Contribute to balance, wire in accordance: Emergence of backpropagation from a simple, bio-plausible neuroplasticity rule

Xinhao Fan[1*]   **Shreesh P Mysore**[1,2*]
[1]Department of Neuroscience
[2]Department of Psychological and Brain Sciences
*Corresponding author
Johns Hopkins University
Baltimore, MD, 21218
{xfan20, mysore}@jhu.edu

Abstract

Over the past several decades, backpropagation (BP) has played a critical role in the advancement of machine learning and remains a core method in numerous computational applications. It is also utilized extensively in comparative studies of biological and artificial neural network representations. Despite its widespread use, the implementation of BP in the brain remains elusive, and its biological plausibility is often questioned due to inherent issues such as the need for symmetry of weights between forward and backward connections, and the requirement of distinct forward and backward phases of computation. Here, we introduce a novel neuroplasticity rule that offers a potential mechanism for implementing BP in the brain. Similar in general form to the classical Hebbian rule, this rule is based on the core principles of maintaining the balance of excitatory and inhibitory inputs as well as on retrograde signaling, and operates over three progressively slower timescales: neural firing, retrograde signaling, and neural plasticity. We hypothesize that each neuron possesses an internal state, termed credit, in addition to its firing rate. After achieving equilibrium in firing rates, neurons receive credits based on their contribution to the E-I balance of postsynaptic neurons through retrograde signaling. As the network's credit distribution stabilizes, connections from those presynaptic neurons are strengthened that significantly contribute to the balance of postsynaptic neurons. We demonstrate mathematically that our learning rule precisely replicates BP in layered neural networks without any approximations. Simulations on artificial neural networks reveal that this rule induces varying community structures in networks, depending on the learning rate. This simple theoretical framework presents a biologically plausible implementation of BP, with testable assumptions and predictions that may be evaluated through biological experiments.

1   Introduction

Backpropagation (BP) is the most widely used supervised learning algorithm for training artificial neural networks. At its core, it updates the weights of a structured neural network in terms of the gradient of an error function with respect to these weights, doing so via the application of the Leibniz chain rule. BP has achieved remarkable success in diverse fields including image [1] and speech recognition [2], natural language processing [3], financial prediction [4], and medical diagnosis [5]. Despite its effectiveness within machine learning, a major criticism of BP has been its lack

of biological plausibility. Two key issues have been raised in the literature in this context: (1) the 'weight symmetry' problem, and (2) the 'update locking' problem. The weight symmetry problem refers to the requirement that the feedback connectivity matrix needs to be the exact transpose of the feedforward connectivity matrix in order to propagate error gradients accurately. The time-locking problem refers to BP's two distinct phases of feedforward and feedback information processing: respectively, the forward-flowing updates of unit activation states (or firing rates) across the layers of the network, followed by the backward-flowing updates to weights between successive pairs of layers. These two serially implemented stages of information flow necessarily require that certain parts of the network are transiently 'locked', i.e., can be updated only after other parts, resulting in 'locking'. As a result, learning cannot occur 'online', or simultaneuously, as the network is actively engaged in task-solving through feedforward processing. Researchers have explored various strategies to address both problems with varying degrees of success.

With respect to the weight symmetry problem, *one strategy* has been to relax the requirement for an exact transpose of the feedforward connectivity matrix. The feedback alignment (FA) algorithm [6] does this by starting with random feedback connectivity, that is shown to gradually align with the transpose of the feedforward matrix. The sign symmetry (SS) framework [7] goes further, requiring that only the signs of the feedback connections match that of the transposed feedforward matrix. The use of random connectivity was later shown to impair performance on larger-scale problems [8, 9]. In response,the weight mirror (WM) mechanism [10] employed specific neural circuits to approximate the feedforward connectivity more accurately. However, it remains uncertain whether these algorithms can perform as well as BP on diverse tasks, and the biological mechanisms for such complex feedback tuning are largely unknown. *Another strategy* proposes entirely different feedback computation methods from that in the original BP framework. Direct/Indirect Feedback Alignment (DFA/IFA) [11] replaces layer-to-layer feedback connectivity with direct feedback connections from the output layer to the hidden layers, eliminating the need for layer-by-layer propagation. Global Error Vector Broadcasting (GEVB) [12] simplifies the feedback signal to a single vectorized error signal permeated across the entire network. Target Propagation (TP), a distinct family of solutions [13, 14, 15, 16, 9, 17], propagates the correct activation to each hidden neuron rather than the error, thereby removing the necessity for symmetry. Additionally, the concept of a separate teacher network has been proposed, wherein the teacher network tunes the feedback to be symmetrical to the feedforward connectivity. This idea has been framed within a reinforcement learning context in node perturbation (NP) [18], and as a controller in deep feedback control (DFC) [19] and least control principle (LCP) [20]. However, the newly proposed feedback computations in most models lack biological validation, and drastically redesigning BP feedback could lose many of its beneficial properties. A *third strategy* has been to directly invoke a biological mechanism that permits synapses to transmit information backward along the axon to the soma, namely retrograde signaling. As the same synapse is used in both directions, symmetry is inherently maintained. Although some research has explored this avenue [21], the fundamental limitation perceived with retrograde signaling in the context of BP is that it occurs on a much slower timescale compared to neural firing [22, 23]. Nonetheless, given the established biological feasibility of retrograde signaling, novel strategies that could achieve congruence between its slow timescale and the rapid timescale of feedforward activity are a potentially promising (but overlooked) route to resolving the weight symmetry problem.

With respect to addressing the update locking problem, significant efforts have been made to unify the feedforward and feedback phases into a single phase. Predictive Coding (PC) [24, 25] attempts to tackle this issue by introducing a second type of neuron that encodes prediction errors and influences network updates. The error neurons are only active when errors are present, allowing PC to automatically differentiate between phases. Similarly, Equilibrium Propagation (EP) [26, 27, 28] recognizes phases by describing the system through a unified rule of minimizing Hopfield energy, with and without the output layer clamped by an external signal. LCP [20] integrates error assignment and system firing dynamics into a single phase of optimal control within a control theory framework. Direct Random Target Propagation (DRTP) [17] replaces backward error propagation with forward target activity propagation, thus maintaining only two types of forward propagation within the model. Additionally, other studies attempt to leverage more biologically plausible mechanisms such as spike-timing-dependent plasticity (STDP) to resolve this issue [29]. Broadly speaking, however, there does not yet exist a widely accepted approach of rule unification for the resolution of the update locking problem. Novel strategies that, instead, use distinct rules operating simultaneously on different timescales are a potentially promising (but unexplored) alternative.

In addition to the widely investigated weight symmetry and update locking problems, there is a third potential concern regarding the biological implausibility of BP, namely (3) the 'derivative computation' problem: how can neurons compute derivatives? Although not the main focus of research in this field, this question has been addressed implicitly in many frameworks, with partial success. In energy-based models such as PC and EP, derivatives are introduced through the first principle of energy minimization, where the network evolves according to the gradient of energy [24, 25, 26, 27]. For more biologically realistic neurons, derivatives are embedded in the differential dynamics governing neural activity [29]. Alternatively, derivatives can be computed by a teacher network acting as a reinforcement learning agent [18] or a PID controller [19]. Some frameworks avoid the direct computation of derivatives by replacing it with an algebraic computation. For instance, in GEVB, specific forms of activation nonlinearity are assumed [12], while in contrastive Hebbian learning (CHL) [30, 31] and generalized recirculation (GeneRec) [32, 33], the use, solely, of multiplication is a limiting case of BP, resulting in a Hebbian-like rule. The TP family goes a step further, replacing the derivative of the loss function with a target signal [13, 14, 15, 16, 9, 17], thereby eschewing altogether the need for the derivative computation in its setup.

