


RESEARCH ARTICLE

The calcitron: A simple neuron model that implements many learning rules via the calcium control hypothesis

Toviah Moldwin^{1*}, Li Shay Azran^{1,2}, Idan Segev^{1,3}

1 Edmond and Lily Safra Center for Brain Sciences, The Hebrew University of Jerusalem, Jerusalem, Israel, **2** Department of Brain Sciences, Weizmann Institute of Science, Rehovot, Israel, **3** Department of Neurobiology, The Hebrew University of Jerusalem, Jerusalem, Israel

* Toviah.moldwin@mail.huji.ac.il



OPEN ACCESS

Citation: Moldwin T, Azran LS, Segev I (2025) The calcitron: A simple neuron model that implements many learning rules via the calcium control hypothesis. PLoS Comput Biol 21(1): e1012754. <https://doi.org/10.1371/journal.pcbi.1012754>

Editor: Boris S. Gutkin, École Normale Supérieure, Collège de France, CNRS, FRANCE

Received: May 19, 2024

Accepted: December 27, 2024

Published: January 29, 2025

Copyright: © 2025 Moldwin et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data availability statement: Code for this project can be found at <https://github.com/tmoldwin/Calcitron>.

Funding: This work was supported by a grant awarded to I.S. by the ONR, Grant #: 3013007336, and by the Gatsby Charitable Foundation. The funders played no role in the study design, data collection and analysis,

Abstract

Theoretical neuroscientists and machine learning researchers have proposed a variety of learning rules to enable artificial neural networks to effectively perform both supervised and unsupervised learning tasks. It is not always clear, however, how these theoretically-derived rules relate to biological mechanisms of plasticity in the brain, or how these different rules might be mechanistically implemented in different contexts and brain regions. This study shows that the calcium control hypothesis, which relates synaptic plasticity in the brain to the calcium concentration ($[Ca^{2+}]$) in dendritic spines, can produce a diverse array of learning rules. We propose a simple, perceptron-like neuron model, the calcitron, that has four sources of $[Ca^{2+}]$: *local* (following the activation of an excitatory synapse and confined to that synapse), *heterosynaptic* (resulting from the activity of other synapses), *postsynaptic spike-dependent*, and *supervisor-dependent*. We demonstrate that by modulating the plasticity thresholds and calcium influx from each calcium source, we can reproduce a wide range of learning and plasticity protocols, such as Hebbian and anti-Hebbian learning, frequency-dependent plasticity, and unsupervised recognition of frequently repeating input patterns. Moreover, by devising simple neural circuits to provide supervisory signals, we show how the calcitron can implement homeostatic plasticity, perceptron learning, and BTSP-inspired one-shot learning. Our study bridges the gap between theoretical learning algorithms and their biological counterparts, not only replicating established learning paradigms but also introducing novel rules.

Author summary

Researchers have developed various learning rules for artificial neural networks, but it is unclear how these rules relate to the brain's natural processes. This study focuses on the calcium control hypothesis, which links changes in brain connections to calcium levels in neurons. The researchers created a simple neuron model that includes four sources of calcium and showed that by adjusting these, the model can mimic different types of learning, like recognizing patterns or learning from single events. This study helps

decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

connect theoretical learning models with how the brain might actually work, offering insights into both established and new learning mechanisms.

Introduction

Artificial neural networks (ANNs) have demonstrated a remarkable ability to solve both supervised and unsupervised learning tasks [1]. Because ANNs are inspired by a simple model of biological neurons and their synapses [2], theoretical neuroscientists have used ANNs to explore questions about the brain [3]. ANNs learn to solve a wide variety of problems by making use of “learning rules” whereby connection strengths between network nodes are modified [4–6]. These learning rules are inspired by the biological phenomenon of synaptic plasticity, i.e., the observation that synapses between neurons in the brain often undergo experience- or stimulation-dependent modifications.

Although algorithms for learning in deep multilayer ANNs, such as backpropagation, [7], are designed to solve optimization problems and not necessarily intended to be biologically plausible [8], the single-neuron versions of these algorithms are often sufficiently simple to be implemented in biology via plasticity mechanisms. Moreover, there is some evidence for the existence of these single-neuron learning algorithms in the brain, such as perceptron learning in the cerebellum [9] and Hebbian learning in the cortex [10]. In recent years, there have been attempts to extend these single-neuron learning rules in more biologically plausible ways at the network level [11–13]. However, it is often not apparent how biological mechanisms of synaptic plasticity can produce the types of learning rules used in ANNs.

One of the dominant theories for how long-term synaptic plasticity operates in the brain is the calcium control hypothesis [14–19]. The calcium control hypothesis states that the magnitude and direction of change of synaptic strength is mediated by the intracellular calcium concentration ($[Ca^{2+}]$) at the synapse. In the classic version of the calcium control hypothesis, at low levels of $[Ca^{2+}]$, the synapse is unaffected, at medium levels of $[Ca^{2+}]$, the synapse is depressed, and at high levels of $[Ca^{2+}]$, the synapse is potentiated.

