Stability by gating plasticity in recurrent neural networks

Katharina A. Wilmes[1,†], Claudia Clopath[1]

1 Bioengineering Department, Imperial College London, SW72AZ, London, UK

† present address: Department of Physiology, Universität Bern, Bühlplatz 5, 3012 Bern

September 10, 2020

Abstract

With Hebbian learning 'who fires together wires together', well-known problems arise. On the one hand, plasticity can lead to unstable network dynamics, manifesting as run-away activity or silence. On the other hand, plasticity can erase or overwrite stored memories. Unstable dynamics can partly be addressed with homeostatic plasticity mechanisms. Unfortunately, the time constants of homeostatic mechanisms required in network models are much shorter than what has been measured experimentally.

Here, we propose that homeostatic time constants can be slow if plasticity is gated. We investigate how the gating of plasticity influences the stability of network activity and stored memories. We use plastic balanced spiking neural networks consisting of excitatory neurons with a somatic and a dendritic compartment (which resemble cortical pyramidal cells in their firing properties), and inhibitory neurons targeting those compartments. We compare how different factors such as excitability, learning rate, and inhibition can lift the requirements for the critical time constant of homeostatic plasticity. We specifically investigate how gating of dendritic versus somatic plasticity allows for different amounts of weight changes in networks with the same critical homeostatic time constant. We suggest that the striking compartmentalisation of pyramidal cells and their inhibitory inputs enable large synaptic changes at the dendrite while maintaining network stability. We additionally show that spatially restricted plasticity in a subpopulation of the network improves stability. Finally, we compare how the different gates affect the stability of memories in the network.

Keywords Synaptic Plasticity · Neuromodulation · Gating · Recurrent Network · Spiking neurons · Balanced network

1 Introduction

Hebbian plasticity is considered to be the neural hallmark for learning and memory. It enables the formation of cell assemblies as it strengthens connections between cells with correlated activity. On the downside, correlations between cells are increased even further with Hebbian plasticity. Theoretically, such a positive feedback loop leads to undesired unstable runaway activity (Abbott and Nelson, 2000). Cortical cells, however, fire at low rates in an asynchronous irregular manner. It is therefore unclear how neural activity in the functioning brain remains stable despite Hebbian plasticity. To resolve this dilemma, it has been suggested that homeostatic processes keep the network activity stable (Turrigiano and Nelson, 2004). Homeostatic processes, such as homeostatic scaling (Turrigiano et al., 1998; Turrigiano, 2017; Desai et al., 2002; Goel and Lee, 2007; Keck et al., 2013; Maffei and Turrigiano, 2008) or inhibitory plasticity (Woodin and Maffei, 2010; Goel and Lee, 2007; Kuhlman et al., 2013; Keck et al., 2011; Chen et al., 2011; Vogels et al., 2011; van Versendaal and Levelt, 2016; Li et al., 2014), counteract increases in the network activity, but it has been proposed that they might be insufficient to keep the network activity stable for the following reason: these processes operate on a timescale of hours or days (Keck et al., 2017; Kaneko et al., 2008a; Greenhill et al., 2015; Kaneko et al., 2008b; Watt and Desai, 2010), but theoretical models require homeostatic mechanisms that act on the same timescale as Hebbian plasticity or faster (Zenke et al., 2013, 2017; Toyoizumi et al., 2014; Litwin-Kumar et al., 2016; Miller and MacKay, 1994; van Rossum et al., 2000). Zenke et al. (2017), therefore, proposed that there must be a fast compensatory mechanism. Such a mechanism could modulate plasticity itself (Naumann and Sprekeler, 2020). Models requiring fast mechanisms typically assume that plasticity is continuously happening (e.g. Litwin-Kumar et al., 2016; Zenke et al., 2013). In contrast, in the brain plasticity is highly regulated by different neuromodulators (Couey et al., 2007; Bissière et al., 2003; Lin et al., 2003; Pawlak et al., 2010; Seol et al., 2007; Zhang et al., 2009; Brzosko et al., 2015), astrocytes (Valtcheva and Venance, 2016), and inhibitory interneurons (Artinian and Lacaille, 2018; Kuhlman et al., 2013; Letzkus et al., 2011). These different regulators of plasticity can slow down, speed up, gate, or flip plasticity. They differ in their temporal and spatial precision and hence enable rigorous plasticity control. We, therefore, explore how different regulators or gates affect plasticity and stability. Specifically, we study using computational modelling whether gating of plasticity can lower the requirements for time constants of homeostatic plasticity to values that are more in line with experimentally observed homeostatic processes without strongly impairing plasticity of synaptic connections.

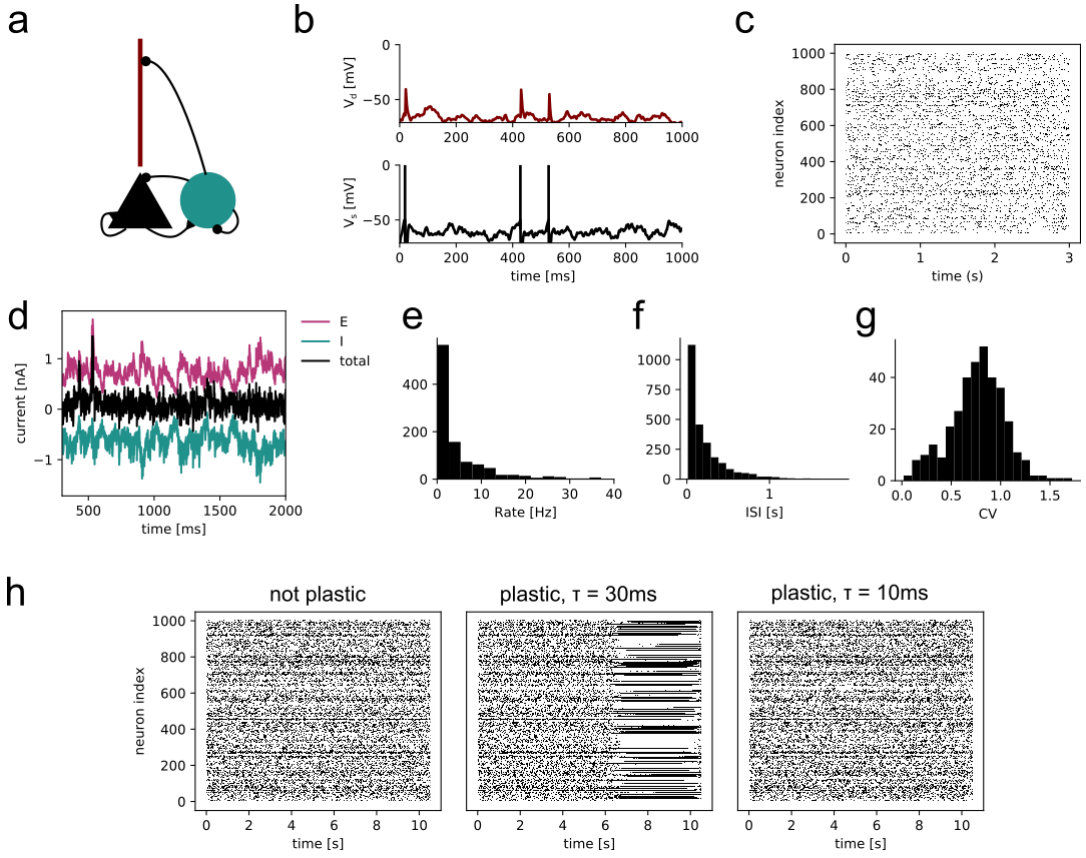

Figure 1: a: The network consisted of 1000 recurrently connected 2-compartment pyramidal cells (triangle and stick), and 250 recurrently connected inhibitory cells (circle). b: Somatic (black) and dendritic (red) voltage traces from one example pyramidal cell. c: Raster plot of excitatory cell activity in the network. d: Example currents from one example pyramidal cell. It receives large E (magenta) and I (cyan) currents which cancel on average (black). e: Distribution of excitatory firing rates. f: distribution of excitatory interspike intervals. g: Distribution of coefficient of variation (CV) of the interspike intervals. e-g indicate that the network is in a balanced state. h: Raster plots of excitatory network activity in a network without plasticity (left), with plasticity and a homeostatic time constant τ =30ms, and with plasticity and a homeostatic time constant τ =10ms.

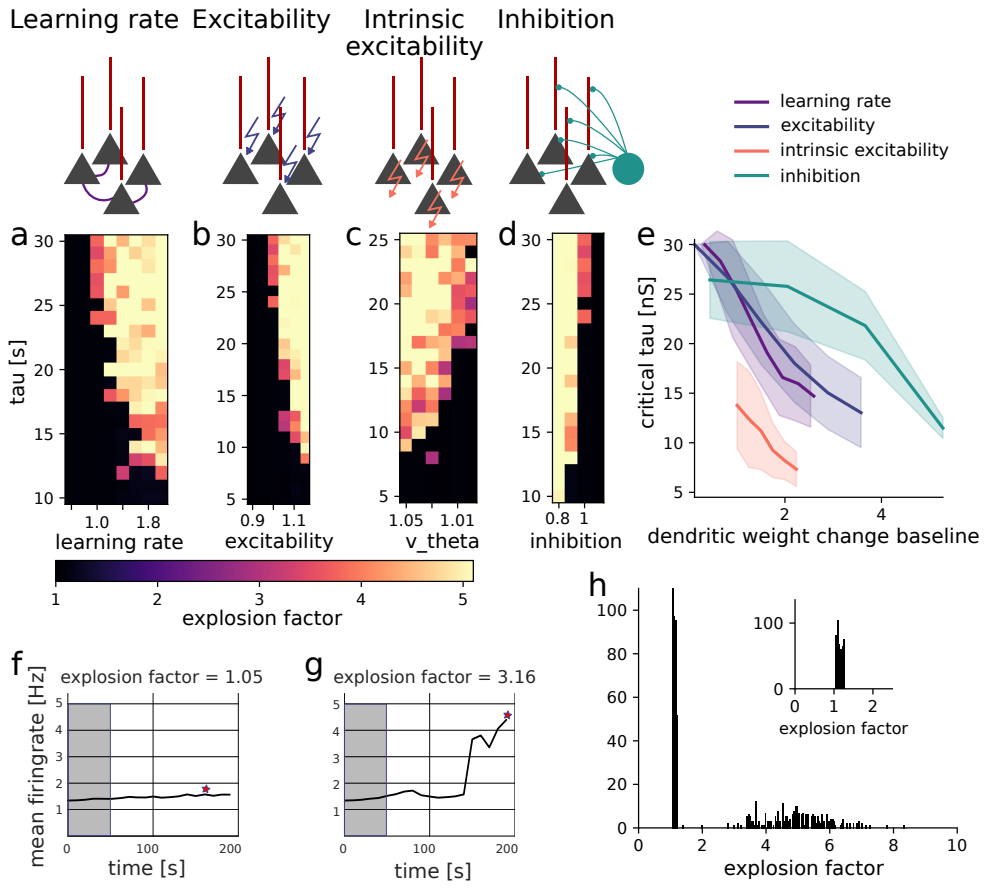

Figure 2: **Gating plasticity loosens the requirements for fast homeostatic processes.** a-d: Explosion factor as a function of homeostatic time constant τ and the respective gate. a: learning rate, b: excitability, c: intrinsic excitability (spiking threshold), d: inhibition. e: comparison of the critical homeostatic time constant τ_{crit} for different gates, plotted as a function of baseline dendritic weight change to allow for comparison. f-g: Illustration of the explosion factor. The star indicates the maximum firing rate of each simulation that was taken for the measurement of the explosion factor. The grey area denotes the reference firing rate at the beginning of the simulation, which was taken to calculate the explosion factor. f: Example network simulation, where the firing rate does not explode (with explosion factor 1.05). g: Example network simulation, where the firing rate explodes (with explosion factor 3.16). h: Distribution of explosion factors. Inset: zoom into the x-axis.

