 5. *Neuron*, 88(3):539–552, November 2015. doi: 10.1016/j.neuron.2015.09.025.
- [39] Megumi Kaneko, David Stellwagen, Robert C. Malenka, and Michael P. Stryker. Tumor Necrosis Factor- α Mediates One Component of Competitive, Experience-Dependent Plasticity in Developing Visual Cortex. *Neuron*, 58(5):673–680, June 2008. doi: 10.1016/j.neuron.2008.04.023.
- [40] Lu Li, Melanie A. Gainey, Joseph E. Goldbeck, and Daniel E. Feldman. Rapid homeostasis by disinhibition during whisker map plasticity. *PNAS*, 111(4):1616–1621, January 2014. doi: 10.1073/pnas.1312455111.
- [41] Anubhuthi Goel, Bin Jiang, Linda W. Xu, Lihua Song, Alfredo Kirkwood, and Hey-Kyoung Lee. Cross-modal regulation of synaptic AMPA receptors in primary sensory cortices by visual experience. *Nat Neurosci*, 9(8):1001–1003, August 2006. doi: 10.1038/nn1725.
- [42] Larry F Abbott, J A Varela, K Sen, and S B Nelson. Synaptic depression and cortical gain control. *Science*, 275(5297):220–224, January 1997.
- [43] Misha V. Tsodyks and Henry Markram. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *PNAS*, 94(2):719–723, January 1997.
- [44] Misha Tsodyks, Klaus Pawelzik, and Henry Markram. Neural Networks with Dynamic Synapses. *Neural Comput*, 10(4):821–835, 1998. doi: 10.1162/089976698300017502.

- [45] T. V. P. Bliss and T. Lømo. Long-Lasting Potentiation of Synaptic Transmission in the Dentate Area of the Anaesthetized Rabbit Following Stimulation of the Perforant Path. *J Physiol*, 232(2):331–356, July 1973.
- [46] Robert C. Malenka and and Roger A. Nicoll. Long-Term Potentiation—A Decade of Progress? *Science*, 285(5435):1870–1874, September 1999. doi: 10.1126/science.285.5435.1870.
- [47] John Lisman. Long-term potentiation: outstanding questions and attempted synthesis. *Phil Trans R Soc Lond B*, 358(1432):829–842, April 2003. doi: 10.1098/rstb.2002.1242.
- [48] Uwe Frey and Richard G. M. Morris. Synaptic tagging: implications for late maintenance of hippocampal long-term potentiation. *Trends Neurosci*, 21(5):181–188, May 1998. doi: 10.1016/S0166-2236(97)01189-2.
- [49] Klaus G. Reymann and Julietta U. Frey. The late maintenance of hippocampal LTP: Requirements, phases, ‘synaptic tagging’, ‘late-associativity’ and implications. *Neuropharmacology*, 52(1):24–40, January 2007. doi: 10.1016/j.neuropharm.2006.07.026.
- [50] Roger L. Redondo and Richard G. M. Morris. Making memories last: the synaptic tagging and capture hypothesis. *Nat Rev Neurosci*, 12(1):17–30, January 2011. doi: 10.1038/nnr2963.
- [51] Stefano Fusi, Patrick J. Drew, and Larry F. Abbott. Cascade models of synaptically stored memories. *Neuron*, 45(4):599–611, February 2005. doi: 10.1016/j.neuron.2005.02.001.
- [52] Subhaneil Lahiri and Surya Ganguli. A memory frontier for complex synapses. In *Advances in Neural Information Processing Systems*, volume 26, pages 1034–1042, Tahoe, USA, 2013. Curran Associates, Inc.
- [53] Rui Ponte Costa, Robert C. Froemke, Per Jesper Sjöström, and Mark C. W. van Rossum. Unified pre- and postsynaptic long-term plasticity enables reliable and flexible learning. *eLife Sciences*, page e09457, August 2015. doi: 10.7554/eLife.09457.
