Breaking Balance: Encoding local error signals in perturbations of excitation-inhibition balance

Julian Rossbroich[1,2] and **Friedemann Zenke**[1,2,*]

[1]Friedrich Miescher Institute for Biomedical Research, Basel, Switzerland
[2]Faculty of Science, University of Basel, Switzerland

[*]Corresponding author: friedemann.zenke@fmi.ch

Abstract: Efficient learning algorithms such as backpropagation, predictive coding, and self-supervised learning require instructive learning signals to influence synaptic plasticity locally at specific neurons. This requirement is inconsistent with classic Hebbian theories and neuromodulation involving a non-specific third factor, thus raising the question of how neurobiology can achieve efficient learning. Here, we propose a simple and biologically plausible solution: local deviations from precise excitation/inhibition (E/I) balance encode error signals that instruct synaptic plasticity. Using a computational model derived from an adaptive control theory framework, we demonstrate that breaking E/I balance through targeted feedback to inhibitory interneurons, can produce neuron- or assembly-specific error signals that enable learning in multi-layer networks. Simulations reveal that such a balance-controlled plasticity mechanism is consistent with phenomenological local plasticity models while enabling online learning in hierarchically organized networks. Furthermore, we demonstrate that this framework is consistent with key features of disinhibitory microcircuit dynamics during in-vivo learning experiments. These results suggest that the brain may exploit E/I balance not only for stability, but also as a substrate for error-driven learning.

Keywords: Synaptic plasticity · Excitation-inhibition balance · Inhibitory interneurons · Neuronal microcircuits · Credit assignment · Adaptive control

Introduction

A central question in neuroscience is how the brain learns complex tasks across distributed networks of interacting brain areas. Machine learning algorithms that solve complex tasks require local error signals—neuron-specific cues that determine the sign of synaptic changes. Similarly, bio-plausible credit assignment algorithms, which determine which synapses to modify, and how, to improve performance on behavioral tasks, explicitly require local error signals that control the sign of plasticity [1, 2]. Moreover, neuron-specific errors, or teaching signals, are also required in self-supervised learning [3, 4] and predictive coding models [5–11]. While local error signals are fundamental to learning in hierarchical networks, the biological mechanisms for computing local errors in the brain remain elusive [2, 12, 13].

Existing models assume that neurons separate error signals from ongoing activity by utilizing segregated dendritic compartments (Fig. 1; [6–9, 11, 14–16]) or by alternating between distinct temporal phases [3, 7, 9, 14, 17–19]. However, direct experimental evidence for electrotonic segregation required for such compartmentalized schemes is limited to certain neuron types, such as cortical Layer 5 pyramidal cells [16], and, in some cases, inconsistent with phenomenological plasticity models [20–27], which raises questions about their generality.

In contrast, E/I balance is a pervasive phenomenon widely observed in cortical circuits, whereby neurons maintain a tight balance between excitatory and inhibitory currents across stimuli and time [28–34]. While balance is considered essential to support network stability, efficient coding, and precise spike timing [35–37], its overall function remains elusive. Importantly, inhibition has been shown to influence synaptic plasticity [38–45] and learning outcomes [46–54], suggesting a potential role for E/I balance in learning beyond mere stabilization [55–57]. Nevertheless, existing work typically considered inhibition, or the lack thereof, as a permissive signal for synaptic plasticity.

This article demonstrates how E/I balance affords neuron-specific error encoding, allowing it to act as an instructive signal that could support powerful learning algorithms in the brain for solving complex

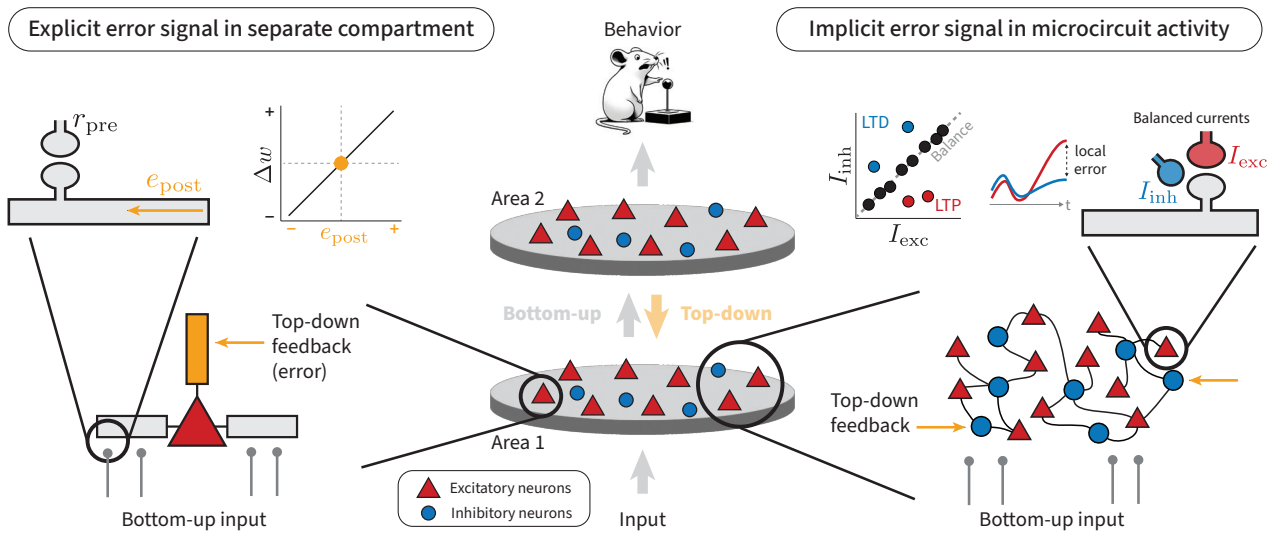

Figure 1. Efficient learning algorithms require local error signals that determine the sign of synaptic plasticity. How synaptic plasticity is coordinated across neuronal circuits across the brain to tune individual synapses for specific computational tasks (middle) remains a fundamental question in neuroscience. Efficient learning algorithms require local neuron-specific error signals that dictate the sign of synaptic weight changes. Most theoretically motivated models of biologically plausible credit assignment and predictive coding assume that such error signals are explicitly represented in separate temporal phases of neuronal activity or distinct neuronal compartments (left). However, such compartmentalization is often at odds with neurophysiology and inconsistent with phenomenological plasticity rules. Here we develop an alternative model that sidesteps these issues by encoding error signals in deviations from E/I balance in local E/I assemblies (right).

tasks. Specifically, we introduce a theoretical framework in which local error signals are encoded implicitly at the microcircuit-level as controlled deviations from E/I balance (Fig. 1). We show how top-down feedback pathways with neuron-specific targeting can give rise to such systematic deviations in E/I balance and, combined with a suitable credit assignment strategy, orchestrate learning in multi-layer networks. Importantly, encoding local error signals in E/I balance constitutes a versatile circuit mechanism that reconciles powerful learning algorithms with established Hebbian principles, without the need for electronically separate dendritic compartments or distinct learning phases. By integrating error encoding with the notion of E/I balance, our model connects theories of biologically plausible credit assignment, predictive coding, and self-supervised learning with experimentally observed plasticity mechanisms, providing a fresh view into how the brain may orchestrate synaptic plasticity across large distributed neural networks to solve complex tasks.

Results

To study whether deviations from precise E/I balance [35] can encode target-specific error signals, we consider an E/I assembly as the fundamental computational unit (Fig. 2a; Methods). We further distinguish between two types of distinct inhibitory interneuron populations (Supplementary Fig. S1). Type 1 interneurons provide feedforward inhibition and contribute to neuronal selectivity. In contrast, Type 2 interneurons are reciprocally connected and co-tuned with excitatory neurons resulting in a precise E/I balance (Fig. 2b). Signatures of such co-tuning exist in neurobiology [28, 30, 31, 33, 58] and numerous computational studies demonstrated how precise balance can emerge through suitable Hebbian plasticity mechanisms [25, 57, 59–62].

We model excitatory neurons as leaky integrators with membrane potential u_i and rectified linear output firing rate $r_i = (u_i)_+$ of neuron i. For slowly changing input currents, the membrane potential is approximately

$$u_i(t) \approx R\left(I_i^{\text{exc}}(t) - I_i^{\text{inh1}}(t) - I_i^{\text{inh2}}(t)\right) \quad , \tag{1}$$

where R is the input resistance. We speak of neuron-specific "precise balance" [35] when the net feedforward currents are proportionally matched by recurrent Type 2 inhibition:

$$\alpha\left(\bar{I}_i^{\text{exc}}(t) - \bar{I}_i^{\text{inh1}}(t)\right)_+ = I_i^{\text{inh2}}(t) \tag{2}$$

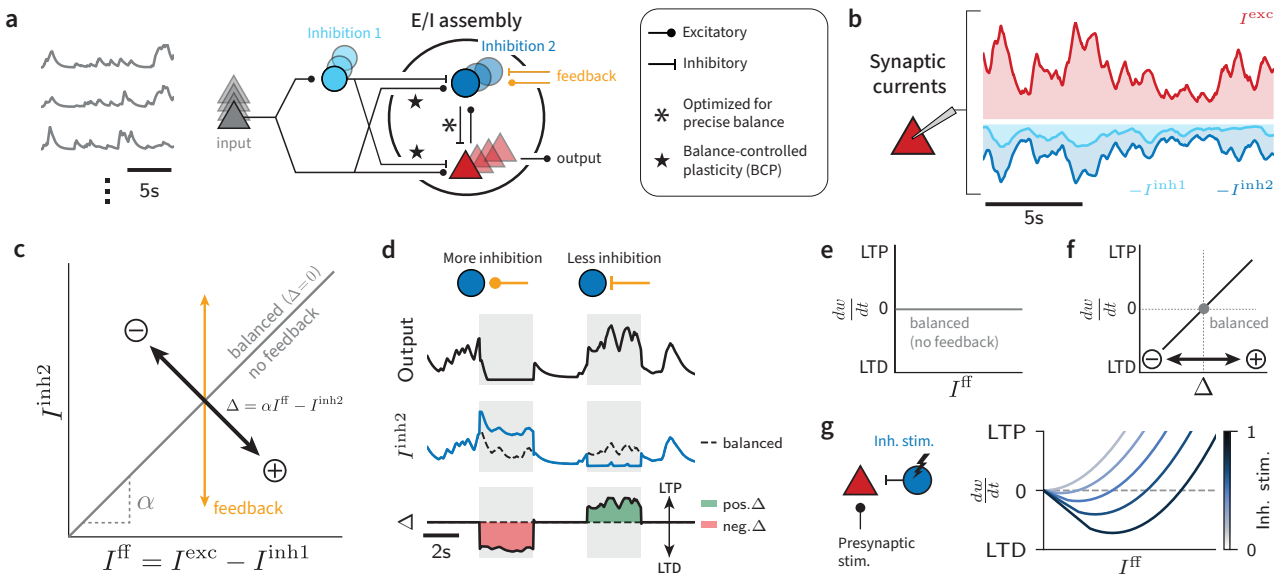

Figure 2. Deviations from E/I balance can encode assembly-specific error signals. (a) Schematic representation of an E/I assembly receiving excitatory input and feed-forward inhibition (Type 1). All input synapses are subject to balance-controlled plasticity (BCP). Within the assembly, excitatory and Type 2 inhibitory neurons are reciprocally connected, maintaining a precise E/I balance. Top-down feedback projections (orange) targeting Type 2 interneurons modulate the E/I balance. **(b)** Example synaptic currents into one excitatory neuron showing excitatory, Type 1 inhibitory and Type 2 inhibitory currents. **(c)** Illustration of E/I assembly dynamics and the effect of top-down feedback. In the absence of feedback, Type 2 inhibitory currents are proportional to the total feed-forward currents $I^{ff} = I^{exc} - I^{inh1}$ due to precise balance, allowing feed-forward currents to change neuronal firing rates while ensuring that the balance error $\Delta = 0$ at all times. Top-down feedback into Type 2 neurons (orange) causes deviations from the balanced regime, leading to changes in firing rate and non-zero Δ. **(d)** Example time series of excitatory output firing rate (top), Type 2 inhibitory current (middle), and balance error Δ (bottom) for positive and negative feedback into Type 2 interneurons. Initially the neuronal input is precisely balanced resulting in $\Delta \approx 0$. Additional top-down input to Type 2 neurons suppresses firing while increasing inhibitory current resulting in a negative error encoding ($\Delta < 0$). Conversely, Type 2 mediated dis-inhibition results in a positive error encoding ($\Delta > 0$). **(e)** Illustration of the BCP rule in the absence of top-down feedback. Regardless of feed-forward input strength (I^{ff}), no plasticity is induced. **(f)** Schematic of the weight change $\frac{dw}{dt}$ as a function of Δ. Sign and magnitude of synaptic plasticity are directly proportional to Δ. **(g)** BCP in a simulated patch-clamp paradigm (left), in which recurrent Type 2 inhibition is not driven by excitatory activity but is under experimental control. Synaptic weight change $\frac{dw}{dt}$ as a function of presynpatic stimulation for different values of Type 2 inhibition (right). In this case, BCP resembles a Hebbian learning rule in which the amount of inhibitory current controls the plasticity threshold.

---

where $\alpha < 1$ is a positive parameter and $\bar{I}^x$ denotes a short-term moving average over the respective synaptic current (Methods). We further define the balance error Δ_i of neuron i as follows:

$$\Delta_i(t) \equiv \alpha \left(\bar{I}_i^{exc}(t) - \bar{I}_i^{inh1}(t) \right)_+ - I_i^{inh2}(t) \tag{3}$$

To test whether Δ could serve error encoding, we simulated an assembly of excitatory and inhibitory neurons in which the connectivity between Type 2 inhibitory neurons and excitatory neurons was optimized to ensure $\Delta_i \approx 0$ for any input (Fig. 2a,b; Methods). Moreover, Type 2 interneurons received top-down feedback, a common connectivity motif observed in neurobiology [63], causing changes in excitatory firing rates *and* deviations from precise balance reflected in non-zero Δ (Fig. 2c). Specifically, more Type 2 inhibition due to excitatory feedback suppressed excitatory firing rates and generated negative Δ (Fig. 2d) Conversely, Type 2 disinhibition caused increased excitatory firing rates and positive Δ. Thus Δ's dynamics are consistent with a bidirectional error signal driven by top-down feedback to inhibitory interneurons.

To see whether the local errors Δ_i could serve as an error signal for error-driven learning, we have to make assumptions about the underlying learning algorithm. We now focus on the general class of local-error-correcting learning algorithms which includes backpropagation of error (backprop) and predictive-coding based learning algorithms as specific limiting cases [5, 6, 9, 11]. In accordance with the circuit introduced above (cf. Fig 2a), the required feedback is communicated through top-down control inputs to

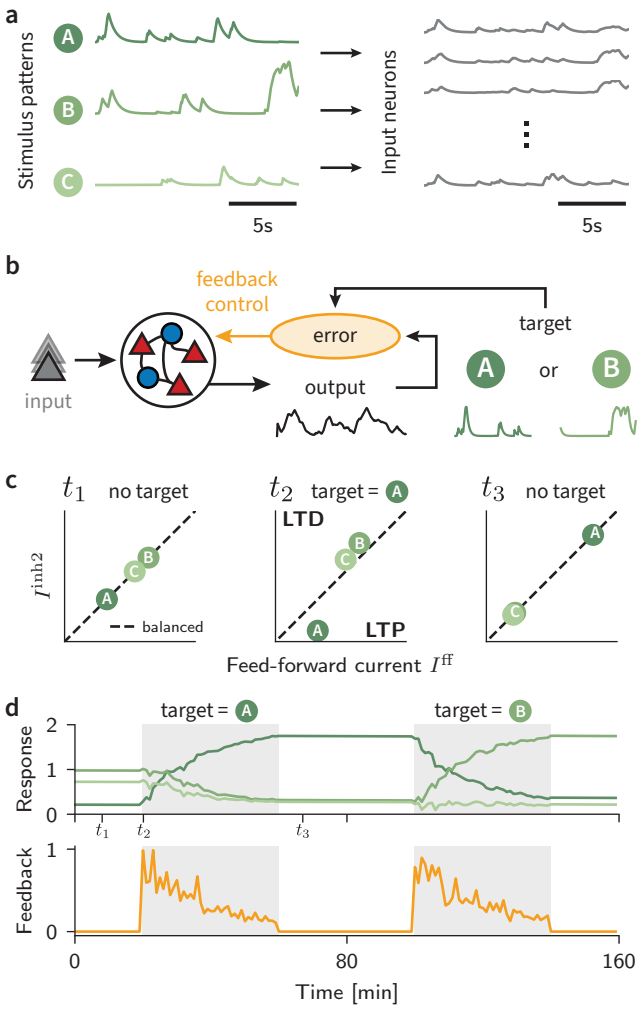

Figure 3. BCP supports plausible target-based learning through feedback control. **(a)** Schematic of the input paradigm. Three latent stimulus timeseries (A, B, C) are mixed linearly to yield 20 model inputs. **(b)** Sketch of the E/I assembly model. The assembly receives bipolar inputs through feedforward synapses subject to balance-controlled plasticity (BCP). The firing rate output of the excitatory assembly is compared to a given target signal. Deviations from the target result in a bidirectional error. Feedback control connections are set up such that their input to Type 2 interneurons reduces this error (Methods). **(c)** Feedforward input versus recurrent inhibition (Type 2) for the three different stimuli in isolation, evaluated at three timepoints of learning indicated in panel (d). Without a target signal (left), feedback control signals are zero and the inputs are precisely balanced (cf. Eq. (2)). When a target is given (middle), the balance is tipped in favor of LTP for the target pattern, whereas the two non-target patterns are suppressed. After learning (right), inputs are precisely balanced again, but the evoked feed-forward currents of each pattern changed due to synaptic plasticity. **(d)** Assembly response to the three target stimuli in the open-loop setting (top) and the strength of the feedback signal in the closed-loop setting (bottom) over time. Initially, there is no target and the response to all stimuli remains constant. When a target comes on (shaded regions), a feedback signal immediately corrects for any output error. Feedback signals decay slowly over time due to ongoing BCP which restructures the feedforward weights such that less feedback is required to produce the target output.

Type 2 interneurons. Given these constraints, the learning rule can be derived from an adaptive control theory framework (Methods). It is given by:

$$\frac{dw_j}{dt} = \eta \bar{x}_j(t)\Delta_i(t) \tag{4}$$

where w_j is the synaptic weight from presynaptic neuron j, η is a small positive learning rate and $\bar{x}_j$ is a presynaptic activity trace. Because Δ_i corresponds to the balance error introduced above, we refer to Equation (4) as the balance-controlled plasticity (BCP) learning rule. In contrast to Hebbian learning rules, plasticity changes in BCP are independent of the total feed-forward drive into the neuron or the corresponding excitatory firing rate (Fig. 2e) as long as synaptic currents are precisely balanced. Instead, negative $\Delta < 0$ results in long-term depression (LTD) while $\Delta > 0$ triggers long-term potentiation (LTP) independent of the magnitude of postsynaptic firing rate (Fig. 2f). For feed-forward Type 1 inhibitory synapses, BCP operates with reversed polarity: negative $\Delta < 0$ induces inhibitory LTP while positive $\Delta > 0$ leads to inhibitory LTD. Importantly, for an individual neuron under tonic Type 2 inhibition, the BCP rule resembles a classic Hebbian plasticity rule whose plasticity threshold depends on Type 2 inhibition [56, 64] (Fig. 2g). Thus, BCP reconciles existing phenomenological plasticity models with the notion of error-driven synaptic plasticity through assembly-specific error signals.