This description of the relationship between calcium concentration and plasticity fits with experimental evidence from hippocampus and cortex [18,20]. In cerebellar Purkinje cells however, the calcium thresholds for potentiation and depression seem to be reversed; medium concentrations of $[Ca^{2+}]$ cause potentiation and high concentrations lead to depression [21]. While many downstream molecular mechanisms (most notably CaMKII and calcineurin) are involved in mediating plasticity [22–24], the calcium control hypothesis stands as one of the most parsimonious and effective theories for how synaptic plasticity works in the brain. It remains to be explained, however, how the calcium control hypothesis for plasticity can produce *learning rules* akin to those used in ANNs for solving tasks in an iterative manner.

In this work, we aim to bridge the gap between learning rules in artificial neurons and the biological mechanisms of synaptic plasticity. We propose a simple, perceptron-like threshold-linear neuron model, the *calcitron*, that has four potential sources of calcium, including one local (synapse-specific) source and three global (common to all synapses) sources. By adjusting the amount of calcium obtained from each calcium source and the calcium thresholds for plasticity, we can reproduce a wide range of learning and plasticity protocols, such as Hebbian and anti-Hebbian learning, frequency-dependent plasticity, and unsupervised recognition of frequently repeating input patterns. Moreover, by devising simple neural circuits to provide supervisory signals, we show how the calcitron can implement homeostatic plasticity,

perceptron learning, and BTSP-inspired one-shot learning. We thereby demonstrate that calcium control of synaptic plasticity can be a highly versatile mechanism which enables neurons to implement many different “programs” for modifying their synapses and storing information.

Results

The calcitron model

The calcitron is a simple neuron model, akin to a McCulloch and Pitts (M&P) neuron, or perceptron, which applies a transfer function to the weighted sum of its inputs. Formally we have:

$$\hat{y} = g \left(\sum_{i=1}^N w_i x_i + b \right) \quad (1)$$

where $\hat{y}$ is the output of the neuron, g is the transfer function, w_i is the weight of synapse i , x_i is the input to synapse i , N is the total number of synapses and b is a bias term. Depending on the particular use case, the transfer function g can be a simple threshold nonlinearity (e.g., a sign function), a sigmoid, a linear function, or any of the standard activation functions used for artificial neural networks, so the set of possible outputs $\hat{y}$ is contingent on the choice of activation function. For the purposes of this work, the weights w_i and inputs x_i of the calcitron are restricted to be non-negative to maintain fidelity to the experimental literature of calcium control hypothesis which mostly focuses on excitatory synapses [14–19]. In some contexts, the inputs x_i will be restricted to be binary (0 or 1); in other contexts, the inputs can be any positive real value (i.e., a rate model). The synaptic weights are also bounded between a minimum and maximum strength (w_{min} and w_{max} , respectively) due to the nature of the calcium-based plasticity rule; this will be discussed further below. Because the weights and inputs are restricted in our model to be excitatory, the bias term b will generally be negative or 0, thus representing the aggregate inhibitory input to the neuron.

The calcitron has four sources of calcium, based on biological mechanisms of calcium influx at dendritic spines (Fig 1A–1C). The first source of calcium is the local calcium due to the presynaptic input at each synapse, C_{local}^i . The local calcium at synapse i is defined as:

$$C_{local}^i = \alpha x_i \quad (2)$$

where α is a non-negative coefficient that determines the marginal increase in spine calcium for a unitary increase in input magnitude. Biologically, C_{local} can be thought of as the calcium that enters a dendritic spine through its NMDA channels during excitatory synaptic stimulation. We note that according to the above formulation, the local calcium does *not* depend on the synaptic weight, w_i , only on the synaptic input, x_i . This decision is motivated by both biological and computational considerations. Biologically, synaptic strength is a consequence of AMPA receptor conductance [22] while plasticity-inducing calcium influx occurs primarily via NMDA receptors. Thus, while changing the synaptic weight (i.e., the number of AMPA receptors) will influence the somatic depolarization observed for a given presynaptic input, it will not necessarily change the calcium influx via the NMDA receptors. (It is true that the NMDA receptor’s conductance is also voltage-dependent [25,26], so the increase of AMPA conductance of the synapse, as well as the weighted input from other synapses, can indirectly affect the calcium influx by depolarizing the neuron. We model the aggregate calcium influx due to input-dependent depolarization as heterosynaptic calcium, see next paragraph and