2 Results

2.1 Balanced spiking neural network with 2-compartment pyramidal cells.

To study how different modulators of plasticity affect stability and plasticity, we built a balanced recurrent network of 1000 excitatory pyramidal cells (E) and 250 inhibitory cells (I, Fig. 1). To explore somatic and dendritic gating separately, we modelled the pyramidal cells with two compartments, one for the soma and one for the dendrite (Fig. 1a,b, Naud and Sprekeler, 2018). Before implementing plasticity in our model, we made sure that the network is in the asynchronous irregular regime (Fig. 1 c,e,f,g), due to a balance between excitation and inhibition. That is, strong excitatory recurrent inputs were balanced by strong inhibitory feedback (inhibition-stabilized regime, Tsodyks et al., 1997). On the single-cell level, this is reflected in large excitatory and inhibitory currents, which cancel each other on average (Fig. 1d).

To test the effect of plasticity in our network, we added a standard triplet STDP rule (Pfister and Gerstner, 2006; Clopath et al., 2010; Bono et al., 2017) to the excitatory connections. As this form of plasticity is Hebbian, it can lead to an explosion of activity in recurrent networks (Keck et al., 2017; Abbott and Nelson, 2000). To keep the activity of the network in the balanced state despite ongoing plasticity, we included homeostatic plasticity (Keck et al., 2017; Abbott and Nelson, 2000). Following previous work (Bienenstock et al., 1982; Pfister and Gerstner, 2006; Clopath et al., 2010; Zenke et al., 2013), the homeostatic process in our network monitored the postsynaptic firing rate and adjusted long-term depression (LTD) to keep the neurons at their target firing rate. The time constant τ of this homeostatic process is critical for stability as it determines how quickly the homeostatic process reacts to changes in firing rate. If τ is too large, the homeostatic plasticity cannot compensate for the correlation-based weight changes and the network activity explodes (Fig. 1h middle). When τ is sufficiently small, the homeostatic plasticity maintains stability (Fig. 1h right). The time constants required for stability in models are orders of magnitude smaller than those of homeostatic processes in the brain. Therefore, we explored in our model whether gating plasticity can allow for larger homeostatic time constants, to reconcile stability constraints with experimental data (Zenke et al., 2017; Turrigiano et al., 1998).

2.2 Gating plasticity loosens the requirements for fast homeostatic processes.

We next explored how different forms of plasticity modulation affect the required homeostatic time constant in our model. The first gate we explored was the *excitability* of individual neurons to model the fact that neuromodulation can change the size or the duration of excitatory postsynaptic potentials (EPSPs) (Rasmusson, 2000). In our two-compartment leaky-integrate-and-fire neuron model, we modelled excitability as a factor $\boldsymbol{\gamma}$ multiplied to the excitatory synaptic currents (Eq. 1,2, the superscript d indicates the dendritic variable). The second gate was *inhibition*. We modelled its modulation by changing the inhibitory conductance $\boldsymbol{g_I}$ (Eq. 12). The somatic and dendritic voltage (V_s and V_d respectively) was therefore modelled as

$$C_s\frac{dV_s}{dt} = -g_l(V_s - E_l) - \boldsymbol{\gamma} g_E V_s - \boldsymbol{g_I}(V_s - E_I) + \lambda(g_s\frac{1}{1+\exp(-\frac{V_d-E_d}{D_d})} + \omega^s) \tag{1}$$

$$C_d\frac{dV_d}{dt} = -(V_d - E_l) - \boldsymbol{\gamma^d} g_E^d V_d - \boldsymbol{g_I^d}(V_d - E_I) + g_d\frac{1}{1+\exp(-\frac{V_d-E_d}{D_d})} + c_d K(t) + \omega^d \tag{2}$$

where $C_{s/d}$ is the somatic/dendritic capacitance, $g_{I/E}$ is the inhibitory/excitatory conductance (index d indicates the dendritic variables), E_l is the reversal potential, g_s is the coupling from the dendrite to the soma, $\frac{1}{1+\exp(-\frac{V_d-E_d}{dd})}$ is a nonlinear term for the dendritic calcium dynamics (see Methods), $\omega_{ad/s}$ is the somatic/dendritic adaptation variable, K the kernel for the back-propagating action potential current with amplitude c_d (Naud and Sprekeler, 2018). λ ensures that the somato-dendritic coupling and adaptation are the same as in the model of Naud and Sprekeler (2018) (see Methods).

The third gate we explored was *intrinsic excitability*, which we modelled by a change in the spiking threshold v_θ. These three gates modulate plasticity indirectly by modulating the activity of the network. Finally, the last gate was *learning rate*, which modulates the synaptic weight changes directly, and was modelled as a factor $\boldsymbol{\eta}$ in the weight update. Formally, the somatic synaptic weight w_{ij} and dendritic synaptic weight w_{ij}^d from neuron j to neuron i changed as:

$$\frac{dw_{ij}}{dt} = \boldsymbol{\eta} w_0 A^+ z_j^+(t) z_i^{slow}(t-\epsilon) S_i(t) \tag{3}$$

$$- \boldsymbol{\eta} w_0 A_i^-(t) z_i^-(t) S_j(t) \tag{4}$$

$$\frac{dw_{ij}^d}{dt} = \boldsymbol{\eta^d} w_0 A^+ z_j^+(t) z_i^{slow}(t-\epsilon) S_i(t) \tag{5}$$

$$- (\boldsymbol{\eta^d} w_0 (A_i^-(t) z_i^-(t) + A^{Ca}(v_d > \theta_{Ca})) - \alpha) S_j(t) \tag{6}$$

where w_0 is the initial weight, $A^{+/-}$ is the amount of potentiation/depression constants for the triplet rule, A^{Ca} is the potentiation constant for the Ca^{2+} spike-dependent potentiation, $S_{i/j}$ is the post-/presynaptic spike train, z_j^+ is the presynaptic trace, z_i^- is the postsynaptic trace, z_i^{slow} is the postsynaptic trace with slow time constant. $t-\epsilon$ denotes that the value of the trace is taken before the action potential, which happened at time t. v_d is the dendritic membrane potential. θ_{Ca} is the threshold for Ca^{2+} spike-dependent plasticity (the term $v_d > \theta_{Ca}$ takes values 1 or 0 depending on whether v_d is above the threshold θ_{Ca}) and α is transmitter-induced depression.

To quantify how gating affects stability, we first defined the explosion factor measure to quantify whether the network is stable (explosion factor close to 1, Fig. 2f) or explodes (explosion factor > 1.5, Fig. 2g). The threshold of 1.5 for a network to be defined as exploding was based on the bimodal distribution of explosion factors (Fig. 2h).

We started by varying learning rate in both the somatic and the dendritic compartment (Fig. 2a). Expectedly, we found that with a low learning rate and a large homeostatic time constant τ, the network was stable (the black region in Fig. 2a). For higher learning rates, the network activity exploded already at low values of τ. This is expected as a higher learning rate increases the rate of synaptic change, which compromises the stability of the network. We defined the largest τ at which the network was still stable as the *critical homeostatic time constant* τ_{crit} (Fig. 2a). A decrease in learning rate increased this critical time constant τ_{crit}. Similarly, a decrease in excitability also increased the τ_{crit} (Fig. 2b, c). This is because increasing excitability increases the overall activity in the network, which in turn increases the amount of plasticity. An increase in inhibition on the contrary had the opposite effect on the critical time constant τ_{crit} (Fig. 2d). Increasing inhibitory inputs decreases firing rates in the network which improves network stability. Therefore, homeostatic mechanisms for network stability can be slower when excitability and learning rate are downregulated or when inhibition is upregulated.

Although these effects were to be expected, qualitatively, our computational model allowed us to compare them quantitatively. We next characterised the different gates by comparing their effects on τ_{crit}. To compare gates despite their different scales, we defined a common variable. That is, we plotted τ_{crit} as a function of the total dendritic weight change happening in a stable network (with a τ of 5 ms, see Methods, Fig. 2d). This analysis revealed that excitability and learning rate affect the critical time constant τ_{crit} in a different way than inhibition. τ_{crit} increases supralinearly as a function of the total dendritic weight change for the excitability and learning rate gates, whereas it increases sublinearly for the inhibition gate. To conclude, all gates can improve network stability and lower the requirements for the time constant of homeostatic mechanisms. However, excitability and learning rate have a larger modulating gain than inhibition.