Despite the numerous alternative frameworks that have been proposed, BP remains the most powerful learning algorithm for a diverse range of tasks and datasets, and is still the most widely used method for machine learning applications and developments. Some alterative frameworks can mathematically reduce to BP, but they require additional assumptions, such as a weak teaching signal in EP [27] or weak feedback connections in CHL [31], which lack real-world relevance and are merely mathematical constraints. Furthermore, many frameworks attempting to solve the bio-plausibility problem of BP introduce additional requirements that may themselves be biologically implausible. Examples include the specific structure of neural circuits in DFA/IFA, WM, and PC [11, 10, 25], as well as the widespread existence of teacher networks in NP and DFC [18, 19]. Consequently,an enduring open question remains, "what is a learning rule implementable by the brain that possesses the features of, and exhibits the achievements realized by, BP?"

Here, we propose a new framework for learning in neural networks based on the foundational principle that neurons pursue excitatory-inhibitory (E-I) balance in their inputs. This framework operates on three progressively slower timescales: neural firing for computation, retrograde signaling for credit assignment, and neural plasticity for weight updates. On the one hand, this formulation allows the system to automatically detect when to update weights, thereby resolving the (2) update locking problem. On the other, it addresses the (1) weight symmetry problem by embedding the slower time-scale retrograde signaling naturally within its three-timescale operation. Separately, by hypothesizing the existence of an internal mechanism in each neuron that measures deviations from its preferred state of E-I balance, we derive a three-factor Hebbian-like neuroplasticity rule that does not rely on derivatives, thereby resolving the (3) derivative computation problem. We prove that in layered neural networks, this framework can be mathematically reduced to BP, *without any approximations*. Additionally, our framework makes several experimentally testable predictions.

2 The Model

2.1 E-I balance as the central principle

The core conceptual foundation of our neuroplasticity framework is the maintenance of E-I balance of neurons. E-I balance refers to the equilibrium between excitatory and inhibitory inputs to neurons. Maintaining this balance has been shown to help direct neuronal firing toward the linear response zone, which is essential for efficient information processing and representation [34, 35, 36]. Although multiple biological mechanisms have been shown to contribute to maintenance of E-I balance, including homeostatic regulation of synaptic strengths [37, 38], alterations in the intrinsic excitability of neurons [39], and heterosynaptic plasticity [40], learning rules have, as yet, not incorporated this requirement directly into their formulations. Our model fills this gap.

2.2 Neural firing

We utilize an all-to-all connected neural network to establish our model (Figure 1A), which provides a high degree of generality. The firing activities of neurons are represented as an n-dimensional vector $\mathbf{x}(t'')$, where the i^{th} element $x_i(t'')$ reflects the firing activity of the i^{th} neuron at time step t''. The firing activity depends on the nonlinear activation function σ and the inputs to each neuron

$\mathbf{s}(t'') = \mathbf{W}\mathbf{x}(t'')$, where $\mathbf{W}$ is the connectivity matrix. The element w_{ij} in $\mathbf{W}$ indicates the strength of the synapse received by neuron i from neuron j. Consequently, the network dynamics and steady state (fixed points) can be described as follows:

$$\mathbf{x}(t''+1) = \sigma(\mathbf{W}\mathbf{x}(t'')), \quad \mathbf{x} = \sigma(\mathbf{W}\mathbf{x}) \tag{1}$$

The time t here is used specifically for the dynamics of the firing activity, and is in the range of milliseconds.

Figure 1: The model. (A) All-to-all connectivity in an ANN. (B) Bell-shaped function quantifying E-I balance of a neuron. (C) Credit distribution from postsynaptic neuron i to presynaptic neuron j (hypothesized to occur via retrograde signaling mechanisms, orange arrow).

2.3 Neural credits

In addition to the firing rate, we hypothesize that each neuron encodes a second quantity that we refer to as *credit*. Credits of neurons are represented as an n-dimensional vector $\mathbf{c}(t')$, where the i^{th} element $c_i(t')$ reflects the amount of credit possessed by the i^{th} neuron at time step t'.

The magnitude of this quantity for each neuron represents the extent to which it has contributed to the E-I balance of its postsynaptic neurons. In turn, the degree of E-I balance for a postsynaptic neuron is measured as the deviation of the neuron's net input from the range that produces a linear response, and is quantified by a *balance function* f: a bell-shaped function which peaks at the balanced set-point of net input, and decreases toward zero when the net input deviates from the set-point (Figure 1B). With this framework in place, the credit g_{ji} distributed from postsynaptic neuron i to presynaptic neuron j is governed by the following rule (Figure 1C):

$$g_{ji} = c_i f(s_i) w_{ij} \tag{2}$$

where c_i is the amount of credit possessed by the postsynaptic neuron i; $s_i = \sum_k w_{ik} x_k$ is neuron i's input and $f(s_i)$ represents how balanced it is.

This credit redistribution rule encapsulates three underlying relationships: **(i)** postsynaptic neurons with more credits give more credits to their presynaptic neurons; **(ii)** postsynaptic neurons that are more balanced give more credits to their presynaptic neurons; and **(iii)** presynaptic neurons receive more credits if the synaptic strength is stronger. From a biological perspective, we propose that neuronal credits are encoded via biomolecules related to retrograde messengers, and distributed via retrograde signaling from postsynaptic to presynaptic neurons (see also section 5).

The dynamics of credit redistribution for each presynaptic neuron j can be described as $c_j(t'+1) = \sum_i g_{ji}(t')$, with the summation occuring over all of its postsynaptic neurons i. Consequently, the network dynamics of credit redistribution and the corresponding fixed points can be described as:

$$\mathbf{c}(t'+1) = \mathbf{J}\mathbf{c}(t') = \begin{bmatrix} f(s_1)w_{11} & f(s_2)w_{21} & \cdots & f(s_n)w_{n1} \\ f(s_1)w_{12} & f(s_2)w_{22} & \cdots & f(s_n)w_{n2} \\ \vdots & \vdots & \ddots & \vdots \\ f(s_1)w_{1n} & f(s_2)w_{2n} & \cdots & f(s_n)w_{nn} \end{bmatrix} \begin{bmatrix} c_1(t') \\ c_2(t') \\ \vdots \\ c_n(t') \end{bmatrix}, \quad \mathbf{c} = \mathbf{J}\mathbf{c} \tag{3}$$

The matrix $\mathbf{J}$ summarizes these dynamics. Here, t' is the time index specifically for the dynamics of credit redistribution, which is independent from the time index in neural firing dynamics t''. The time for credit redistribution is expected to be in the range of a second or a few seconds, based on the known biological features of retrograde signaling.

2.4 Neuroplasticity rule

With the neural firing activity and credit values defined, we propose a new three-factor plasticity rule as an extension to the two-factor Hebbian rule. In it, both the firing activity and the contribution to network E-I balance are taken into account for weight updating, which we mathematically formalize:

$$\Delta w_{ij} = c_i f(s_i) x_j, \quad \text{i.e.} \quad \mathbf{W}(t+1) = (\mathbf{c} \odot f(\mathbf{s})) \mathbf{x}^T + \mathbf{W}(t) \tag{4}$$

with the $\odot$ symbol denoting point-wise multiplication.

The three factors in this learning rule encapsulate the following relationships: **(a)** If the postsynaptic neuron has more credits, the synapse will be strengthened more; **(b)** If the postsynaptic neuron is more balanced, the synapse will be strengthened more. **(c)** If the presynaptic neuron has a higher firing rate, the synapse will be strengthened more. Separately, the timing for weight updates in equation 4 is expected to operate in the range of seconds to minutes, on stable values of both firing rates x_j as well as of credit values c_i.