Fig 1. The calcitron model and calcium-based plasticity rules. (A) Sources of Ca^{2+} at the synapse. Local glutamate release from an activated presynaptic axon binds to an NMDA receptor in the postsynaptic dendritic spine, enabling local Ca^{2+} influx. Depolarization of the neuron opens voltage-gated calcium channels (VGCCs), enabling calcium influx from global signals. (Glutamate also binds to AMPA receptors, enabling Na^+ influx, and depolarization also affects NMDAR conductance.) (B) Possible sources of Ca^{2+} influx in a neuron. Ca^{2+} can enter due to presynaptic input (C_{local}^i), heterosynaptically-induced depolarization of VGCCs (C_{het}), the backpropagating action potential (C_{BAP}) or a supervisory signal, such as a calcium plateau induced by input to the apical tuft (C_{sprv}) (Neuron image courtesy of Dean Geck and MICrONS Consortium [80]). (C) The four Ca^{2+} sources in a point neuron model. Each Ca^{2+} source is associated with a respective coefficient ($\alpha, \beta, \gamma, \delta$) determining how much Ca^{2+} comes from each source. (D) The calcium control hypothesis. $[\text{Ca}^{2+}]$ below θ_D induces no change, $[\text{Ca}^{2+}]$ between θ_D and θ_P induces depression, and $[\text{Ca}^{2+}]$ above θ_P induces potentiation. (E) Weight change as a function of calcium in the linear version of Ca^{2+} -based plasticity, as in (D), shown as phase plane. Magnitude of weight change is independent of current weight. Blue indicates depression, red indicates potentiation, white indicates no change. (F) Step stimulus to show the plastic effect of different levels of $[\text{Ca}^{2+}]$. $[\text{Ca}^{2+}]$ is either raised to a depressive level ($\theta_D \leq [\text{Ca}^{2+}] \leq \theta_P$, blue line) or to a potentative level ($[\text{Ca}^{2+}] > \theta_P$, red line) for several timesteps, then reduced to 0. S and E refer to the start and end of the calcium step. (G) Dynamics of the linear rule in response to the step stimulus from (F). Synaptic weights increase or decrease linearly in response to the potentative or depressive levels of calcium (red and blue traces, respectively), then remain stable after calcium is turned off. (H) Fixed points (black) and learning rates (pink) in the asymptotic fixed point – learning rate (FPLR) version of the calcium control hypothesis. (I) Weight change as a function of $[\text{Ca}^{2+}]$ for different values of the present synaptic weight. Darker colors indicate higher weights. (J) Phase plane of weight changes for the FPLR rule. (K) Stimulus to demonstrate FPLR rule, identical to F. (L) Dynamics of the FPLR rule. Synaptic weights potentiate or depress asymptotically toward the potentative ($w_{\text{max}}, 1$) or depressive ($w_{\text{min}}, 0$) fixed point.

<https://doi.org/10.1371/journal.pcbi.1012754.g001>

Discussion. The weight-independence of the local calcium influx also has the computational advantage of avoiding feedback loops – if calcium influx was weight-dependent, potentiating or depressing a synapse would change the synapse's sensitivity to plasticity protocols.

The second source of intracellular calcium in the calcitron is calcium that globally enters dendritic spines due to the aggregate activity of all nearby synaptic inputs, resulting in the activation of voltage-gated calcium channels (VGCCs) in regions of the dendrite that are sufficiently depolarized [19,27] (and by increasing the conductance of NMDA at active synapses via NMDA's voltage dependence, see Discussion). We call this calcium source heterosynaptic calcium, or C_{het} . This calcium is responsible for heterosynaptic plasticity, i.e., plasticity that can be induced at non-activated synapses by presynaptic stimulation at other nearby synapses. While heterosynaptic plasticity is spatially sensitive, for simplicity we assume that synaptic activity is distributed uniformly on the neuron and we can thus approximate this calcium, C_{het} , as being proportional to the aggregate activity from all synaptic inputs. We thus have

$$C_{het} = \beta * \left(\sum_{i=1}^N w_i x_i \right) \quad (3)$$

where β is a coefficient that determines how much calcium enters each spine due to the overall depolarization of the dendritic membrane and x_i and w_i are as in Eq. (1).

The third source of calcium, C_{BAP} , comes from the backpropagating action potential, or BAP. When a neuron fires an action potential, the axonal/somatic spike backpropagates to the dendrites and depolarizes the dendritic membrane [28], which can globally activate voltage-gated calcium channels at all spines. We model this as:

$$C_{BAP} = \gamma \hat{y} \quad (4)$$

Where γ is a coefficient that determines the amount of calcium that enters the postsynaptic spines due to each spike. (Here we ignore timing effects of the postsynaptic spike relative to the timing of the input; for the purposes of the calcitron we assume that synaptic input and the spike it generates happen within a single time step. This assumption prevents the calcitron from implementing spike-timing dependent plasticity [STDP [29]]; a more detailed version of the calcitron with temporal dynamics would be necessary to capture STDP. See our previous work, [30], as well as [14,16] for modeling STDP via calcium-dependent plasticity.)

The fourth source of calcium, C_{sprv} , comes from an external supervisor, denoted as Z . Z may be binary or positive real-valued. In hippocampal pyramidal neurons, a likely candidate for this supervisory signal is strong input to the apical tuft, which can induce bursts of spikes at the soma, potentially leading to global calcium influx at VGCCs at basal dendrites [31–34]. A similar calcium-based supervisory scheme exists in cerebellar Purkinje neurons, where strong input to the Purkinje neurons from climbing fibers can induce long-term depression of synapses between presynaptic parallel fibers and the postsynaptic Purkinje cell [35]. (In our model, Z does not contribute directly to the neuron's output $\hat{y}$, because although the supervisory signal often does depolarize the neuron, this depolarization is treated as incidental to the plasticity induction, rather than as part of the neuron's input-output function. In other words, we assume that downstream neurons only care about “output spikes” rather than “plasticity plateaus/bursts”). The calcium influx due to this supervisory signal is defined as:

$$C_{sprv} = \delta Z \quad (5)$$

Where δ is the coefficient determining the amount of calcium that comes from the supervising signal. The total calcium per dendritic spine, C_{total}^i , is the sum of these four calcium sources (Fig 1C):

$$C_{total}^i = C_{local}^i + C_{het} + C_{BAP} + C_{sprv} = \alpha x_i + \beta * \left(\sum_{j=1}^N w_j x_j \right) + \gamma \hat{y} + \delta Z \quad (6)$$

We note that because the calcium sources in the calcitron are related to current sources (e.g., the presynaptic input or postsynaptic spike), the calcitron bears some similarity to the BCM plasticity rule [36].