BCP supports plausible target-based learning through feedback control.

Next we sought to understand whether BCP supports target-based learning [2, 9, 11]. To address this question, we designed a learning task in which assembly neurons needed to become selective to distinct input stimuli according to top-down feedback. Assembly neurons received temporally dynamic inputs from a population of 20 input neurons collectively encoding three continuous latent input signals (Fig. 3a,b; Methods). The connectivity between Type 2 inhibitory and excitatory neurons in the assembly was optimized to ensure precise E/I balance (Methods). We assessed the circuit's selectivity in an open-loop setting (i.e. with-

out top-down feedback) by presenting each stimulus in isolation. Before learning, assembly neurons maintained precise balance and showed nonselective responses to all stimuli (Fig. 3c). To test whether feedback could bias the assembly toward one of the input stimuli, we added feedback control to Type 2 interneurons proportional to a supervised target selectivity (Fig. 3b; Methods). In this closed-loop setting, feedback induced stimulus-specific balance perturbations when the response of the assembly did not align with the target, causing the assembly to respond strongly to a specific input signal and weakly to the other two. For instance, when we set A as the target, feedback caused disinhibition during stimulus A which encoded a positive error. Conversely, increasing Type 2 inhibition during non-target stimuli B and C encoded negative balance errors (Fig. 3c). As a result, the assembly's selectivity and its associated open-loop response shifted toward A, while the required magnitude of the feedback control signal decreased monotonically (Fig. 3d). Importantly, this change in selectivity persisted in the absence of top-down feedback. To assess if these plasticity changes were reversible, we subsequently switched the output target to Stimulus B, causing a gradual loss of selectivity for Stimulus A and an acquired selective responses to B (Fig. 3d). These results show that top-down feedback can control plasticity through specific changes in E/I balance and enable target-based learning.

BCP enables online learning in multi-layer networks

We now wondered whether top-down control of BCP could serve as a possible circuit mechanism capable of controlling the sign of plasticity in multi-layer networks. Since many learning algorithms require top-down feedback and local error signals controlling plasticity—irrespective of the specific objective function they optimize and which credit assignment strategy they use—we focused on supervised learning with deep feedback control [9]. Unless mentioned otherwise, we compute optimal top-down feedback signals from the network Jacobian (Methods), which requires full knowledge of the forward weights. Although this credit assignment strategy is considered biologically implausible, this choice allowed us to focus on the mechanism for encoding local error signals without having to commit to any specific biologically plausible approximation of backprop. With these thoughts in mind, we constructed a network model in which 20 periodic input units project to a hidden layer consisting of 240 excitatory and 60 Type 2 inhibitory neurons organized into 20 distinct E/I assemblies (Methods). The assembly output projected to two output neurons whose activity we interpreted as X-Y coordinates. The learning task was to map the periodic inputs to a predefined X-Y target trajectory corresponding to the outline of a turtle (Fig. 4a,b). The network output did not resemble the target prior to training (Fig. 4b). As before, we initialized synaptic weights to ensure E/I balance in the absence of feedback (Fig. 4c). We let input-to-hidden synapses evolve according to the BCP learning rule (cf. Eq. (4)), while output synapses learned according to the standard delta rule ([65]; Methods). In the closed-loop setup, Type 2 inhibitory neurons in the hidden layer received assembly-specific feedback control signals proportional to the output error that conveyed credit information. The resulting E/I balance deviations Δ collectively improved the match between the network output and the target trajectory (Fig. 4b,c), but showed only a weak correlation with neuronal firing rates on a single-neuron level (Fig. 4d,e). Because Δ_i determines the sign of plasticity in a neuron-specific manner, this behavior is different from classic Hebbian learning rules in which the sign of plasticity is usually correlated with postsynaptic firing rate [20–27] or by a global third-factor as in three-factor learning rules [66, 67].

Next we switched on BCP and simulated the network for 360 minutes simulated time, corresponding to 4320 trajectory iterations, while periodically monitoring network performance in the open loop setting (Methods). BCP caused continuous synaptic weight changes in the hidden layer for which LTD and LTP were approximately balanced across synapses (Fig. 4f) and monotonically improved the match of the output to the target (Fig. 4g,h). At the end of the simulation, the network output trajectory closely matched the target (Fig. 4j) with an $R^2 = 0.996$.

To test whether top-down feedback was necessary to solve the task, we repeated the same simulation without feedback, thereby restricting plasticity to output synapses. The resulting network failed to solve the task ($R^2 = 0.428$; Fig. 4k). Similarly, when we removed the hidden layer entirely by directly mapping input nodes into the readout units, the network was unable to solve the task ($R^2 = 0.531$). Lastly, to check whether the specificity of feedback connections was critical, we implemented random feedback projections [1, 68], either targeting entire assemblies or individual Type 2 interneurons, and found that networks could still successfully solve the task (Supplementary Fig. S2). These results demonstrate that BCP in combination with E/I assemblies and targeted feedback signals derived from an adaptive control theory framework supports

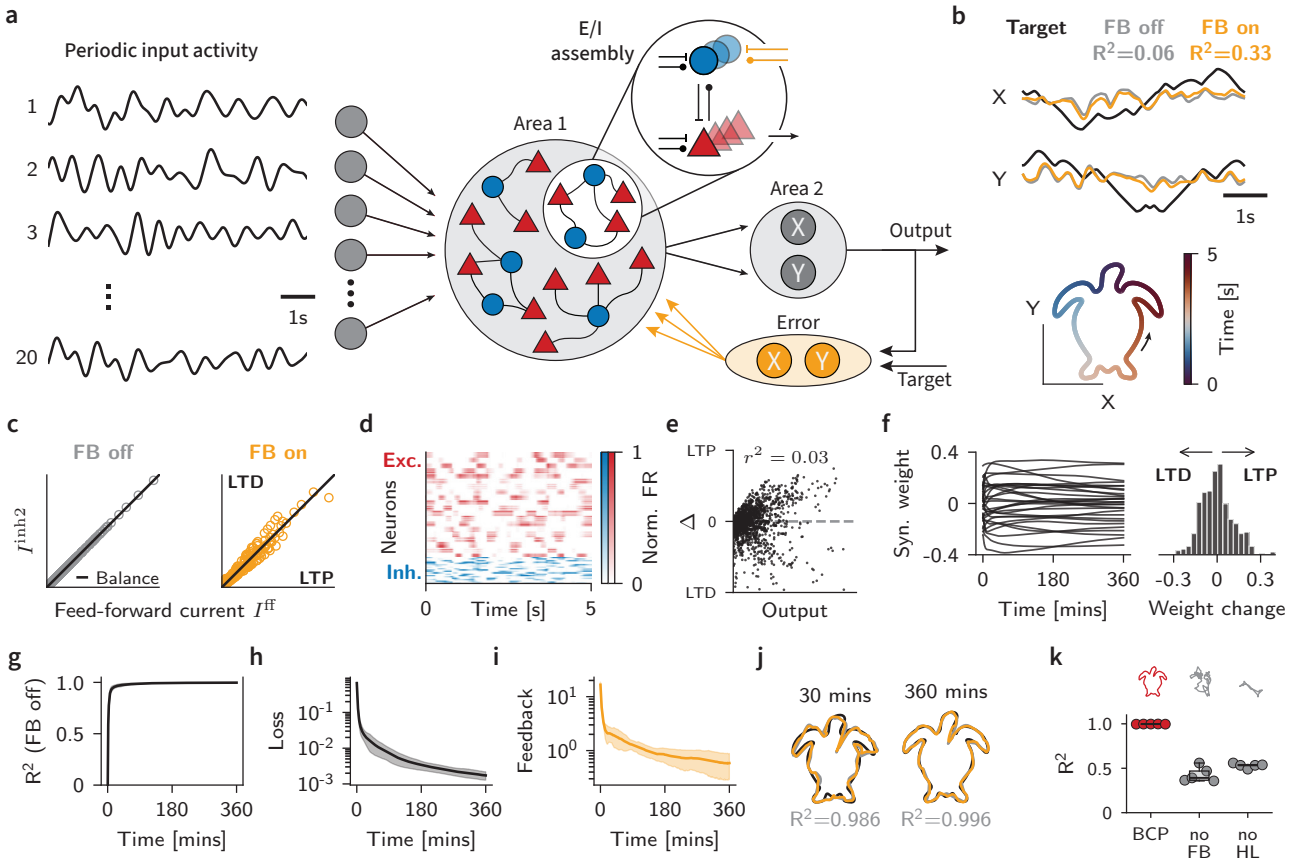

Figure 4. BCP enables online learning in multi-layer networks through feedback-guided plasticity. **(a)** Schematic of input stimulus and network architecture. 20 periodic inputs serve as inputs to a hidden layer organized into 20 E/I assemblies. Hidden layer neurons are connected to two output neurons encoding X-Y coordinates. Inhibitory neurons receive continuous feedback from a top-down controller proportional to the output error. **(b)** Target trajectory as a function of time (black) versus network output without (gray) and with top-down feedback (orange). Feedback nudges the output toward the target trajectory, resulting in increased R^2. **(c)** EI balance for each neuron without (left) and with feedback (right). Each point corresponds to the synaptic currents into one excitatory neuron in the hidden layer averaged over 1s. Feedback induces bidirectional deviations from E/I balance that encode error signals. **(d)** Heat map of hidden-neuron firing rates over one trajectory before learning. **(e)** Synaptic weight changes as a function of postsynaptic firing rate. Each dot represents the postsynaptic firing rate and the corresponding neuron-specific Δ averaged over 1s. **(f)** Synaptic weight evolution of representative input-to-hidden weights (left). Distribution of weight final changes (right). **(g)** R^2 during training. **(h)** Training loss over training (mean ± SD, n=5 networks). **(i)** Feedback magnitude as a function of time. **(j)** Output trajectory after 30 minutes and 6h respectively, demonstrating precise target matching. **(k)** Performance comparison (R^2) across three conditions: BCP with feedback, no feedback (No FB), and no hidden layer (No HL). Control conditions fail to learn the task, showing that assembly specific feedback is necessary (mean ± SD, n=5 networks).

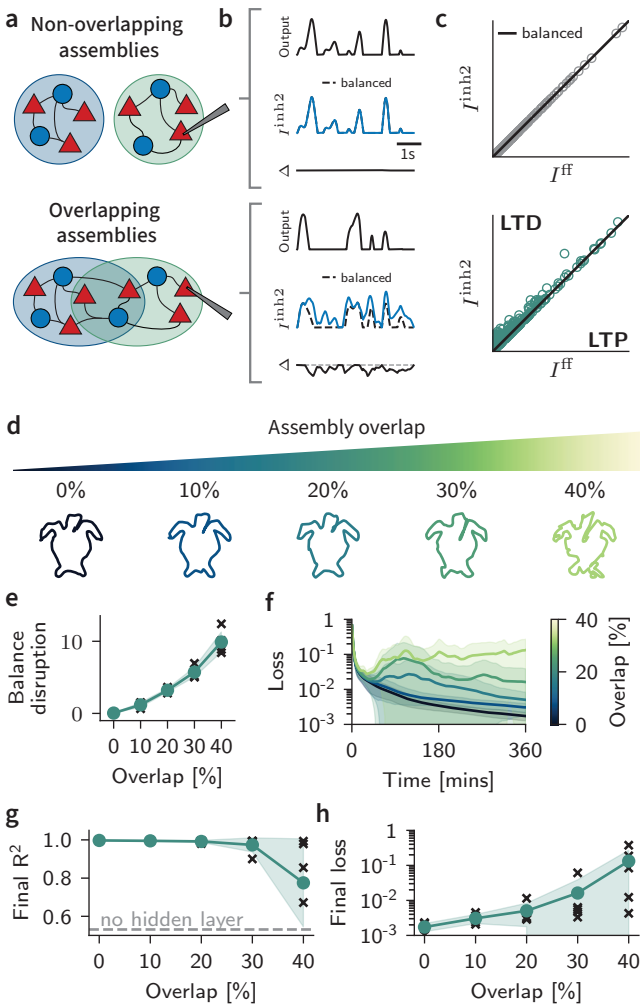

Figure 5. Learning is robust to moderate overlap in E/I assemblies. **(a)** Schematic illustration of non-overlapping (top) and overlapping (bottom) E/I assemblies. **(b)** Representative firing rate output, Type 2 inhibitory currents, and balance error Δ traces over time from non-overlapping (top) and overlapping (bottom) assemblies. Overlap induces deviations from the E/I balance target. **(c)** Scatter plot of E/I balance from the neuronal population in non-overlapping (top) and overlapping (bottom) networks. Each point represents average currents over one second for one hidden layer excitatory neuron. Overlap induces systematic deviations toward excess inhibition. **(d)** Output trajectories of networks with different levels of overlap trained on the trajectory matching task (see Fig.4). **(e)** Balance disruption, quantified as $\sum_i |\Delta_i|$, as a function of overlap for individual networks (crosses) and mean (green; $n = 5$). **(f)** Temporal evolution of training loss for different degrees of overlap. More overlap leads to unstable learning and worse performance. **(g)** Final R^2 after 360 minutes of training on the trajectory matching task. **(h)** Same as (g) for the loss.

learning in multi-layer networks.

Learning is robust to moderate overlap in E/I assemblies

Thus far, we had considered idealized non-overlapping assemblies in the hidden layer in which each excitatory and each Type 2 inhibitory neuron belonged to a single assembly. However, it seems unlikely that such a clear assembly separation exists in neurobiology. Therefore, we sought to understand how robust the findings are to assemblies with overlap [69–71]. To address this question, we systematically varied the degree of overlap (Fig. 5a) by randomly assigning both excitatory and inhibitory neurons to additional assemblies (Methods). We characterize the degree of overlap as the ratio of additional assembly memberships to total assembly memberships.

We first examined how assembly overlap affected E/I balance in hidden layer neurons. We found that assembly overlap induced a systematic disruption of the precise E/I balance, predominantly toward excess inhibition (Fig. 5b,c). This result can be understood because overlap causes excitatory neurons to receive input from and provide feedback to multiple inhibitory microcircuits. Consequently, the systematic disruptions induced by E/I assembly overlap might impair task performance when using E/I balance deviations as error signals for learning.

To test this hypothesis, we quantified the magnitude of balance disruptions for different degrees of assembly overlap (Fig. 5d; Methods). Doing so revealed that the degree of systematic balance disruption scaled monotonically with increasing assembly overlap as expected (Fig. 5e). We then trained these networks on the trajectory matching task as before (cf. Fig. 4) and measured final loss and R^2. While all networks showed initial decreases in task loss, some configurations exceeding 10% overlap exhibited unstable learning dynamics (Fig. 5f). Notably, this effect had a strong initialization dependence as some networks with 40% overlap achieved comparable performance to non-overlapping architectures, while others showed severe impairment. Overall, we observed a graceful degradation of final performance with increasing assembly overlap (Fig. 5g,h). These results demonstrate that BCP learning is robust to moderate deviations

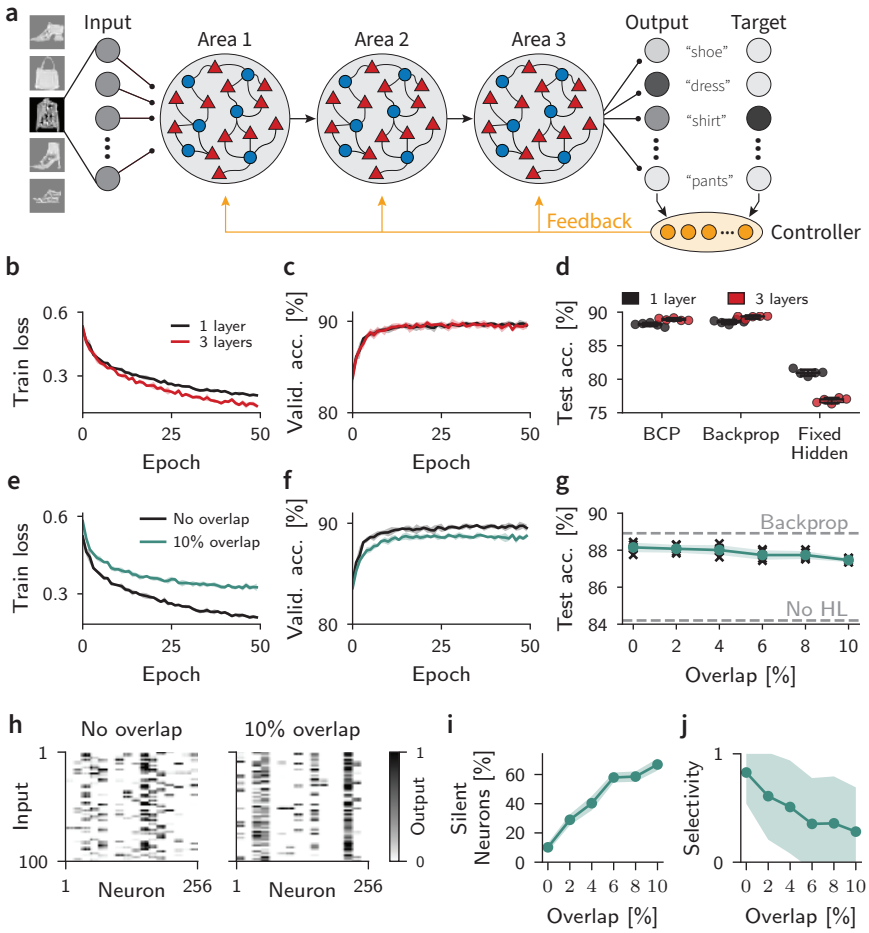

Figure 6. BCP enables training of deep networks on a computer vision benchmark. **(a)** Schematic of network for Fashion MNIST classification. Each hidden layer contains 2048 excitatory and 512 inhibitory neurons organized into 128 assemblies with assembly-specific feedback. **(b)** Evolution of training loss for networks with one or three hidden layers. **(c)** Evolution of validation accuracy during training. **(d)** Final test accuracy (n = 5). **(e,f)** Same as (b,c) for one hidden layer with 0% and 10% assembly overlap. **(g)** Final test accuracy as a function of assembly overlap. **(h)** Post-training activity of 256 representative hidden layer excitatory neurons in a network without overlap (left) and with 10% overlap (right). **(i)** Percentage of silent neurons after training as a function of assembly overlap. **(j)** Neuronal selectivity measured as lifetime sparseness post training as a function of assembly overlap. In all panels, shaded regions and error bars correspond to mean ± SD across n = 5 networks trained with different random seeds.

from an ideal E/I assembly structure, whereas increasing levels of overlap degrades learning performance.