- [54] Lorric Ziegler, Friedemann Zenke, David B. Kastner, and Wulfram Gerstner. Synaptic Consolidation: From Synapses to Behavioral Modeling. *J Neurosci*, 35(3):1319–1334, January 2015. doi: 10.1523/JNEUROSCI.3989-14.2015.
- [55] Stephen Grossberg. On learning of spatiotemporal patterns by networks with ordered sensory and motor components. 1. Excitatory components of the cerebellum. *Stud Appl Math*, pages 105–132, 1969.
- [56] David J Willshaw and Christoph Von Der Malsburg. How patterned neural connections can be set up by self-organization. *Proc R Soc Lond B Biol Sci*, 194(1117):431–445, 1976.

- [57] Erkki Oja. Simplified neuron model as a principal component analyzer. *J Math Biol*, 15(3): 267–273, 1982. doi: 10.1007/BF00275687.
- [58] Richard Kempter, Wulfram Gerstner, and J. van Hemmen. Hebbian learning and spiking neurons. *Phys Rev E*, 59(4):4498–4514, April 1999. doi: 10.1103/PhysRevE.59.4498.
- [59] Stefano Fusi, Mario Annunziato, Davide Badoni, Andrea Salamon, and Daniel J. Amit. Spike-Driven Synaptic Plasticity: Theory, Simulation, VLSI Implementation. *Neural Comput*, 12(10):2227–2258, 2000. doi: 10.1162/089976600300014917.
- [60] Walter Senn, Henry Markram, and Misha Tsodyks. An Algorithm for Modifying Neurotransmitter Release Probability Based on Pre- and Postsynaptic Spike Timing. *Neural Comput*, 13(1):35–67, January 2001. doi: 10.1162/089976601300014628.
- [61] Harel Z. Shouval, Mark F. Bear, and Leon N. Cooper. A Unified Model of NMDA Receptor-Dependent Bidirectional Synaptic Plasticity. *Proc Natl Acad Sci U S A*, 99(16):10831–10836, August 2002. doi: 10.1073/pnas.152343099.
- [62] Jonathan E. Rubin, Richard C. Gerkin, Guo-Qiang Bi, and Carson C. Chow. Calcium Time Course as a Signal for Spike-Timing-Dependent Plasticity. *J Neurophysiol*, 93(5):2600–2613, May 2005. doi: 10.1152/jn.00803.2004.
- [63] Claudia Clopath and Wulfram Gerstner. Voltage and spike timing interact in STDP – a unified model. *Front Synaptic Neurosci*, 2:25, 2010. doi: 10.3389/fnsyn.2010.00025.
- [64] Abigail Morrison, Markus Diesmann, and Wulfram Gerstner. Phenomenological models of synaptic plasticity based on spike timing. *Biol Cybern*, 98(6):459–478, June 2008. doi: 10.1007/s00422-008-0233-1.
- [65] Michael Graupner and Nicolas Brunel. Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. *Proc Natl Acad Sci U S A*, 109(10):3991–3996, March 2012. doi: 10.1073/pnas.1109359109.
- [66] G. S. Lynch, T. Dunwiddie, and V. Gribkoff. Heterosynaptic depression: a postsynaptic correlate of long-term potentiation. *Nature*, 266(5604):737–739, April 1977. doi: 10.1038/266737a0.
- [67] W. B. Levy and O. Steward. Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience*, 8(4):791–797, April 1983. doi: 10.1016/0306-4522(83)90010-6.
- [68] Henry Markram, Joachim Lübke, Michael Frotscher, and Bert Sakmann. Regulation of Synaptic Efficacy by Coincidence of Postsynaptic APs and EPSPs. *Science*, 275(5297):213–215, January 1997. doi: 10.1126/science.275.5297.213.

- [69] Guo-Qiang Bi and Mu-Ming Poo. Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. *J Neurosci*, 18(24):10464–10472, December 1998.