An equivalent way to write Eq. (6) is in terms of local and global calcium sources:

$$C_{total}^i = C_{local}^i + C_{global} \quad (7)$$

where:

$$C_{global} = C_{het} + C_{BAP} + C_{sprv} \quad (8)$$

This formulation emphasizes that global calcium signals from the total feed-forward depolarization, backpropagating action potential, and supervisor are broadcast equally to all synapses, thus the local calcium is needed to break the symmetry between different synapses at each time step.

To implement any particular learning rule, it is usually only necessary to use a subset of calcium sources. Mathematically, this is accomplished by setting the coefficients of the unused calcium sources to 0. For example, if we are interested in Hebbian plasticity (see below) which depends only on the presynaptic and postsynaptic spikes, we set $\beta = 0$ and $\delta = 0$ so we don't have to worry about heterosynaptic plasticity or supervisory signals.

Calcium-based plasticity for the calcitron

At each discrete time step t of the calcitron's operation, a vector of inputs $\mathbf{x} = [x_1, x_2, \dots, x_N]$ is presented to the calcitron, and the calcitron produces an output $\hat{y}$ (Eq. (1)). The calcitron calculates the calcium concentration per spine, C_{total}^i , from the inputs (x), weights (w), output ($\hat{y}$), and the supervising signal (Z) at that time step (Eq. (6)). The calcium concentration at each dendritic spine (C_{total}^i) is used to determine the magnitude and direction of plastic change (if any) for that synapse's weight (w_i) at the next time step.

To implement calcium-based plastic changes to each synapse, we consider two versions of a calcium-based plasticity rule: a linear version and an asymptotic version. The latter rule is also referred to as the fixed point – learning rate (FPLR) rule. Although we will use the FPLR rule for most of the simulations in this study, we will first introduce the linear rule as it is simpler and provides a point of entry to the FPLR rule.

The linear version of calcium-based plasticity (Fig 1D–1G) changes a synaptic weight by a fixed amount at each time step depending on the amount of calcium present at the synapse at that time step. Formally, following the notation of [16] this may be described as

$$\Delta w_i(t) = w_i(t+1) - w_i(t) = \Omega(C_{total}^i(t)) \quad (9)$$

where $\Omega(C)$ is any function from positive-valued calcium concentrations to positive or negative changes in synaptic weight at each time step. As a simple representation of the classic calcium control paradigm observed in hippocampal and cortical cells, we can choose $\Omega(C)$ to be a step function with two thresholds: the depression threshold, θ_D , and the potentiation threshold, θ_P , where $\theta_D < \theta_P$ (Fig 1D and 1E). Our plasticity function $\Omega(C)$ returns 0 when the calcium is below θ_D , returns $-\eta_D$ when the calcium is between θ_D and θ_P (depressing the synapse by η_D units per time step), and returns η_P when the calcium is above the θ_P (potentiating the synapse by η_P units per time step). Formally this can be written as:

$$\Omega(C) = \begin{cases} 0, & C < \theta_D \\ -\eta_D, & \theta_D \leq C \leq \theta_P \\ \eta_P, & C > \theta_P \end{cases} \quad (10)$$

In the above formulation of the plasticity rule, synaptic weights can become arbitrarily large or small provided a sufficient number of plasticity events or a sufficiently long calcium plateau. Moreover, nothing in the linear rule prevents synapses from changing signs, i.e., from inhibitory to excitatory or vice versa (Fig 1F and 1G). This is not particularly biologically realistic, as synapses generally do not change sign from excitatory to inhibitory, and synaptic strengths are not observed to become arbitrarily large (and there are also biophysical limitations on the maximum possible depolarization that can be achieved from a single synaptic input).

In this work, we thus employ a different rule, inspired by the rules of Shouval [16] and Graupner [14], wherein synaptic weights are modified asymptotically toward some fixed value, depending on the calcium concentration (Fig 1H–1L).

$$\Delta w_i(t) = \eta(C_{total}^i(t)) * (F(C_{total}^i(t)) - w_i(t)) \quad (11)$$

This is known as the fixed point – learning rate (FPLR) rule, as the rule is specified by defining the asymptotic fixed points of the weights, $F(C)$, and rate of synaptic weight modification, $\eta(C)$, as a function of the calcium concentration (Fig 1H) [30]. The learning rates $\eta(C)$ define the fraction of the difference between the present weight and the target fixed point $F(C)$ that is traversed at each time step, resulting in asymptotic plasticity dynamics for a given level of calcium at a particular synapse (Fig 1L). A standard two-threshold calcium control function would have the following structure:

$$F(C) = \begin{cases} w_{drift}, & C < \theta_D \\ w_{min}, & \theta_D \leq C \leq \theta_P \\ w_{max}, & C > \theta_P \end{cases} \quad (12)$$

And

$$\eta(C) = \begin{cases} \eta_{drift}, & C < \theta_D \\ \eta_D, & \theta_D \leq C \leq \theta_P \\ \eta_P, & C > \theta_P \end{cases} \quad (13)$$