Neurons that contribute to balance, wire in accordance. The product term $c_i f(s_i)$, also shared in equation 2, represents the effective credit a postsynaptic neuron can distribute to other neurons. This is because a higher credit value possessed by a postsynaptic neuron reflects a greater potential to distribute credits to its presynaptic neurons, and a higher degree of E-I balance suggests a higher propensity to allocate those credits. Consequently, establishing connections with neurons that have higher effective credits increases the likelihood of receiving more credits in the future. Separately, a higher firing rate of a presynaptic neuron leads to the strengthening of all its synapses, which is consistent with Hebbian plasticity. Taken together, presynaptic neurons that fire will strengthen their connections to postsynaptic neurons that contribute significantly to network balance and are well-balanced themselves. We encapsulate this rule with the phrase "Neurons that contribute to balance, wire in accordance," mirroring Hebb's principle: "Neurons that fire together, wire together."

2.5 Timescales of the model

Our model operates on three sequentially slower timescales: firing activity on the millisecond scale, credit redistribution on the seconds scale, and neural plasticity from seconds to minutes. This is depicted graphically in the raster plot of figure 2A, where each update mechanism is shown at its respective timescale using distinct time bin sizes for reference. When the system undergoes significant changes (for instance, owing to the arrival of a new input), the amplitude of curve is high. As the system stabilizes and converges, the amplitude gradually decreases to zero. Due to the hierarchical nature of these different timescales, we posit that the time it takes for one level to converge is much shorter than the update time interval of the subsequent level. Consequently, most changes due to credit update (yellow shadow under the curve) occur when network firing has already converged, and most weight changes (green shadow) occur when credit redistribution has converged. In other words, equation 3 is based on the equilibrium state of equation 1, while equation 4 is based on the equilibrium state of equation 3.

For learning to occur, the neuroplasticity rule (equation 4) requires both firing activity $\mathbf{x}$ and credit distribution $\mathbf{c}$ to be non-zero. However, these values can naturally converge to a zero state in the absence of external sensory signals, as $\mathbf{0} = \mathbf{J0}$ and $\mathbf{0} = \sigma(\mathbf{W0})$ for many activation functions such as *tanh* or *shifted sigmoid*. To address this, we propose that two sets of "clampings" are necessary for learning to take place. One subgroup of neurons needs to have their firing activity clamped to ensure the network converges to a non-zero state, while a different subgroup of neurons needs to have their credit value clamped. We refer to the former as input neurons and the corresponding clamping as *input clamping*, and the latter as output neurons and *output clamping*. We note that the term "clamping", here, is used to mean sensory inputs or error/reward signals from other neural circuits occurring for some duration of time (see also Discussion).

The time windows for input clamping (signal 1) and output clamping (signal 2) are illustrated in figure 2A. An overlap between these two time windows is necessary because, outside this overlap, either $\mathbf{x}$ or $\mathbf{c}$ will quickly converge to zero. Additionally, the overlap must be long enough for neuroplasticity to occur. Therefore, the durations of both input and output clamping should align with the neuroplasticity timescale, which is from seconds to longer. As a result of these multiple timescales, the system automatically distinguishes between the prediction phase (only input clamping) and the learning phase (overlapped input and output clamping), thereby avoiding the need to explicitly introduce distinct

phases. From a biological perspective, we hypothesize that these multiple hierarchical timescales are orchestrated by neural oscillations (see also section 5).

Figure 2: Three timescales of network operation. (A) Each row represents updates to the corresponding quantity; timescales are progressively faster from top to bottom. For the top three rows, each vertical line represents one timebin, and the curve represents the amplitude of the rate of updates. The area under the curve thus represents the total amount of change made to different quantities. As the quantity converges to steady-state, the corresponding curve value decreases to zero. Learning occurs with input clamping on neural firing (signal 1, blue) and output clamping on credit values (signal 2, yellow), as long as there is a sufficiently long period of overlap. (B,C) Input and output clamping in an all-to-all connected ANN, and a layered ANN, respectively.

2.6   Summary of the model

In summary, our new framework operates following three update rules: **(1)** Neural firing dynamics at the fastest timescale: $\mathbf{x}(t'' + 1) = \sigma(\mathbf{W}\mathbf{x}(t''))$; **(2)** Credit redistribution dynamics at the intermediate timescale: $\mathbf{c}(t' + 1) = \mathbf{J}\mathbf{c}(t')$; **(3)** Neural plasticity dynamics at the slowest timescale: $\mathbf{W}(t + 1) = (\mathbf{c} \odot f(\mathbf{s}))\mathbf{x}^T + \mathbf{W}(t)$

3   Model reduces *exactly* to backpropagation in layered neural networks

Here, we demonstrate the link between our proposed neuroplasticity framework and BP. We will first provide an intuitive understanding of how the chain rule emerges from our model setup, using derivations in an all-to-all connected neural network (Figure 2B). We will then present a theorem proving the equivalence of both methods in layered neural networks.

Assuming the credit values for the network start with the initial value $\mathbf{c}(0)$, and following the credit redistribution rule in equation 2, the elements for the credits at the next step $\mathbf{c}(1)$ are therefore:

$$\mathbf{c}(1) = \mathbf{J}\mathbf{c}(0), \quad \text{i.e.} \quad c_i(1) = \sum_j f(s_j)w_{ji}c_j(0) \tag{5}$$

Following this line of thinking, the subsequent credit distribution can also be derived:

$$c_i(2) = \sum_j f(s_j)w_{ji}c_j(1) = \sum_{j,k} f(s_j)w_{ji}f(s_k)w_{kj}c_k(0) \tag{6}$$

$$c_i(3) = \sum_{j,k,l} f(s_j)w_{ji}f(s_k)w_{kj}f(s_l)w_{lk}c_l(0) \quad \text{......} \tag{7}$$

where an interleaved pattern of w and f emerges. For simplicity, we approximate $\mathbf{c}(2)$ as the final credit distribution at equilibrium. Consequently, the neuroplasticity update based on this distribution, following equation 4, is:

$$\Delta w_{ij} = \sum_{k,l} f(s_k)w_{ki}f(s_l)w_{lk}c_l(0)f(s_i)x_j = \sum_{k,l} c_l(0)f(s_l)w_{lk}f(s_k)w_{ki}f(s_i)x_j \tag{8}$$

To further illustrate the connection between our model and gradient descent, we need to specify the form of the balance function f. This brings us to the *key idea of linking to BP*, which relies on the chain rule and gradient computations, with our framework, which involves simple algebraic computations. Specifically, we propose that the derivative of the activation function serves as the balance function, i.e., $\partial\sigma/\partial s = f$ (Figure 3).

The derivative of the activation function is a natural choice for the balance function because, for monotonically increasing activation functions that saturate at the extremes, its derivative typically forms a bell-shaped curve centered around the linear response zone. For instance, the derivative of the sigmoid function exhibits this property (Figure 3). For other activation functions like ReLU, the derivative can also display an appropriate peak structure if the positive side is bounded due to the biological constraint of a saturated firing rate, making it a suitable balance function.