BCP enables training of deep networks.

So far we focused on a simple trajectory matching task in a network with a single hidden layer. To test whether BCP would also be conducive for learning for more complex classification tasks in multi-layer networks, we trained networks with either one or three hidden layers, each comprising 2048 excitatory and 512 inhibitory neurons organized into 128 assemblies (Fig. 6a), on Fashion MNIST, a ten-way classification task (Methods). In all networks, the output layer contained ten neurons corresponding to the ten classes. As before, we ensured assembly-specific feedback weights tuned to convey top-down feedback signals from a proportional controller at the network output (Methods). Given the static nature of Fashion MNIST inputs, we presented each input image until neuronal dynamics converged before updating synaptic weights based on the balance error Δ. We trained all networks for 50 epochs while monitoring performance on a held-out validation dataset and measured final performance on the test set.

The training loss decreased during training of all networks (Fig. 6b) resulting in increasing classification accuracy (Fig. 6c). Networks with one hidden layer achieved a mean test accuracy of $88.1 \pm 0.3\%$, while networks with three hidden layers reached a mean test accuracy of $88.7 \pm 0.3\%$. For comparison we trained comparable networks using backpropagation which achieved comparable performance. In contrast, networks in which we trained the readout weights but not the hidden-layer weights performed substantially worse, indicating that feedback-driven plasticity in hidden layers was required for high task performance (Fig. 6d). These results demonstrate that BCP can successfully train deep networks to competitive performance levels on a standard machine learning benchmark.

To assess whether the resilience of BCP to assembly overlap extends to more complex learning scenarios, we trained single-hidden-layer networks with overlapping assemblies on the Fashion MNIST task. Networks with 10% overlap exhibited slower learning trajectories and converged to lower performance levels

compared to non-overlapping networks (Fig. 6e,f). To quantify the relationship between assembly overlap and final performance, we trained networks with varying degrees of overlap (0-10%) and evaluated their test accuracy. We found that that final performance decreased slightly with increasing overlap percentage (Fig. 6g), consistent with our findings from the trajectory matching task.

To understand how assembly overlap affected network dynamics and learning performance, we analyzed the activity of hidden neurons in response to Fashion MNIST inputs after training. In networks with 10% overlap, we observed that several neurons failed to respond to any input, indicating that the network did not utilize these neurons for classification (Fig. 6h). We quantified the number of silent neurons and found that the percentage of these silent neurons increased with assembly overlap (Fig. 6i). While only $10.2 \pm 3.2\%$ of neurons were silent in non-overlapping networks, this proportion rose to $66.8 \pm 4.1\%$ in networks with 10% overlap.

To investigate how this increase in silent neurons affected the representation of task-relevant information in the network, we calculated the lifetime sparseness of hidden layer neurons across different inputs. This analysis revealed that the selectivity of the remaining active neurons decreased with assembly overlap (Fig. 6j). Together, these results suggest that assembly overlap impairs learning through two mechanisms: by reducing the number of neurons available for encoding task-relevant information and by decreasing the selectivity of the remaining active neurons.

BCP captures signatures of *in-vivo* interneuron activity during learning

A plethora of experimental studies suggests that inhibitory circuits play essential roles in gating synaptic plasticity and learning [43, 45, 46, 48, 52, 53, 72–79]. More specifically, several studies implicated top-down activation of vasoactive intestinal peptide (VIP) interneurons with learning. We wanted to understand to what extend BCP could explain the results by Krabbe et al. [46] who showed that VIP interneuron activation in the Basolateral Amygdala (BLA) is required for learning in an auditory fear conditioning experiment. In this experiment, freely moving mice were exposed to ten conditioned stimulus (CS)-unconditioned stimulus (US) pairings over the course of two days while the activity of VIP interneurons was assessed via two-photon Ca2+ imaging through a head-fixed miniature microscope (Fig. 7a). The authors found that VIP interneurons responded strongly to the US in naive animals and that this response reduced over the course of ten pairings, accompanying the acquisition of a behavioral freezing response to the CS (Fig. 7b). To test whether BCP could account for these findings, we used it to train E/I assembly networks on the same task (Fig. 7c). In the model, the CS corresponded to sparse noisy input and the network was trained to output the freezing probability (Methods). We modeled the US as a negative valence target for the network output and implemented proportional error feedback through a dis-inhibitory circuit motif mediated by Type 2 interneurons (Methods). We simulated ten CS-US pairings and recorded the amount of US-driven dis-inhibitory feedback as well as freezing responses as in the experiment. The model output acquired a reliable freezing response which was accompanied by a reduction in US-driven dis-inhibitory feedback (Fig. 7d) similar to the experiment. In the experiment, optogenetically inhibiting the US response of VIP interneurons impaired learning (Fig. 7e). To test whether dis-inhibitory feedback is also required for learning in the model, we inhibited the VIP population by reducing the magnitude of the control signal by a factor of ten. This manipulation led to a stark reduction in freezing probability similar to the experimental data (Fig. 7f). Thus, our model can explain VIP interneuron activity during learning, suggesting a putative role for these inhibitory microcircuits in encoding top-down error signals in the BLA.

To understand whether BCP could also capture interneuron dynamics in the cortex during learning, we trained a multi-layer network model with BCP on a motor learning task (Fig. 8a) known to require cholinergic top-down feedback to VIP-somatostatin (SOM) dis-inhibitory microcircuits in primary motor cortex (M1) [47]. Networks learned to generate a motor output resembling a lever push in response to an auditory cue over the course of 2000 training iterations (Fig. 8b,c), as in the experiment by Ren et al. [47] (Fig. 8d). Importantly, Type 2 neurons in the model showed bi-directional activity changes over learning, with an overall trend towards increased activity (Fig. 8e), consistent with activity changes in SOM interneurons in the experiment (Fig. 8f). In contrast, the amount of top-down feedback to almost all E/I assemblies in the model decreased over learning (Fig. 8g), closely resembling the experimental time course of VIP interneuron activity in the experiment (Fig. 8h). Thus, the evolution of activity of Type 2 interneurons and top-down feedback during learning in the BCP framework captures the activity changes observed in SOM and VIP interneurons in M1, respectively, suggesting a role for these inhibitory microcircuits in local error

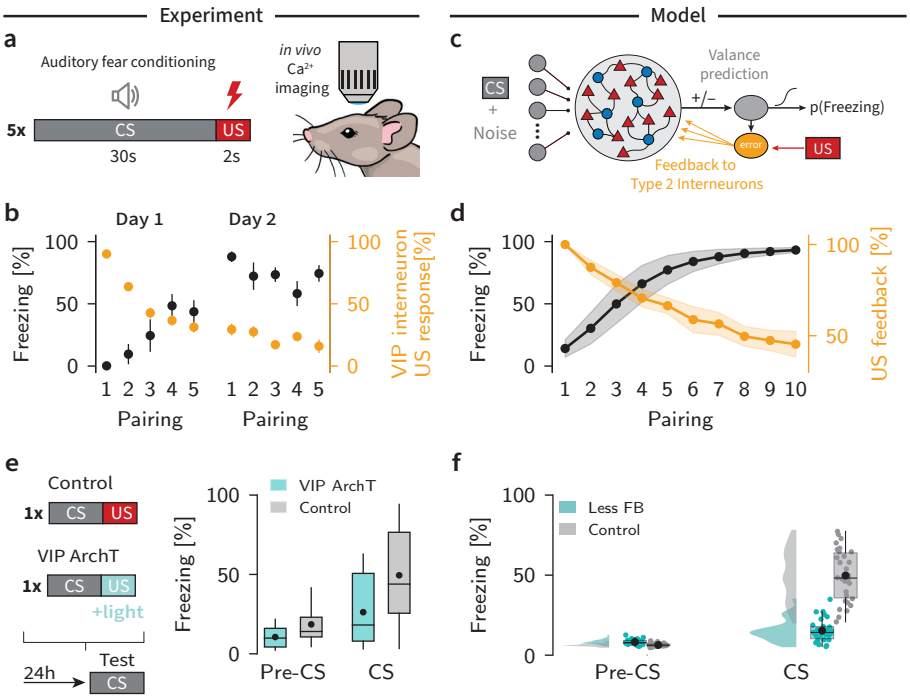

Figure 7. BCP captures signatures of fear learning in the rodent BLA. (a) Schematic of the fear conditioning experiment by Krabbe et al. [46]. Freely moving mice were exposed to five pairings of CS (tone) and US (footshock) while the activity of VIP interneurons in the BLA was recorded using a genetically encoded Ca2+ indicator. Learning was assessed by recording the animal's freezing response to the CS presentation. **(b)** Time course of freezing responses and VIP neuron activation over sessions in the experiment [46]. Mice gradually acquired a freezing response to the CS over 10 pairings over 2 days (black). Learning was accompanied by decreasing VIP interneuron responses to the US (orange). Data extracted from [46]. **(c)** Schematic of the BLA network model trained with BCP. The output of the network serves as a valance prediction, which is translated through a sigmoidal activation function into freezing probability. Error signals are computed by comparing valence prediction to a target value associated with the US and fed back to the network through a dis-inhibitory circuit (orange; Methods). **(d)** Timecourse of freezing and VIP feedback during fear learning in the model. Learning corresponds to increasing feezing probaility (black) accompanied by decreasing US-driven dis-inhibitory feedback (orange)mirroring key elements of the experiment (cf. (b)). Shaded areas correspond to the standard deviation ($n = 30$ simulation runs). **(e)** Left: Schematic of the experiment [46] in which the response of VIP interneurons to the US was optogenetically decreased during CS-US pairing (VIP ArchT condition). Right: Freezing behavior before and after fear conditioning indicates that US-driven VIP interneuron activation during pairing is necessary for learning. Data extracted from [46]. **(f)** Freezing probability in the model before (left) and following learning (right) for simulted ArchT and the control condition. Reducing the amount of US-driven feedback in the model leads to impaired learning, which is consistent with the experimental findings (cf. (e)).

encoding.

Discussion

In this article, we investigated a circuit-level mechanism for error-driven learning based on deviations from precise E/I balance. We introduced balance-controlled plasticity (BCP), a plasticity model that assumes that neurons and synapses can detect both the amount and source of inhibitory input in a cell-type-specific manner. We showed that BCP supports online learning in multi-layer networks when suitable top-down feedback influences local microcircuits. Finally, models in which feedback targets an inhibitory subpopulation naturally replicate *in-vivo* activity changes of SOM and VIP interneurons during learning [46, 47].

Neuron-specific errors that influence plasticity are essential for many plasticity algorithms—most notably the ones solving the spatial credit assignment problem. How the brain accomplishes such error specificity remains an open question [2, 13]. Recent models proposed mechanisms such as separate temporal phases [14] for feed-forward vs error encoding or argued that electrotonically segregated dendrites are ideally suited for integrating error or target signals [7, 9, 11, 80–83]. However, such a separation of feed-forward activity and error encoding is often ad-hoc, not always easy to reconcile with neuronal data, and raises conceptual questions. On important question is how error signals can be sensed at synapses targeting the soma

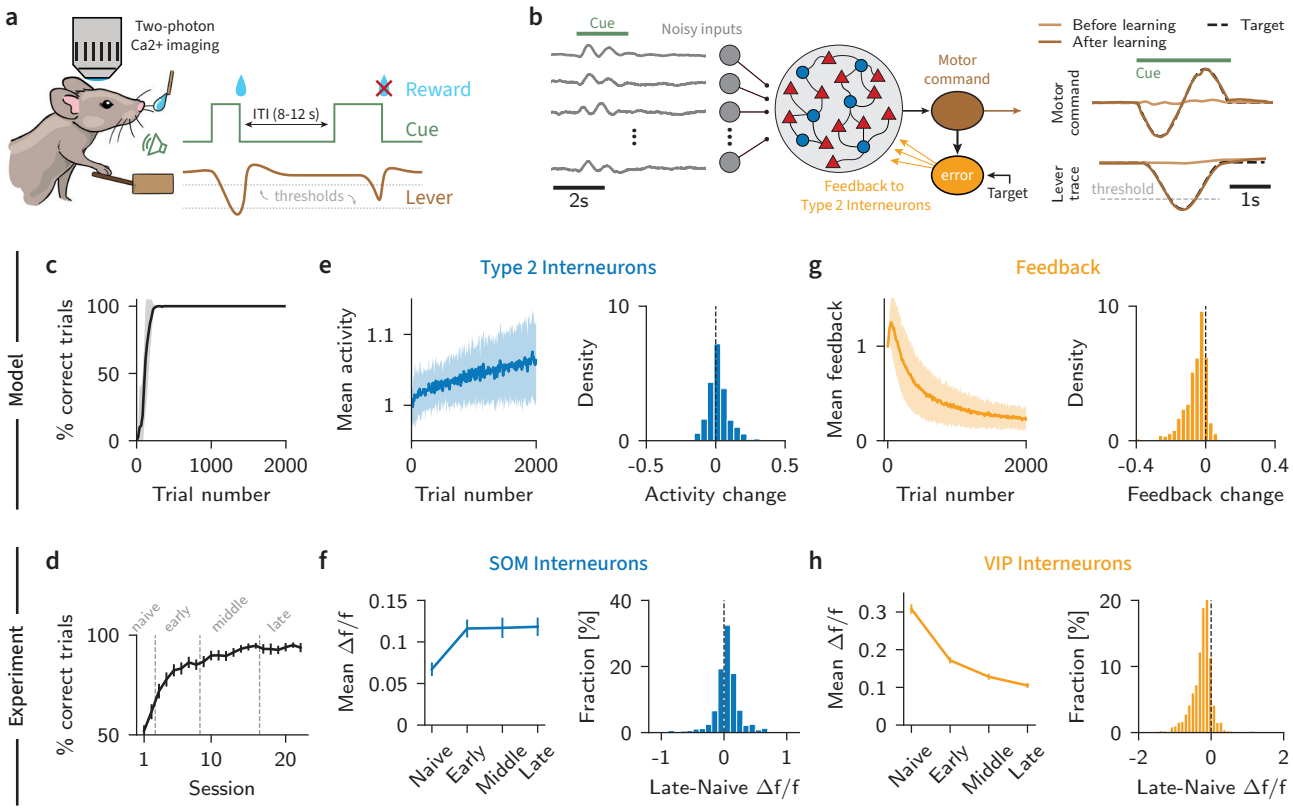

Figure 8. BCP captures activity changes of M1 interneurons during motor learning. **(a)** Schematic of the motor learning experiment by Ren et al. [47]. Mice were trained to push and pull a lever in response to an auditory cue to obtain a water reward over 22 sessions spread across multiple days. In all sessions, the activity of VIP and SOM interneurons in M1 was recorded using two-photon Ca2+ imaging. **(b)** Schematic of the M1 network model trained with BCP. In each trial, the model receives noisy cue-driven inputs that are processed by a hidden layer with $n = 40$ E/I Assemblies. A single output neuron was trained to follow the simulated time course of a push-pull movement in response to the cue. **(c)** Evolution of network performance during task learning. The shaded region indicates one standard deviation ($n = 10$ networks). **(d)** Performance of mice across training sessions. Data extracted from [47]. **(e)** Left: Evolution of mean Type 2 interneuron activity in the model during motor output across trials. Right: Distribution of learning-induced Type 2 interneuron activity changes in the model. **(f)** Left: Mean activity of SOM interneurons in M1 during movement at different stages of learning. Right: Distribution of activity changes across SOM interneurons in M1. Data extracted from [47]. **(g)** As in panel **(d)**, displaying the magnitude of top-down feedback in the network during motor output. **(h)** As in panel **(e)**, for VIP interneurons. Data extracted from [47].

or basal dendrites if they are encoded in the apical tuft dendrite? While some models compellingly argue that an intricate burst multiplexing code based on nonlinear interactions of apical and basal dendritic compartments constitute a plausible mechanism that circumvents this problem [15, 84], these models chiefly apply to Layer 5 pyramidal cells, which possess the required well-separated apical and basal dendritic compartments and exhibit bursting activity consistent with the notion of multiplexing [16]. However, this reliance leaves open the question of how other neurons, for instance in cortical Layers 2/3 or in subcortical brain areas, that do not meet these criteria, accomplish credit assignment.

In this article we have proposed one possible alternative. Since E/I currents are often precisely balanced, i.e., correlated, this create a null-space in which neuronal activity does not change. BCP's core idea is to use this null-space for error coding. Consequently, BCP does not rely on any anatomical specialization or compartmentalization. Instead, it requires mechanisms to retain and break precise E/I balance in an orchestrated manner. While some of the models discussed above also posit a form of balance in their apical dendrite [7, 15, 84], our work shows that we can dispense with this dendritic compartment and effectively externalize it into a local recurrent microcircuit, i.e., an E/I assembly. What we gain through this seemingly small change is that the resulting learning rules are more closely aligned with phenomenological voltage-dependent plasticity models. A direct consequence is that E/I assemblies effectively replace

individual neurons as the computational units in this model.

To enable learning in multi-layer networks, BCP needs to be accompanied by a credit assignment strategy. Here, our work draws from recent models that use strong top-down feedback to influence neuronal activity and plasticity [5, 11, 81, 82, 85–87]. These models can be derived from a minimization of control principle in which the error signal emerges from the difference between *open-loop* (in the absence of top-down feedback) and *closed-loop* activity (with feedback). A central challenge for these models remains that there is little evidence for explicit open and closed-loop phases in neurobiology. Previous models solved this issue by, once more, asserting segregated dendrites [5, 81, 82, 87] or dedicated error neurons [11] based on similar mathematical ideas [10]. Other recent work demonstrated that the need for separate phases or compartmentalization can be relaxed when temporally delayed top-down feedback induces firing rate changes detectable by an spike-timing-dependent plasticity (STDP) rule [88]. In contrast, BCP allows reformulating the problem by mapping the open- versus closed-loop contributions onto a distinct coding axis related to deviations from E/I balance. Here, the expected inhibitory current under precise balance encodes the *open-loop* activity, while the actual inhibitory input reflects the *closed-loop* state. This separation gives each neuron implicit access to its own *open-loop* activity without requiring additional circuitry or compartmentalization.

Importantly, mapping error signals to E/I balance instead of putative anatomical compartments also results in learning rules that more closely align with phenomenological plasticity models. We showed that BCP resembles a Hebbian plasticity rule with a dynamic plasticity threshold controlled by inhibition (cf. Fig. 2h). While plasticity thresholds are central to classic models of Hebbian learning [89, 90] and phenomenological plasticity rules [23, 24, 26, 91, 92], most learning rules do not depend explicitly on inhibition. Inhibition-dependent learning has been investigated in recent theoretical studies, which implicated it in receptive field formation and stabilization of learning [55–57, 64]. While these studies shared BCP's core principles of E/I balance and an explicit inhibition dependence of plasticity at excitatory synapses, they did not study the implications of implicit error encoding for credit assignment in multi-layer networks.