This means that synaptic weights with pre-depressive calcium concentrations (C_{total}^i) eventually drift toward a “neutral” state of w_{drift} at a rate of η_{drift} , synapses with a depressive calcium concentration will depress towards w_{min} at a rate of η_D , and synapses with a potentiative C_{total}^i will be potentiated towards w_{max} at a rate of η_P . In practice, for the purpose of this work, we will neglect the pre-depressive drift, i.e., the tendency of synapses to slowly drift back to baseline in the presence of very low levels of calcium ($C < \theta_D$). In the FPLR framework, this is accomplished by setting $\eta_{drift} = 0$. Also for simplicity, we will generally set the learning rate for depression to be equal to the learning rate for potentiation, i.e., $\eta_D = \eta_P$. We note however that in biology, LTD requires a more prolonged and sustained calcium pulse than LTP [37]. To replicate this aspect of LTD, one can set $\eta_D < \eta_P$, which is the standard method of handling the prolonged duration requirement in calcium control models [14,16,30].

Calcitron leaning in a binary model: Hebbian, Anti-Hebbian, and other pre-post learning rules

As simple illustration of the ability of the calcitron to implement learning rules, we consider Hebbian learning. There are a variety of different versions of the Hebb rule, and we will describe several of them here. As a simple mnemonic to describe these different rules, we will

use a 3-letter code, where the first letter refers to what happens at an active synapse when there is *no* postsynaptic spike (pre only), the second letter refers to what happens at *inactive* synapses when there *is* a postsynaptic spike (post only), and the third letter refers to what happens at *active* synapses when there *is* a postsynaptic spike (both). In any of the three positions, there might be an N (no change), D (depression) or P (potentiation). For example, the acronym DNP would refer to a rule where presynaptic input alone induces depression at the active synapse, no change occurs to inactive synapses when there is a postsynaptic spike, and active synapses are potentiated when there is a postsynaptic spike.

One simple formulation of the Hebb rule is “neurons that fire together, wire together”. In other words, presynaptic inputs that were active at the same time as a postsynaptic spike are potentiated, otherwise, no plasticity occurs (NNP). In the calcitron, this rule can be implemented by adjusting the coefficients for $C_{pre}(\alpha)$ and $C_{post}(\gamma)$. (As noted above, β and δ will be set to 0 for all Hebb-like rules.) If we assume binary inputs and outputs (i.e., x_i and $\hat{y}$ are either 0 or 1), this can be accomplished by setting α and γ in Eq. (6) to both be below θ_D , while enforcing that $\alpha + \gamma$ is above θ_P . This ensures that an active synapse not accompanied by a spike does not change (because $\alpha * 1 + \gamma * 0 < \theta_D$), nor does an inactive synapse that is accompanied by a spike (because $\alpha * 0 + \gamma * 1 < \theta_D$), but an active synapse accompanied by a spike will potentiate (because $\alpha * 1 + \gamma * 1 > \theta_P$) (Fig 2A1).

Another version of the Hebbian learning rule can be stated as “fire together, wire together; out of sync, lose your link” (DDP). In other words, synapses that were active at the same time as a postsynaptic spike are potentiated, as before, but now we penalize (via depression) synapses that were active at a time when there was no postsynaptic spike, as well as synapses that were inactive at a time when a postsynaptic spike did occur. In the calcitron, this is accomplished by setting α and γ to individually be in the depressive region (between θ_D and θ_P), which penalizes ‘out of sync’ synapses, while still maintaining $\alpha + \gamma$ is above θ to enforce “fire together wire together” behavior (Fig 2A2).

It is also possible to obtain an anti-Hebbian “fire together, lose your link” (NND) plasticity rule by setting α and γ to individually be below θ_D , while enforcing that $\theta_D < \alpha + \gamma < \theta_P$. However, it is *not* possible, using the standard plasticity thresholds, to get an anti-Hebbian rule that rewards out-of-sync synapses, i.e., “fire together, lose your link, out of sync wire together” (PPD), as a synapse that fires synchronously with the postsynaptic spike will always have a higher C_{total}^i than a synapse that fires asynchronously, and thus cannot potentiate. In other words, if $\theta_D < \alpha + \gamma < \theta_P$, necessarily it holds that $\alpha < \theta_P$ and $\gamma < \theta_P$ (Fig 2A3).

However, if the plasticity thresholds are reversed, i.e. $\theta_P < \theta_D$, as in Purkinje neurons, we can get “fire together, lose your link; out of sync, wire together” (PPD) plasticity. We set α and γ to both be in the potentiation region (between θ_P and θ_D), while enforcing that $\alpha + \gamma$ is above θ_D (Fig 2A4).

As a simple demonstration of these Hebbian and anti-Hebbian rules, we created a random sequence of binary input patterns (Fig 2B1–2B4) consisting of 10 synaptic input lines (i.e., $N = 10$) and presented this sequence to four different calcitrons, each implementing one of four rules described above by applying different calcium source coefficients and plasticity thresholds. We then compare how the different Calcitrons yield different plasticity dynamics (Fig 2C–2F).