Consequently, the neuroplasticity update becomes:

$$\Delta w_{ij} = \sum_{k,l} c_l(0)\frac{\partial\sigma(s_l)}{\partial s_l}w_{lk}\frac{\partial\sigma(s_k)}{\partial s_k}w_{ki}\frac{\partial\sigma(s_i)}{\partial s_i}x_j \tag{9}$$

Since the neural firing dynamics have already converged to equilibrium by the time neuroplasticity occurs, the firing rates satisfy $\mathbf{x} = \sigma(\mathbf{W}\mathbf{x})$. Therefore, we have $\frac{\partial\sigma(s_l)}{\partial s_l}w_{lk} = \frac{\partial x_l}{\partial x_k}$ and $\frac{\partial\sigma(s_i)}{\partial s_i}x_j = \frac{\partial x_i}{\partial w_{ij}}$. The update then becomes:

$$\Delta w_{ij} = \sum_{k,l} c_l(0)\frac{\partial x_l}{\partial x_k}\frac{\partial x_k}{\partial x_i}\frac{\partial x_i}{\partial w_{ij}} \tag{10}$$

Figure 3: Choosing the balance function (orange curve) as the derivative of the activation function (black curve).

In this form, we intuitively demonstrate how the credit information in the l^{th} neuron can propagate to the k^{th} neuron, then to the i^{th} neuron, and finally reach the synapse w_{ij}. Although this derivation is informal, it reveals a connection between the timesteps of credit redistribution and the steps of the chain rule in BP. Therefore, by clamping the credit values of certain neurons with loss-related information, we can propagate the error signal through the neural network similarly to BP.

To formalize the connection we've outlined, we present a theorem that proves the equivalence of BP and our learning framework in a layered neural network (Figure 2C):

Theorem. *Consider an m-layer neural network with neurons $\{\mathbf{x}^1, \mathbf{x}^2, \ldots, \mathbf{x}^m\}$ and connectivities $\{\mathbf{W}^1, \mathbf{W}^2, \ldots, \mathbf{W}^{m-1}\}$, governed by $\mathbf{x}^i = \sigma(\mathbf{W}^{i-1}\mathbf{x}^{i-1})$. Let the balance function be defined as $f = \partial\sigma/\partial s$. If the input layer's firing activity is clamped at $\mathbf{x}^1$ and the output layer's credit distribution is clamped at $-\partial L/\partial\mathbf{x}^m$, then $\Delta\mathbf{W} = -\partial L/\partial\mathbf{W}$.*

(For the detailed proof, please refer to the appendix.) This theorem indicates that simultaneous clamping of the input and output layers in a layered neural network can facilitate gradient descent on a loss function L that is equivalent to BP. At first glance, the output clamping signal, $-\partial L/\partial \mathbf{x}^m$, seems to necessitate explicit derivative computation by neurons. However, this can often be simplified to algebraic computation, depending on the form of the loss function. For example, with the squared loss $L = \frac{1}{2}\|\mathbf{x}^m - \mathbf{y}\|^2$, the gradient $-\partial L/\partial \mathbf{x}^m$ simplifies to $\mathbf{y} - \mathbf{x}^m$. Therefore, the output clamping signal is merely the difference between the target activity $\mathbf{y}$ and the current activity $\mathbf{x}^m$. This difference computation is biologically plausible, as it can be implemented by other neural circuits.

4   Simulation results: Learning rates dictate connectome patterns

Next, we present simulations of our proposed framework. Given that the computational power of BP, to which our model reduces exactly, has been validated extensively, our simulation focuses not on the performance of our model on task datasets but on its predictions about neural connectomes. We investigate the classes of connectivity structures that can emerge from our neuroplasticity rule.

We note four key points about the set up of our simulation algorithm. *First*, for the purposes of this exploration, the simulation algorithm makes one simplifying assumption In the model dynamics, the input/output clamping signal is replaced with a random Gaussian noise term, ϵ. This substitution serves the dual purposes of eliminating the need for specific input-loss pairs from open task datasets, and preventing the system from being trapped indefinitely in a trivial state. *Second*, to ensure the existence of a non-trivial equilibrium states, the first 10% of the rows of the matrix $\mathbf{J}$ were fixed to be a diagonal matrix. This can be interpreted as output clamping for that initial 10% of neurons, thereby guaranteeing the existence of a non-trivial equilibrium state via clamping. *Third*, in the wiring of biological networks, pruning of weak synaptic connections, a mechanism of structural plasticity, co-occurs with classical mechanisms of synaptic plasticity. Therefore, for the purposes of our exploration of neural connectivity structures, we pair our new synaptic plasticity rule with pruning. This encourages sparsity and the emergence of structured patterns of connectivity. *Finally,* a learning rate factor, r, is added to Equation 4, which serves as a key free parameter of interest.

The pseudocode for our simulation is as follows:

Algorithm 1 Weight Update and Pruning

1: Initialize weight matrix $\mathbf{W}$, firing activity vector $\mathbf{x} \in \mathbb{R}^n$, credit vector $\mathbf{c} \in \mathbb{R}^n$
2: Set learning rate r, pruning probability p
3: Define activation function σ, balance function f, noise distribution ϵ
4: **for** $i = 1$ to max_iteration **do**
5:     $\mathbf{x} \leftarrow \sigma(\mathbf{W}\mathbf{x}) + \epsilon$ ▷ Update firing activity vector $\mathbf{x}$ with noise
6:     $\mathbf{c} \leftarrow \mathbf{J}\mathbf{c} + \epsilon$ ▷ Update credit vector $\mathbf{c}$ with noise
7:     $\mathbf{W} \leftarrow \mathbf{W} + r \cdot (\mathbf{c} \odot f(\mathbf{s}))\mathbf{x}^\top$ ▷ Update weight matrix $\mathbf{W}$
8:     Prune $\mathbf{W}$ by setting the p proportion of smallest $|w_{ij}|$ to zero   ▷ Encourage sparsity
9: **end for**

We simulated an all-to-all connected neural network with 300 neurons. After running our simulation, we employed the Leiden algorithm [41] to reorder the neurons in the search of community structure within this network. We discovered that with sufficiently strong pruning, different learning rates led to the emergence of distinct connectivity patterns in the network. At a low learning rate, a block-diagonal-like structure emerged (Figure 4A), indicating stronger within-module connectivity. Conversely, at a high learning rate, we observed prominent off-diagonal block connectivity (Figure 4B), suggesting stronger across-module connectivity. The computational requirements for this simulation are lightweight, allowing it to be run on a personal computer within minutes (Configuration used for simulation: Intel Core i7-14700F CPU, 32GB RAM, NVIDIA RTX 4070 GPU).

5   Biological plausibility of our neuroplasticity framework

Here, we address the biological correspondence of the key aspects of our framework.

Credit and Neurotrophic factors. We introduced a novel factor in our model, namely, the credit value c. Just as neural firing rates are taken to encode x, it is important to clarify what neural

Figure 4: Simulation. Network connectivity structures with low (A) versus high (B) learning rates.

feature may encode c, or possibly the product of the credit and balance value, cf. We propose that neurotrophic factors (NTFs) could satisfy this role. NTFs, which are retrograde messengers, can be involved in retrograde credit redistribution processes. Moreover, NTFs are critical for neuronal survival, growth, maintenance, and repair [42, 43, 44], for neural plasticity [45, 46], and for mediating E-I balance [47, 48]. NTFs are, therefore, plausible candidates to encode c or cf.

Backward credit assignment and Retrograde Signaling. Retrograde signaling involves the release of a retrograde messenger by a postsynaptic dendrite or cell body, which then travels backward across a chemical synapse to bind to the axon terminal of a presynaptic neuron [49]. Different types of retrograde messengers have been identified experimentally: gaseous types such as nitric oxide (NO) [50, 51], endocannabinoids such as anandamide [52, 53, 54], neurotrophic factors such as neurotrophin [55, 56] and nerve growth factor [43], prostaglandins [57, 58], and classical neurotransmitters [59, 60]. Although retrograde signaling systems can vary in the physical spread of their action (affecting synapses in an input-specific way versus affecting millions of synapses) and in the speed of their reuptake and degradation [53, 60, 61, 62, 63], their time-of-action is usually within seconds to minutes [64, 65, 66], consistent with model requirements.