2.3 Learning in dendrites facilitates synaptic changes while maintaining network stability.

The increase in the critical time constant by gating comes at the cost of the lack of plasticity (measured as the total dendritic weight change, Fig. 2e). However, pyramidal neurons consist of a soma and a complex ramified structure of

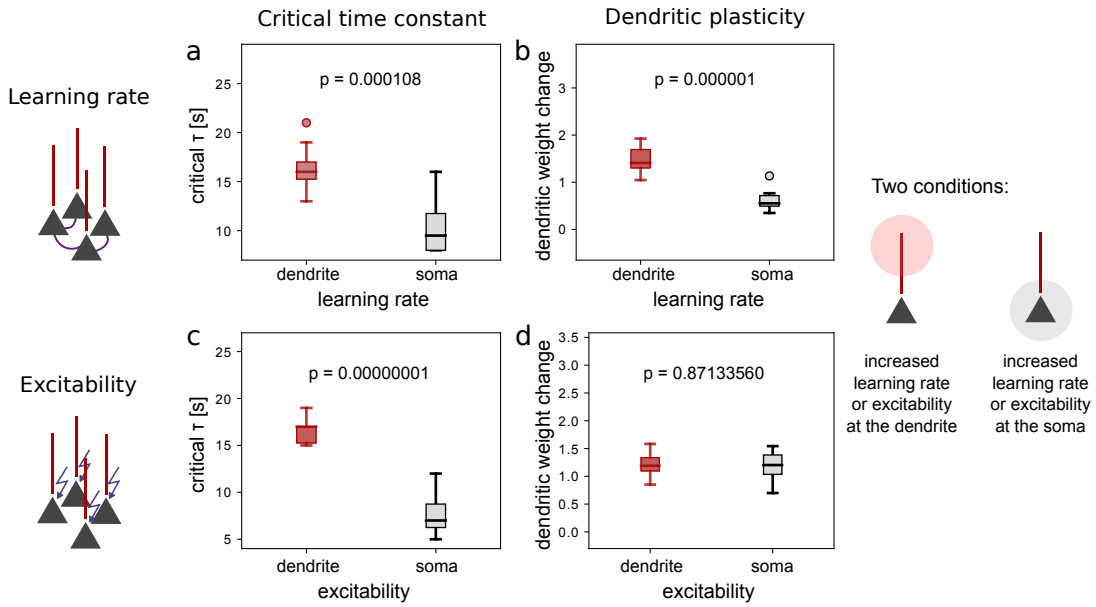

Figure 3: **Learning in dendrites helps mitigate the plasticity-stability dilemma.** a: Distribution of critical homeostatic time constants for a two-fold increase in learning rate in the dendrite (red) and in the soma (black). b: Distribution of dendritic weight changes for a two-fold increase in learning rate in the dendrite (red) and in the soma (black). c: Distribution of critical homeostatic time constants for a 15% increase in excitability in the dendrite (red) and in the soma (black). d: Distrbution of dendritic weight changes for a 15% increase in excitability in the dendrite (red) and in the soma (black). p-values were obtained by using the two-sample student's t-test.

dendrites. Interestingly, the majority of excitatory synapses are located on dendrites, electrotonically distant from the soma. Inspired by these observations, we hypothesised that the anatomy of pyramidal cells could enable both plasticity of dendritic synapses and stable somatic activity at the same time. We, therefore, increased the learning rate and the excitability separately for the soma and the dendrite and compared their impact on τ_{crit}.

We found that increasing plasticity (by increasing learning rate or excitability) in the dendrite compromised the critical time constant τ_{crit} less than in the soma (Fig. 3). τ_{crit} was significantly larger for a two-fold increase in learning rate in the dendrite than for the same increase in learning rate in the soma (Fig. 3a). Moreover, modulating learning rate only in the dendrite allowed for significantly higher dendritic weight changes at a larger critical time constant (Fig. 3b). Increasing excitability by 15% in the dendrite led to a significantly larger τ_{crit} than increasing excitability by 15% in the soma, while there was no difference in dendritic plasticity between the two conditions. Note that we chose a two-fold increase in learning rate and a 15% increase in excitability as these changes lower τ_{crit} by more than 50% (maximum explored values in Fig. 2e). In summary, by opening the gates for plasticity exclusively in the dendrite, the network can afford slower homeostatic mechanisms, higher network stability, while allowing the same or a higher amount of plasticity as when the gate is open at the soma.

2.4 Spatially precise neuromodulation enables learning while keeping network activity stable.

Neuromodulators were typically thought of as global and diffuse (Schwarz et al., 2015). However, neuromodulatory projections could in principle precisely target specific cell types and subpopulations, depending on their projective field and the receptor channels expressed in their targets. Specific neuromodulation (Totah et al., 2018) could enable plasticity locally when learning requires synaptic adjustments only in a subset of neurons. To test how local neuromodulation affects the critical time constant τ_{crit}, we opened the gate for plasticity in only a subpopulation (one-fourth of the neurons) in the network and compared it to opening plasticity in the entire network.

We found that spatially confined plasticity had a much lower impact on the critical time constant than global plasticity. Both an increase in learning rate and an increase in excitability lowered the critical time constant substantially (Fig. 4 c,d black). Opening these gates in only one-fourth of the neurons lowered the time constant significantly less than opening the gates in the entire network. Therefore, spatially confined neuromodulation has advantages for network stability beyond enabling precise control.

2.5 Protection of memories by gating plasticity.

In addition to the impact on the stability of network activity, plasticity interferes with the stability of memories. In a plastic network, neural activity patterns lead to synaptic changes, which could overwrite memories that were previously stored in the synaptic connections. Especially when resources are limited, forming new memories can come at the expense

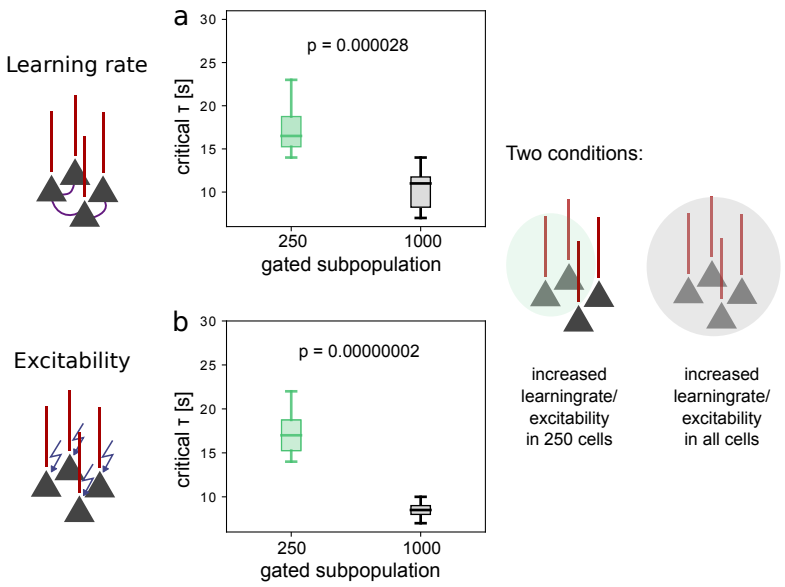

Figure 4: **Spatially precise neuromodulation enables learning while keeping network activity stable.** a: Distribution of critical homeostatic time constants for a two-fold increase in learning rate in a subpopulation of excitatory cells (green) and in the entire network (black). b: Distribution of critical homeostatic time constants for a 15% increase in excitability in a subpopulation of excitatory cells (green) and in the entire network (black). p-values were obtained by using the two-sample student's t-test.

of old ones. By gating plasticity, memories can in principle be protected. For example, trivially, by switching off plasticity, old memories are protected. Interestingly, the gates we considered here are not simple all-or-nothing gates but can be continuously modulated. We, therefore, investigated how their modulation affects the maintenance and encoding of memories.

Before gating plasticity, we first tested how a memory interferes with a previously stored memory in a plastic network. We added inhibitory plasticity to ensure network stability during memory formation and heterosynaptic depression to introduce synaptic competition (see Methods). First, we showed pattern P1 to the network (Fig. 5a). Afterwards, neurons activated by P1 formed a neural ensemble E1 by increasing their connectivity (Fig. 5a top left). Then, we showed a pattern P2 to the network. P2 is similar to P1 and hence activated a group of neurons E2 that overlapped with the previously formed ensemble. Neurons activated by P2 increased their connectivity. Because the patterns overlapped, synapses were increased at the expense of connections from the old memory (Fig. 5a bottom). Therefore, the new memory interfered with the old memory (Fig. 5 top right). We defined the difference between the mean connection strength of the P1 neurons after memory formation of P1 and the mean connection strength of the P1 neurons at the end of the simulation - after P2 has been learned - as the *breakdown* of the memory.

To test the effect of gating on the protection of memories, we applied the different gates after the first memory is formed. Reducing learning rate to 0 after the first memory is formed trivially protects the memory from being overwritten by the second pattern (Fig. 5b), as this blocked further weight changes. It is less clear how a change in excitability or inhibition affect the storage of the memory. Unlike the learning rate, we cannot modulate excitability or inhibition to their extremes without silencing neural activity. On the one hand, because they decrease neural firing rates, these gates could protect memories by reducing weight changes. On the other hand, by decreasing firing rates, they could also increase LTD as experimentally, low firing rates promote more LTD than LTP (Dudek and Bear, 1992; Sjöström et al., 2001).

We found that a reduction in excitability could indeed protect the memory (Fig. 5c) without permanently silencing the network (Fig. A.1c). With a lower excitability, the stimulated neural ensemble, E2, fired at a lower rate (compare E2 in Fig. A.1g and Fig. A.1e), and weights within E2, including those projecting to the overlapping ensemble, O, potentiated less (compare Fig. 5c bottom with Fig. 5a bottom). The new ensemble, E2, hence competed to a lesser extent with the old memory, leading to less memory decay due to heterosynaptic depression.

Similar to reduced excitability, increased inhibition could also protect the memory (Fig. 5c) as it reduced firing rates in the network (Fig. A.1d,h). Notably, the inhibitory plasticity in the network additionally protected the previously formed memory, as it led to increased inhibition of the old memory ensemble E1 (Fig. A.1i). This further reduced potentiation of synapses from the new ensemble, E2, to the overlapping ensemble, O.

Because decreased excitability and increased inhibition lower the firing rate of the network, we asked whether this low firing rate induces LTD and hence counteracts the protection of the old memory ensemble. We found that there was no increased LTD within the old memory ensemble due to low firing rates. First, the non-overlapping population of the old ensemble fired at a very low rate (E1-O in Fig. A.1g,h). There was hence little depression from the non-overlapping population to the overlapping one, as depression happens upon presynaptic spiking. Second, the memory breakdown was weaker at lower excitability and higher inhibition, i.e. at lower firing rates (Fig. 5f,g). The old memory was hence mostly at risk due to heterosynaptic depression. In line with this, the memory breakdown correlated with the maximum mean strength of synaptic connections to the overlapping ensemble O during pattern P2 (Fig. A.1j-l).