Finally, our work suggests a reinterpretation of the basic computational units in biological neural networks. Specifically, E/I assemblies are the computational units in our model rather than individual neurons. This interpretation opens up exciting parallels to recent models on efficient coding [36, 93], low-rank networks [94, 95], and work showing that low-dimensional error signals are often sufficient for learning complex functions in deep networks [96]. Because the number of E/I assemblies puts an upper bound on the effective dimensionality of the neuronal manifold, we expect that brain areas encoding high-dimensional information comprise many small assemblies rather than a few large assemblies. Interestingly, there is indirect evidence for such E/I assembly structures, for instance, in the area Dp, the homologue of the olfactory cortex in zebrafish [37, 97].

Putative bio-chemical mechanisms of BCP. Our model is built on the assumption that neurons can distinguish between excitatory and specific types of inhibitory input. Fundamentally, BCP's balance condition (cf. Eq. (2)) can be interpreted more broadly as a balance between LTP-inducing and inhibition-controlled LTD-inducing plasticity mechanisms. While the concrete mechanism is unknown, several molecular pathways could be involved in mediating such balance control. First, inhibitory currents can precisely regulate local dendritic voltage [98], thereby controlling Ca^{2+} influx through NMDA channels [99] known to modulate plasticity [100]. Additionally, inhibition can directly modulate Ca^{2+} permeability of both NMDA receptors and voltage-gated Ca^{2+} channels through non-ionotropic GABA-B receptor signaling [101–103]. Another conceivable pathway extending beyond calcium signaling may rely on inhibitory inputs regulating local kinases and phosphatases involved in synaptic plasticity [104, 105]. While the model formulation in this article assumes these processes operate at synaptic timescales, biological plasticity mechanisms likely incorporate multiple timescales of balanced processes [106, 107]. Thus, the effects of E/I balance on plasticity may manifest more slowly than synaptic currents, suggesting that plasticity-relevant processes could still cancel even without millisecond-precision of the E/I balance [35].

Limitations. We made several simplifying assumptions. First, we assumed that top-down feedback weights were optimized to deliver relevant credit signals to their target assemblies (except in Supplementary Figure S2, where we used random feedback). The need for appropriate feedback signals is a fundamental

challenge shared by all biologically plausible credit assignment strategies [2, 13]. Still several studies have shown that this limitation can be mitigated by learning the feedback weights [9, 108–113] or by redesigning feedback pathways [114]. In this article, we did not focus on any particular candidate feedback mechanism. Instead we presented an effective error encoding mechanism which could work with most, if not all, biologically plausible credit assignment models that traditionally rely on separate compartments for receiving such error signals.

In this article we focused on feedback to Type 2 interneurons which provided recurrent precisely balanced inhibition. This choice was inspired by the prominent dis-inhibitory circuit motifs found in many brain areas [115] and involved in gating of plasticity [116]. However, it is important to note that our adaptive control theory framework is agnostic to how E/I balance is established and how it is influenced through feedback. Hence other circuit models and feedback pathways are possible. For example, BCP would yield a similar workable model in which feed-forward inhibition is responsible for precise balance and targeted by feedback (cf. Fig. S1). Similarly, feedback pathways relying entirely on excitatory neurons are possible too. As long as E/I balance is modulated in an orchestrated way, we expect BCP to work.

Another limitation of our work is that while our model requires E/I assemblies, it remains agnostic about how they are established and maintained during learning. In this article, we artificially enforced E/I assemblies for simplicity. In neurobiology, the required assembly connectivity would necessarily have to be self-organized during development and learning. Fortunately, several studies have shown how inhibitory plasticity can establish co-tuned E/I assemblies through decorrelation of excitatory activity [59–61, 117–121]. Still, maintaining E/I balance while using balance errors for learning poses a chicken-egg problem. Solving this problem may ultimately require cycling between different plasticity states. A promising candidate could be our wake-sleep cycle. For instance, balance errors would serve error-driven learning during wakefulness. In contrast, one of the roles of sleep may be to re-establish E/I balance [122], a notion that also seems consistent with work showing that inhibitory engrams emerge over prolonged periods *after* memory storage [123]. Beyond maintaining E/I balance, inhibitory plasticity could serve complementary functions in learning that we did not consider in our model. While we have considered circuits in which top-down feedback transiently modulates inhibitory activity to induce plasticity, other computational work explored networks in which top-down feedback drives cell-type specific inhibitory plasticity [124] that could allow for prolonged control over excitatory plasticity even after top-down feedback subsides.

Moreover, BCP is a simplified plasticity model which does not yield classic plasticity curves associated with low postsynaptic activity result in LTD whereas high activity levels yield LTP in the absence of inhibition. This choice was intentional for ease of analysis, but is not a fundamental limitation of the model. Importantly, we have shown previously that learning rules with an explicit inhibition dependence in which error signals are decoded through a linear approximation of the inverse of the neuronal transfer function [125] naturally recover classic plasticity curves while also providing rapid compensatory control of Hebbian plasticity through recurrent inhibition alone [126, 127].

Finally, our model does not *require* electrically well-separated dendritic compartments to encode error signals. However, beyond any doubt dendritic compartmentalization fulfills important computational functions [128–130]. Curiously, recent experiments have shown that E/I balance is maintained at the level of local dendritic segments [131], and that some interneurons form synapses on individual dendritic spines [98]. While we have only considered simple neuron models without dendritic compartments in this article, in future work we would like to extend the BCP framework to compartment-specific error signals through precise E/I balance that is maintained at the level of individual dendrites and explore whether such compartment-specific error encoding offers computational advantages beyond neuron-wide error signals.

Experimental predictions. Our study makes several testable predictions. First, the BCP learning rule predicts that inhibitory interneuron activity should be modulated by top-down feedback during learning and thereby effect both excitatory activity *and* plasticity. While there is ample evidence for inhibitory gating of plasticity both *in vitro* [39–45, 53] and *in vivo* [46–49, 52, 54, 76, 77], most experiments did not separately control inhibition and neuronal firing rates because they are typically anti-correlated. To isolate balance effects from firing rate changes, we propose a closed-loop experimental learning paradigm. In this setup, stimulus strength or excitatory neuronal input is dynamically adjusted to maintain constant target firing rates in an excitatory population. At the same time recurrent inhibitory activity is modulated, e.g., optoge-

netically, to break E/I balance. Under these conditions, BCP predicts observable plasticity at input synapses despite stable output firing rates. In addition, the direction and magnitude of plasticity should be a function of inhibitory activity.

Second, our model posits a dual role for inhibition, suggesting that inhibitory plasticity rules should differ based on the presynaptic population's circuit function. In our model, we distinguish Type 1 inhibition, which acts as the negative part of a neuron's receptive field, and Type 2 inhibition, which establishes precise balance. BCP predicts that Type 2 (primarily recurrent) inhibitory synapses should exhibit co-dependent plasticity to maintain balance [57], while the sign of plasticity at Type 1 (primarily feed-forward) inhibitory synapses should be governed by deviations from E/I balance. This prediction aligns with the diverse forms of inhibitory STDP observed experimentally [120, 132]. Importantly, the relationship between connectivity motifs and circuit function does not have to be rigid—some circuits may employ feed-forward inhibition to maintain balance [25, 133–135] while using lateral inhibition between different assemblies to shape receptive fields (c.f. Supplementary Fig. S1).

Finally, our model assumes that specific inhibitory interneuron subpopulations are modulated by top-down feedback during learning consistent with feedback control. This assumption is also a prediction. Specifically, we expect that such inhibition-mediated feedback responding to external stimuli early in learning can predict both behavioral and neuronal responses that emerge later in response to the same stimuli. Several studies have presented evidence for such (dis-)inhibitory feedback control [46, 49, 74, 78, 136] that guides learning and population activity. However, whether this top-down control predicts plasticity at the single neuron level remains unresolved.

In summary, our work introduces a biologically grounded mechanism by which neurons can encode local error signals through deviations from precise excitation-inhibition balance. This balance-controlled plasticity enables online credit assignment without requiring specialized compartments or phases, unifying normative principles with observed cortical microcircuit dynamics. By linking synaptic plasticity to transient deviations in E/I balance, our model offers a testable framework for how neuronal circuits may solve complex learning tasks using local information.

Acknowledgements

We thank all members of the Zenke Lab for comments and discussions throughout this project. Thanks to Peter Buttaroni for providing helpful illustrations for figures in this manuscript. This project was supported by the Swiss National Science Foundation (Grant Number PCEFP3_202981), EU's Horizon Europe Research and Innovation Programme (Grant Agreement No. 101070374, CONVOLVE) funded through SERI (Ref. 1131–52302), and the Novartis Research Foundation.

Author contributions

Julian Rossbroich: Formal analysis, Investigation, Methodology, Software, Visualization, Writing – original draft, Writing – review and editing; Friedemann Zenke: Conceptualization, Funding acquisition, Methodology, Supervision, Writing – original draft, Writing – review and editing.

Declaration of interests

The authors declare no competing interests.

Data and code availability

The code to reproduce all modeling results reported in this article is available at https://github.com/fmi-basel/balance-controlled-plasticity

Supplemental information

Supplementary Material document. Figures S1 & S2, Tables S1-S5 and supplementary methods.

References

[1] T. P. Lillicrap, D. Cownden, D. B. Tweed, and C. J. Akerman. "Random Synaptic Feedback Weights Support Error Backpropagation for Deep Learning". In: *Nature Communications* 7 (2016), p. 13276.

[2] T. P. Lillicrap, A. Santoro, L. Marris, C. J. Akerman, and G. Hinton. "Backpropagation and the Brain". In: *Nature Reviews Neuroscience* 21 (2020), pp. 335–346.

[3] B. Illing, J. Ventura, G. Bellec, and W. Gerstner. "Local Plasticity Rules Can Learn Deep Representations Using Self-Supervised Contrastive Predictions". In: *Advances in Neural Information Processing Systems*. Vol. 34. Curran Associates, Inc., 2021, pp. 30365–30379.

[4] K. K. Nejad, P. Anastasiades, L. Hertäg, and R. P. Costa. "Self-supervised predictive learning accounts for cortical layer-specificity". en. In: *Nature Communications* 16.1 (2025), p. 6178. ISSN: 2041-1723. DOI: 10.1038/s41467-025-61399-5.

[5] A. Gilra and W. Gerstner. "Predicting Non-Linear Dynamics by Stable Local Learning in a Recurrent Spiking Neural Network". In: *eLife* 6 (2017). Ed. by P. Latham, e28295. ISSN: 2050-084X. DOI: 10.7554/eLife.28295.

[6] J. C. R. Whittington and R. Bogacz. "An Approximation of the Error Backpropagation Algorithm in a Predictive Coding Network with Local Hebbian Synaptic Plasticity". In: *Neural computation* 29.5 (2017), pp. 1229–1262. ISSN: 0899-7667, 1530-888X. DOI: 10.1162/NECO_a_00949. PMID: 28333583.

[7] J. Sacramento, R. Ponte Costa, Y. Bengio, and W. Senn. "Dendritic Cortical Microcircuits Approximate the Backpropagation Algorithm". In: *Advances in Neural Information Processing Systems 31*. Ed. by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett. Curran Associates, Inc., 2018, pp. 8721–8732.

[8] J. C. R. Whittington and R. Bogacz. "Theories of Error Back-Propagation in the Brain". In: *Trends in cognitive sciences* 23.3 (2019), pp. 235–250. ISSN: 1364-6613, 1879-307X. DOI: 10.1016/j.tics.2018.12.005. PMID: 30704969.

[9] A. Meulemans, M. Tristany Farinha, J. Garcia Ordonez, P. Vilimelis Aceituno, J. Sacramento, and B. F. Grewe. "Credit Assignment in Neural Networks through Deep Feedback Control". In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, and J. W. Vaughan. Vol. 34. Curran Associates, Inc., 2021, pp. 4674–4687.

[10] F. A. Mikulasch, L. Rudelt, M. Wibral, and V. Priesemann. "Where Is the Error? Hierarchical Predictive Coding through Dendritic Error Computation". In: *Trends in Neurosciences* 46.1 (2023), pp. 45–59. ISSN: 1878-108X. DOI: 10.1016/j.tins.2022.09.007. PMID: 36577388.

[11] Y. Song, B. Millidge, T. Salvatori, T. Lukasiewicz, Z. Xu, and R. Bogacz. "Inferring Neural Activity before Plasticity as a Foundation for Learning beyond Backpropagation". In: *Nature Neuroscience* 27.2 (2024), pp. 348–358. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/s41593-023-01514-1.

[12] F. Crick. "The Recent Excitement about Neural Networks". In: *Nature* 337.6203 (1989), pp. 129–132. ISSN: 0028-0836. DOI: 10.1038/337129a0. PMID: 2911347.

[13] B. A. Richards, T. P. Lillicrap, P. Beaudoin, Y. Bengio, R. Bogacz, A. Christensen, C. Clopath, R. P. Costa, A. de Berker, S. Ganguli, C. J. Gillon, D. Hafner, A. Kepecs, N. Kriegeskorte, P. Latham, G. W. Lindsay, K. D. Miller, R. Naud, C. C. Pack, P. Poirazi, P. Roelfsema, J. Sacramento, A. Saxe, B. Scellier, A. C. Schapiro, W. Senn, G. Wayne, D. Yamins, F. Zenke, J. Zylberberg, D. Therien, and K. P. Kording. "A Deep Learning Framework for Neuroscience". In: *Nature neuroscience* 22.11 (2019), pp. 1761–1770. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/s41593-019-0520-2. PMID: 31659335.

[14] J. Guerguiev, T. P. Lillicrap, and B. A. Richards. "Towards Deep Learning with Segregated Dendrites". In: *eLife* 6 (2017). ISSN: 2050-084X. DOI: 10.7554/eLife.22901. PMID: 29205151.

[15] A. Payeur, J. Guerguiev, F. Zenke, B. A. Richards, and R. Naud. "Burst-Dependent Synaptic Plasticity Can Coordinate Learning in Hierarchical Circuits". In: *Nature neuroscience* (2021). ISSN: 1097-6256, 1546-1726. DOI: 10.1038/s41593-021-00857-x. PMID: 33986551.

[16] M. E. Larkum, J. J. Zhu, and B. Sakmann. "A New Cellular Mechanism for Coupling Inputs Arriving at Different Cortical Layers". In: *Nature* 398.6725 (1999), pp. 338–341. ISSN: 0028-0836. DOI: 10.1038/18686. PMID: 10192334.

[17] B. Scellier and Y. Bengio. "Equilibrium Propagation: Bridging the Gap between Energy-Based Models and Backpropagation". In: *Frontiers in computational neuroscience* 11 (2017), p. 24. ISSN: 1662-5188. DOI: 10.3389/fncom.2017.00024. PMID: 28522969.

[18] R. Urbanczik and W. Senn. "Learning by the Dendritic Prediction of Somatic Spiking". In: *Neuron* 81.3 (2014), pp. 521–528. ISSN: 0896-6273, 1097-4199. DOI: 10.1016/j.neuron.2013.11.030. PMID: 24507189.

[19] A. Laborieux and F. Zenke. "Holomorphic Equilibrium Propagation Computes Exact Gradients Through Finite Size Oscillations". In: *Advances in Neural Information Processing Systems*. Ed. by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh. Vol. 35. Curran Associates, Inc., 2022, pp. 12950–12963.

[20] S. Song, K. D. Miller, and L. F. Abbott. "Competitive Hebbian Learning through Spike-Timing-Dependent Synaptic Plasticity". In: *Nature Neuroscience* 3.9 (2000), pp. 919–926. ISSN: 1546-1726. DOI: 10.1038/78829.

[21] W. Gerstner and W. M. Kistler. "Mathematical Formulations of Hebbian Learning". In: *Biological cybernetics* 87.5–6 (2002), pp. 404–415. ISSN: 0340-1200. DOI: 10.1007/s00422-002-0353-y. PMID: 12461630.

[22] H. Z. Shouval, M. F. Bear, and L. N. Cooper. "A Unified Model of NMDA Receptor-Dependent Bidirectional Synaptic Plasticity". In: *Proc Natl Acad Sci U S A* 99.16 (2002), pp. 10831–10836. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.152343099.

[23] J.-P. Pfister and W. Gerstner. "Triplets of Spikes in a Model of Spike Timing-Dependent Plasticity". In: *J Neurosci* 26.38 (2006), pp. 9673–9682. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.1425-06.2006.

[24] C. Clopath, L. Büsing, E. Vasilaki, and W. Gerstner. "Connectivity Reflects Coding: A Model of Voltage-Based STDP with Homeostasis". In: *Nat Neurosci* 13.3 (2010), pp. 344–352. ISSN: 1097-6256. DOI: 10.1038/nn.2479.

[25] T. P. Vogels, H. Sprekeler, F. Zenke, C. Clopath, and W. Gerstner. "Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks". In: *Science* 334.6062 (2011), pp. 1569–1573. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1211095. PMID: 22075724.

[26] M. Graupner and N. Brunel. "Calcium-Based Plasticity Model Explains Sensitivity of Synaptic Changes to Spike Pattern, Rate, and Dendritic Location". In: *Proc Natl Acad Sci U S A* 109.10 (2012), pp. 3991–3996. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.1109359109.

[27] F. Zenke, E. J. Agnes, and W. Gerstner. "Diverse Synaptic Plasticity Mechanisms Orchestrated to Form and Retrieve Memories in Spiking Neural Networks". In: *Nature communications* 6 (2015), p. 6922. ISSN: 2041-1723. DOI: 10.1038/ncomms7922. PMID: 25897632.

[28] P. Rupprecht and R. W. Friedrich. "Precise Synaptic Balance in the Zebrafish Homolog of Olfactory Cortex". In: *Neuron* 100.3 (2018), 669–683.e5.

[29] M. Xue, B. V. Atallah, and M. Scanziani. "Equalizing Excitation–Inhibition Ratios across Visual Cortical Neurons". In: *Nature* 511.7511 (2014), pp. 596–600. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature13321.

[30] M. Okun and I. Lampl. "Instantaneous Correlation of Excitation and Inhibition during Ongoing and Sensory-Evoked Activities". In: *Nature Neuroscience* 11.5 (2008), pp. 535–537. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/nn.2105.

[31] M. Wehr and A. M. Zador. "Balanced Inhibition Underlies Tuning and Sharpens Spike Timing in Auditory Cortex". In: *Nature* 426.6965 (2003), pp. 442–446.

[32] A. Y. Y. Tan, L. I. Zhang, M. M. Merzenich, and C. E. Schreiner. "Tone-Evoked Excitatory and Inhibitory Synaptic Conductances of Primary Auditory Cortex Neurons". In: *Journal of Neurophysiology* 92.1 (2004), pp. 630–643. ISSN: 0022-3077, 1522-1598. DOI: 10.1152/jn.01020.2003.