- [70] Per Jesper Sjöström, Gina G Turrigiano, and Sacha B Nelson. Rate, Timing, and Cooperativity Jointly Determine Cortical Synaptic Plasticity. *Neuron*, 32(6):1149–1164, December 2001. doi: 10.1016/S0896-6273(01)00542-6.
- [71] T. H. Brown, A. M. Zador, Z. F. Mainen, and B. J. Claiborne. Hebbian modifications in hippocampal neurons. In M. Baudry and J. L. Davis, editors, *Long-term potentiation.*, pages 357–389. MIT Press, Cambridge, London, 1991.
- [72] Wulfram Gerstner, Werner M Kistler, Richard Naud, and Liam Paninski. *Neuronal dynamics: from single neurons to networks and models of cognition*. Cambridge University Press, Cambridge, 2014.
- [73] Nicolas Frémaux and Wulfram Gerstner. Neuromodulated Spike-Timing-Dependent Plasticity, and Theory of Three-Factor Learning Rules. *Front Neural Circuits*, page 85, 2016. doi: 10.3389/fncir.2015.00085.
- [74] Wulfram Gerstner and Werner M. Kistler. Mathematical formulations of Hebbian learning. *Biol Cybern*, 87(5-6):404–415, December 2002. doi: 10.1007/s00422-002-0353-y.
- [75] W. C. Abraham and G. V. Goddard. Asymmetric relationships between homosynaptic long-term potentiation and heterosynaptic long-term depression. *Nature*, 305(5936):717–719, October 1983. doi: 10.1038/305717a0.
- [76] Thomas H. Brown and Sumantra Chattarji. Hebbian Synaptic Plasticity: Evolution of the Contemporary Concept. In Professor Eytan Domany, Professor Dr J. Leo van Hemmen, and Professor Klaus Schulten, editors, *Models of Neural Networks*, Physics of Neural Networks, pages 287–314. Springer New York, 1994.
- [77] Maxim Volgushev, Leon L. Voronin, Marina Chistiakova, and Wolf Singer. Relations Between Long-term Synaptic Modifications and Paired-pulse Interactions in the Rat Neocortex. *European Journal of Neuroscience*, 9(8):1656–1665, August 1997. doi: 10.1111/j.1460-9568.1997.tb01523.x.
- [78] Craig H. Bailey, Maurizio Giustetto, Yan-You Huang, Robert D. Hawkins, and Eric R. Kandel. Is Heterosynaptic modulation essential for stabilizing hebbian plasticity and memory. *Nat Rev Neurosci*, 1(1):11–20, October 2000. doi: 10.1038/35036191.

- [79] Paul A Salin, Robert C Malenka, and Roger A Nicoll. Cyclic AMP Mediates a Presynaptic Form of LTP at Cerebellar Parallel Fiber Synapses. *Neuron*, 16(4):797–803, April 1996. doi: 10.1016/S0896-6273(00)80099-9.
- [80] Varda Lev-Ram, Scott T. Wong, Daniel R. Storm, and Roger Y. Tsien. A new form of cerebellar long-term potentiation is postsynaptic and depends on nitric oxide but not cAMP. *PNAS*, 99(12):8389–8393, June 2002. doi: 10.1073/pnas.122206399.
- [81] Hyung-Bae Kwon and Bernardo L. Sabatini. Glutamate induces de novo growth of functional spines in developing cortex. *Nature*, 474(7349):100–104, June 2011. doi: 10.1038/nature09986.
- [82] Jonathan Rubin, Daniel D. Lee, and H. Sompolinsky. Equilibrium Properties of Temporally Asymmetric Hebbian Plasticity. *Phys. Rev. Lett.*, 86(2):364–367, January 2001. doi: 10.1103/PhysRevLett.86.364.
- [83] R. Gütiğ, R. Aharonov, S. Rotter, and Haim Sompolinsky. Learning Input Correlations through Nonlinear Temporally Asymmetric Hebbian Plasticity. *J Neurosci*, 23(9):3697–3714, May 2003.
- [84] Karl Johan Aström and Richard M. Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press, April 2010.