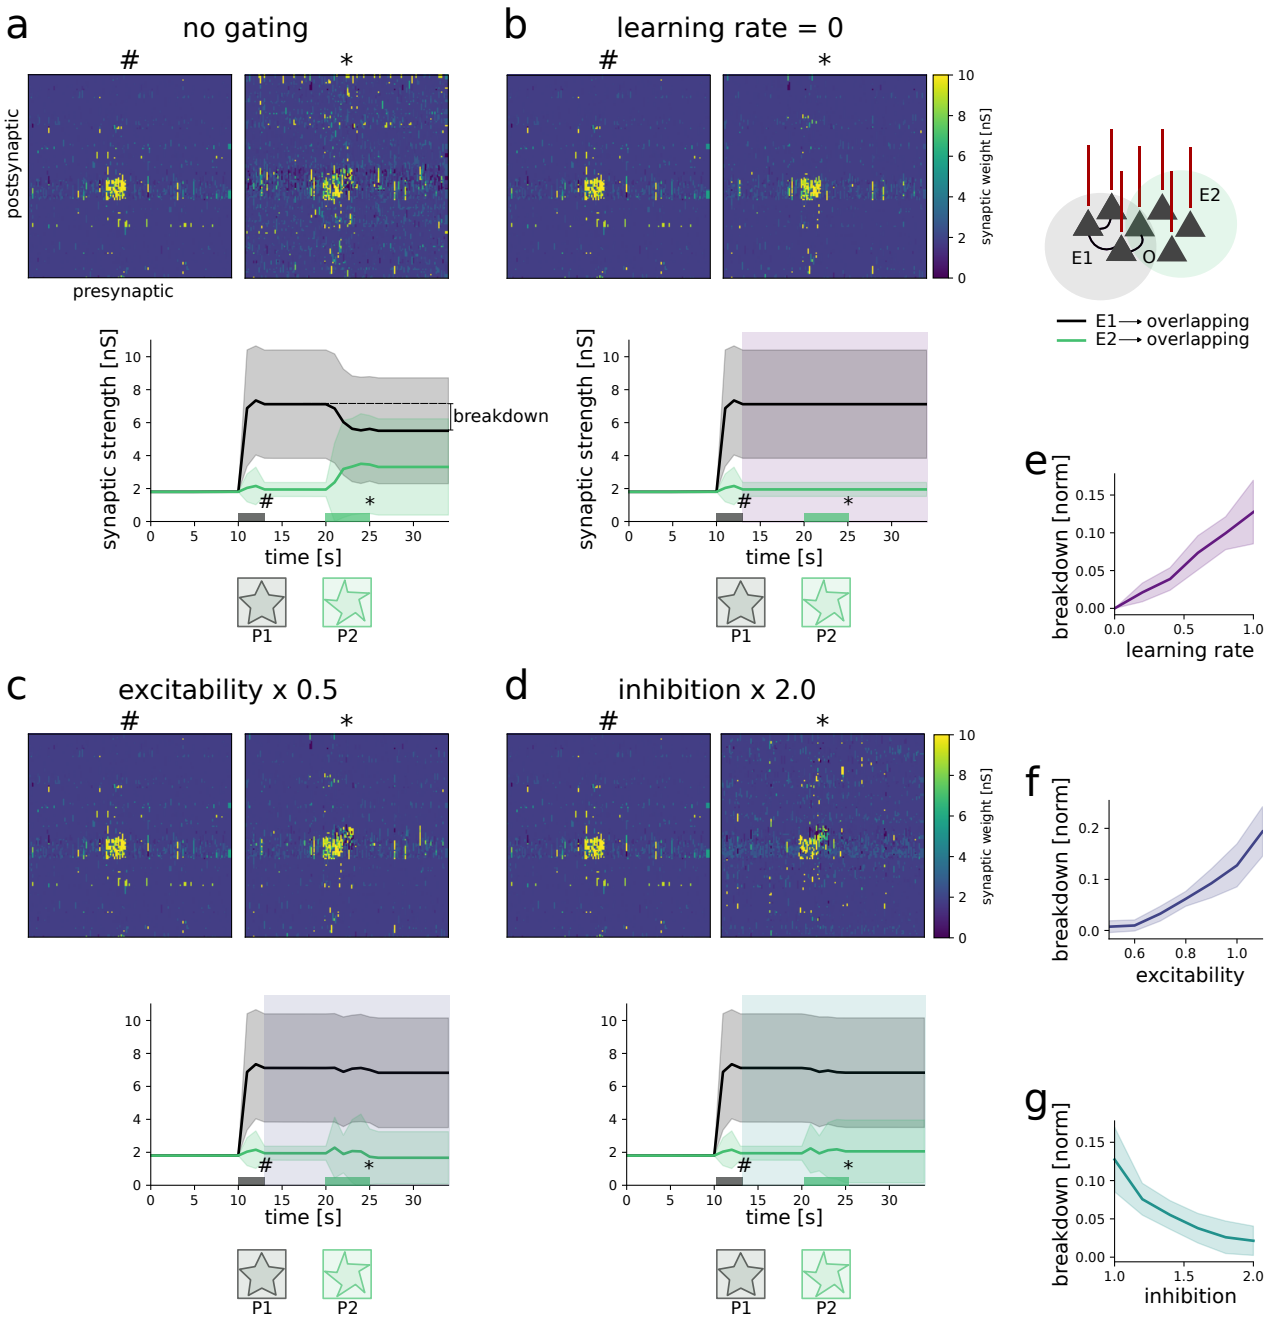

Figure 5: **Protection of memories by gating plasticity.** a-d: After 10 seconds, we show pattern 1 (P1, grey) to the network for 3 seconds. After a gap of 7 seconds, we show pattern 2 (P2, green) to the network for 5 seconds. Top panels: excitatory weight matrix at two time points of the simulation. Bottom panel: mean synaptic weight from ensemble 1 (E1) to the overlapping region (grey) and from ensemble 2 (E2) to the overlapping region over time as an indication of the strength of the memory of P1 and P2, respectively. a: no gating. b: after P1 is learned, the learning rate is set to 0 (denoted by blue flag). c: after P2 is learned, excitability is reduced by 50%. d: after P2 is learned, inhibition is doubled. e-g: breakdown of the memory (see a) as a function of learning rate (e), excitability (f), and inhibition (g).

When we continuously modulated the gates, we found that the breakdown of the memory increased nonlinearly with both increasing excitability (Fig. 5f) and decreasing inhibition (Fig. 5g), and increased linearly with learning rate (Fig. 5e).

To conclude, all gates can protect memories. Learning rate can be modulated independent of network activity and hence act as a switch for plasticity. Although excitability and inhibition do not modulate plasticity separately, they can both protect memories by reducing activity and weight changes without silencing network activity.

3 Discussion

We explored the effect of different gating mechanisms on the stability of neural networks. Using a balanced spiking neural network with 2-compartment pyramidal cells, we showed how gating of plasticity loosens the requirements for fast homeostatic processes. We found that excitability, learning rate, and inhibition affect the critical time constant in different ways. Interestingly, the network was more tolerant towards weight changes, when plasticity gates were opened in the dendrite versus the soma. Plasticity in dendrites thereby could facilitate learning without compromising the stability of the network. . We also showed that spatially precise neuromodulation lifts the critical time constant and thereby could locally enable learning while keeping network activity stable. Finally, gating plasticity can protect memories from being overwritten in addition to keeping the network activity stable.

Examples for the different modulators

Plasticity is highly modulated and gated (Bear and Singer, 1986; Pedrosa and Clopath, 2017; Couey et al., 2007; Bissière et al., 2003; Lin et al., 2003; Pawlak et al., 2010; Seol et al., 2007; Zhang et al., 2009; Brzosko et al., 2015). In this paper, we explored different such modulations of plasticity. First, inhibitory cell types, which target soma and dendrites of excitatory cells, can modulate plasticity. It has been shown that disinhibition - the inhibition of inhibitory cells - promotes learning (Letzkus et al., 2011; Kuhlman et al., 2013; Clopath et al., 2016). Dendritic inhibition can influence plasticity directly by affecting depolarizing events in the dendrite, such as back-propagating action potentials and calcium spikes (Larkum et al., 1999; Wilmes et al., 2016). Somatic inhibition can modulate plasticity indirectly by decreasing the firing rate of the neuron, as synaptic weight changes depend on neural activity. The modulation of plasticity via inhibition can be both 1) fast, because interneurons can be switched on and off quickly, and 2) local, because they can be precisely targeted by cholinergic (Woody and Gruen, 1987; Metherate et al., 1992; Xiang et al., 1998; Froemke et al., 2007, 2013) and noradrenergic neuromodulation (Kuo and Trussell, 2011; Martins and Froemke, 2015).

In addition, neuromodulators influence plasticity by regulating neural excitability (acetylcholine and noradrenaline (Xiang et al., 1998; Rasmusson, 2000; Joshi et al., 2016)). For example, acetylcholine binds to muscarinic receptors, which activate a cascade that leads to a decreased permeability of potassium channels (Rasmusson, 2000). This prolongs the duration of EPSPs and thereby increases excitability.

A similar form of neuromodulation is achieved by presynaptic inhibition. A recent theoretical study showed that presynaptic inhibition can act as a fast modulator of plasticity to stabilize network activity (Naumann and Sprekeler, 2020). They showed that presynaptic inhibition is an attractive control mechanism as it depends on network activity and therefore provides a gain control loop. Similar to excitability in our model, the analysis in Naumann and Sprekeler (2020) shows a supralinear relationship between presynaptic inhibition strength and the critical homeostatic time constant (Fig. 2).

Finally, because many forms of plasticity are NMDAR-dependent (Dudek and Bear, 1992), a modulation of NMDA channels could affect plasticity directly. NMDA channel permeability can be modulated by D-serine, the origin of which is debated (Wolosker et al., 2016), although it was initially thought to be synthesised by astrocytes (Wolosker et al., 2016; Henneberger et al., 2010). Such a direct modulation of plasticity would correspond to modulation of learning rate in our model. A different form of learning rate modulation could be achieved by dendritic inhibition, which is precisely timed to not affect the integration of EPSPs from the dendrite to the soma (Wilmes et al., 2016).

Hence, plasticity is highly gated and modulated. Depending on the form of neuromodulation, the effect on plasticity can be precisely timed and spatially confined.

Predictions from the model

We showed that larger synaptic changes are tolerated in dendrites than in the soma for the same critical time constant (Fig. 3). Therefore, our model predicts that more weight changes should be seen in dendrites. The weaker the dendrite and the soma are coupled, the larger becomes the advantage of the separate dendritic compartment. Hence, we predict that neurons with electrotonically more separate dendrites can undergo more dendritic plasticity.