[33] J. S. Anderson, M. Carandini, and D. Ferster. "Orientation Tuning of Input Conductance, Excitation, and Inhibition in Cat Primary Visual Cortex". In: *Journal of Neurophysiology* 84.2 (1, 2000), pp. 909–926. ISSN: 0022-3077, 1522-1598. DOI: 10.1152/jn.2000.84.2.909.

[34] B. Haider, A. Duque, A. R. Hasenstaub, and D. A. McCormick. "Neocortical Network Activity In Vivo Is Generated through a Dynamic Balance of Excitation and Inhibition". In: *Journal of Neuroscience* 26.17 (26, 2006), pp. 4535–4545. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.5297-05.2006. PMID: 16641233.

[35] G. Hennequin, E. J. Agnes, and T. P. Vogels. "Inhibitory Plasticity: Balance, Control, and Codependence". In: *Annual Review of Neuroscience* 40 (Volume 40, 2017 25, 2017), pp. 557–579. ISSN: 0147-006X, 1545-4126. DOI: 10.1146/annurev-neuro-072116-031005.

[36] S. Denève and C. K. Machens. "Efficient Codes and Balanced Networks". In: *Nature Neuroscience* 19.3 (2016), pp. 375–382.

[37] C. Meissner-Bernard, B. Jenkins, P. Rupprecht, E. A. Bouldoires, F. Zenke, R. W. Friedrich, and T. Frank. "Computational Functions of Precisely Balanced Neuronal Microcircuits in an Olfactory Memory Network". In: *Cell Reports* 44.3 (25, 2025). ISSN: 2211-1247. DOI: 10.1016/j.celrep.2025.115330. PMID: 39985769.

[38] M. J. Higley. "Localized GABAergic Inhibition of Dendritic Ca(2+) Signalling". In: *Nature reviews. Neuroscience* 15.9 (2014), pp. 567–572. ISSN: 1471-003X, 1471-0048. DOI: 10.1038/nrn3803. PMID: 25116141.

[39] H. Hasuo and T. Akasu. "Activation of Inhibitory Pathways Suppresses the Induction of Long-Term Potentiation in Neurons of the Rat Lateral Septal Nucleus". In: *Neuroscience* 105.2 (2001), pp. 343–352. ISSN: 0306-4522. DOI: 10.1016/s0306-4522(01)00195-6. PMID: 11672602.

[40] R. M. Douglas, G. V. Goddard, and M. Riives. "Inhibitory Modulation of Long-Term Potentiation: Evidence for a Postsynaptic Locus of Control". In: *Brain research* 240.2 (1982), pp. 259–272. ISSN: 0006-8993. DOI: 10.1016/0006-8993(82)90221-9. PMID: 7104688.

[41] V. C. Kotak, A. Mirallave, T. M. Mowery, and D. H. Sanes. "GABAergic Inhibition Gates Excitatory LTP in Perirhinal Cortex". In: *Hippocampus* 27.12 (2017), pp. 1217–1223. ISSN: 1050-9631, 1098-1063. DOI: 10.1002/hipo.22799. PMID: 28881444.

[42] L. Wang and A. Maffei. "Inhibitory Plasticity Dictates the Sign of Plasticity at Excitatory Synapses". In: *The Journal of neuroscience* 34.4 (2014), pp. 1083–1093. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.4711-13.2014. PMID: 24453301.

[43] V. Paille, E. Fino, K. Du, T. Morera-Herreras, S. Perez, J. H. Kotaleski, and L. Venance. "GABAergic Circuits Control Spike-Timing-Dependent Plasticity". In: *The Journal of neuroscience* 33.22 (2013), pp. 9353–9363. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.5796-12.2013. PMID: 23719804.

[44] T. Hayama, J. Noguchi, S. Watanabe, N. Takahashi, A. Hayashi-Takagi, G. C. R. Ellis-Davies, M. Matsuzaki, and H. Kasai. "GABA Promotes the Competitive Selection of Dendritic Spines by Controlling Local Ca2+ Signaling". In: *Nature neuroscience* 16.10 (2013), pp. 1409–1416. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/nn.3496. PMID: 23974706.

[45] M. Canto-Bustos, F. K. Friason, C. Bassi, and A.-M. M. Oswald. "Disinhibitory Circuitry Gates Associative Synaptic Plasticity in Olfactory Cortex". In: *The Journal of neuroscience* 42.14 (2022), pp. 2942–2950. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.1369-21.2021. PMID: 35181596.

[46] S. Krabbe, E. Paradiso, S. d'Aquin, Y. Bitterman, J. Courtin, C. Xu, K. Yonehara, M. Markovic, C. Müller, T. Eichlisberger, J. Gründemann, F. Ferraguti, and A. Lüthi. "Adaptive Disinhibitory Gating by VIP Interneurons Permits Associative Learning". In: *Nature neuroscience* 22.11 (2019), pp. 1834–1843. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/s41593-019-0508-y. PMID: 31636447.

[47] C. Ren, K. Peng, R. Yang, W. Liu, C. Liu, and T. Komiyama. "Global and Subtype-Specific Modulation of Cortical Inhibitory Neurons Regulated by Acetylcholine during Motor Learning". In: *Neuron* 110.14 (20, 2022), 2334–2350.e8. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2022.04.031.

[48] A. Adler, R. Zhao, M. E. Shin, R. Yasuda, and W.-B. Gan. "Somatostatin-Expressing Interneurons Enable and Maintain Learning-Dependent Sequential Activation of Pyramidal Neurons". In: *Neuron* 102.1 (2019), 202–216.e7. ISSN: 08966273. DOI: 10.1016/j.neuron.2019.01.036.

[49] M. Neubrandt, N. Lenkey, and K. Vervaeke. *VIP Interneurons Control Hippocampal Place Cell Remapping through Transient Disinhibition in Novel Environments*. 1, 2025. DOI: 10.1101/2025.02.01.636072. URL: https://www.biorxiv.org/content/10.1101/2025.02.01.636072v1 (visited on 04/11/2025). Pre-published.

[50] F. Donato, S. B. Rompani, and P. Caroni. "Parvalbumin-Expressing Basket-Cell Network Plasticity Induced by Experience Regulates Adult Learning". In: *Nature* 504.7479 (2013), pp. 272–276. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature12866. PMID: 24336286.

[51] R. C. Froemke, M. M. Merzenich, and C. E. Schreiner. "A Synaptic Memory Trace for Cortical Receptive Field Plasticity". In: *Nature* 450.7168 (2007), pp. 425–429. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature06289.

[52] J. J. Letzkus, S. B. E. Wolff, E. M. M. Meyer, P. Tovote, J. Courtin, C. Herry, and A. Lüthi. "A Disinhibitory Microcircuit for Associative Fear Learning in the Auditory Cortex". In: *Nature* 480.7377 (2011), pp. 331–335. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature10674. PMID: 22158104.

[53] L. E. Williams and A. Holtmaat. "Higher-Order Thalamocortical Inputs Gate Synaptic Long-Term Potentiation via Disinhibition". In: *Neuron* 101.1 (2019), 91–102.e4. ISSN: 0896-6273, 1097-4199. DOI: 10.1016/j.neuron.2018.10.049. PMID: 30472077.

[54] S. Pagès, N. Chenouard, R. Chéreau, V. Kouskoff, F. Gambino, and A. Holtmaat. "An Increase in Dendritic Plateau Potentials Is Associated with Experience-Dependent Cortical Map Reorganization". In: *Proceedings of the National Academy of Sciences of the United States of America* 118.9 (2021). ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.2024920118. PMID: 33619110.

[55] C. Clopath, T. P. Vogels, R. C. Froemke, and H. Sprekeler. *Receptive Field Formation by Interacting Excitatory and Inhibitory Synaptic Plasticity*. 29, 2016. DOI: 10.1101/066589. URL: https://www.biorxiv.org/content/10.1101/066589v1 (visited on 02/17/2025). Pre-published.

[56] C. Miehl and J. Gjorgjieva. "Stability and Learning in Excitatory Synapses by Nonlinear Inhibitory Plasticity". In: *PLoS computational biology* 18.12 (2022), e1010682. ISSN: 1553-734X, 1553-7358. DOI: 10.1371/journal.pcbi.1010682. PMID: 36459503.

[57] E. J. Agnes and T. P. Vogels. "Co-Dependent Excitatory and Inhibitory Plasticity Accounts for Quick, Stable and Long-Lasting Memories in Biological Networks". In: *Nature Neuroscience* 27.5 (2024), pp. 964–974. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/s41593-024-01597-4.

[58] W. Yang and Q.-Q. Sun. "Circuit-Specific and Neuronal Subcellular-Wide E-I Balance in Cortical Pyramidal Cells". In: *Scientific Reports* 8.1 (5, 2018), p. 3971. ISSN: 2045-2322. DOI: 10.1038/s41598-018-22314-9.

[59] S. Eckmann, E. J. Young, and J. Gjorgjieva. "Synapse-Type-Specific Competitive Hebbian Learning Forms Functional Recurrent Networks". In: *Proceedings of the National Academy of Sciences of the United States of America* 121.25 (18, 2024), e2305326121. ISSN: 1091-6490. DOI: 10.1073/pnas.2305326121. PMID: 38870059.

[60] P. D. King, J. Zylberberg, and M. R. DeWeese. "Inhibitory Interneurons Decorrelate Excitatory Cells to Drive Sparse Code Formation in a Spiking Model of V1". In: *Journal of Neuroscience* 33.13 (2013), pp. 5475–5485.

[61] W. Brendel, R. Bourdoukan, P. Vertechi, C. K. Machens, and S. Denève. "Learning to Represent Signals Spike by Spike". In: *PLOS Computational Biology* 16.3 (16, 2020), e1007692. ISSN: 1553-7358. DOI: 10.1371/journal.pcbi.1007692.

[62] O. Mackwood, L. B. Naumann, and H. Sprekeler. "Learning Excitatory-Inhibitory Neuronal Assemblies in Recurrent Networks". In: *eLife* 10 (26, 2021), e59715. ISSN: 2050-084X. DOI: 10.7554/eLife.59715.

[63] A. Kepecs and G. Fishell. "Interneuron Cell Types Are Fit to Function". In: *Nature* 505.7483 (2014), pp. 318–326. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature12983. PMID: 24429630.

[64] D. Drix, V. V. Hafner, and M. Schmuker. "Sparse Coding with a Somato-Dendritic Rule". In: *Neural Netw.* (2020).

[65] F. Rosenblatt. *Principles of Neurodynamics: Perceptions and the Theory of Brain Mechanisms*. Washington, DC: Spartan, 1962.

[66] N. Frémaux and W. Gerstner. "Neuromodulated Spike-Timing-Dependent Plasticity, and Theory of Three-Factor Learning Rules". In: *Frontiers in neural circuits* 9 (2016), pp. 85–85. ISSN: 1662-5110. DOI: 10.3389/fncir.2015.00085.

[67] Ł. Kuśmierz, T. Isomura, and T. Toyoizumi. "Learning with Three Factors: Modulating Hebbian Plasticity with Errors". In: *Current opinion in neurobiology* 46 (2017), pp. 170–177. ISSN: 0959-4388, 1873-6882. DOI: 10.1016/j.conb.2017.08.020. PMID: 28918313.

[68] A. Nøkland. *Direct Feedback Alignment Provides Learning in Deep Neural Networks*. 2016. arXiv: 1609.01596 [stat.ML]. URL: https://arxiv.org/abs/1609.01596. Pre-published.

[69] E. Russo, N. Becker, A. P. F. Domanski, T. Howe, K. Freud, D. Durstewitz, and M. W. Jones. "Integration of Rate and Phase Codes by Hippocampal Cell-Assemblies Supports Flexible Encoding of Spatiotemporal Context". In: *Nature Communications* 15.1 (22, 2024), p. 8880. ISSN: 2041-1723. DOI: 10.1038/s41467-024-52988-x.

[70] M. Bocchio, A. Vorobyev, S. Sadeh, S. Brustlein, R. Dard, S. Reichinnek, V. Emiliani, A. Baude, C. Clopath, and R. Cossart. "Functional Networks of Inhibitory Neurons Orchestrate Synchrony in the Hippocampus". In: *PLOS Biology* 22.10 (14, 2024), e3002837. ISSN: 1545-7885. DOI: 10.1371/journal.pbio.3002837.

[71] G. Buzsáki. "Neural Syntax: Cell Assemblies, Synapsembles, and Readers". In: *Neuron* 68.3 (4, 2010), pp. 362–385. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2010.09.023.

[72] S. Furutachi, A. D. Franklin, A. M. Aldea, T. D. Mrsic-Flogel, and S. B. Hofer. "Cooperative Thalamocortical Circuit Mechanism for Sensory Prediction Errors". In: *Nature* 633.8029 (2024), pp. 398–406. ISSN: 1476-4687. DOI: 10.1038/s41586-024-07851-w.

[73] Y. Fu, M. Kaneko, Y. Tang, A. Alvarez-Buylla, and M. P. Stryker. "A Cortical Disinhibitory Circuit for Enhancing Adult Plasticity". In: *eLife* 4 (27, 2015), e05558. ISSN: 2050-084X. DOI: 10.7554/eLife.05558.

[74] A. Schroeder, M. B. Pardi, J. Keijser, T. Dalmay, A. I. Groisman, E. M. Schuman, H. Sprekeler, and J. J. Letzkus. "Inhibitory Top-down Projections from Zona Incerta Mediate Neocortical Memory". In: *Neuron* 111.5 (2023), 727–738.e8. ISSN: 08966273. DOI: 10.1016/j.neuron.2022.12.010.

[75] S. X. Chen, A. N. Kim, A. J. Peters, and T. Komiyama. "Subtype-Specific Plasticity of Inhibitory Circuits in Motor Cortex during Motor Learning". In: *Nature Neuroscience* 18.8 (2015), pp. 1109–1115.

[76] J. Basu, J. D. Zaremba, S. K. Cheung, F. L. Hitti, B. V. Zemelman, A. Losonczy, and S. A. Siegelbaum. "Gating of Hippocampal Activity, Plasticity, and Memory by Entorhinal Cortex Long-Range Inhibition". In: *Science* 351.6269 (8, 2016), aaa5694. DOI: 10.1126/science.aaa5694.

[77] S. V. Rolotti, M. S. Ahmed, M. Szoboszlay, T. Geiller, A. Negrean, H. Blockus, K. C. Gonzalez, F. T. Sparks, A. S. Solis Canales, A. L. Tuttman, D. S. Peterka, B. V. Zemelman, F. Polleux, and A. Losonczy. "Local Feedback Inhibition Tightly Controls Rapid Formation of Hippocampal Place Fields". In: *Neuron* 110.5 (2, 2022), 783–794.e6. ISSN: 1097-4199. DOI: 10.1016/j.neuron.2021.12.003. PMID: 34990571.

[78] G. F. Turi, W.-K. Li, S. Chavlis, I. Pandi, J. O'Hare, J. B. Priestley, A. D. Grosmark, Z. Liao, M. Ladow, J. F. Zhang, B. V. Zemelman, P. Poirazi, and A. Losonczy. "Vasoactive Intestinal Polypeptide-Expressing Interneurons in the Hippocampus Support Goal-Oriented Spatial Learning". In: *Neuron* 101.6 (2019), 1150–1165.e8. ISSN: 08966273. DOI: 10.1016/j.neuron.2019.01.009.

[79] S. B. E. Wolff, J. Gründemann, P. Tovote, S. Krabbe, G. A. Jacobson, C. Müller, C. Herry, I. Ehrlich, R. W. Friedrich, J. J. Letzkus, and A. Lüthi. "Amygdala Interneuron Subtypes Control Fear Learning through Disinhibition". In: *Nature* 509.7501 (2014), pp. 453–458. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature13258. PMID: 24814341.

[80] K. P. Körding and P. König. "Supervised and Unsupervised Learning with Two Sites of Synaptic Integration". In: *Journal of Computational Neuroscience* 11.3 (1, 2001), pp. 207–215. ISSN: 1573-6873. DOI: 10.1023/A:1013776130161.

[81] A. Meulemans, N. Zucchet, S. Kobayashi, J. von Oswald, and J. Sacramento. "The Least-Control Principle for Local Learning at Equilibrium". In: *Advances in Neural Information Processing Systems*. Ed. by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh. Vol. 35. Curran Associates, Inc., 2022, pp. 33603–33617.

[82] W. Senn, D. Dold, A. F. Kungl, B. Ellenberger, J. Jordan, Y. Bengio, J. Sacramento, and M. A. Petrovici. "A Neuronal Least-Action Principle for Real-Time Learning in Cortical Circuits". In: *eLife* 12 (20, 2024). Ed. by S. Ostojic and P. Poirazi, RP89674. ISSN: 2050-084X. DOI: 10.7554/eLife.89674.

[83] P. V. Aceituno, S. de Haan, R. Loidl, and B. F. Grewe. *Challenging Backpropagation: Evidence for Target Learning in the Cortex*. 29, 2024. DOI: 10.1101/2024.04.10.588837. URL: https://www.biorxiv.org/content/10.1101/2024.04.10.588837v4 (visited on 04/29/2025). Pre-published.

[84] W. Greedy, H. W. Zhu, J. Pemberton, J. Mellor, and R. Ponte Costa. "Single-Phase Deep Learning in Cortico-Cortical Networks". In: *Advances in Neural Information Processing Systems*. Ed. by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh. Vol. 35. Curran Associates, Inc., 2022, pp. 24213–24225.

[85] A. Alemi, C. Machens, S. Deneve, and J.-J. Slotine. "Learning Nonlinear Dynamics in Efficient, Balanced Spiking Networks Using Local Plasticity Rules". In: *Proceedings of the AAAI Conference on Artificial Intelligence* 32.1 (2018). ISSN: 2374-3468. DOI: 10.1609/aaai.v32i1.11320.

[86] B. Podlaski and C. K. Machens. "Biological Credit Assignment through Dynamic Inversion of Feedforward Networks". In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin. Vol. 33. Curran Associates, Inc., 2020, pp. 10065–10076.

[87] A. Meulemans, M. T. Farinha, M. R. Cervera, J. Sacramento, and B. F. Grewe. "Minimizing Control for Credit Assignment with Strong Feedback". In: *Proceedings of the 39th International Conference on Machine Learning*. Ed. by K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato. Vol. 162. Proceedings of Machine Learning Research. PMLR, 17, 2022, pp. 15458–15483.

[88] P. V. Aceituno, M. T. Farinha, R. Loidl, and B. F. Grewe. "Learning Cortical Hierarchies with Temporal Hebbian Updates". In: *bioRxiv* (2023). DOI: 10.1101/2023.01.02.522459.

[89] E. L. Bienenstock, L. N. Cooper, and P. W. Munro. "Theory for the Development of Neuron Selectivity: Orientation Specificity and Binocular Interaction in Visual Cortex". In: *The Journal of neuroscience* 2.1 (1982), pp. 32–48. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.02-01-00032.1982. PMID: 7054394.