In the “fire together, wire together” calcitron (NNP, 2A1–2F1), synaptic weights can only potentiate or stay stable, they can never depress. Every so often, just by chance, one of the random input patterns will be sufficiently large to generate a spike (2C1, red-headed stems). The calcium from this spike combined with the calcium at active synapses yields potentiation at the active synapses (2D1–2E1). This creates a positive feedback loop: when a synapse is potentiated, that makes the neuron more likely to elicit a spike




Fig 2. Four kinds of Hebbian and anti-Hebbian learning using Ca^{2+} . (A1–A4) Different versions of Hebbian and anti-Hebbian learning rules are implemented by setting the respective coefficients (α and γ in Eq. (6)) for the local (C_{local}^i) and backpropagating spike-dependent (C_{BAP}) [Ca^{2+}]. For each rule, we show the direction of

plastic change (indicated by the letters above the bars: “N”: no change, “D”: depression, “P”: potentiation) for three different conditions: a synapse *with* active presynaptic input in the *absence* of a postsynaptic spike (“pre”), a synapse *without* local input in the *presence* of a postsynaptic spike (“post”) and at a synapse *with* active presynaptic input *and* a postsynaptic spike (“both”). The total synaptic $[Ca^{2+}]$ ($C_{total}^{[Ca^{2+}]}$) for each condition is the sum of the local input-dependent $[Ca^{2+}]$ ($C_{local}^{[Ca^{2+}]}$, green) and the spike-dependent $[Ca^{2+}]$ ($C_{BAP}^{[Ca^{2+}]}$, pink). (When there is neither local input or a postsynaptic spike, the expected $[Ca^{2+}]$ is 0) (B1–B4) For each learning rule from (A1–A4), 10 random binary inputs (black: active, white: inactive) are presented to each synapse at each time step. (Inputs are identical for all learning rules). (C1–C4) Sum of weighted inputs at each time step for each learning rule shown in A1–A4 respectively. Dotted horizontal line indicates the spike threshold ($-b$ from Eq. (6)). Outputs that are above the threshold (produce a postsynaptic spike) are indicated by a red circle. (D1–D4) $[Ca^{2+}]$ per synapse for each time step for the 4 learning rules shown in A1–A4 respectively. (E1–E4) Bar codes indicating occurrence of potentiation (“P”, red), depression (“D”, blue) or no change (“N”, white) shown in A1–A4 respectively. (F1–F4) Synaptic weights over the course of the simulation for A1–A4 cases respectively.

<https://doi.org/10.1371/journal.pcbi.1012754.g002>

whenever that synapse participates in a pattern (because the synapse provides a greater contribution to the overall input to the neuron). Eventually, this results in all synaptic weights becoming strongly potentiated (2F1) and the calcitron spikes in response to almost all input patterns.

In the “fire together, wire together; out of sync, lose your link” calcitron (DDP, 2A2–2F2), the potentiation at active synapses whenever there is a spike is counterbalanced by depression at inactive synapses when a spike occurs, as well as at active synapses when no spike occurs. As such, the synaptic weights do not become overly strong and the neuron doesn’t spike as aggressively.

In the anti-Hebbian “fire together, lose your link” calcitron (NND, 2A3–2F3), synaptic weights can only depress or stay stable, they can never potentiate. As such, initially, whenever the random input elicits a spike, the synapses that were active at that time step depress. However, here we have a negative feedback loop: once a synapse depresses, it contributes less to the neuron’s overall voltage, which makes the neuron less likely to spike, so eventually the synaptic weights stop depressing.

Finally, in the anti-Hebbian “fire together, lose your link; out of sync, wire together” calcitron (PPD, 2A4–2F4), we again have a balance between depression and potentiation for synchronized and unsynchronized activity, so we again observe more moderate changes in the synaptic weights over longer time horizons.

Importantly, it is possible to implement many more pre-post rules with calcium than just the standard Hebbian and anti-Hebbian rules. For example, we could have a “fire together wire together” rule that penalizes synapses that are active when there is no postsynaptic spike, but doesn’t penalize inactive synapses at the time when there is a postsynaptic spike (DNP, [Fig 3B7](#)).

From a combinatoric standpoint, one can imagine 27 possible pre-post learning rules, because there are three possible synaptic scenarios (pre only, post only, pre and post) and there are three possible outcomes for each case (potentiate, depress, nothing). (If we allow for synapses to undergo plasticity when there is neither presynaptic nor postsynaptic activity, there are 81 possibilities, but we assume that synapses are stable in the absence of any activity.)

However, the two-threshold calcium control hypothesis prohibits some of these scenarios, because the $[Ca^{2+}]$ from the “both” (pre and post) scenario must be the sum of the $[Ca^{2+}]$ from the “pre only” and “post only” scenarios, and $[Ca^{2+}]$ is always non-negative. This imposes a constraint on the plasticity rules that can be implemented by the calcitron. For example, it would be impossible to implement PPN, because if the presynaptic and postsynaptic calcium are both above the potentiation threshold, their sum must necessarily also be above the potentiation threshold. If we assume the standard order of the plasticity thresholds, i.e., that $\theta_D < \theta_P$, we are left with 14 potential pre-post rules that can be implemented in the calcitron ([Fig 3](#)). (We note that if the plasticity thresholds are reversed, as in Purkinje neurons, there are also 14 implementable plasticity rules, some of which overlap with the rules implementable