If plasticity is gated in space and time, i.e. synaptic changes are only locally permitted in limited periods of time, then we would observe that the total amount of synaptic change is not constant, but varies in time and space. The amount of synaptic change averaged over longer periods of time may be constant. When taking averages over shorter periods, we predict that the amount of synaptic change varies significantly over time.

Our model shows that the gates differ in their impact on the critical homeostatic time constant (Fig. 2e). We found that, for inhibition, the critical time constant decreases sublinearly as a function of the resulting increased dendritic weight change. For excitability and learning rate, however, the critical time constant decreases supralinearly. Our model, therefore, predicts that gating plasticity with inhibition allows for a larger critical time constant than gating plasticity with excitability or learning rate. We predict that when inhibition and excitability are separately modulated in an experiment, that the network will lose stability earlier with a change in excitability than with a change in inhibition.

We found that the gates also differ in their ability to protect memories (Fig. 5e-g). Learning rate is the only gate, which can completely switch off plasticity to protect the memory. The memory breakdown increased supralinearly with a change in inhibition or excitability, whereas it increased linearly with a change in learning rate. Our model hence predicts that memories break down earlier when inhibition or excitability are modulated than with modulation of learning rate.

Limitations

Our model provides a comparison between different gating mechanisms. We simulated a balanced spiking network undergoing spontaneous activity to allow for the comparison of the different plasticity gates. A network, which is externally stimulated may have additional requirements for the homeostatic time constant.

In our model, we used one form of homeostatic plasticity, which adjusts LTD based on the postsynaptic firing rate. There are, however, different forms of homeostatic plasticity such as inhibitory plasticity (Woodin and Maffei, 2010) and synaptic scaling (Turrigiano et al., 1998). Inhibitory plasticity also requires fast homeostatic mechanisms when plasticity is not gated (Zenke et al., 2017; Litwin-Kumar et al., 2016). We expect that the gates studied here will similarly lift the requirements for the time scale of inhibitory plasticity. It will be interesting to explore the effects of the inhibitory gate on a homeostatic mechanism, which depends on inhibitory plasticity.

Conclusion

In summary, our study using balanced spiking neural networks with 2-compartment pyramidal cells shows how different forms of gating plasticity increase the stability of neural networks in the presence of plasticity. Our results suggest an important role for dendrites as they can undergo synaptic plasticity with minor effects on network stability. Our results also imply that gating via neuromodulation should be locally restricted, supporting the recent finding that neuromodulation may be more specific than initially thought (Totah et al., 2018). Finally, diverse gating mechanisms can protect memories, even if they affect network activity, such as excitability or inhibition.

4 Materials and Methods

Balanced network

We built a recurrent neural network model with N_E =1000 excitatory (E) and N_I =250 inhibitory (I) cells. Both E and I cells received excitatory inputs from a pool of 1000 Poisson processes with a firing rate of 2 Hz and with a connection probability of p=10%. All neurons were randomly connected. Excitatory cells receive excitatory and inhibitory synapses on both their soma and their dendrite. The connection probability is 10% for all connections except from excitatory cells to excitatory cell's somata. The connection probability for those connections is 9% to account for the fact that the cells also receive inputs on their dendrites in the two-compartment model. The connection strength of the synapses is chosen such that the network is balanced (see Table 4).

2-compartment pyramidal cell model

For the excitatory population, we used a 2-compartment integrate and fire pyramidal cell model with spike-triggered adaptation, adapted from the model by Naud and Sprekeler (2018) which was originally fitted to data from layer 5 pyramidal cells. It has two coupled membrane equations, one for the soma (V_s, Eq. 1), one for the dendrite (V_d, Eq. 2), modelled as (for clarity we repeat the equations from the main text):

$$C_s\frac{dV_s}{dt} = -g_l(V_s - E_l) - \gamma g_E V_s - g_I(V_s - E_I) + \lambda(g_s\frac{1}{1+\exp(-\frac{V_d-E_d}{D_d})} + \omega^s) \tag{1}$$

$$C_d\frac{dV_d}{dt} = -(V_d - E_l) - \gamma^d g_E^d V_d - g_I^d(V_d - E_I) + g_d\frac{1}{1+\exp(-\frac{V_d-E_d}{D_d})} + c_d K(t) + \omega^d \tag{2}$$

where $C_{s/d}$ is the somatic/dendritic capacitance, $g_{I/E}$ is the inhibitory/excitatory conductance (index d indicates the dendritic variables), E_l is the reversal potential, γ is excitability, $\omega^{d/s}$ is the somatic/dendritic adaptation variable. When the soma spikes, the dendrite receives a back-propagating action potential after a delay of 0.5 ms, which is modelled as a 2 ms long current pulse (defined by rectangular kernel $K(t)$) with amplitude $c_d = 2600pA$. With $\hat{t}_s$ as the time of the last somatic spike, $K(t)$ is defined as

$$K(t) = \begin{cases} 1 & \text{if } \hat{t}_s + 0.5\text{ms} \leq t \leq \hat{t} + 2.5\text{ms} \\ 0 & \text{otherwise.} \end{cases}$$

The dendrite has a nonlinear (sigmoidal) term $\frac{1}{1+\exp(-\frac{V_d-E_d}{dd})}$ corresponding to the activation of dendritic calcium channels. E_d determines the voltage at which the threshold will be reached and Dd determines the slope of the nonlinear function. The nonlinear dynamics are controlled locally by g_d and are also transmitted to the soma with a coupling factor g_s, such that the soma bursts. The factor λ ensures that the somato-dendritic coupling and adaptation are the same as in the model of Naud and Sprekeler (2018), where the somatic capacitance was 370 pF (we used C_s=200pF). The somatic adaptation variable is modelled as

$$\frac{d\omega^s}{dt} = -\omega^s/\tau_\omega^s + b_\omega^s S_i(t) \tag{7}$$

where b_ω^s is the strength of spike-triggered adaptation and τ_ω^s is the recovery time scale. The dendritic adaptation variable is written as

$$\tau_\omega^d \frac{d\omega^d}{dt} = -\omega^d + a_\omega^d (V_d - E_l) \tag{8}$$

where a_ω^d is the strength of subthreshold adaptation and τ_ω^d is the recovery time scale.

For the inhibitory population, we used a single-compartment leaky integrate-and-fire neuron model, which membrane potential V evolves according to:

$$C_s \frac{dV}{dt} = -g_l(V - E_l) - g_E V \cdot -g_I(V - E_I) \tag{9}$$

For all neurons, excitatory and inhibitory conductances, g_E and g_I respectively, are increased by the synaptic weight w_{iE}/w_{iI}, depending on their type i upon a spike event in a presynaptic excitatory or inhibitory neuron with spike train $S_j(t)$, and decay exponentially with time constants τ_E and τ_I, respectively:

$$\frac{dg_E}{dt} = -\frac{g_E}{\tau_E} + \sum_k w_{ij} S_j(t) \tag{10}$$

$$\frac{dg_I}{dt} = -\frac{g_I}{\tau_I} + \sum_k w_{ij} S_j(t). \tag{11}$$

Both excitatory and inhibitory neurons had a refractory period of 8.3 ms (chosen according to the network model from Zenke et al. (2013). Initial membrane potentials for V_s and V were sampled from a Gaussian distribution with μ = -70mV and σ = 10mV to prevent that all neurons spike at the same time at the beginning of a simulation. V_d was set to -70mV initially.

Plasticity

Synapses from neuron j targeting the somatic compartment of neuron i change their synaptic weight w_{ij} according to the triplet rule (Pfister and Gerstner, 2006). For clarity we repeat the same equation as in the main text:

$$\begin{aligned}\frac{dw_{ij}}{dt} =& \eta w_0 A^+ z_j^+(t) z_i^{slow}(t-\epsilon) S_i(t) \\ & - \eta w_0 A_i^-(t) z_i^-(t) S_j(t)\end{aligned} \tag{3}$$

where w_0 is the initial weight, $A^{+/-}$ is the amplitude of potentiation/depression (the depression one is time dependent, see below), $S_{i/j}$ is the post-/presynaptic spike train, z_j^+ is the presynaptic trace, z_i^- is the postsynaptic trace, z_i^{slow} is the postsynaptic trace with a slower time constant. ϵ denotes a small fraction of time such that $t - \epsilon$ indicates that the value of the trace is taken before the time point of the action potential t. Parameters were chosen as in (Zenke et al., 2013).

Synapses from neuron j to neuron i targeting the dendritic compartment change their synaptic weight w_{ij}^d according to the same triplet rule with the back-propagating action potential (bAP) as the postsynaptic spike and an additional Ca-spike-dependent potentiation at the time of a presynaptic spike.

$$\begin{aligned}\frac{dw_{ij}^d}{dt} =& \eta^d w_0 A^+ z_j^+(t) z_i^{slow}(t-\epsilon) S_i^{\mathrm{bAP}}(t) \\ & - (\eta^d w_0 (A_i^-(t) z_i^-(t) + A^{Ca}(v_d > \theta_{Ca})) - \alpha) S_j(t)\end{aligned} \tag{5}$$

Here, the timing of the back-propagating action potential in the dendrite is used to update the post-synaptic traces z_i^- and z_i^{slow} and $S_i^{\mathrm{bAP}}(t)$ is the postsynaptic train of back-propagating action potentials. A back-propagating action potential is detected if three conditions are met: (1) the dendritic membrane potential v_d exceeds a threshold of -50 mV, and (2) there was a somatic spike within the last 3 ms, and (3) there was no backpropagating action potential within the last 5.8 ms (to account for the refractory period). Synapses are potentiated by a constant amount A^{Ca} when the presynaptic cell fires and the postsynaptic dendritic membrane potential v_d exceeds a threshold θ_{Ca} of -40mV. The term $v_d > \theta_{Ca}$ takes a value of 1 when the threshold is crossed and is 0 otherwise. Synapses are depressed by a constant amount α for each presynaptic spike (transmitter-induced plasticity).

The pre- and postsynaptic traces are defined as:

$$\frac{dz_j^+}{dt} = -\frac{z_j^+}{\tau^+} + S_j(t) \tag{12}$$

$$\frac{dz_i^-}{dt} = -\frac{z_i^-}{\tau^-} + S_i(t) \tag{13}$$

$$\frac{dz_i^{\mathrm{slow}}}{dt} = -\frac{z_i^{\mathrm{slow}}}{\tau^{\mathrm{slow}}} + S_i(t) \tag{14}$$

where τ^+, τ^-, and τ^{slow} are the time constants with which the traces decay. Both somatic and dendritic excitatory synapses are limited by a maximum synaptic weight $w_{\max}$ = 10 nS.