[90] E. Oja. "A Simplified Neuron Model as a Principal Component Analyzer". In: *Journal of mathematical biology* 15.3 (1982), pp. 267–273. ISSN: 0303-6812. DOI: 10.1007/bf00275687. PMID: 7153672.

[91] S. Lim, J. L. McKee, L. Woloszyn, Y. Amit, D. J. Freedman, D. L. Sheinberg, and N. Brunel. "Inferring Learning Rules from Distributions of Firing Rates in Cortical Neurons". In: *Nature neuroscience* 18.12 (2015), pp. 1804–1810. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/nn.4158. PMID: 26523643.

[92] J. Gjorgjieva, C. Clopath, J. Audet, and J.-P. Pfister. "A Triplet Spike-Timing–Dependent Plasticity Model Generalizes the Bienenstock–Cooper–Munro Rule to Higher-Order Spatiotemporal Correlations". In: *Proceedings of the National Academy of Sciences* 108.48 (29, 2011), pp. 19383–19388. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.1105933108.

[93] M. Boerlin, C. K. Machens, and S. Denève. "Predictive Coding of Dynamical Variables in Balanced Spiking Networks". In: *Plos Computational Biology* 9.11 (2013), e1003258.

[94] F. Mastrogiuseppe and S. Ostojic. "Linking Connectivity, Dynamics and Computations in Low-Rank Recurrent Neural Networks". In: *Neuron* 99.3 (2018), 609–623.e29. ISSN: 08966273. DOI: 10.1016/j.neuron.2018.07.003. arXiv: 1711.09672 [q-bio].

[95] B. DePasquale, D. Sussillo, L. F. Abbott, and M. M. Churchland. "The Centrality of Population-Level Factors to Network Computation Is Demonstrated by a Versatile Approach for Training Spiking Networks". In: *Neuron* 111.5 (1, 2023), 631–649.e10. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2022.12.007.

[96] M. Hanut and J. Kadmon. *Training Large Neural Networks With Low-Dimensional Error Feedback*. 20, 2025. DOI: 10.48550/arXiv.2502.20580. arXiv: 2502.20580 [cs]. URL: http://arxiv.org/abs/2502.20580 (visited on 04/09/2025). Pre-published.

[97] C. Meissner-Bernard, F. Zenke, and R. W. Friedrich. "Geometry and Dynamics of Representations in a Precisely Balanced Memory Network Related to Olfactory Cortex". In: *eLife* 13 (13, 2025). Ed. by J. Gjorgjieva and P. Poirazi, RP96303. ISSN: 2050-084X. DOI: 10.7554/eLife.96303.

[98] C. Q. Chiu, G. Lur, T. M. Morse, N. T. Carnevale, G. C. R. Ellis-Davies, and M. J. Higley. "Compartmentalization of GABAergic Inhibition by Dendritic Spines". In: *Science (New York, N.Y.)* 340.6133 (10, 2013), pp. 759–762. ISSN: 1095-9203. DOI: 10.1126/science.1234274. PMID: 23661763.

[99] D. B. Dorman, J. Jędrzejewska-Szmek, and K. T. Blackwell. "Inhibition Enhances Spatially-Specific Calcium Encoding of Synaptic Input Patterns in a Biologically Constrained Model". In: *eLife* 7 (25, 2018). Ed. by M. B. Kennedy and E. Marder, e38588. ISSN: 2050-084X. DOI: 10.7554/eLife.38588.

[100] C. Lüscher and R. C. Malenka. "NMDA Receptor-Dependent Long-Term Potentiation and Long-Term Depression (LTP/LTD)". In: *Cold Spring Harbor Perspectives in Biology* 4.6 (2012), a005710. ISSN: 1943-0264. DOI: 10.1101/cshperspect.a005710. PMID: 22510460.

[101] J. R. Chalifoux and A. G. Carter. "GABAB Receptors Modulate NMDA Receptor Calcium Signals in Dendritic Spines". In: *Neuron* 66.1 (15, 2010), pp. 101–113. ISSN: 1097-4199. DOI: 10.1016/j.neuron.2010.03.012. PMID: 20399732.

[102] J. R. Chalifoux and A. G. Carter. "GABAB Receptor Modulation of Synaptic Function". In: *Current Opinion in Neurobiology* 21.2 (2011), pp. 339–344. ISSN: 1873-6882. DOI: 10.1016/j.conb.2011.02.004. PMID: 21376567.

[103] J. R. Chalifoux and A. G. Carter. "GABAB Receptor Modulation of Voltage-Sensitive Calcium Channels in Spines and Dendrites". In: *Journal of Neuroscience* 31.11 (16, 2011), pp. 4221–4232. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.4561-10.2011. PMID: 21411663.

[104] K. Foley, C. McKee, A. C. Nairn, and H. Xia. "Regulation of Synaptic Transmission and Plasticity by Protein Phosphatase 1". In: *The Journal of Neuroscience* 41.14 (7, 2021), pp. 3040–3050. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.2026-20.2021. PMID: 33827970.

[105] C. Bernard, D. Exposito-Alonso, M. Selten, S. Sanalidou, A. Hanusz-Godoy, A. Aguilera, F. Hamid, F. Oozeer, P. Maeso, L. Allison, M. Russell, R. A. Fleck, B. Rico, and O. Marín. "Cortical Wiring by Synapse Type–Specific Control of Local Protein Synthesis". In: *Science* 378.6622 (25, 2022), eabm7466. DOI: 10.1126/science.abm7466.

[106] C. Clopath, L. Ziegler, E. Vasilaki, L. Büsing, and W. Gerstner. "Tag-Trigger-Consolidation: A Model of Early and Late Long-Term-Potentiation and Depression". In: *PLoS Computational Biology* 4.12 (26, 2008), e1000248. ISSN: 1553-734X. DOI: 10.1371/journal.pcbi.1000248. PMID: 19112486.

[107] E. P. Huang. "Synaptic Plasticity: Going through Phases with LTP". In: *Current Biology* 8.10 (7, 1998), R350–R352. ISSN: 0960-9822. DOI: 10.1016/S0960-9822(98)70219-2.

[108] J. Kolen and J. Pollack. "Backpropagation without Weight Transport". In: *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*. 1994 IEEE International Conference on Neural Networks (ICNN'94). Vol. 3. Orlando, FL, USA: IEEE, 1994, pp. 1375–1380. ISBN: 978-0-7803-1901-1. DOI: 10.1109/ICNN.1994.374486.

[109] M. Akrout, C. Wilson, P. Humphreys, T. Lillicrap, and D. B. Tweed. "Deep Learning without Weight Transport". In: *Advances in Neural Information Processing Systems*. Vol. 32. Curran Associates, Inc., 2019.

[110] Y. Amit. "Deep Learning With Asymmetric Connections and Hebbian Updates". In: *Frontiers in Computational Neuroscience* 13 (4, 2019). ISSN: 1662-5188. DOI: 10.3389/fncom.2019.00018.

[111] J. Lindsey and A. Litwin-Kumar. "Learning to Learn with Feedback and Local Plasticity". In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin. Vol. 33. Curran Associates, Inc., 2020, pp. 21213–21223.

[112] A. Laborieux and F. Zenke. *Improving Equilibrium Propagation without Weight Symmetry through Jacobian Homeostasis*. 8, 2024. DOI: 10.48550/arXiv.2309.02214. arXiv: 2309.02214 [cs]. URL: http://arxiv.org/abs/2309.02214 (visited on 04/14/2025). Pre-published.

[113] K. Max, L. Kriener, G. Pineda García, T. Nowotny, I. Jaras, W. Senn, and M. A. Petrovici. "Learning Efficient Backprojections across Cortical Hierarchies in Real Time". In: *Nature Machine Intelligence* 6.6 (2024), pp. 619–630. ISSN: 2522-5839. DOI: 10.1038/s42256-024-00845-3.

[114] L. Ji-An and M. K. Benna. *Deep Learning without Weight Symmetry*. 31, 2024. DOI: 10.48550/arXiv.2405.20594. arXiv: 2405.20594 [cs]. URL: http://arxiv.org/abs/2405.20594 (visited on 02/17/2025). Pre-published.

[115] C. K. Pfeffer, M. Xue, M. He, Z. J. Huang, and M. Scanziani. "Inhibition of Inhibition in Visual Cortex: The Logic of Connections between Molecularly Distinct Interneurons". In: *Nature Neuroscience* 16.8 (2013), pp. 1068–1076. ISSN: 1546-1726. DOI: 10.1038/nn.3446.

[116] J. J. Letzkus, S. B. E. Wolff, and A. Lüthi. "Disinhibition, a Circuit Mechanism for Associative Learning and Memory". In: *Neuron* 88.2 (2015), pp. 264–276. ISSN: 0896-6273, 1097-4199. DOI: 10.1016/j.neuron.2015.09.024. PMID: 26494276.

[117] P. Földiák. "Forming Sparse Representations by Local Anti-Hebbian Learning". In: *Biological Cybernetics* 64.2 (1990), pp. 165–170.

[118] J. Zylberberg, J. T. Murphy, and M. R. DeWeese. "A Sparse Coding Model with Synaptically Local Plasticity and Spiking Neurons Can Account for the Diverse Shapes of V1 Simple Cell Receptive Fields". In: *Plos Computational Biology* 7.10 (2011), e1002250.

[119] A. Litwin-Kumar and B. Doiron. "Formation and Maintenance of Neuronal Assemblies through Synaptic Plasticity". In: *Nature Communications* 5 (14, 2014), p. 5319. ISSN: 2041-1723. DOI: 10.1038/ncomms6319. PMID: 25395015.

[120] F. Lagzi, M. C. Bustos, A.-M. Oswald, and B. Doiron. *Assembly Formation Is Stabilized by Parvalbumin Neurons and Accelerated by Somatostatin Neurons*. 7, 2021. DOI: 10.1101/2021.09.06.459211. URL: https://www.biorxiv.org/content/10.1101/2021.09.06.459211v1 (visited on 04/30/2025). Pre-published.

[121] F. A. Mikulasch, L. Rudelt, and V. Priesemann. "Local Dendritic Balance Enables Learning of Efficient Representations in Networks of Spiking Neurons". In: *Proceedings of the National Academy of Sciences* 118.50 (14, 2021), e2021925118. DOI: 10.1073/pnas.2021925118.

[122] M. C. D. Bridi, F.-J. Zong, X. Min, N. Luo, T. Tran, J. Qiu, D. Severin, X.-T. Zhang, G. Wang, Z.-J. Zhu, K.-W. He, and A. Kirkwood. "Daily Oscillation of the Excitation-Inhibition Balance in Visual Cortical Circuits". In: *Neuron* 105.4 (19, 2020), 621–629.e4. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2019.11.011.

[123] H. C. Barron, T. P. Vogels, T. E. Behrens, and M. Ramaswami. "Inhibitory Engrams in Perception and Memory". In: *Proceedings of the National Academy of Sciences of the United States of America* 114.26 (27, 2017), pp. 6666–6674. ISSN: 1091-6490. DOI: 10.1073/pnas.1701812114. PMID: 28611219.

[124] K. A. Wilmes and C. Clopath. "Inhibitory Microcircuits for Top-down Plasticity of Sensory Representations". In: *Nature Communications* 10.1 (7, 2019), p. 5055. ISSN: 2041-1723. DOI: 10.1038/s41467-019-12972-2.

[125] J. Rossbroich and F. Zenke. “Dis-Inhibitory Neuronal Circuits Can Control the Sign of Synaptic Plasticity”. In: *Neural Information Processing Systems* abs/2310.19614 (2023). DOI: 10.48550/arXiv.2310.19614. eprint: 2310.19614.

[126] F. Zenke and W. Gerstner. “Hebbian Plasticity Requires Compensatory Processes on Multiple Timescales”. In: *Philosophical Transactions of the Royal Society B* 372.1715 (2017), p. 20160259.

[127] F. Zenke, G. Hennequin, and W. Gerstner. “Synaptic Plasticity in Neural Networks Needs Homeostasis with a Fast Rate Detector”. In: *Plos Computational Biology* 9.11 (2013), e1003330.

[128] P. Poirazi and B. W. Mel. “Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue”. In: *Neuron* 29.3 (1, 2001), pp. 779–796. ISSN: 0896-6273. DOI: 10.1016/S0896-6273(01)00252-5. PMID: 11301036.

[129] M. Pagkalos, R. Makarov, and P. Poirazi. “Leveraging Dendritic Properties to Advance Machine Learning and Neuro-Inspired Computing”. In: *Current Opinion in Neurobiology* 85 (1, 2024), p. 102853. ISSN: 0959-4388. DOI: 10.1016/j.conb.2024.102853.

[130] S. Chavlis and P. Poirazi. “Dendrites Endow Artificial Neural Networks with Accurate, Robust and Parameter-Efficient Learning”. In: *Nature Communications* 16.1 (22, 2025), p. 943. ISSN: 2041-1723. DOI: 10.1038/s41467-025-56297-9.

[131] D. M. Iascone, Y. Li, U. Sümbül, M. Doron, H. Chen, V. Andreu, F. Goudy, H. Blockus, L. F. Abbott, I. Segev, H. Peng, and F. Polleux. “Whole-Neuron Synaptic Mapping Reveals Spatially Precise Excitatory/Inhibitory Balance Limiting Dendritic and Somatic Spiking”. In: *Neuron* 106.4 (2020), 566–578.e8. ISSN: 08966273. DOI: 10.1016/j.neuron.2020.02.015.

[132] Y. K. Wu, C. Miehl, and J. Gjorgjieva. “Regulation of Circuit Organization and Function through Inhibitory Synaptic Plasticity”. In: *Trends in Neurosciences* 45.12 (1, 2022), pp. 884–898. ISSN: 0166-2236, 1878-108X. DOI: 10.1016/j.tins.2022.10.006. PMID: 36404455.

[133] T. P. Vogels and L. F. Abbott. “Gating Multiple Signals through Detailed Balance of Excitation and Inhibition in Spiking Networks”. In: *Nature Neuroscience* 12.4 (2009), pp. 483–491. ISSN: 1546-1726. DOI: 10.1038/nn.2276. PMID: 19305402.

[134] J. Kremkow, A. Aertsen, and A. Kumar. “Gating of Signal Propagation in Spiking Neural Networks by Balanced and Correlated Excitation and Inhibition”. In: *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 30.47 (24, 2010), pp. 15760–15768. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.3874-10.2010. PMID: 21106815.

[135] F. I. Kleberg, T. Fukai, and M. Gilson. “Excitatory and Inhibitory STDP Jointly Tune Feedforward Neural Circuits to Selectively Propagate Correlated Spiking Activity”. In: *Frontiers in Computational Neuroscience* 8 (2014), p. 53. ISSN: 1662-5188. DOI: 10.3389/fncom.2014.00053. PMID: 24847242.

[136] A. Silbaugh, K. P. Koster, and C. Hansel. *Cerebellum Instructs Plasticity in the Mouse Primary Somatosensory Cortex*. 21, 2025. DOI: 10.1101/2025.04.18.649578. URL: https://www.biorxiv.org/content/10.1101/2025.04.18.649578v1 (visited on 04/21/2025). Pre-published.

[137] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. *JAX: Composable Transformations of Python+NumPy Programs*. Version 0.3.13. 2018.

[138] J. Heek, A. Levskaya, A. Oliver, M. Ritter, B. Rondepierre, A. Steiner, and M. van Zee. *Flax: A Neural Network Library and Ecosystem for JAX*. Version 0.6.9. 2023.

[139] P. Kidger. “On Neural Differential Equations”. PhD thesis. University of Oxford, 2021.

[140] *TensorFlow Datasets, A Collection of Ready-to-Use Datasets*.

[141] C. Tsitouras. “Runge–Kutta Pairs of Order 5(4) Satisfying Only the First Column Simplifying Assumption”. In: *Computers & mathematics with applications* 62.2 (2011), pp. 770–775. ISSN: 0898-1221. DOI: 10.1016/j.camwa.2011.06.002.

[142] X. Zhang, W. Guan, T. Yang, A. Furlan, X. Xiao, K. Yu, X. An, W. Galbavy, C. Ramakrishnan, K. Deisseroth, K. Ritola, A. Hantman, M. He, Z. Josh Huang, and B. Li. "Genetically Identified Amygdala–Striatal Circuits for Valence-Specific Behaviors". In: *Nature Neuroscience* 24.11 (2021), pp. 1586–1600. ISSN: 1546-1726. DOI: 10.1038/s41593-021-00927-0.

[143] G. Hennequin, T. P. Vogels, and W. Gerstner. "Optimal Control of Transient Dynamics in Balanced Networks Supports Generation of Complex Movements". In: *Neuron* 82.6 (18, 2014), pp. 1394–1406. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2014.04.045. PMID: 24945778.

Methods

Neuronal dynamics

Throughout this article, we simulated networks of N_{E} excitatory and $N_{\mathrm{I}} = {}^{N_{\mathrm{E}}}/_{4}$ Type 2 inhibitory neurons with reciprocal connections. We write u_i^{E} for the membrane potential of excitatory neuron i and u_i^{I} for the i-th inhibitory neuron. These membrane potentials evolved according to the following differential equations

$$\tau_{\mathrm{E}} \frac{du_i^{\mathrm{E}}}{dt} = -u_i^{\mathrm{E}} + R\left(I_i^{\mathrm{exc}} - I_i^{\mathrm{inh1}} - I_i^{\mathrm{inh2}}\right), \tag{5}$$

$$\tau_{\mathrm{I}} \frac{du_i^{\mathrm{I}}}{dt} = -u_i^{\mathrm{I}} + R\left(I_i^{\mathrm{exc}} - I_i^{\mathrm{inh1}} + I_i^{\mathrm{exc,rec}} + I_i^{\mathrm{fb}}\right), \tag{6}$$

where I_i^{exc} and I_i^{inh1} denote the excitatory and Type 1 inhibitory currents originating from external input populations, I_i^{inh2} are Type 2 inhibitory currents, $I_i^{\mathrm{exc,rec}}$ are recurrent excitatory currents and I_i^{fb} are top-down feedback currents to Type 2 inhibitory interneurons that can be either excitatory or inhibitory. Finally, τ_{E} and τ_{I} are the membrane time constants for excitatory and inhibitory neurons, respectively. For ease of notation we modeled synaptic currents in natural units assuming an input resistance $R = 1$.