Multiple timescales of operation and Neural oscillations. The brain operates on hierarchical timescales with nested dynamics [67, 68]. Neural oscillations, the rhythmic patterns of activity in the nervous system, reflect this hierarchy by spanning a wide range of frequencies. In the mammalian forebrain, for example, networks exhibit oscillatory bands from 0.05 Hz to 500 Hz [69]. It has been suggested that neural oscillations are crucial for coordinating activity across timescales, providing phase information [70, 71], and in aiding functional integration [72]. Neural oscillations are also involved in neuroplasticity [73, 74] and are closely related to E-I balance [75, 76].

Input/output clamping and Associative learning. Our model requires clamping some input neurons with input signals and output neurons with loss-signal information. For the network to update weights, the input clamping period must overlap sufficiently with the output clamping period to allow plasticity. The longer the overlap, the more updates occur, resembling the temporal contiguity feature in associative learning for both classical and operant conditioning. The notion of input/output clamping is, therefore, congruent with typical scenarios of animals interacting with their environments.

6 Discussion

In this work, we proposed a simple neuroplasticity rule rooted in the principle of E-I balance, and with a functional form reminiscent of the classical Hebbian rule. Strikingly, we show that in layered neural networks, it is isomorphic to the BP algorithm under the sole assumption that the functional metric quantifying the E-I balance of a neuron be chosen to have the functional form of the derivative of its activation function. Moreover, the model is biologically plausible under two further assumptions: hierarchical timescales, and retrograde signaling, both of which have well-established biological mechanisms in neuroscience. In the framework of this model, BP can be interpreted as a simple consequence of individual neurons striving to maintain E-I balance. Our simulations demonstrate that this synaptic plasticity rule, when paired with structural plasticity mechanisms like pruning, can lead to the self-organization of different network structures in a fully connected network.

Our model makes three practical predictions that are experimentally testable. *First*, the amplitude of retrograde signaling should be highest for balanced neurons, but low for neurons with saturated firing rates or those close to being silent. This is based on the equation $g_{ji} = c_i f(s_i) w_{ij}$, where the balance value f approaches zero as a neuron becomes more unbalanced. *Second*, the ratio of the amplitude of synaptic change to the amplitude of retrograde signal should be proportional to the firing rate of the

presynaptic neuron and inversely proportional to synaptic strength. This prediction is derived from the ratio of the credit redistribution rule to the neuroplasticity rule, leading to $\frac{\Delta w_{ij}}{g_{ji}} = \frac{x_j}{w_{ij}}$. ***Third***, during early stages of neural development, associated with a greater degree of plasticity and therefore higher learning rate, the brain is more likely to form across-module connectivity. By contrast, in later stages, associated with lower learning rates, within-module connectivity is more likely.

Our model also reveals an intriguing link between BP and associative learning. First, the core formulation of our learning rule (that reduces exactly to BP) is Hebbian-like, promoting parallels with associative learning. Second, the clamping of input and loss onto the "input" and "output" neurons for overlapping durations are reminiscent of the temporal contiguity needed in associative learning paradigms. In this sense, BP can be viewed as associative learning in a broader context.

The model we have proposed here possesses a potential limitation. It is the assumption that credit redistribution will converge for the network updating rule to function. Our current mathematical framework does not guarantee the existence of a non-trivial fixed point in the dynamical system for credit redistribution. The matrix $\mathbf{J}$ in Equation 5 does not necessarily have an eigenvalue of one, meaning it could result in convergence to a zero vector or divergence to infinity. We currently circumvent this issue through the plausible means of clamping the credit values of some neurons (technically, even of just one neuron) to a constant. With such clamping, multiple (≥ 1) rows of the matrix $\mathbf{J}$ will be replaced by an identity matrix, ensuring the existence of an eigenvalue of one. Consequently, an intriguing observation is that the network's stability may depend on external credit signals and that spontaneous credit dynamics could lead to instability. Long-term sensory deprivation, which causes cognitive instability [77, 78] and affects connectivity and E-I balance in cortical circuits [79], is one biological phenomenon that supports this observation. More generally, then, the stability of this model will benefit from further exploration and systematic analysis.

Acknowledgements

This work was supported by funding from a JHU Discovery Award co-funded with the One Neuro Initiative (SPM), and from NIH grant 2R01EY027718 (SPM). We thank Hokin Deng (JHU) for making helpful suggestions about the title of the manuscript.

References

[1] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.

[2] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al. Deep speech 2: End-to-end speech recognition in english and mandarin. In *International conference on machine learning*, pages 173–182. PMLR, 2016.

[3] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[4] Omer Berat Sezer, Mehmet Ugur Gudelek, and Ahmet Murat Ozbayoglu. Financial time series forecasting with deep learning: A systematic literature review: 2005–2019. *Applied soft computing*, 90:106181, 2020.

[5] Pranav Rajpurkar, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis Langlotz, Katie Shpanskaya, et al. Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *arXiv preprint arXiv:1711.05225*, 2017.

[6] Timothy P Lillicrap, Daniel Cownden, Douglas B Tweed, and Colin J Akerman. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, 7(1):13276, 2016.

[7] Qianli Liao, Joel Leibo, and Tomaso Poggio. How important is weight symmetry in back-propagation? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.

[8] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015.

[9] Sergey Bartunov, Adam Santoro, Blake Richards, Luke Marris, Geoffrey E Hinton, and Timothy Lillicrap. Assessing the scalability of biologically-motivated deep learning algorithms and architectures. *Advances in neural information processing systems*, 31, 2018.

[10] Mohamed Akrout, Collin Wilson, Peter Humphreys, Timothy Lillicrap, and Douglas B Tweed. Deep learning without weight transport. *Advances in neural information processing systems*, 32, 2019.

[11] Arild Nøkland. Direct feedback alignment provides learning in deep neural networks. *Advances in neural information processing systems*, 29, 2016.

[12] David Clark, LF Abbott, and SueYeon Chung. Credit assignment through broadcasting a global error vector. *Advances in Neural Information Processing Systems*, 34:10053–10066, 2021.

[13] Yann Le Cun. Learning process in an asymmetric threshold network. In *Disordered systems and biological organization*, pages 233–240. Springer, 1986.

[14] Yoshua Bengio. How auto-encoders could provide credit assignment in deep networks via target propagation. *arXiv preprint arXiv:1407.7906*, 2014.

[15] Alexander Meulemans, Francesco Carzaniga, Johan Suykens, João Sacramento, and Benjamin F Grewe. A theoretical framework for target propagation. *Advances in Neural Information Processing Systems*, 33:20024–20036, 2020.

[16] Dong-Hyun Lee, Saizheng Zhang, Asja Fischer, and Yoshua Bengio. Difference target propagation. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2015, Porto, Portugal, September 7-11, 2015, Proceedings, Part I 15*, pages 498–515. Springer, 2015.

[17] Charlotte Frenkel, Martin Lefebvre, and David Bol. Learning without feedback: Fixed random learning signals allow for feedforward training of deep neural networks. *Frontiers in neuroscience*, 15:629892, 2021.

[18] Benjamin James Lansdell, Prashanth Ravi Prakash, and Konrad Paul Kording. Learning to solve the credit assignment problem. *arXiv preprint arXiv:1906.00889*, 2019.

[19] Alexander Meulemans, Matilde Tristany Farinha, Javier García Ordóñez, Pau Vilimelis Aceituno, João Sacramento, and Benjamin F Grewe. Credit assignment in neural networks through deep feedback control. *Advances in Neural Information Processing Systems*, 34:4674–4687, 2021.