Homeostatic Plasticity

The depression amplitude A_i^- for all synapses onto neuron i is a function of a moving average of neuron i's activity $\bar{s}_i$:

$$A_i^-(t) = \frac{A^+\tau^+\tau^{\text{slow}}}{\tau^-\kappa}\bar{s}_i^2 \tag{15}$$

where κ is the target firing rate, A^+, τ^+, τ^- and τ^{slow} are variables from the triplet STDP rule and $\bar{s}_i$ is the low-pass filtered spike train:

$$\tau\frac{d\bar{s}_i}{dt} = -\bar{s}_i + S_i(t) \tag{16}$$

with τ defining the time constant of the homeostatic plasticity.

Explosion factor

We quantify the stability with the explosion factor EF. We calculate it as follows:

$$EF = \frac{r_{max}}{r_{baseline}} \tag{17}$$

where r_{max} is the maximum population firing rate within the duration of the simulation and $r_{baseline}$ is the population firing rate averaged over the first 50 ms of the simulation. Therefore, an explosion factor close to 1 indicates that the network activity is stable. The distribution of explosion factors was bimodal with a sharp peak close to 1 and a broader distribution of larger EFs (Fig. 2h). We defined a threshold separating those two modes, which defines whether the network is stable or explodes:

$$\text{network is } \begin{cases} \text{stable} & \text{if } EF \leq 1.5 \\ \text{unstable} & \text{if } EF > 1.5. \end{cases}$$

Critical time constant τ_{crit}

For each value of the gate, we calculated the maximum τ for which the network was stable. It additionally had to be smaller than the minimum τ for which the network was unstable.

Baseline dendritic weight change

For each gating value, we calculated the sum of all weight changes in the dendrite in a 200 s simulation with a τ of 5 ms.

Statistical analyses

To test for significance in Fig. 3 and Fig. 4, we used the two-sample student's t-test.

Simulation

For Fig. 1, we simulated the network for 10s without plasticity. Simulations to calculate the explosion factor (for Figs. 2-4) were run for 200 seconds. We simulated an initial warm-up phase for 3 τ seconds without plasticity to calculate the average population firing rate for the balanced network. We used the average population firing rate of the last 2 seconds of the warm-up phase to set the target firing rate κ in our model. We then switch on plasticity. All simulations were run at a timestep of 0.1 ms. For the plots in Figs. 1-3, each condition was simulated with 10 different seeds.

Memory network (Fig. 5)

We used the described network and added plasticity on inhibitory to excitatory synapses and a competition mechanism for postsynaptic weights.

Inhibitory Plasticity Synapses from inhibitory neuron j to excitatory neuron i change their weight w_{ij}^I according to Vogels et al. (2011)

$$\frac{dw_{ij}^I}{dt} = \eta^I z_j^I(t)S_i(t) \tag{18}$$

$$+ (\eta^I z_i^I(t) - \alpha^I)S_j(t) \tag{19}$$

where η^I is the inhibitory learning rate, $z_j^I(t)/z_i^I(t)$ is the pre/post-synaptic trace, $S_j(t)/S_i(t)$ is the pre/post-synaptic spike train, and $\alpha^I = 2\kappa\tau_{iSTDP}$ determines the amount of transmitter-induced depression Vogels et al. (2013). Plastic inhibitory weights are limited by a maximum weight $w_{\text{max}}^I = 100nS$. The pre/post-synaptic traces $z_j^I(t)/z_i^I(t)$ are written as

$$\frac{dz_j^I}{dt} = -\frac{z_j^I}{\tau_{\text{iSTDP}}} + S_j(t) \tag{20}$$

$$\frac{dz_i^I}{dt} = -\frac{z_i^I}{\tau_{\text{iSTDP}}} + S_i(t). \tag{21}$$

$$\tag{22}$$

Heterosynaptic depression If the sum of postsynaptic weights in the soma or the dendrite exceeds a maximum $1.5pN_E w_0$ (hard bound), all somatic/dendritic synaptic weights are scaled down equally by the average synaptic weight change of the postsynaptic compartment in the current time step, that is the total somatic/dendritic weight change divided by the number of incoming somatic/dendritic synapses.

Stimulation protocol We first simulated the non-plastic network for 5 s to calculate the steady state population firing rate. We then introduced plasticity with a target firing rate κ of the measured population firing rate and simulated the network with plasticity for another 5 s. Then, we activated pattern P1 for 3 s by injecting current into neurons with indices 400 to 499 (ensemble E1). Afterwards, we changed the gating variable under investigation (except in Fig. 5a, where we did not apply any gating). After a stimulation pause of 7 s, we activated pattern P2 for 5 s by injecting current into neurons with indices 450 to 549 (ensemble E2). We continued the simulation for further 10 s without stimulation.

Parameters

Table 1: **Parameters of the network.**

Parameter	Value
N_E	1000
N_I	250
N_{Poisson}	1000
λ_{Poisson}	2 Hz
p	0.1
w_{EP}	1.6 nS
w_{IP}	1.6 nS
w_{EE}	1.8 nS
w^d_{EE}	1.8 nS
w_{IE}	4.0 nS
w_{II}	6.0 nS
w_{EI}	8.0 nS
w^d_{EI}	4.0 nS

Table 2: **Parameters of the neuron model.**

Parameter	Value
g_l	10.0 nS
E_l(leak)	-70 mV
E_I(Inhibitory)	-80 mV
v_{theta}	-50.0 mV
v^I_{theta}	-50 mV
τ_m	20 ms
τ^I_m	10 ms
C_s	200 pF
C_d	170 pF
c_d	2600 pA
a^d_ω	-13 nS
b^s_ω	-200 pA
τ^s_ω	100 ms
τ^d_ω	30 ms
τ^s	16 ms
τ^d	7 ms
g_s	1300 pA
g_d	1200 pA
E_d	-38 mV
D_d	6 mV
λ	0.54

Table 3: **Parameters of the plasticity.**

Parameter	Value
A^+	6.5e-3
A^{Ca}	7.2e-2
α	1e-4
θ_{bAP}	-50 mV
θ_{Ca}	-40 mV
τ^+	16.8 ms
τ^-	33.7 ms
τ^{slow}	114 ms
$w_{\max}$	10 nS
τ_{iSTDP}	20 ms
$w^I_{\max}$	100 nS

References

L. F. Abbott and Sacha B. Nelson. Synaptic plasticity: taming the beast. *Nature Neuroscience*, 3(11):1178–1183, 2000. doi: 10.1038/81453. URL `https://doi.org/10.1038/81453`.

Julien Artinian and Jean-Claude Lacaille. Disinhibition in learning and memory circuits: New vistas for somatostatin interneurons and long-term synaptic plasticity. *Brain Research Bulletin*, 141:20 – 26, 2018. ISSN 0361-9230. doi: https://doi.org/10.1016/j.brainresbull.2017.11.012. URL `http://www.sciencedirect.com/science/article/pii/S0361923017305130`. Memory Mechanisms in Health and Disease.

Mark F. Bear and Wolf Singer. Modulation of visual cortical plasticity by acetylcholine and noradrenaline. *Nature*, 320: 172, Mar 1986. URL `https://doi.org/10.1038/320172a0`.

EL Bienenstock, LN Cooper, and PW Munro. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, 2(1):32–48, 1982. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.02-01-00032.1982. URL `http://www.jneurosci.org/content/2/1/32`.

Stephanie Bissière, Yann Humeau, and Andreas Lüthi. Dopamine gates ltp induction in lateral amygdala by suppressing feedforward inhibition. *Nature Neuroscience*, 6(6):587–592, 2003. doi: 10.1038/nn1058. URL `https://doi.org/10.1038/nn1058`.

Jacopo Bono, Katharina A Wilmes, and Claudia Clopath. Modelling plasticity in dendrites: from single cells to networks. *Current Opinion in Neurobiology*, 46:136 – 141, 2017. ISSN 0959-4388. doi: https://doi.org/10.1016/j.conb.2017.08.013. URL `http://www.sciencedirect.com/science/article/pii/S0959438816302665`.

Zuzanna Brzosko, Wolfram Schultz, and Ole Paulsen. Retroactive modulation of spike timing-dependent plasticity by dopamine. *eLife*, 4:e09685, oct 2015. ISSN 2050-084X. doi: 10.7554/eLife.09685.

Jerry L Chen, Walter C Lin, Jae Won Cha, Peter T So, Yoshiyuki Kubota, and Elly Nedivi. Structural basis for the role of inhibition in facilitating adult brain plasticity. *Nature Neuroscience*, 14(5):587–594, 2011. doi: 10.1038/nn.2799. URL `https://doi.org/10.1038/nn.2799`.

Claudia Clopath, Lars Büsing, Eleni Vasilaki, and Wulfram Gerstner. Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nature Neuroscience*, 13:344, Jan 2010. URL `https://doi.org/10.1038/nn.2479`. Article.

Claudia Clopath, Tim P. Vogels, Robert C. Froemke, and Henning Sprekeler. Receptive field formation by interacting excitatory and inhibitory synaptic plasticity. *bioRxiv*, 2016. doi: 10.1101/066589. URL `https://www.biorxiv.org/content/early/2016/07/29/066589`.

Jonathan J. Couey, Rhiannon M. Meredith, Sabine Spijker, Rogier B. Poorthuis, August B. Smit, Arjen B. Brussaard, and Huibert D. Mansvelder. Distributed network actions by nicotine increase the threshold for spike-timing-dependent plasticity in prefrontal cortex. *Neuron*, 54(1):73 – 87, 2007. ISSN 0896-6273. doi: https://doi.org/10.1016/j.neuron.2007.03.006. URL `http://www.sciencedirect.com/science/article/pii/S0896627307001845`.

Niraj S. Desai, Robert H. Cudmore, Sacha B. Nelson, and Gina G. Turrigiano. Critical periods for experience-dependent synaptic scaling in visual cortex. *Nature Neuroscience*, 5(8):783–789, 2002. doi: 10.1038/nn878. URL `https://doi.org/10.1038/nn878`.

S M Dudek and M F Bear. Homosynaptic long-term depression in area ca1 of hippocampus and effects of n-methyl-d-aspartate receptor blockade. *Proceedings of the National Academy of Sciences of the United States of America*, 89(10): 4363–4367, 05 1992. doi: 10.1073/pnas.89.10.4363. URL `https://pubmed.ncbi.nlm.nih.gov/1350090`.

Robert C. Froemke, Michael M. Merzenich, and Christoph E. Schreiner. A synaptic memory trace for cortical receptive field plasticity. *Nature*, 450(7168):425–429, 2007. doi: 10.1038/nature06289. URL https://doi.org/10.1038/nature06289.