Synaptic connections. Each input current was modeled as a weighted sum over presynaptic activities. In general, we denote the synaptic weights from presynaptic neuron j in population A to postsynaptic neuron i in population B as $w_{ij}^{a\to b}$. Neurons received excitatory and Type 1 (feed-forward) inhibitory inputs from a population of N_{X} neurons with activities $\{x_j\}_{j=1}^{N_{\mathrm{X}}}$. For ease of analysis, we assume instantaneous Type 1 feed-forward inhibition. In other words, we treated Type 1 inhibitory neurons as part of the input population and decomposed the input weights into excitatory and inhibitory components

$$w_{ij}^{\mathrm{X}\to b,+} \quad \text{and} \quad w_{ij}^{\mathrm{X}\to b,-}, \quad b \in \{\mathrm{E},\mathrm{I}\}, \tag{7}$$

to denote the excitatory and inhibitory weights, respectively. Accordingly, the dynamics become

$$\tau_{\mathrm{E}} \frac{du_i^{\mathrm{E}}}{dt} = -u_i^{\mathrm{E}} + \sum_{j=1}^{N_{\mathrm{X}}} \left[w_{ij}^{\mathrm{X}\to\mathrm{E},+} - w_{ij}^{\mathrm{X}\to\mathrm{E},-} \right] x_j - \sum_{j=1}^{N_{\mathrm{I}}} w_{ij}^{\mathrm{I}\to\mathrm{E}} r_j^{\mathrm{I}}, \tag{8}$$

$$\tau_{\mathrm{I}} \frac{du_i^{\mathrm{I}}}{dt} = -u_i^{\mathrm{I}} + \sum_{j=1}^{N_{\mathrm{X}}} \left[w_{ij}^{\mathrm{X}\to\mathrm{I},+} - w_{ij}^{\mathrm{X}\to\mathrm{I},-} \right] x_j + \sum_{j=1}^{N_{\mathrm{E}}} w_{ij}^{\mathrm{E}\to\mathrm{I}} r_j^{\mathrm{E}} + I_i^{\mathrm{fb}}, \tag{9}$$

where r_j^{E} and r_j^{I} denote instantaneous firing rates that were computed as

$$r_j^{\mathrm{E}} = \phi(u_j^{\mathrm{E}}), \quad r_j^{\mathrm{I}} = \phi(u_j^{\mathrm{I}}), \tag{10}$$

with the rectifying nonlinearity $\phi(u) = \max\{0, u\}$. To simplify further analysis, we combine the positive and negative feedforward weights into a single weight matrix as follows:

$$w_{ij}^{\mathrm{X}\to\mathrm{E}} \equiv w_{ij}^{\mathrm{X}\to\mathrm{E},+} - w_{ij}^{\mathrm{X}\to\mathrm{E},-}, \quad w_{ij}^{\mathrm{X}\to\mathrm{I}} \equiv w_{ij}^{\mathrm{X}\to\mathrm{I},+} - w_{ij}^{\mathrm{X}\to\mathrm{I},-}. \tag{11}$$

With these definitions, the dynamics of the E/I assembly populations simplify to

$$\tau_{\mathrm{E}} \frac{du_i^{\mathrm{E}}}{dt} = -u_i^{\mathrm{E}} + \sum_{j=1}^{N_{\mathrm{X}}} w_{ij}^{\mathrm{X}\to\mathrm{E}} x_j - \sum_{j=1}^{N_{\mathrm{I}}} w_{ij}^{\mathrm{I}\to\mathrm{E}} r_j^{\mathrm{I}}, \tag{12}$$

$$\tau_{\mathrm{I}} \frac{du_i^{\mathrm{I}}}{dt} = -u_i^{\mathrm{I}} + \sum_{j=1}^{N_{\mathrm{X}}} w_{ij}^{\mathrm{X}\to\mathrm{I}} x_j + \sum_{j=1}^{N_{\mathrm{E}}} w_{ij}^{\mathrm{E}\to\mathrm{I}} r_j^{\mathrm{E}} + I_i^{\mathrm{fb}}, \tag{13}$$

or, in vectorized notation,

$$\tau_{\mathrm{E}} \frac{d\mathbf{u}^{\mathrm{E}}}{dt} = -\mathbf{u}^{\mathrm{E}} + \mathbf{W}^{\mathrm{X}\to\mathrm{E}}\mathbf{x} - \mathbf{W}^{\mathrm{I}\to\mathrm{E}}\mathbf{r}^{\mathrm{I}}, \tag{14}$$

$$\tau_{\mathrm{I}} \frac{d\mathbf{u}^{\mathrm{I}}}{dt} = -\mathbf{u}^{\mathrm{I}} + \mathbf{W}^{\mathrm{X}\to\mathrm{I}}\mathbf{x} + \mathbf{W}^{\mathrm{E}\to\mathrm{I}}\mathbf{r}^{\mathrm{E}} + \mathbf{I}^{\mathrm{fb}}. \tag{15}$$

Readout neurons. Unless otherwise stated, we modeled the output of the network as a group of N_{out} linear readout neurons with the following dynamics:

$$\tau_{\text{out}} \frac{du_i^{\text{out}}}{dt} = -u_i^{\text{out}} + \sum_{j=1}^{N_{\text{E}}} w_{ij}^{\text{E}\to\text{out}} r_j^{\text{E}} \tag{16}$$

where τ_{out} is the readout time constant and number of readout neurons N_{out} depends on the specific task (see below).

E/I Assembly structure

Neuronal E/I assemblies are the basic computational element in our model. The population of excitatory and Type 2 inhibitory neurons was divided into N_{A} assemblies, each containing N_{A}^{E} excitatory neurons and N_{A}^{I} inhibitory neurons. The membership of neurons to these assemblies was specified by the membership matrices $\mathbf{M}^{\text{E}}$ and $\mathbf{M}^{\text{I}}$. In particular, for both membership matrices, the element

$$m_{ij} = \begin{cases} 0, & \text{if neuron } i \text{ is not part of assembly } j, \\ > 0, & \text{if neuron } i \text{ is part of assembly } j, \end{cases}$$

with the positive value quantifying the relative contribution of neuron i to assembly j.

Within-assembly excitatory synaptic weights. A neuron's assembly memberships determine its recurrent connections within the network. We considered fixed E → I weights so that the recurrent weight from excitatory neuron j to Type 2 inhibitory neuron i is given by

$$w_{ij}^{\text{E}\to\text{I}} = \frac{1}{N_{\text{A}}^{\text{E}}} \sum_{k=1}^{N_{\text{A}}} m_{jk}^{\text{E}} m_{ik}^{\text{I}}, \tag{17}$$

where m_{ik}^{a} denotes the membership of neuron i in population $a \in \{\text{E}, \text{I}\}$ to assembly k. This led to a connectivity structure,in which each inhibitory neuron receives recurrent input from excitatory neurons belonging to the same assembly, with the connection strength determined by the degree of membership of both neurons in that assembly.

Parameterization of feed-forward synaptic weights. Our model requires co-tuned E/I assemblies which we assumed as given, for instance, as the outcome of developmental or learning processes [59, 60, 62, 118]. Specifically, we assumed that neurons belonging to the same assembly share the same feed-forward tuning. Consequently, feed-forward weights to assemblies were the learnable parameters

$$\mathbf{V}^{\text{X}\to\text{A}} \in \mathbb{R}^{N_{\text{A}} \times N_{\text{X}}}, \tag{18}$$

which map inputs to assembly space, in conjunction with the respective membership matrices. Specifically, this parameterization is expressed as

$$w_{ij}^{\text{X}\to\text{E}} = \sum_{k=1}^{N_{\text{A}}} m_{ik}^{\text{E}} v_{kj}^{\text{X}\to\text{A}}, \quad w_{ij}^{\text{X}\to\text{I}} = \sum_{k=1}^{N_{\text{A}}} m_{ik}^{\text{I}} v_{kj}^{\text{X}\to\text{A}}. \tag{19}$$

Similarly, we assumed that network readout neurons receive assembly-specific connections given by the learnable parameters

$$\mathbf{V}^{\text{A}\to\text{out}} \in \mathbb{R}^{N_{\text{out}} \times N_{\text{A}}}, \tag{20}$$

that determine the readout weights

$$w_{ij}^{\text{E}\to\text{out}} = \sum_{k=1}^{N_{\text{A}}} m_{jk}^{\text{E}} v_{ik}^{\text{A}\to\text{out}}. \tag{21}$$

Initialization of assembly structure. In simulations with a single assembly (Figs. 2 and 3) or networks with a non-overlapping assembly structure (Figs. 4, 6, 7, 8), each neuron was assigned exclusively to one assembly. To that end, we partitioned the excitatory and inhibitory populations into N_A assemblies such that the numbers of excitatory and inhibitory neurons per assembly are given by

$$N_\mathrm{A}^\mathrm{E} = \frac{N_\mathrm{E}}{N_\mathrm{A}}, \quad N_\mathrm{A}^\mathrm{I} = \frac{N_\mathrm{I}}{N_\mathrm{A}}. \tag{22}$$

For each population $a \in \{\mathrm{E}, \mathrm{I}\}$, with N_a denoting the number of neurons in population a, we defined a binary membership matrix

$$\mathbf{B}^a \in \{0,1\}^{N_a \times N_\mathrm{A}},$$

where

$$b_{ij}^a = \begin{cases} 1, & \text{if neuron } i \text{ is assigned to assembly } j, \\ 0, & \text{otherwise.} \end{cases}$$

In the non-overlapping case, each neuron was a member of exactly one assembly. To account for heterogeneous contributions within an assembly, we created a weight matrix

$$\mathbf{L}^a \in \mathbb{R}^{N_a \times N_\mathrm{A}},$$

with entries independently sampled from a lognormal distribution with parameters μ_memb and σ_memb. The final membership matrix was then defined via the elementwise product,

$$\mathbf{M}^a = \mathbf{B}^a \circ \mathbf{L}^a, \quad a \in \{\mathrm{E}, \mathrm{I}\}. \tag{23}$$

Initialization of overlapping assemblies. In networks with overlapping assemblies (Figs. 5 and 6), we started from the binary matrices $\mathbf{B}^a$ defined above for $a \in \{\mathrm{E}, \mathrm{I}\}$. We then introduced additional memberships by randomly setting extra entries to 1, yielding an augmented binary matrix $\mathbf{B}_\mathrm{ov}^a$. Since neurons may now have multiple memberships, we normalized each row of $\mathbf{B}_\mathrm{ov}^a$ to ensure that the total membership weight per neuron equals one:

$$\tilde{b}_{ij}^a = \frac{b_{\mathrm{ov},ij}^a}{\sum_{j=1}^{N_\mathrm{A}} b_{\mathrm{ov},ij}^a}, \quad i = 1, \ldots, N_a. \tag{24}$$

As before, we generated a weight matrix

$$\mathbf{L}^a \in \mathbb{R}^{N_a \times N_\mathrm{A}},$$

with lognormally distributed entries and defined the final membership matrix as

$$\mathbf{M}^a = \tilde{\mathbf{B}}^a \circ \mathbf{L}^a, \quad a \in \{\mathrm{E}, \mathrm{I}\}. \tag{25}$$

Quantification of assembly overlap. To quantify the degree of assembly overlap, we considered the ratio of additional entries in the augmented binary matrices $\mathbf{B}_\mathrm{ov}^a$ relative to the total number of entries in these matrices. In particular, we define the overlap ratio as

$$\text{Overlap} = \frac{\sum_{i=1}^{N_a} \sum_{j=1}^{N_\mathrm{A}} \left(b_{\mathrm{ov},ij}^a - b_{ij}^a \right)}{\sum_{i=1}^{N_a} \sum_{j=1}^{N_\mathrm{A}} b_{\mathrm{ov},ij}^a}, \tag{26}$$

which ranges from 0 (no overlap) to values approaching 1 (very high overlap). In the main text, the overlap ratio is expressed as a percentage.

Optimization of Type 2 inhibitory connections for E/I balance Throughout this paper, we required that E/I assemblies operate in a precise E/I balance regime (cf. Eq. (2)), in which Type 2 inhibitory currents maintain a proportional relationship with the net feed-forward activity. In contrast to previous computational studies that investigated how E/I balance can be established and homeostatically maintained through inhibitory plasticity mechanisms [25, 56, 57, 59], we chose a mechanism-agnostic approach in this article. To that end, we directly optimized the Type 2 inhibitory I → E synaptic weights to satisfy the balance condition in equation (2) according to

$$\min_{\mathbf{W}^{\mathrm{I}\to\mathrm{E}}} \mathscr{L}_{\text{balance}} \quad \text{with} \quad \mathscr{L}_{\text{balance}} = \sum_i \|\Delta_i\|^2, \tag{27}$$

using standard gradient descent methods (see Supplementary Material Section 1.1). Minimizing the balance error Δ for each neuron at the time of network initialization ensured precise balance regardless of the network inputs that persisted during learning without the need for continuous optimization.

Top-down error feedback

In our model, top-down feedback to Type 2 interneurons provides an adaptive control input that changes activity to minimize a loss at the network output [5, 81, 87]. We modeled the top-down currents $\mathbf{I}^{\mathrm{fb}}$ as feedback control currents from a controller population $\mathbf{c}$ through excitatory or inhibitory feedback afferents with weights $\mathbf{Q}$. Importantly, our framework makes little assumptions about the actual dynamics of the controller, its implementation, and the nature of the feedback weights, as long as the control currents lead to reducing the loss function (see also [9, 81]). In this article, we considered simple leaky proportional-integral controllers that follow the dynamics

$$\mathbf{c}(t) = k_p \mathbf{e}(t) + k_i \mathbf{c}^{\mathrm{int}}(t) \quad , \quad \tau_c \frac{d\mathbf{c}^{\mathrm{int}}}{dt} = -\mathbf{c}^{\mathrm{int}} + \mathbf{e}(t) \tag{28}$$

where $\mathbf{e}(t) = -\nabla_{\mathbf{u}^{\mathrm{out}}} \mathscr{L}(\mathbf{u}^{\mathrm{out}})$ is the magnitude of the error at the network output, k_p is the proportional gain, k_i is the integral gain, and τ_c is the time constant of the integral controller.

Feedback weights. The feedback weights $\mathbf{Q}$ describe the effect of the controller on the inhibitory interneurons. In this article, we only considered direct linear feedback from the controller to the inhibitory interneurons, which could be interpreted as direct feedback projections from another brain area. However, other forms of feedback involving local mechanisms or non-linear pathways are possible [81]. Following previous work [81, 87] we chose a direct linear feedback mapping that approximates the network Jacobian (see Supplementary Material Section 1.2), such that

$$\mathbf{Q} = \mathbf{D}^{\mathrm{I}} \left(\mathbf{M}^{\mathrm{I}}\right)^T \left(\mathbf{V}^{\mathrm{A}\to\mathrm{out}}\right)^T \tag{29}$$

where $\mathbf{D}^{\mathrm{I}}$ is a diagonal matrix containing the derivative of the rectified nonlinearity for each inhibitory neuron so that $\mathbf{D}^{\mathrm{I}}_{ii} = \phi'(u^{\mathrm{I}}_i)$. The inclusion of this derivative term ensured that top-down feedback only targeted active inhibitory neurons that operate in the linear regime.

Balance controlled synaptic plasticity

Given the E/I assembly structure and top-down feedback connectivity discussed above, the BCP learning rule can be derived from a minimization of control objective function (Supplementary Material, Section 1.3). In brief, the idea is that learning should reduce the amount of control the controller has to exert to minimize the output error. Given linearized Type 2 inhibitory dynamics, we obtain the BCP learning rule for feed-forward weights $\mathbf{V}^{\mathrm{X}\to\mathrm{A}}$ as the negative gradient of the least control loss function $\mathscr{H} = ||\mathbf{W}^{\mathrm{I}\to\mathrm{E}}\mathbf{Q}\,\mathbf{c}||^2$:

$$\nabla_{\mathbf{V}^{\mathrm{X}\to\mathrm{A}}} \mathscr{H}(\mathbf{V}^{\mathrm{X}\to\mathrm{A}}) = -\left(\frac{\partial f(\mathbf{u}^{\mathrm{E}}, \mathbf{V}^{\mathrm{X}\to\mathrm{A}})}{\partial \mathbf{V}^{\mathrm{X}\to\mathrm{A}}}\right)^T \left(\mathbf{W}^{\mathrm{I}\to\mathrm{E}}\mathbf{Q}\mathbf{c}\right), \tag{30}$$

which yields the following gradient-based update rule for the feed-forward parameters $\mathbf{V}^{\mathrm{X}\to\mathrm{A}}$:

$$\Delta v_{kj} = \frac{\eta_{\mathrm{ff}}}{N^{\mathrm{E}}_{\mathrm{A}}} \sum_i^{N^{\mathrm{E}}} m^{\mathrm{E}}_{ik}\, x_j \left(\alpha \left(I^{\mathrm{exc}}_i - I^{\mathrm{inh1}}_i\right)_+ - I^{\mathrm{inh2}}_i\right), \tag{31}$$

Online learning. Importantly, the derivation from a least control objective assumes that neuronal dynamics have converged to an equilibrium state for a given input x_j. While this is possible for static inputs (Fig. 6), during online learning with dynamic input stimuli (Figs. 3, 4, 5, 7 and 8), this condition is not formally met. That is, the temporal dynamics of Type 2 inhibitory neurons and the recurrent E/I assemblies cause a temporal delay of Type 2 inhibition with respect to feed-forward input currents. To compensate for this delay in the continuous-time version BCP, we implemented short-term filtering of feed-forward synaptic currents and presynaptic activity. The resulting learning rule is:

$$\frac{dv_{kj}}{dt} = \frac{\eta_{\text{ff}}}{N_{\text{A}}^{\text{E}}} \sum_{i}^{N^{\text{E}}} m_{ik}^{\text{E}} \, \bar{x}_j \left(\alpha \left(\bar{I}_i^{\text{exc}} - \bar{I}_i^{\text{inh1}} \right)_+ - I_i^{\text{inh2}} \right), \tag{32}$$

where short-term moving averages of the feed-forward currents were implemented as leaky integrators

$$\tau_{\text{adjusted}} \frac{d\bar{I}_i}{dt} = -\bar{I}_i + I_i. \tag{33}$$

Here we choose the filtering time constant simply as

$$\tau_{\text{adjusted}} = \min(\tau_E, \tau_I), \tag{34}$$

which matches the fastest network mode under slow inputs (see Supplementary Material 1.4).

Similarly, the presynaptic activity is filtered to obtain the moving average $\bar{x_j}$ using a leaky integrator with time constant τ_{pre}:

$$\tau_{\text{pre}} \frac{d\bar{x}_j}{dt} = -\bar{x}_j + x_j. \tag{35}$$

This temporal filtering of both synaptic currents and presynaptic inputs reduces the temporal mismatch between Type 2 inhibitory currents and feed-forward inputs, thereby stabilizing online learning.