[20] Alexander Meulemans, Nicolas Zucchet, Seijin Kobayashi, Johannes Von Oswald, and João Sacramento. The least-control principle for local learning at equilibrium. *Advances in Neural Information Processing Systems*, 35:33603–33617, 2022.

[21] Kenneth D Harris. Stability of the fittest: organizing learning through retroaxonal signals. *Trends in neurosciences*, 31(3):130–136, 2008.

[22] Timothy P Lillicrap, Adam Santoro, Luke Marris, Colin J Akerman, and Geoffrey Hinton. Backpropagation and the brain. *Nature Reviews Neuroscience*, 21(6):335–346, 2020.

[23] Emin Oztas et al. Neuronal tracing. *Neuroanatomy*, 2(2):5, 2003.

[24] Karl Friston. A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological sciences*, 360(1456):815–836, 2005.

[25] James CR Whittington and Rafal Bogacz. An approximation of the error backpropagation algorithm in a predictive coding network with local hebbian synaptic plasticity. *Neural computation*, 29(5):1229–1262, 2017.

[26] Benjamin Scellier and Yoshua Bengio. Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. *Frontiers in computational neuroscience*, 11:24, 2017.

[27] Benjamin Scellier and Yoshua Bengio. Equivalence of equilibrium propagation and recurrent backpropagation. *Neural computation*, 31(2):312–329, 2019.

[28] Axel Laborieux and Friedemann Zenke. Holomorphic equilibrium propagation computes exact gradients through finite size oscillations. *Advances in Neural Information Processing Systems*, 35:12950–12963, 2022.

[29] Alexandre Payeur, Jordan Guerguiev, Friedemann Zenke, Blake A Richards, and Richard Naud. Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nature neuroscience*, 24(7):1010–1019, 2021.

[30] Javier R Movellan. Contrastive hebbian learning in the continuous hopfield model. In *Connectionist models*, pages 10–17. Elsevier, 1991.

[31] Xiaohui Xie and H Sebastian Seung. Equivalence of backpropagation and contrastive hebbian learning in a layered network. *Neural computation*, 15(2):441–454, 2003.

[32] Geoffrey E Hinton and James McClelland. Learning representations by recirculation. In *Neural information processing systems*, 1987.

[33] Randall C O'Reilly. Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural computation*, 8(5):895–938, 1996.

[34] Carl Van Vreeswijk and Haim Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724–1726, 1996.

[35] Sophie Denève and Christian K Machens. Efficient codes and balanced networks. *Nature neuroscience*, 19(3):375–382, 2016.

[36] Shanglin Zhou and Yuguo Yu. Synaptic ei balance underlies efficient neural coding. *Frontiers in neuroscience*, 12:307227, 2018.

[37] Gina G Turrigiano and Sacha B Nelson. Homeostatic plasticity in the developing nervous system. *Nature reviews neuroscience*, 5(2):97–107, 2004.

[38] Gina G Turrigiano. The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell*, 135(3):422–435, 2008.

[39] Niraj S Desai, Lana C Rutherford, and Gina G Turrigiano. Plasticity in the intrinsic excitability of cortical pyramidal neurons. *Nature neuroscience*, 2(6):515–520, 1999.

[40] Rachel E Field, James A D¡¯amour, Robin Tremblay, Christoph Miehl, Bernardo Rudy, Julijana Gjorgjieva, and Robert C Froemke. Heterosynaptic plasticity determines the set point for cortical excitatory-inhibitory balance. *Neuron*, 106(5):842–854, 2020.

[41] Vincent A Traag, Ludo Waltman, and Nees Jan Van Eck. From louvain to leiden: guaranteeing well-connected communities. *Scientific reports*, 9(1):5233, 2019.

[42] Alun Milward Davies. The neurotrophic hypothesis: where does it stand? *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 351(1338):389–394, 1996.

[43] Eric C Yuen, Charles L Howe, Yiwen Li, David M Holtzman, and William C Mobley. Nerve growth factor and the neurotrophic factor hypothesis. *Brain and Development*, 18(5):362–368, 1996.

[44] Lorne M Mendell. Neurotrophin action on sensory neurons in adults: an extension of the neurotrophic hypothesis. *Pain*, 82:S127–S132, 1999.

[45] Hans Thoenen. Neurotrophins and neuronal plasticity. *Science*, 270(5236):593–598, 1995.

[46] A Kimberley McAllister, Lawrence C Katz, and Donald C Lo. Neurotrophins and synaptic plasticity. *Annual review of neuroscience*, 22(1):295–318, 1999.

[47] Lana C Rutherford, Andrew DeWan, Holly M Lauer, and Gina G Turrigiano. Brain-derived neurotrophic factor mediates the activity-dependent regulation of inhibition in neocortical cultures. *Journal of Neuroscience*, 17(12):4527–4535, 1997.

[48] Niraj S Desai, Lana C Rutherford, and Gina G Turrigiano. Bdnf regulates the intrinsic excitability of cortical neurons. *Learning & Memory*, 6(3):284–291, 1999.

[49] Huizhong W Tao and Mu-ming Poo. Retrograde signaling at central synapses. *Proceedings of the National Academy of Sciences*, 98(20):11009–11015, 2001.

[50] Thomas J O'dell, Robert D Hawkins, Eric R Kandel, and Ottavio Arancio. Tests of the roles of two diffusible substances in long-term potentiation: evidence for nitric oxide as a possible early retrograde messenger. *Proceedings of the National Academy of Sciences*, 88(24):11285–11289, 1991.

[51] Peter L Malen and Paul F Chapman. Nitric oxide facilitates long-term potentiation, but not long-term depression. *Journal of Neuroscience*, 17(7):2645–2651, 1997.

[52] Bradley E Alger. Retrograde signaling in the regulation of synaptic transmission: focus on endocannabinoids. *Progress in neurobiology*, 68(4):247–286, 2002.

[53] Rachel I Wilson and Roger A Nicoll. Endogenous cannabinoids mediate retrograde signalling at hippocampal synapses. *Nature*, 410(6828):588–592, 2001.

[54] Anatol C Kreitzer and Wade G Regehr. Retrograde signaling by endocannabinoids. *Current opinion in neurobiology*, 12(3):324–330, 2002.

[55] Mu-ming Poo. Neurotrophins as synaptic modulators. *Nature reviews neuroscience*, 2(1):24–32, 2001.

[56] David D Ginty and Rosalind A Segal. Retrograde neurotrophin signaling: Trk-ing along the axon. *Current opinion in neurobiology*, 12(3):268–274, 2002.

[57] Nan Sang, Jian Zhang, Victor Marcheselli, Nicolas G Bazan, and Chu Chen. Postsynaptically synthesized prostaglandin e2 (pge2) modulates hippocampal synaptic transmission via a presynaptic pge2 ep2 receptor. *Journal of Neuroscience*, 25(43):9858–9870, 2005.

[58] Jie Li, Elizabeth Serafin, and Mark L Baccei. Prostaglandin signaling governs spike timing-dependent plasticity at sensory synapses onto mouse spinal projection neurons. *Journal of Neuroscience*, 38(30):6628–6639, 2018.

[59] Ian C Duguid and Trevor G Smart. Retrograde activation of presynaptic nmda receptors enhances gaba release at cerebellar interneuron–purkinje cell synapses. *Nature neuroscience*, 7(5):525–533, 2004.

[60] Wade G Regehr, Megan R Carey, and Aaron R Best. Activity-dependent regulation of synapses by retrograde messengers. *Neuron*, 63(2):154–170, 2009.