Fig 3. Possible plasticity rules for presynaptic input- and spike-dependent calcium. (A) First scenario for calcium thresholds, where the depressive region is larger than the pre-depressive region, i.e., $\theta_p - \theta_D > \theta_D$. Setting different coefficient values (α and γ in Eq. (6)) for the local (C_{local}^i) and backpropagating spike-dependent (C_{BAP}) [Ca^{2+}] can lead to thirteen possible learning rules in the case of binary input and output. Vertical lines indicate the values of α that would be needed to induce depression (blue line) or potentiation (red line) with presynaptic input alone, horizontal lines indicate the values of γ that would be required to induce plasticity with a postsynaptic spike alone, and diagonal lines indicate the values of α and γ that would induce plasticity at activated synapses in the presence of a postsynaptic spike. Asterisk indicates rule (DDD) that can't be implemented under the alternative threshold scenario from panel (C). (B1–B13) Each of the 13 regions from panel (A) represented as a bar plot. (C) Second calcium-threshold scenario, where the pre-depressive region is larger than the depressive region, i.e., $\theta_D > \theta_p - \theta_D$. Asterisk indicates rule (NNP) that can't be implemented under the first threshold scenario from (A). (D) Bar plot for the NNP rule from panel (C). See S1 Fig for reversed plasticity thresholds, i.e., when $\theta_D > \theta_p$.

<https://doi.org/10.1371/journal.pcbi.1012754.g003>

via the standard threshold order (e.g., NNN) and some of which are different (e.g., PPD, as described above. See S1 Fig). Some rules, like PPN, cannot be implemented irrespective of the threshold order.)

Moreover, for any given set of plasticity thresholds θ_p and θ_D , only 13 out of 14 plasticity rules may be implemented by adjusting the coefficients for the local (C_{local}^i) and backpropagating spike-dependent (C_{BAP}) (α and γ in Eq. (6)), depending on whether the thresholds fall into one of two scenarios. In the first scenario (Fig 3A and 3B13), the depressive region is larger than pre-depressive region, i.e., $\theta_p - \theta_D > \theta_D$. In this scenario, it is impossible to implement classic no-penalty Hebbian plasticity (NNP). This is because if the pre- and postsynaptic calcium are both below the depressive threshold, as required for this form of plasticity, we have that $C_{local}^i + C_{BAP} < 2\theta_D < \theta_p$, so potentiation is impossible. On the other hand, if the pre-depressive region is larger than the depressive region, i.e., $\theta_D > \theta_p - \theta_D$ (Fig 3C and 3D), then it is possible to implement this no-penalty Hebbian plasticity (NNP, Fig 3D). However, a rule where all pre-post scenarios produce depression (DDD), which is possible in the first

scenario (Fig 3B8), would be impossible in this second scenario, because we would have that $C_{local}^i + C_{BAP} > 2\theta_D > \theta_P$. In other words, in this scenario, if C_{local}^i and C_{BAP} individually are sufficient to produce depression, their sum necessarily produces potentiation. See also [14], who perform a similar exploration in the context of STDP.

Frequency-dependent pre-post learning rules in a rate model

In the previous section, we assumed binary inputs and outputs (i.e., x_i and $\hat{y}$ are either 0 or 1). If we instead consider a rate model, where x_i and $\hat{y}$ represent the firing rates of the presynaptic and postsynaptic neuron, respectively, it is possible to implement frequency-dependent plasticity whose outcome depends on both the pre-synaptic and postsynaptic firing rate (Fig 4A). Experimentally, we note that presynaptic-only frequency-dependent plasticity has been observed, where low frequency stimulation causes depression and high frequency stimulation causes potentiation [38–40]. To replicate classical (presynaptic only) frequency-dependent plasticity, we set all coefficients other than α to 0. The direction of plasticity at any given synapse will depend on the input strength, i.e., the value of x_i . According to Eqs. 6




Fig 4. Frequency-dependent pre- and post-synaptic plasticity in rate-based models. (A) Calcium-dependent plasticity in a rate model where $\alpha, \gamma > 0$. In the absence of postsynaptic spikes, sufficiently strong presynaptic inputs (x_i) alone can generate plasticity (green bars), postsynaptic firing alone ($\hat{y}$) can induce plasticity even at inactive synapses (pink bars), and the combination of presynaptic input and postsynaptic spiking can sum to induce plasticity. (B1–B5) $[Ca^{2+}]$ (binned into regions of no change (white), depressive (blue) or potentiative (red)) as a function of presynaptic (x-axis) and postsynaptic (y-axis) firing rate. Each panel has a different value for α ($[Ca^{2+}]$ per presynaptic spike) and γ ($[Ca^{2+}]$ per postsynaptic spike). Values for θ_P and θ_D as in A.

<https://doi.org/10.1371/journal.pcbi.1012754.g004>

and 12, if x_i is between $\frac{\theta_D}{\alpha}$ and $\frac{\theta_P}{\alpha}$, w_i depresses, if x_i is larger than $\frac{\theta_P}{\alpha}$, w_i potentiates, otherwise there is no change (Fig 4B1).

If we choose nonzero values for both α and γ , the direction of plasticity at each synapse will depend on both the rate of its presynaptic input x_i as well as the output rate $\hat{y}$. By choosing different values for α and γ , it is possible to emphasize the effect of the presynaptic input versus the postsynaptic output on the $[Ca^{2+}]$, and consequently the plasticity (Fig 4B2–4B4).