- [85] Vincent Delattre, Daniel Keller, Matthew Perich, Henry Markram, and Eilif Benjamin Muller. Network-timing-dependent plasticity. *Front. Cell. Neurosci.*, page 220, 2015. doi: 10.3389/fncel.2015.00220.
- [86] Won Chan Oh, Laxmi Kumar Parajuli, and Karen Zito. Heterosynaptic Structural Plasticity on Local Dendritic Segments of Hippocampal CA1 Neurons. *Cell Reports*, 10(2):162–169, January 2015. doi: 10.1016/j.celrep.2014.12.016.
- [87] A. Artola, S. Bröcher, and W. Singer. Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature*, 347(6288):69–72, September 1990. doi: 10.1038/347069a0.
- [88] Leon N Cooper, Nathan Intrator, Brian S Blais, and Harel Z Shouval. *Theory of Cortical Plasticity*. World Scientific, New Jersey, April 2004.
- [89] Daniel Harnack, Miha Pelko, Antoine Chaillet, Yacine Chitour, and Mark C.W. van Rossum. Stability of Neuronal Networks with Homeostatic Regulation. *PLoS Comput Biol*, 11(7): e1004357, July 2015. doi: 10.1371/journal.pcbi.1004357.
- [90] Pierre Yger and Matthieu Gilson. Models of Metaplasticity: A Review of Concepts. *Front. Comput. Neurosci*, page 138, 2015. doi: 10.3389/fncom.2015.00138.

- [91] Nicolas Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J Comput Neurosci*, 8(3):183–208, May 2000.
- [92] Tim P. Vogels, Kanaka Rajan, and Larry F. Abbott. Neural network dynamics. *Annu Rev Neurosci*, 28:357–376, January 2005. doi: 10.1146/annurev.neuro.28.061604.135637.
- [93] A Artola and W Singer. Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. *Trends Neurosci*, 16(11):480–487, November 1993.
- [94] Pierre Yger and Kenneth D. Harris. The Convallis Rule for Unsupervised Learning in Cortical Networks. *PLoS Comput Biol*, 9(10):e1003272, October 2013. doi: 10.1371/journal.pcbi.1003272.
- [95] Melanie A. Gainey and Daniel E. Feldman. Multiple shared mechanisms for homeostatic plasticity in rodent somatosensory and visual cortex. *Phil Trans R Soc B*, (accepted), 2016.
- [96] D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *J Physiol*, 160(1):106–154.2, January 1962.
- [97] Nikos K Logothetis and David L Sheinberg. Visual Object Recognition. *Annual Review of Neuroscience*, 19(1):577–621, 1996. doi: 10.1146/annurev.ne.19.030196.003045.
- [98] C. P. Hung, G. Kreiman, T. Poggio, and J. J. DiCarlo. Fast Readout of Object Identity from Macaque Inferior Temporal Cortex. *Science*, 310:863 – 866, 2005.
- [99] R. Quian Quiroga, L. Reddy, G. Kreiman, C. Koch, and I. Fried. Invariant visual representation by single neurons in the human brain. *Nature*, 435(7045):1102–1107, June 2005. doi: 10.1038/nature03687.
- [100] Yasushi Miyashita. Neural correlate of visual associative long-term memory in the primate temporal cortex. *Nature*, 335(27):817–820, October 1988.
- [101] S. Funahashi, C. J Bruce, and P. S Goldman-Rakic. Mnemonic Coding of Visual Space in the Monkey’s Dorsolateral Prefrontal Cortex. *J Neurophysiol*, 61(2):331–349, February 1989.
- [102] Fraser A. W. Wilson, Séamas P. Ó Scalaidhe, and Patricia S. Goldman-Rakic. Dissociation of Object and Spatial Processing Domains in Primate Prefrontal Cortex. *Science*, 260(5116):1955–1958, June 1993.
- [103] Friedemann Zenke, Everton J. Agnes, and Wulfram Gerstner. Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nat Commun*, 6, April 2015. doi: doi:10.1038/ncomms7922.