Robert C Froemke, Ioana Carcea, Alison J Barker, Kexin Yuan, Bryan A Seybold, Ana Raquel O Martins, Natalya Zaika, Hannah Bernstein, Megan Wachs, Philip A Levis, Daniel B Polley, Michael M Merzenich, and Christoph E Schreiner. Long-term modification of cortical synapses improves sensory perception. *Nature Neuroscience*, 16(1):79–88, 2013. doi: 10.1038/nn.3274. URL https://doi.org/10.1038/nn.3274.

Anubhuthi Goel and Hey-Kyoung Lee. Persistence of experience-induced homeostatic synaptic plasticity through adulthood in superficial layers of mouse visual cortex. *Journal of Neuroscience*, 27(25):6692–6700, 2007. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.5038-06.2007. URL https://www.jneurosci.org/content/27/25/6692.

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, 2020/08/19 2015. doi: 10.1016/j.neuron.2015.09.025. URL https://doi.org/10.1016/j.neuron.2015.09.025.

Christian Henneberger, Thomas Papouin, Stéphane H. R. Oliet, and Dmitri A. Rusakov. Long-term potentiation depends on release of d-serine from astrocytes. *Nature*, 463(7278):232–236, 2010. doi: 10.1038/nature08673. URL https://doi.org/10.1038/nature08673.

Ankur Joshi, Bopanna I. Kalappa, Charles T. Anderson, and Thanos Tzounopoulos. Cell-specific cholinergic modulation of excitability of layer 5b principal neurons in mouse auditory cortex. *Journal of Neuroscience*, 36(32):8487–8499, 2016. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.0780-16.2016. URL https://www.jneurosci.org/content/36/32/8487.

Megumi Kaneko, Jessica L Hanover, Pamela M England, and Michael P Stryker. Trkb kinase is required for recovery, but not loss, of cortical responses following monocular deprivation. *Nature Neuroscience*, 11(4):497–504, 2008a. doi: 10.1038/nn2068. URL https://doi.org/10.1038/nn2068.

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, 2020/08/19 2008b. doi: 10.1016/j.neuron.2008.04.023. URL https://doi.org/10.1016/j.neuron.2008.04.023.

Tara Keck, Volker Scheuss, R. Irene Jacobsen, Corette J. Wierenga, Ulf T. Eysel, Tobias Bonhoeffer, and Mark Hübener. Loss of sensory input causes rapid structural changes of inhibitory neurons in adult mouse visual cortex. *Neuron*, 71(5): 869–882, 2020/08/19 2011. doi: 10.1016/j.neuron.2011.06.034. URL https://doi.org/10.1016/j.neuron.2011.06.034.

Tara Keck, Georg B. Keller, R. Irene Jacobsen, Ulf T. Eysel, Tobias Bonhoeffer, and Mark Hübener. Synaptic scaling and homeostatic plasticity in the mouse visual cortex in vivo. *Neuron*, 80(2):327–334, 2020/08/19 2013. doi: 10.1016/j.neuron.2013.08.018. URL https://doi.org/10.1016/j.neuron.2013.08.018.

Tara Keck, Taro Toyoizumi, Lu Chen, Brent Doiron, Daniel E. Feldman, Kevin Fox, Wulfram Gerstner, Philip G. Haydon, Mark Hübener, Hey-Kyoung Lee, John E. Lisman, Tobias Rose, Frank Sengpiel, David Stellwagen, Michael P. Stryker, Gina G. Turrigiano, and Mark C. van Rossum. Integrating hebbian and homeostatic plasticity: the current state of the field and future research directions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1715): 20160158, 2017. doi: 10.1098/rstb.2016.0158. URL https://royalsocietypublishing.org/doi/abs/10.1098/rstb.2016.0158.

Sandra J. Kuhlman, Nicholas D. Olivas, Elaine Tring, Taruna Ikrar, Xiangmin Xu, and Joshua T. Trachtenberg. A disinhibitory microcircuit initiates critical-period plasticity in the visual cortex. *Nature*, 501(7468):543–546, 2013. doi: 10.1038/nature12485. URL https://doi.org/10.1038/nature12485.

Sidney P Kuo and Laurence O Trussell. Spontaneous spiking and synaptic depression underlie noradrenergic control of feed-forward inhibition. *Neuron*, 71(2):306–318, 07 2011. doi: 10.1016/j.neuron.2011.05.039. URL https://pubmed.ncbi.nlm.nih.gov/21791289.

M E Larkum, K M Kaiser, and B Sakmann. Calcium electrogenesis in distal apical dendrites of layer 5 pyramidal cells at a critical frequency of back-propagating action potentials. *Proceedings of the National Academy of Sciences of the United States of America*, 96(25):14600–4, December 1999. ISSN 0027-8424. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=24482&tool=pmcentrez&rendertype=abstract.

Johannes J. Letzkus, Steffen B. E. Wolff, Elisabeth M. M. Meyer, Philip Tovote, Julien Courtin, Cyril Herry, and Andreas Lüthi. A disinhibitory microcircuit for associative fear learning in the auditory cortex. *Nature*, 480(7377):331–335, December 2011. ISSN 0028–0836. doi: 10.1038/nature10674. URL http://www.nature.com/doifinder/10.1038/nature10674.

Lu Li, Melanie A. Gainey, Joseph E. Goldbeck, and Daniel E. Feldman. Rapid homeostasis by disinhibition during whisker map plasticity. *Proceedings of the National Academy of Sciences*, 111(4):1616–1621, 2014. ISSN 0027-8424. doi: 10.1073/pnas.1312455111. URL https://www.pnas.org/content/111/4/1616.

Yi-Wen Lin, Ming-Yuan Min, Tsai-Hsien Chiu, and Hsiu-Wen Yang. Enhancement of associative long-term potentiation by activation of -adrenergic receptors at ca1 synapses in rat hippocampal slices. *Journal of Neuroscience*, 23(10): 4173–4181, 2003. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.23-10-04173.2003. URL https://www.jneurosci.org/content/23/10/4173.

Ashok Litwin-Kumar, Robert Rosenbaum, and Brent Doiron. Inhibitory stabilization and visual coding in cortical circuits with multiple interneuron subtypes. *Journal of Neurophysiology*, 115(3):1399–1409, 2016. doi: 10.1152/jn.00732.2015. URL https://doi.org/10.1152/jn.00732.2015.

Arianna Maffei and Gina G. Turrigiano. Multiple modes of network homeostasis in visual cortical layer 2/3. *Journal of Neuroscience*, 28(17):4377–4384, 2008. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.5298-07.2008. URL https://www.jneurosci.org/content/28/17/4377.

Ana Raquel O Martins and Robert C Froemke. Coordinated forms of noradrenergic plasticity in the locus coeruleus and primary auditory cortex. *Nature neuroscience*, 18(10):1483–1492, 10 2015. doi: 10.1038/nn.4090. URL https://pubmed.ncbi.nlm.nih.gov/26301326.

R Metherate, CL Cox, and JH Ashe. Cellular bases of neocortical activation: modulation of neural oscillations by the nucleus basalis and endogenous acetylcholine. *Journal of Neuroscience*, 12(12):4701–4711, 1992. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.12-12-04701.1992. URL https://www.jneurosci.org/content/12/12/4701.

Kenneth D. Miller and David J. C. MacKay. The role of constraints in hebbian learning. *Neural Computation*, 6(1): 100–126, 1994. doi: 10.1162/neco.1994.6.1.100.

Richard Naud and Henning Sprekeler. Sparse bursts optimize information transmission in a multiplexed neural code. *Proceedings of the National Academy of Sciences*, 115(27):E6329–E6338, 2018. ISSN 0027-8424. doi: 10.1073/pnas.1720995115. URL https://www.pnas.org/content/115/27/E6329.

Laura Bella Naumann and Henning Sprekeler. Presynaptic inhibition rapidly stabilises recurrent excitation in the face of plasticity. *bioRxiv*, 2020. doi: 10.1101/2020.02.11.944082. URL https://www.biorxiv.org/content/early/2020/02/12/2020.02.11.944082.

Verena Pawlak, Jeffery Wickens, Alfredo Kirkwood, and Jason Kerr. Timing is not everything: Neuromodulation opens the stdp gate. *Frontiers in Synaptic Neuroscience*, 2:146, 2010. ISSN 1663-3563. doi: 10.3389/fnsyn.2010.00146. URL https://www.frontiersin.org/article/10.3389/fnsyn.2010.00146.

Victor Pedrosa and Claudia Clopath. The role of neuromodulators in cortical plasticity. a computational perspective. *Frontiers in Synaptic Neuroscience*, 8:38, 2017. ISSN 1663-3563. doi: 10.3389/fnsyn.2016.00038. URL https://www.frontiersin.org/article/10.3389/fnsyn.2016.00038.

Jean-Pascal Pfister and Wulfram Gerstner. Triplets of spikes in a model of spike timing-dependent plasticity. *Journal of Neuroscience*, 26(38):9673–9682, 2006. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.1425-06.2006. URL https://www.jneurosci.org/content/26/38/9673.

D. D Rasmusson. The role of acetylcholine in cortical synaptic plasticity. *Behavioural Brain Research*, 115(2):205–218, 2000. doi: https://doi.org/10.1016/S0166-4328(00)00259-X. URL http://www.sciencedirect.com/science/article/pii/S016643280000259X.

Lindsay A. Schwarz, Kazunari Miyamichi, Xiaojing J. Gao, Kevin T. Beier, Brandon Weissbourd, Katherine E. DeLoach, Jing Ren, Sandy Ibanes, Robert C. Malenka, Eric J. Kremer, and Liqun Luo. Viral-genetic tracing of the input–output organization of a central noradrenaline circuit. *Nature*, 524(7563):88–92, 2015. doi: 10.1038/nature14600. URL https://doi.org/10.1038/nature14600.

Geun Hee Seol, Jokubas Ziburkus, ShiYong Huang, Lihua Song, In Tae Kim, Kogo Takamiya, Richard L. Huganir, Hey-Kyoung Lee, and Alfredo Kirkwood. Neuromodulators control the polarity of spike-timing-dependent synaptic plasticity. *Neuron*, 55(6):919 – 929, 2007. ISSN 0896-6273. doi: https://doi.org/10.1016/j.neuron.2007.08.013. URL http://www.sciencedirect.com/science/article/pii/S0896627307006265.