[61] Stephen C Harward, Nathan G Hedrick, Charles E Hall, Paula Parra-Bueno, Teresa A Milner, Enhui Pan, Tal Laviv, Barbara L Hempstead, Ryohei Yasuda, and James O McNamara. Autocrine bdnf–trkb signalling within a single dendritic spine. *Nature*, 538(7623):99–103, 2016.

[62] J Wood and J Garthwaite. Models of the diffusional spread of nitric oxide: implications for neural nitric oxide signalling and its pharmacological properties. *Neuropharmacology*, 33(11):1235–1244, 1994.

[63] Andrew Philippides, Phil Husbands, and Michael O'Shea. Four-dimensional neuronal signaling by nitric oxide: a computational analysis. *Journal of Neuroscience*, 20(3):1199–1207, 2000.

[64] Boris D Heifets and Pablo E Castillo. Endocannabinoid signaling and long-term synaptic plasticity. *Annual review of physiology*, 71:283–306, 2009.

[65] Hyungju Park and Mu-ming Poo. Neurotrophin regulation of neural circuit development and function. *Nature Reviews Neuroscience*, 14(1):7–23, 2013.

[66] Gerard P Ahern, Vitaly A Klyachko, and Meyer B Jackson. cgmp and s-nitrosylation: two routes for modulation of neuronal excitability by no. *Trends in neurosciences*, 25(10):510–517, 2002.

[67] Harris S Kaplan, Oriana Salazar Thula, Niklas Khoss, and Manuel Zimmer. Nested neuronal dynamics orchestrate a behavioral hierarchy across timescales. *Neuron*, 105(3):562–576, 2020.

[68] Annemarie Wolff, Nareg Berberian, Mehrshad Golesorkhi, Javier Gomez-Pilar, Federico Zilio, and Georg Northoff. Intrinsic neural timescales: temporal integration and segregation. *Trends in cognitive sciences*, 26(2):159–173, 2022.

[69] Gyorgy Buzsaki and Andreas Draguhn. Neuronal oscillations in cortical networks. *science*, 304(5679):1926–1929, 2004.

[70] Ryan T Canolty and Robert T Knight. The functional role of cross-frequency coupling. *Trends in cognitive sciences*, 14(11):506–515, 2010.

[71] John E Lisman and Ole Jensen. The theta-gamma neural code. *Neuron*, 77(6):1002–1016, 2013.

[72] Satu Palva and J Matias Palva. Roles of brain criticality and multiscale oscillations in temporal predictions for sensorimotor processing. *Trends in neurosciences*, 41(10):729–743, 2018.

[73] György Buzsáki. Theta oscillations in the hippocampus. *Neuron*, 33(3):325–340, 2002.

[74] Mircea Steriade and Igor Timofeev. Neuronal plasticity in thalamocortical networks during sleep and waking oscillations. *Neuron*, 37(4):563–576, 2003.

[75] Bassam V Atallah and Massimo Scanziani. Instantaneous modulation of gamma oscillation frequency by balancing excitation with inhibition. *Neuron*, 62(4):566–577, 2009.

[76] Jeffry S Isaacson and Massimo Scanziani. How inhibition shapes cortical activity. *Neuron*, 72(2):231–243, 2011.

[77] William Harold Bexton, Woodburn Heron, and Thomas H Scott. Effects of decreased variation in the sensory environment. *Canadian Journal of Psychology/Revue canadienne de psychologie*, 8(2):70, 1954.

[78] Philip Solomon, P Herbert Leiderman, Jack Mendelson, and Donald Wexler. Sensory deprivation: A review. *American Journal of Psychiatry*, 114(4):357–363, 1957.

[79] Arianna Maffei and Alfredo Fontanini. Network homeostasis: a matter of coordination. *Current opinion in neurobiology*, 19(2):168–173, 2009.

7 Appendix

7.1 Link to Backpropagation: Proof for layered neural network

Theorem. *Consider an m-layer neural network with neurons $\{\mathbf{x}^1, \mathbf{x}^2, \ldots, \mathbf{x}^m\}$ and connectivities $\{\mathbf{W}^1, \mathbf{W}^2, \ldots, \mathbf{W}^{m-1}\}$, governed by $\mathbf{x}^i = \sigma(\mathbf{W}^{i-1}\mathbf{x}^{i-1})$. Let the balance function be defined as $f = \partial\sigma/\partial s$. If the input layer's firing activity is clamped at $\mathbf{x}^1$ and the output layer's credit distribution is clamped at $-\partial L/\partial \mathbf{x}^m$, then $\Delta\mathbf{W} = -\partial L/\partial\mathbf{W}$.*

Proof. We consider a typical layered neural network: the network has m layers, with n_i neurons in the i^{th} layer, and their activities are represented as a column vector $\mathbf{x}^i = (x_1^i, x_2^i, \ldots, x_{n_i}^i)$. The overall activity of the neural network is the concatenation of all layers, $\mathbf{x} = (\mathbf{x}^1, \mathbf{x}^2, \ldots, \mathbf{x}^m)$. Each layer $\mathbf{x}^i$ receives input solely from the previous layer $\mathbf{x}^{i-1}$ through the connectivity matrix $\mathbf{W}^{i-1}$. Consequently, the firing dynamics of this layered network can be described as:

$$\mathbf{x}(t'' + 1) = \sigma(\mathbf{W}\mathbf{x}(t'')) = \sigma(\begin{bmatrix} \mathbf{0} & & & & \\ \mathbf{W}^1 & \mathbf{0} & & & \\ & \mathbf{W}^2 & \mathbf{0} & & \\ & & \ddots & \ddots & \\ & & & \mathbf{W}^{m-1} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{x}^1(t'') \\ \mathbf{x}^2(t'') \\ \mathbf{x}^3(t'') \\ \vdots \\ \mathbf{x}^{m-1}(t'') \\ \mathbf{x}^m(t'') \end{bmatrix}) \tag{11}$$

For this network, the *neural activity* at equilibrium can be obtained by fixing the value of the input layer, i.e., keeping $\mathbf{x}^1(t'') \equiv \mathbf{x}^1$. For ease of explanation, if the network is initialized with $\mathbf{x} = (\mathbf{x}^1, \mathbf{0}, \mathbf{0}, \ldots)$, and $\mathbf{x}^1$ is kept fixed, reflecting a consistent external input signal, the firing activities will evolve as follows:

$$\begin{bmatrix} \mathbf{x}^1 \\ \mathbf{0} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} \rightarrow \begin{bmatrix} \mathbf{x}^1 \\ \mathbf{x}^2 = \sigma(\mathbf{W}^1\mathbf{x}^1) \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} \rightarrow \begin{bmatrix} \mathbf{x}^1 \\ \mathbf{x}^2 \\ \mathbf{x}^3 = \sigma(\mathbf{W}^2\mathbf{x}^2) \\ \vdots \\ \mathbf{0} \end{bmatrix} \rightarrow \begin{bmatrix} \mathbf{x}^1 \\ \mathbf{x}^2 \\ \mathbf{x}^3 \\ \vdots \\ \mathbf{x}^m = \sigma(\mathbf{W}^{m-1}\mathbf{x}^{m-1}) \end{bmatrix} \tag{12}$$

It can be observed that an m-layered neural network takes $m - 1$ timesteps to reach equilibrium. Non-zero initialization in other layers does not affect the stable state because activities in successive postsynaptic layers will emerge as a result of inputs from their presynaptic layers.