- [104] Werner M. Kistler and J. Leo Hemmen. Modeling synaptic plasticity in conjunction with the timing of pre-and postsynaptic action potentials. *Neural Computation*, 12(2):385–405, 2000.
- [105] Abigail Morrison, Ad Aertsen, and Markus Diesmann. Spike-timing-dependent plasticity in balanced random networks. *Neural Comput*, 19(6):1437–67, June 2007. doi: 10.1162/neco.2007.19.6.1437.
- [106] Guillaume Hennequin, Wulfram Gerstner, and Jean-Pascal Pfister. STDP in Adaptive Neurons Gives Close-To-Optimal Information Transmission. *Front Comput Neurosci*, 4, December 2010. doi: 10.3389/fncom.2010.00143.
- [107] Christian Tetzlaff, Christoph Kolodziejski, Marc Timme, and Florentin Wörgötter. Analysis of synaptic scaling in combination with Hebbian plasticity in several simple networks. *Front Comput Neurosci*, 6:36, 2012. doi: 10.3389/fncom.2012.00036.
- [108] Matthieu Gilson, Anthony N. Burkitt, David B. Grayden, Doreen A. Thomas, and J. Leo Hemmen. Emergence of network structure due to spike-timing-dependent plasticity in recurrent neuronal networks. I. Input selectivity–strengthening correlated input pathways. *Biol Cybern*, 101(2):81–102, June 2009. doi: 10.1007/s00422-009-0319-4.
- [109] R Kempter, W Gerstner, and J L van Hemmen. Intrinsic stabilization of output rates by spike-based Hebbian learning. *Neural Comput*, 13(12):2709–2741, December 2001. doi: 10.1162/089976601317098501.
- [110] Baktash Babadi and L. F. Abbott. Stability and Competition in Multi-spike Models of Spike-Timing Dependent Plasticity. *PLOS Comput Biol*, 12(3):e1004750, March 2016. doi: 10.1371/journal.pcbi.1004750.
- [111] Tim P. Vogels, Henning Sprekeler, Friedemann Zenke, Claudia Clopath, and Wulfram Gerstner. Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks. *Science*, 334(6062):1569–1573, December 2011. doi: 10.1126/science.1211095.
- [112] Tim P Vogels, Robert C Froemke, Nicolas Doyon, Matthieu Gilson, Julie S Haas, Robert Liu, Arianna Maffei, Paul Miller, Corette Wierenga, Melanie A Woodin, Friedemann Zenke, and Henning Sprekeler. Inhibitory Synaptic Plasticity - Spike timing dependence and putative network function. *Front Neural Circuits*, 7(119), 2013. doi: 10.3389/fncir.2013.00119.
- [113] James A. D’Amour and Robert C. Froemke. Inhibitory and Excitatory Spike-Timing-Dependent Plasticity in the Auditory Cortex. *Neuron*, 86(2):514–528, April 2015. doi: 10.1016/j.neuron.2015.03.014.
- [114] Sukbin Lim, Jillian L. McKee, Luke Woloszyn, Yali Amit, David J. Freedman, David L. Sheinberg, and Nicolas Brunel. Inferring learning rules from distributions of firing rates in

- cortical neurons. *Nat Neurosci*, advance online publication, November 2015. doi: 10.1038/nn.4158.
- [115] Lang Wang and Arianna Maffei. Inhibitory Plasticity Dictates the Sign of Plasticity at Excitatory Synapses. *J. Neurosci.*, 34(4):1083–1093, January 2014. doi: 10.1523/JNEUROSCI.4711-13.2014.
 - [116] J. Y. Delgado, J. F. Gomez-Gonzalez, and N. S. Desai. Pyramidal Neuron Conductance State Gates Spike-Timing-Dependent Plasticity. *J Neurosci*, 30(47):15713–15725, November 2010. doi: 10.1523/JNEUROSCI.3068-10.2010.