Per Jesper Sjöström, Gina G Turrigiano, and Sacha B Nelson. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32(6):1149–1164, 2001. ISSN 0896-6273. doi: https://doi.org/10.1016/S0896-6273(01)00542-6. URL http://www.sciencedirect.com/science/article/pii/S0896627301005426.

Nelson K. Totah, Ricardo M. Neves, Stefano Panzeri, Nikos K. Logothetis, and Oxana Eschenko. The locus coeruleus is a complex and differentiated neuromodulatory system. *Neuron*, 99(5):1055 – 1068.e6, 2018. ISSN 0896-6273. doi: https://doi.org/10.1016/j.neuron.2018.07.037. URL http://www.sciencedirect.com/science/article/pii/S0896627318306354.

Taro Toyoizumi, Megumi Kaneko, Michael P. Stryker, and Kenneth D. Miller. Modeling the dynamic interaction of hebbian and homeostatic plasticity. *Neuron*, 84(2):497 – 510, 2014. ISSN 0896-6273. doi: https://doi.org/10.1016/j.neuron.2014.09.036. URL http://www.sciencedirect.com/science/article/pii/S0896627314008940.

Misha Tsodyks, William E Skaggs, Terrence J Sejnowski, and Bruce L Mcnaughton. Paradoxical effects of inhibitory interneurons. *J. Neurosci.*, 17(11):4382–4388, 1997. ISSN 0270-6474.

Gina G. Turrigiano. The dialectic of hebb and homeostasis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1715):20160258, 2017. doi: 10.1098/rstb.2016.0258. URL https://royalsocietypublishing.org/doi/abs/10.1098/rstb.2016.0258.

Gina G. Turrigiano and Sacha B. Nelson. Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience*, 5(2):97–107, 2004. doi: 10.1038/nrn1327. URL https://doi.org/10.1038/nrn1327.

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, 1998. doi: 10.1038/36103. URL https://doi.org/10.1038/36103.

Silvana Valtcheva and Laurent Venance. Astrocytes gate hebbian synaptic plasticity in the striatum. *Nature Communications*, 7:13845–13845, 12 2016. doi: 10.1038/ncomms13845. URL https://pubmed.ncbi.nlm.nih.gov/27996006.

M. C. W. van Rossum, G. Q. Bi, and G. G. Turrigiano. Stable hebbian learning from spike timing-dependent plasticity. *Journal of Neuroscience*, 20(23):8812–8821, 2000. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.20-23-08812.2000. URL https://www.jneurosci.org/content/20/23/8812.

Daniëlle van Versendaal and Christiaan N. Levelt. Inhibitory interneurons in visual cortical plasticity. *Cellular and Molecular Life Sciences*, 73(19):3677–3691, 2016. doi: 10.1007/s00018-016-2264-4. URL https://doi.org/10.1007/s00018-016-2264-4.

T. P. Vogels, H. Sprekeler, F. Zenke, C. Clopath, and W. Gerstner. Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science*, 334(6062):1569–1573, 2011. ISSN 0036-8075. doi: 10.1126/science.1211095. URL https://science.sciencemag.org/content/334/6062/1569.

Tim Vogels, Robert Froemke, Nicolas Doyon, Matthieu Gilson, Julie Haas, Robert Liu, Arianna Maffei, Paul Miller, Corette Wierenga, Melanie Woodin, Friedemann Zenke, and Henning Sprekeler. Inhibitory synaptic plasticity: spike timing-dependence and putative network function. *Frontiers in Neural Circuits*, 7:119, 2013. ISSN 1662-5110. doi: 10.3389/fncir.2013.00119. URL https://www.frontiersin.org/article/10.3389/fncir.2013.00119.

Alanna Watt and Niraj Desai. Homeostatic plasticity and stdp: keeping a neuron's cool in a fluctuating world. *Frontiers in Synaptic Neuroscience*, 2:5, 2010. ISSN 1663-3563. doi: 10.3389/fnsyn.2010.00005. URL https://www.frontiersin.org/article/10.3389/fnsyn.2010.00005.

Katharina Anna Wilmes, Henning Sprekeler, and Susanne Schreiber. Inhibition as a binary switch for excitatory plasticity in pyramidal neurons. *PLoS Computational Biology*, 12(3):e1004768, 2016.

Herman Wolosker, Darrick T Balu, and Joseph T Coyle. The rise and fall of the d-serine-mediated gliotransmission hypothesis. *Trends in neurosciences*, 39(11):712–721, 11 2016. doi: 10.1016/j.tins.2016.09.007. URL https://pubmed.ncbi.nlm.nih.gov/27742076.

M.A. Woodin and A. Maffei, editors. *Inhibitory Synaptic Plasticity*. Springer, 2010.

CD Woody and E Gruen. Acetylcholine reduces net outward currents measured in vivo with single electrode voltage clamp techniques in neurons of the motor cortex of cats. *Brain research*, 424(1):193—198, October 1987. ISSN 0006-8993. doi: 10.1016/0006-8993(87)91210-8. URL https://doi.org/10.1016/0006-8993(87)91210-8.

Zixiu Xiang, John R. Huguenard, and David A. Prince. Cholinergic switching within neocortical inhibitory networks. *Science*, 281(5379):985–988, 1998. ISSN 0036-8075. doi: 10.1126/science.281.5379.985. URL https://science.sciencemag.org/content/281/5379/985.

Friedemann Zenke, Guillaume Hennequin, and Wulfram Gerstner. Synaptic plasticity in neural networks needs homeostasis with a fast rate detector. *PLOS Computational Biology*, 9(11):1–14, 11 2013. doi: 10.1371/journal.pcbi.1003330. URL https://doi.org/10.1371/journal.pcbi.1003330.

Friedemann Zenke, Wulfram Gerstner, and Surya Ganguli. The temporal paradox of hebbian learning and homeostatic plasticity. *Current Opinion in Neurobiology*, 43:166 – 176, 2017. ISSN 0959-4388. doi: https://doi.org/10.1016/j.conb.2017.03.015. URL http://www.sciencedirect.com/science/article/pii/S0959438817300910. Neurobiology of Learning and Plasticity.

Ji-Chuan Zhang, Pak-Ming Lau, and Guo-Qiang Bi. Gain in sensitivity and loss in temporal contrast of stdp by dopaminergic modulation at hippocampal synapses. *Proceedings of the National Academy of Sciences*, 106(31):13028–13033, 2009. ISSN 0027-8424. doi: 10.1073/pnas.0900546106. URL https://www.pnas.org/content/106/31/13028.

Acknowledgments

Funding

This work was supported by DFG 398005926, BBSRC BB/N013956/1, BB/N019008/1, Wellcome Trust 200790/Z/16/Z, Simons Foundation 564408, EPSRC EP/R035806/1.

A Figures

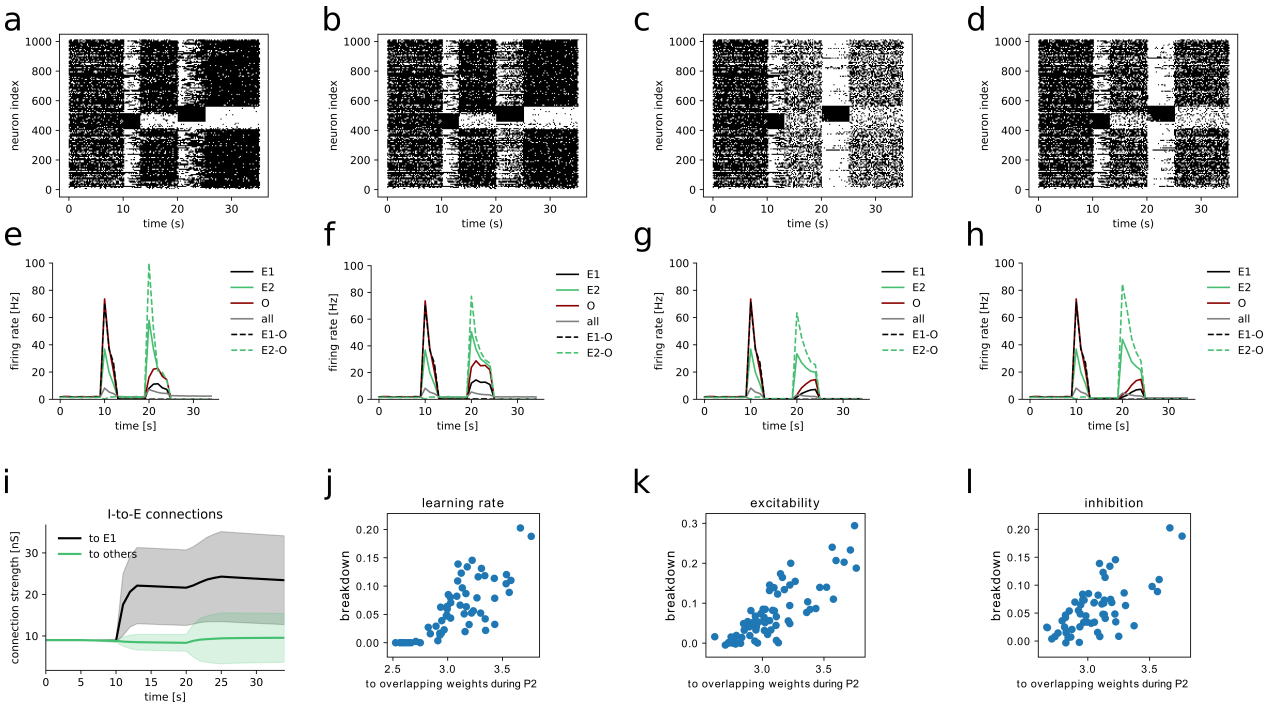

Figure A.1: a-d: Raster plots for the simulation without gating (a), with gated learning rate (b), with gated excitability (c) and with gated inhibition (d) as in Fig. 5. e-h: Firing rates over time of different subpopulations: E1: memory ensemble E1 ($x \in E1$), E2: memory ensemble E2 ($x \in E1$), O: overlapping ensemble O ($x \in (E1 \cap E2)$, all: all excitatory neurons in the network, E1-O: ensemble E1 excluding the overlapping ensemble O ($x \in (E1 \setminus O)$), E2-O: ensemble 2 excluding the overlapping ensemble ($x \in (E2 \setminus O)$). i: Connection strengths over time from the inhibitory population to ensemble E1 (black) and from the inhibitory population to all other cells in the ungated network. j-l: Memory breakdown as a function of the maximum during pattern P2 of the mean strength of synaptic connections from all excitatory cells to the overlapping ensemble O for a network with gated learning rate (j), gated excitability (k) and gated inhibition (l).