Next, for *credit distribution* in this m-layered neural network, we denote the credit values for the i^{th} layer as $\mathbf{c}^i$. The entire network's credit distribution can thus be represented as a concatenated column vector $\mathbf{c} = (\mathbf{c}^1, \mathbf{c}^2, \ldots, \mathbf{c}^m)$. If the balance function is chosen as $f = \frac{\partial\sigma(s)}{\partial s}$, the elements in the $\mathbf{J}$ matrix become derivatives between pairs of neurons. Representing this in block matrix form:

$$\mathbf{c}(t' + 1) = \mathbf{J}\mathbf{c}(t') = \begin{bmatrix} \mathbf{0} & (\frac{\partial\mathbf{x}^2}{\partial\mathbf{x}^1})^T & & & \\ & \mathbf{0} & (\frac{\partial\mathbf{x}^3}{\partial\mathbf{x}^2})^T & & \\ & & \ddots & \ddots & \\ & & & \mathbf{0} & (\frac{\partial\mathbf{x}^m}{\partial\mathbf{x}^{m-1}})^T \\ & & & & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{c}^1(t') \\ \mathbf{c}^2(t') \\ \mathbf{c}^3(t') \\ \vdots \\ \mathbf{c}^{m-1}(t') \\ \mathbf{c}^m(t') \end{bmatrix} \tag{13}$$

where the blocks are transposes of Jacobian matrices. For example,

$$(\frac{\partial\mathbf{x}^m}{\partial\mathbf{x}^{m-1}})^T = \begin{bmatrix} \frac{\partial x_1^m}{\partial x_1^{m-1}} & \frac{\partial x_2^m}{\partial x_1^{m-1}} & \cdots & \frac{\partial x_{n_m}^m}{\partial x_1^{m-1}} \\ \frac{\partial x_1^m}{\partial x_2^{m-1}} & \frac{\partial x_2^m}{\partial x_2^{m-1}} & \cdots & \frac{\partial x_{n_m}^m}{\partial x_2^{m-1}} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial x_1^m}{\partial x_{n_{m-1}}^{m-1}} & \frac{\partial x_2^m}{\partial x_{n_{m-1}}^{m-1}} & \cdots & \frac{\partial x_{n_m}^m}{\partial x_{n_{m-1}}^{m-1}} \end{bmatrix} \tag{14}$$

The stable state of the credit distribution can be obtained by fixing the credits at the output layer, reflecting a consistent external signal: $\mathbf{c}^m(t) \equiv \mathbf{c}^m$. For simplicity, if the other layers are initialized with zero credits, the credit distribution will evolve as follows:

$$
\begin{bmatrix} \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{c}^m \end{bmatrix}
\rightarrow
\begin{bmatrix} \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \mathbf{c}^{m-1} = (\frac{\partial \mathbf{x}^m}{\partial \mathbf{x}^{m-1}})^T \mathbf{c}^m \\ \mathbf{c}^m \end{bmatrix}
\rightarrow
\begin{bmatrix} \mathbf{0} \\ \vdots \\ \mathbf{c}^{m-2} = (\frac{\partial \mathbf{x}^{m-1}}{\partial \mathbf{x}^{m-2}})^T \mathbf{c}^{m-1} \\ \mathbf{c}^{m-1} \\ \mathbf{c}^m \end{bmatrix}
\rightarrow
\begin{bmatrix} \mathbf{c}^1 = (\frac{\partial \mathbf{x}^2}{\partial \mathbf{x}^1})^T \mathbf{c}^2 \\ \vdots \\ \mathbf{c}^{m-2} \\ \mathbf{c}^{m-1} \\ \mathbf{c}^m \end{bmatrix}
\tag{15}
$$

Therefore, it takes $m - 1$ timesteps for an m-layered network to reach equilibrium. In this scenario, the converged distribution depends solely on the consistent activity of the output layer. Initializing other layers with non-zero values will not affect the stable distribution.

Finally, we consider *updates to network weights* by the application of our neuroplasticity rule. Since we have specified the balance function f as $\partial\sigma/\partial s$, the original rule becomes:

$$
\Delta w_{ij} = c_i \sigma'(s_i) x_j \tag{16}
$$

Additionally, given the layered network structure, updates for all illegal connections (cross-layer, within-layer, or recurrent connections) are disregarded, and these connections remain unconnected. The update for connections between the k^{th} and $(k+1)^{\text{th}}$ layers is given by:

$$
\Delta \mathbf{W}^k = \sigma'(\mathbf{s}^{k+1}) \odot \mathbf{c}^{k+1} \, (\mathbf{x}^k)^T \tag{17}
$$

where $\mathbf{s}^{k+1}$ is the vector summarizing the inputs for all neurons in layer $k + 1$, $\sigma'(\mathbf{s}^{k+1})$ is the vector of corresponding balance values, and $\odot$ denotes term-wise multiplication.

Therefore, we have:

$$
\Delta w_{ij}^k = c_i^{k+1} \sigma'(s_i^{k+1}) x_j^k = \left((\frac{\partial \mathbf{x}^{k+2}}{\partial \mathbf{x}^{k+1}})^T \mathbf{c}^{k+2} \right)_i \sigma'(s_i^{k+1}) x_j^k \tag{18}
$$

$$
= \left((\frac{\partial \mathbf{x}^{k+2}}{\partial \mathbf{x}^{k+1}})^T (\frac{\partial \mathbf{x}^{k+3}}{\partial \mathbf{x}^{k+2}})^T \mathbf{c}^{k+3} \right)_i \sigma'(s_i^{k+1}) x_j^k \tag{19}
$$

$$
= \left((\frac{\partial \mathbf{x}^{k+2}}{\partial \mathbf{x}^{k+1}})^T (\frac{\partial \mathbf{x}^{k+3}}{\partial \mathbf{x}^{k+2}})^T \dots (\frac{\partial \mathbf{x}^m}{\partial \mathbf{x}^{m-1}})^T \mathbf{c}^m \right)_i \sigma'(s_i^{k+1}) x_j^k \tag{20}
$$

$$
= \left((\mathbf{c}^m)^T \frac{\partial \mathbf{x}^m}{\partial \mathbf{x}^{m-1}} \dots \frac{\partial \mathbf{x}^{k+3}}{\partial \mathbf{x}^{k+2}} \frac{\partial \mathbf{x}^{k+2}}{\partial \mathbf{x}^{k+1}} \right)_i \sigma'(s_i^{k+1}) x_j^k \tag{21}
$$

$$
= \left((\mathbf{c}^m)^T \frac{\partial \mathbf{x}^m}{\partial \mathbf{x}^{m-1}} \dots \frac{\partial \mathbf{x}^{k+3}}{\partial \mathbf{x}^{k+2}} \frac{\partial \mathbf{x}^{k+2}}{\partial \mathbf{x}^{k+1}} \right)_i \frac{\partial x_i^{k+1}}{\partial w_{ij}^k} \tag{22}
$$

If the activity of neurons in the output layer, driven by some external signal, is equal to the derivative of the loss L with respect to the neural firing rate, i.e., $\mathbf{c}^m = -\frac{\partial L}{\partial \mathbf{x}^m}$, then we have:

$$
\Delta w_{ij}^k = \left(-(\frac{\partial L}{\partial \mathbf{x}^m})^T \frac{\partial \mathbf{x}^m}{\partial \mathbf{x}^{m-1}} \dots \frac{\partial \mathbf{x}^{k+3}}{\partial \mathbf{x}^{k+2}} \frac{\partial \mathbf{x}^{k+2}}{\partial \mathbf{x}^{k+1}} \right)_i \frac{\partial x_i^{k+1}}{\partial w_{ij}^k} \tag{23}
$$

$$
= -\frac{\partial L}{\partial x_i^{k+1}} \frac{\partial x_i^{k+1}}{\partial w_{ij}^k} = -\frac{\partial L}{\partial w_{ij}^k} \tag{24}
$$

which is the defining weight update rule of backpropagation.