Plasticity at output neurons. We parameterize the readout weights from excitatory neuron j to readout neuron i as

$$w_{ij}^{\text{E}\to\text{out}} = \sum_{k}^{N_{\text{A}}} v_{ik} m_{jk}^{\text{E}}, \tag{36}$$

where v_{ik} is the element i, k of the readout weight matrix $\mathbf{V}^{\text{A}\to\text{out}} \in \mathbb{R}^{N_{\text{out}} \times N_{\text{A}}}$. To update readout weights during learning of target $\mathbf{y}$, we used a simple Delta rule [65] multiplying the presynaptic activity by the difference between the target output and the actual output:

$$\Delta v_{ik} = \frac{\eta_{\text{out}}}{N_{\text{A}}^{\text{E}}} \sum_{j}^{N_{\text{E}}} m_{jk}^{\text{E}} (y_i - u_i^{\text{out}}), \tag{37}$$

where η_{out} is the learning rate for the readout weights. This learning rule is equivalent to the gradient of the squared error loss function $\mathscr{L}(\mathbf{u}^{\text{out}}) = \frac{1}{2}||\mathbf{y} - \mathbf{u}^{\text{out}}||^2$ with respect to the readout weights.

Simulation details

All numerical simulations were implemented using custom software written in Python 3.11.8 and executed on NVIDIA Quadro RTX 5000 GPUs. Software packages employed for numerical simulations included Jax [137], Flax [138], and Diffrax [139]. The Fashion-MNIST dataset was obtained from the TensorFlow datasets library [140]. The code to reproduce our results is available at https://github.com/fmi-basel/balance-controlled-plasticity. All hyperparameters for numerical simulations can be found in Supplementary Tables S1-S5.

Single-assembly experiments

To investigate the learning dynamics in a single assembly (Figs. 2 and 3), we simulated an assembly comprising $N_{\text{E}} = N_{\text{A}}^{\text{E}} = 16$ excitatory neurons and $N_{\text{I}} = N_{\text{A}}^{\text{I}} = 4$ inhibitory neurons.

Dynamic input stimuli. Assembly neurons received feed-forward input from $N_X = 20$ input neurons, each encoding a linear combination of $N_S = 3$ dynamic stimuli (A, B, and C). The temporal activity of each stimulus was modeled as a rectified Ornstein-Uhlenbeck (OU) process. Specifically, for each stimulus i, we simulated a standard OU process $z_i(t)$ according to

$$\frac{dz_i(t)}{dt} = -\frac{1}{\tau_{\mathrm{OU}}}\Big(z_i(t) - \mu_{\mathrm{OU}}\Big) + \sigma_{\mathrm{OU}}\xi_i(t), \tag{38}$$

where $\xi_i(t)$ is Gaussian white noise with zero mean and unit variance, τ_{OU} is the OU time constant, μ_{OU} is the process mean, and σ_{OU} is the fluctuation magnitude. The effective stimulus activity was obtained by thresholding:

$$s_i^{\mathrm{raw}}(t) = \max\Big(z_i(t) - \theta_{\mathrm{OU}}, 0\Big). \tag{39}$$

To generate smooth input signals, the thresholded activity was low-pass filtered using an exponential filter with time constant τ_{filter}, resulting in the corresponding stimulus dynamics:

$$\tau_{\mathrm{filter}}\frac{ds_i(t)}{dt} = -s_i(t) + s_i^{\mathrm{raw}}(t), \tag{40}$$

with initial condition $s_i(0) = s_i^{\mathrm{raw}}(0)$.

Stimulus encoding. Input neurons encoded the stimuli by forming linear combinations based on a Gaussian tiling of the stimulus space with circular boundaries. The activity of input neuron $x_i(t)$ was given by

$$x_i(t) = \sum_{j=1}^{N_S} w_{ij}^{S\to X}\, s_j(t) + b_i, \tag{41}$$

where $w_{ij}^{S\to X}$ denotes the synaptic weight from stimulus j to input neuron i, and b_i is a bias term drawn from a standard lognormal distribution scaled by b_{scale}. The synaptic weights $\mathbf{W}^{S\to X}$ were determined using a Gaussian tiling scheme. In particular, the weight from stimulus j to input neuron i was given by

$$w_{ij} = \exp\left(-\frac{d_{ij}^2}{2\sigma_X^2}\right), \tag{42}$$

where σ_X is the tuning width parameter and d_{ij} denotes the circular distance between the preferred tuning Λ_i of input neuron i and the stimulus location ξ_j. The circular distance was defined as

$$d_{ij} = \min\Big\{|\Lambda_i - \xi_j|,\ (N_S - 1) - |\Lambda_i - \xi_j|\Big\}, \tag{43}$$

with preferred tunings

$$\Lambda_i = \frac{(i-1)(N_S - 1)}{N_X - 1} \tag{44}$$

and stimulus locations $\xi_j = j - 1$. This construction guarantees that the preferred tunings of input neurons are distributed uniformly over a circular domain.

Weight initialization. The feed-forward weights $\mathbf{V}^{X\to A}$ were initialized from a normal distribution with mean 0.2 and standard deviation $^2/_{N_X}$. The output weights were set as

$$v_{ij}^{\mathrm{A}\to\mathrm{out}} = \frac{1}{N_E}, \quad \forall\, i, j. \tag{45}$$

Numerical integration. Input stimulus activity was generated in 15-s chunks using the forward Euler method with a step size of $\Delta t = 1$ ms. Linear interpolation between time steps was used to produce continuous signals. The network dynamics were then integrated using Tsitouras' 5th order explicit Runge-Kutta method [141] with an adaptive step size. Simulations were run continuously for the duration of each experiment.

Manual feedback perturbations. For the manual feedback perturbation experiments (Fig. 2d), network dynamics were simulated for $T = 15$ s while applying two sets of strong external current perturbations I^{ext} to all inhibitory neurons in the assembly. The dynamics of the inhibitory membrane potential were governed by

$$\tau_{\text{I}} \frac{du_i^{\text{I}}}{dt} = -u_i^{\text{I}} + I_i^{\text{exc,ff}} - I_i^{\text{inh,ff}} + I_i^{\text{exc,rec}} + I^{\text{ext}}. \tag{46}$$

Two perturbation phases were administered, each lasting 3 s and separated by a 3-s interval. In the first phase, a positive perturbation of $I^{\text{ext}} = 10$ was applied, while in the second phase a negative perturbation of $I^{\text{ext}} = -10$ was applied.

Trajectory matching task

For the trajectory matching task (see Figs. 4 and 5), we simulated networks with $N_X = 20$ input neurons, a hidden layer consisting of $N_E = 240$ excitatory and $N_I = 60$ inhibitory neurons arranged in $N_A = 20$ assemblies, and an output layer with $N_{\text{Y}} = 2$ output neurons.

Periodic inputs. Input neurons encoded random linear combinations of 20 sinusoidal functions. The activity of the jth sinusoidal function was defined as

$$s_j(t) = A_j \sin\left(2\pi f_j t\right), \quad f_j = \frac{j}{10}, \tag{47}$$

where A_j is a random amplitude drawn from a uniform distribution between 0.1 and 1.0 and f_j is the frequency of the sinusoidal function, parameterized so that each sinusoidal function completes an integer number of periods every 10 seconds. Each input neuron encoded a linear combination of sinusoidal functions so that the activity of the ith input neuron was given by

$$x_i(t) = \sum_{j=1}^{N_X} w_{ij}^{\text{S}\to\text{X}} s_j(t), \tag{48}$$

where $w_{ij}^{\text{S}\to\text{X}}$ are the elements of a mixing matrix $\mathbf{W}^{\text{S}\to\text{X}}$ drawn from a standard normal distribution.

Target trajectory. The target trajectory was generated from a hand-drawn turtle-shaped vector path in the 2D plane. We obtained the target $\mathbf{y}$ by sampling $T_{\text{traj}}/\Delta t$ uniformly spaced points along the path, where Δt is the simulation time step and $T_{\text{traj}} = 5$ seconds is the total duration of the trajectory. The trajectory was centered at the origin and scaled to lie within $[-1, 1]$ in both the X and Y dimensions.

Parameter initialization. Feed-forward weights $\mathbf{V}^{\text{X}\to\text{A}}$ were initialized from a normal distribution with mean 0 and standard deviation $2/N_X$. Similarly, readout weights $\mathbf{V}^{\text{A}\to\text{out}}$ were initialized from a normal distribution with mean 0 and standard deviation $2/N_A$.

Training. Networks were trained to minimize the mean squared error (MSE) between the network output and the target trajectory. Top-down feedback was applied via a proportional controller with fixed gain k_p through feedback weights $\mathbf{Q}$ defined in equation (29). The network was trained for a total duration of $T =$ 360 minutes, corresponding to 4320 iterations of the trajectory. Synaptic weights were recorded every 20 iterations during training. Numerical integration was performed using the forward Euler method with a time step of $\Delta t = 1$ ms. To prevent runaway LTD due to excess inhibition in the absence of postsynaptic activity in networks with assembly overlap (Fig. 5), we constrained all synaptic weights to the interval $[-1, 1]$ during training.

Assessing model performance. Model performance was evaluated under two conditions: with top-down feedback (closed-loop) and without top-down feedback (open-loop). Closed-loop performance was measured from the recorded network output every 20 iterations. Since open-loop performance cannot be directly observed when top-down feedback is active, we created a copy of the model every 20 iterations, disabled top-down feedback, and recorded the open-loop output for 2 iterations. Performance was quantified

using the MSE and the coefficient of determination (R^2) between the network output and the target trajectory.

Control experiments. Two control conditions were examined: one without top-down feedback and one without a hidden layer. In the no-feedback condition, feedback weights $\mathbf{Q}$ were set to zero, effectively disabling top-down feedback while still optimizing the readout weights to minimize the MSE. In the no-hidden layer condition, the hidden layer was removed and the input neurons were connected directly to the output neurons, with the readout weights trained normally.

Training deep hierarchical networks on Fashion-MNIST

We rescaled all Fashion-MNIST input images to values between 0 and 1 and flattened each image into a 784-dimensional vector that served as the network input. Each hidden layer consisted of $N_{\mathrm{A}} = 128$ E/I assemblies, with each assembly containing $N_{\mathrm{A}}^{\mathrm{E}} = 16$ excitatory and $N_{\mathrm{A}}^{\mathrm{I}} = 4$ inhibitory neurons. The output layer comprised $N_{\mathrm{Y}} = 10$ neurons, one for each class in the Fashion-MNIST dataset. For networks with multiple hidden layers, each hidden layer was treated as a separate E/I network with dynamics defined in equations (14) and (15). The first hidden layer received the image input x, and subsequent layers were connected sequentially, with the output of one layer serving as the input to the next. The final hidden layer was connected to the output layer.

Top-down feedback. Each hidden layer $l \in \{1, 2, \ldots, L\}$ received direct top-down feedback from a controller. Feedback weights followed the mapping defined in equation (29) adjusted for multiple hidden layers, so that the feedback weights to the lth hidden layer were given by

$$\mathbf{Q}_l = \mathbf{D}_l^{\mathrm{I}} \left(\mathbf{M}_l^{\mathrm{I}}\right)^T \frac{\partial \mathbf{u}^{\mathrm{out}}}{\partial \mathbf{r}_l^{\mathrm{A}}}, \tag{49}$$

where $\mathbf{r}_l^{\mathrm{A}} = \left(\mathbf{M}_l^{\mathrm{E}}\right)^T \mathbf{r}_l^{\mathrm{E}}$ represents the activity of each assembly in the lth hidden layer. For networks with three hidden layers, we additionally implemented a layer-wise normalization procedure by dividing each layer's feedback weight matrix $\mathbf{Q}_l$ by its Frobenius norm $\|\mathbf{Q}_l\|_F$. This normalization, while not required for successful training, significantly accelerated the network's convergence to equilibrium, and reduced variance in performance across training runs, resulting in more reliable network behavior.

Classification loss function. For the classification task, we adapted the standard cross-entropy loss function. Conventionally, the output layer employs a Softmax activation to match a one-hot encoded target vector $\mathbf{y}$. However, since softmax outputs can cause issues in feedback control settings [81], we used soft targets instead of one-hot encodings: the target value was set to 0.99 for the correct class and to $^{0.01}/_{N_{\mathrm{Y}}-1}$ for all other classes. In line with previous work [81, 125], we utilized a linear readout layer and incorporated the softmax operation within the loss function:

$$\mathscr{L}_{\text{classification}} = -\sum_{b=1}^{B} \mathbf{y}_b \log\left(\mathrm{Softmax}\left(\mathbf{u}_b^{\mathrm{out}}\right)\right), \tag{50}$$

where B is the batch size and $\mathbf{u}_b^{\mathrm{out}}$ denotes the network output for the bth sample.

Parameter initialization. Network parameters were initialized from normal distributions with mean 0 and standard deviation $^1/_{N_{\mathrm{pre},l}}$, where $N_{\mathrm{pre},l}$ is the number of neurons in the preceding layer.

Training. Networks were trained for 50 epochs using a batch size of $B = 100$. For each sample, the network was first allowed to converge to an open-loop equilibrium state in the absence of top-down feedback. Numerical integration was performed using Tsitouras' 5th order explicit Runge-Kutta method [141] with an adaptive step size, running for up to 2 seconds or until an equilibrium state was reached (defined by an absolute tolerance of 10^{-6}). At the open-loop equilibrium state, feedback weights $\mathbf{Q}$ were computed.

Thereafter, top-down feedback control was enabled, and the network was allowed to converge to a closed-loop equilibrium state using the same integration method. Parameter updates were then computed at the closed-loop equilibrium state using the BCP learning rule in equation (31). We constrained all parameters to the interval $[-1, 1]$ during training to prevent extreme values in the network Jacobian that could arise from excessively large weights. All simulations employed the ADAM optimizer with a learning rate of 0.001.

Backpropagation control experiments. To compare the performance of the BCP learning rule with standard backpropagation, we trained networks of identical architecture using standard backpropagation. To match the number of parameters in the E/I assembly networks, the number of hidden units was set to $N = 128$. As in the BCP networks, all hidden units used a rectified linear activation function. We trained networks with either 1 or 3 hidden layers using the ADAM optimizer (learning rate 0.001, batch size 100) for 50 epochs. In contrast to the BCP networks, these networks employed a softmax output layer and a standard cross-entropy loss function.

Model of learning in BLA during fear conditioning

To model the fear conditioning experiments by Krabbe et al. [46], we simulated networks with $N_X = 20$ input neurons, a hidden layer consisting of $N_E = 160$ excitatory and $N_I = 40$ inhibitory neurons arranged in $N_A = 40$ assemblies, and a single readout neuron that was used to model freezing probability in response to the CS. Networks were trained on a total of ten fear conditioning trials, each trial lasting 20 seconds. During each trial, the CS was presented for a duration of 4 seconds, followed by a US presentation for 2 seconds and an inter-trial interval period of 14 seconds.

CS encoding and inputs. The activity of input neurons was modeled as $x_j(t) = \epsilon_j(t) + r_j(t)$, where $\epsilon_j(t)$ is random noise determined by an OU process with parameters τ_{OU}, μ_{OU}, and σ_{OU}, and $r_j(t)$ represents the filtered response of neuron j to the CS. The response selectivity of inputs to the CS was chosen to generate a sparse encoding. Specifically, we randomly selected 25% of the input neurons to respond to the CS with response strengths q_j drawn from a lognormal distribution and normalized to sum to 1. The raw binary CS presentations were filtered with an alpha kernel with separate rise (τ_{rise}) and decay (τ_{decay}) time constants to produce temporally smoothed responses $r_j(t)$, calculated as:

$$r_j(t) = \int_0^t s_j(\tau) \cdot \frac{\tau - \tau'}{\tau_{\text{rise}}} \exp\left(-\frac{\tau - \tau'}{\tau_{\text{decay}}}\right) d\tau' \tag{51}$$

where $s_j(t)$ is a binary input signal for neuron j that represents CS presentation.

Freezing probability. We used a single readout neuron representing a continuous valence prediction output of the BLA. Following experimental evidence for genetically identifiable subpopulations of BLA neurons for positive and negative valence encoding [142], we used random, fixed readout weights, so that a negative (positive) weight indicated an assembly encoding negative (positive) valence. Freezing probability was modeled as a shifted logistic function operating on the valence prediction:

$$p\left(\text{freezing}|\text{u}_{\text{out}}\right) = \frac{1}{1 + e^{a(-u_{\text{out}} + \theta)}} \tag{52}$$

US-driven feedback and loss function. We modeled the US as an external stimulus that sets a target freezing probability of 1 for the network. During US presentation, a top-down controller as described above provided feedback to all hidden layer Type 2 interneurons. Top-down feedback was set to zero in the absence of the US.

Model of motor learning in M1

To model the motor learning experiments by Ren et al. [47], we trained networks to output a motor command representing a lever push-and-pull movement in response to a sensory cue. We simulated networks

with $N_X = 20$ input neurons, a hidden layer consisting of $N_E = 320$ excitatory and $N_I = 80$ inhibitory neurons arranged in $N_A = 80$ assemblies, and a single readout neuron representing a motor command. Networks were trained for a total of 2000 trials, each lasting 8 seconds. Two seconds into each trial, the cue was presented and the network was tasked to generate the appropriate motor command.

Cue encoding and inputs. The cue signal was designed to capture the transient non-normal amplification dynamics observed in motor cortex [143]. Specifically, we modeled the cue-evoked neuronal response as an alpha kernel with rise and decay time constants τ_{rise} and τ_{decay} that provided a transient temporal envelope. Each neuron's response to the cue within this temporal window was modeled as a linear combination of N_{sines} sinusoidal waves with frequencies randomly sampled between 0.1 and 2.0 Hz and random phases. This was implemented as:

$$c_j(t) = \alpha(t) \cdot \sum_{i=1}^{N_{\text{sines}}} w_{ij} \sin(2\pi f_i t + \phi_i) \tag{53}$$

where w_{ij} are elements of a random weight matrix drawn from a normal distribution and scaled by a factor of 10, and $\alpha(t)$ is the normalized alpha kernel. The resulting cue activity began 2 seconds into each trial and was normalized to have a maximum amplitude of 1. Input neuron activity was modeled as $x_j(t) = c_j(t) + \epsilon_j(t)$, where $c_j(t)$ is the filtered cue input to neuron j and $\epsilon_j(t)$ is random noise determined by an OU process with parameters τ_{OU}, μ_{OU}, and σ_{OU}.

Motor output. The target motor command was defined as a single sinusoidal cycle that began at time $t_{\text{cue}} = 2$ seconds and lasted for a duration of $T_{\text{movement}} = 2$ seconds:

$$m(t) = \begin{cases} -\sin\left(2\pi \frac{t - t_{\text{cue}}}{T_{\text{movement}}}\right) & \text{if } t_{\text{cue}} \le t \le t_{\text{cue}} + T_{\text{movement}} \\ 0 & \text{otherwise} \end{cases} \tag{54}$$

The network output was trained to follow the motor command $m(t)$ given the cue inputs like described previously for the trajectory matching task.

Lever output and correct trials. The lever position was modeled as the temporal integration of the motor command, with an offset correction to ensure the lever position started at zero at cue onset:

$$l(t) = \int_0^t m(t')\, dt' - l(t_{\text{cue}}) \tag{55}$$

Following Ren et al. [47], we classified trials as successful when the lever trajectory crossed a threshold θ_{lever} both downward and upward within 2 seconds after cue onset, corresponding to a complete push-pull movement.