 - [117] Sébastien Royer and Denis Paré. Conservation of total synaptic weight through balanced synaptic depression and potentiation. *Nature*, 422(6931):518–522, April 2003. doi: 10.1038/nature01530.
 - [118] Jen-Yung Chen, Peter Lonjers, Christopher Lee, Marina Chistiakova, Maxim Volgushev, and Maxim Bazhenov. Heterosynaptic Plasticity Prevents Runaway Synaptic Dynamics. *J Neurosci*, 33(40):15915–15929, October 2013. doi: 10.1523/JNEUROSCI.5088-12.2013.
 - [119] Maxim Volgushev, Jen-Yung Chen, Vladimir Ilin, Roman Goz, Marina Chistiakova, and Maxim Bazhenov. Partial Breakdown of Input Specificity of STDP at Individual Synapses Promotes New Learning. *J. Neurosci.*, 36(34):8842–8855, August 2016. doi: 10.1523/JNEUROSCI.0552-16.2016.
 - [120] Wickliffe C. Abraham, Barbara Logan, Amy Wolff, and Lubica Benuskova. “Heterosynaptic” LTD in the Dentate Gyrus of Anesthetized Rat Requires Homosynaptic Activity. *J Neurophysiol*, 98(2):1048–1051, August 2007. doi: 10.1152/jn.00250.2007.
 - [121] Uwe Frey and Richard G. M. Morris. Synaptic tagging and long-term potentiation. *Nature*, 385(6616):533–6, March 1997. doi: 10.1038/385533a0.
 - [122] Friedemann Zenke. *Memory formation and recall in recurrent spiking neural networks*. PhD thesis, École polytechnique fédérale de Lausanne EPFL, Lausanne, Switzerland, 2014.
 - [123] Eve Marder and Jean-Marc Goaillard. Variability, compensation and homeostasis in neuron and network function. *Nat Rev Neurosci*, 7(7):563–574, July 2006. doi: 10.1038/nrn1949.
 - [124] Gina G. Turrigiano. Homeostatic Synaptic Plasticity: Local and Global Mechanisms for Stabilizing Neuronal Function. *Cold Spring Harb Perspect Biol*, 4(1):a005736, January 2012. doi: 10.1101/cshperspect.a005736.
 - [125] Alex Williams, Timothy O’Leary, and Eve Marder. Homeostatic Regulation of Neuronal Excitability. *Scholarpedia*, 8(1):1656, 2013. doi: 10.4249/scholarpedia.1656.

- [126] Graeme W. Davis. Homeostatic Signaling and the Stabilization of Neural Function. *Neuron*, 80(3):718–728, October 2013. doi: 10.1016/j.neuron.2013.09.044.
- [127] Julijana Gjorgjieva, Guillaume Drion, and Eve Marder. Computational implications of biophysical diversity and multiple timescales in neurons and synapses for circuit performance. *Current Opinion in Neurobiology*, 37:44–52, April 2016. doi: 10.1016/j.conb.2015.12.008.
- [128] Yann Sweeney, Jeanette Hellgren Kotaleski, and Matthias H. Hennig. A Diffusive Homeostatic Signal Maintains Neural Heterogeneity and Responsiveness in Cortical Networks. *PLoS Comput Biol*, 11(7):e1004389, July 2015. doi: 10.1371/journal.pcbi.1004389.
- [129] Gaël Daoudal and Dominique Debanne. Long-Term Plasticity of Intrinsic Excitability: Learning Rules and Mechanisms. *Learn Mem*, 10(6):456–465, November 2003. doi: 10.1101/lm.64103.
- [130] Gina Turrigiano. Too Many Cooks? Intrinsic and Synaptic Homeostatic Mechanisms in Cortical Circuit Refinement. *Annu Rev Neurosci*, 34(1):89–103, 2011. doi: 10.1146/annurev-neuro-060909-153238.
- [131] David R. C. House, Justin Elstrott, Eileen Koh, Jason Chung, and Daniel E. Feldman. Parallel Regulation of Feedforward Inhibition and Excitation during Whisker Map Plasticity. *Neuron*, 72(5):819–831, December 2011. doi: 10.1016/j.neuron.2011.09.008.
- [132] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent Component Analysis*. Wiley, New York, 2001.
- [133] Cengiz Pehlevan and Dmitri B. Chklovskii. A Hebbian/Anti-Hebbian Network Derived from Online Non-Negative Matrix Factorization Can Cluster and Discover Sparse Features. *arXiv:1503.00680 [cs, q-bio, stat]*, March 2015.
- [134] Bruno A. Olshausen and David J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, June 1996. doi: 10.1038/381607a0.
- [135] Paul D. King, Joel Zylberberg, and Michael R. DeWeese. Inhibitory Interneurons Decorrelate Excitatory Cells to Drive Sparse Code Formation in a Spiking Model of V1. *J. Neurosci.*, 33(13):5475–5485, March 2013. doi: 10.1523/JNEUROSCI.4188-12.2013.
- [136] D. J. Willshaw, O. P. Bunemann, and H. C. Longuet-Higgins. Non-holographic associative memory. *Nature*, 222:960–962, 1969.
- [137] W. A. Little and Gordon L. Shaw. Analytic study of the memory storage capacity of a neural network. *Math Biosci*, 39(3):281–290, June 1978. doi: 10.1016/0025-5564(78)90058-5.

- [138] John J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci U S A*, 79(8):2554, 1982.
- [139] Daniel J. Amit, Hanoch Gutfreund, and H. Sompolinsky. Storing Infinite Numbers of Patterns in a Spin-Glass Model of Neural Networks. *Phys Rev Lett*, 55(14):1530–1533, September 1985. doi: 10.1103/PhysRevLett.55.1530.
- [140] Daniel J. Amit and Nicolas Brunel. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb Cortex*, 7(3):237–252, May 1997.
- [141] Florence I. Kleberg, Tomoki Fukai, and Matthieu Gilson. Excitatory and inhibitory STDP jointly tune feedforward neural circuits to selectively propagate correlated spiking activity. *Front Comput Neurosci*, 8:53, 2014. doi: 10.3389/fncom.2014.00053.
- [142] Gina G. Turrigiano and Sacha B. Nelson. Homeostatic plasticity in the developing nervous system. *Nat Rev Neurosci*, 5(2):97–107, February 2004. doi: 10.1038/nrn1327.
- [143] Sarah R. Hulme, Owen D. Jones, Clarke R. Raymond, Pankaj Sah, and Wickliffe C. Abraham. Mechanisms of heterosynaptic metaplasticity. *Phil Trans R Soc B*, 369(1633):20130148, January 2014. doi: 10.1098/rstb.2013.0148.
- [144] Qiang Zhou, Huizhong W. Tao, and Mu-Ming Poo. Reversal and Stabilization of Synaptic Modifications in a Developing Visual System. *Science*, 300(5627):1953–1957, June 2003. doi: 10.1126/science.1082212.
- [145] M.e. Gilbert and C.m. Mack. Field potential recordings in dentate gyrus of anesthetized rats: Stability of baseline. *Hippocampus*, 9(3):277–287, January 1999. doi: 10.1002/(SICI)1098-1063(1999)9:3<277::AID-HIPO7>3.0.CO;2-H.
- [146] John Th. Rick and Norton W. Milgram. Instability of Dentate Gyrus Field Potentials in Awake and Anesthetized Rats. *Hippocampus*, 9(3):333–339, January 1999. doi: 10.1002/(SICI)1098-1063(1999)9:3<333::AID-HIPO12>3.0.CO;2-2.
- [147] Zhihua Wu and Yoko Yamaguchi. Conserving total synaptic weight ensures one-trial sequence learning of place fields in the hippocampus. *Neural Networks*, 19(5):547–563, June 2006. doi: 10.1016/j.neunet.2005.06.048.