It is similarly possible to have postsynaptic-only frequency-dependent plasticity, which depends only on the output strength, $\hat{y}$, by setting all coefficients other than γ to 0. Now, the plasticity at all synapses will depend only on the output firing rate, $\hat{y}$. If $\hat{y}$ is between $\frac{\theta_D}{\gamma}$ and $\frac{\theta_P}{\gamma}$, all synapses depress, if $\hat{y}$ is larger than $\frac{\theta_P}{\gamma}$, all synapses potentiate, otherwise no change occurs at any synapse (Fig 4B5).

Because we are using a step function for the learning rate $\eta(C)$ (Eq. (11)), the frequency of the input and output affect only the direction of the synaptic change, not its magnitude. However, if desired, it is possible to implement more biologically realistic frequency-dependent rules where the magnitude of plasticity is more precisely titrated by the $[Ca^{2+}]$ by defining $\eta(C)$ as a soft threshold function (e.g., a sum of sigmoid functions) instead of a step function (see [30]).

Unsupervised learning of repetitive patterns with heterosynaptic plasticity

One task we might want a neuron to perform is to learn to recognize a particular input pattern, i.e., by emitting a spike in response to that pattern while not firing in response to other input patterns. This task is usually implemented as a supervised learning task (such as in the perceptron algorithm, see below), but it is also possible for a neuron to learn to recognize specific patterns in an unsupervised fashion using heterosynaptic plasticity. Instead of directly telling the neuron which inputs should elicit a spike, it is possible to teach a neuron using heterosynaptic plasticity to spike only in response to frequently repeated “signal” patterns, while ignoring sporadic random “noise” patterns (Fig 5A). For this rule, we will only consider presynaptic and heterosynaptic plasticity, so we will adjust the coefficients for $C_{local}(\alpha)$ and $C_{het}(\beta)$; γ and δ will be set to 0.

We assume that both the signal and noise input patterns here have the same sparsity, i.e., that there are always k out of N active synapses at every time step. We enforce that input patterns always heterosynaptically depress non-active synapses by setting β in Eq. (6) such that $\theta_D < \beta k w_{min} < \beta k w_{max} < \theta_P$ and homosynaptically potentiate active synapses by enforcing that $\beta k w_{min} + \alpha > \theta_P$ (Fig 5B).

Every time an input pattern is presented to the neuron, the synapses that are active in that pattern will be potentiated and the synapses that are inactive will be depressed. If $\eta(C) < 1$, the potentiation and depression will occur gradually as the patterns are presented, and each synapse thus retains a “memory” of the recent input history. This creates a sort of competition between input patterns, because each pattern potentiates its active synapses while depressing synapses that were active in other patterns. Patterns that are repeated frequently (the signal patterns) will tend to dominate this competition, as the other (noise) non-repeated patterns will cancel out each other’s plastic influence on the synaptic weights. Over time, the calcitron will tend to strengthen the synapses associated with the frequently repeated signal pattern and depress other synapses, eventually inducing the calcitron to spike only in response to the signal pattern (Fig 5C1–5C5).




Fig 5. Learning to recognize repetitive patterns with heterosynaptic plasticity. (A) A “signal” pattern is presented repeatedly to the neuron interspersed with non-repeating random “noise” patterns of the same sparsity. (B) Within each input pattern (both signal or noise) inactive synapses depress (above θ_D at left) due to the heterosynaptic calcium, whereas active synapses will potentiate from the sum of heterosynaptic calcium, C_{het} , and local calcium C^i_{local} (above θ_P at right) (C1) Signal and noise patterns are presented to the neuron. (S: signal, N: noise). (C2) Spiking output of the calcitron. Black: no spike, Red: spike. An ‘x’ marker indicates incorrect output (e.g., no spike in response to a signal pattern, or a spike in response to a noise pattern), filled circles indicate correct outputs. Note the increase in correct spiking output over time. (C3–C5) Calcium, plasticity, and weights over time respectively as the input patterns in C1 are presented.

<https://doi.org/10.1371/journal.pcbi.1012754.g005>

We note that because this learning rule only depends on the weighted inputs and not the activation function g , the learning rule is similar in structure and function to Oja’s rule for neural learning of principle components [4]. For additional work on plasticity models for unsupervised learning of input patterns, see [41,42].

“One-shot flip-flop” (1SFF) model of behavioral time-scale plasticity (BTSP)

Recent experimental findings in the hippocampus have revealed a novel form of plasticity, known as behavioral time scale plasticity (BTSP) [31,34,33,43]. A mouse running on a treadmill can spontaneously form a place field in a CA1 hippocampal neuron when the soma of the neuron is injected with a strong current, inducing a plateau potential (this also occurs spontaneously in vivo via a supervising signal from the entorhinal cortex, see [32,33]). After a single induction, this plateau potential results in the neuron exhibiting a place field selective to the mouse's location few seconds before or after the time of the plateau potential. Moreover, this place field can be modified; if a second plateau potential is induced while the mouse is at a different location near the first place field, the place field will shift to the new location, thus “overwriting” the first place field [34]. In previous work, we showed how the calcium control hypothesis might be able to explain various aspects of these experimental results [30].

BTSP can be reduced to a more abstract, idealized form of learning. If we ignore the precise temporal dynamics, BTSP can be thought of as a form of supervised one-shot learning, wherein a supervisory signal potentiates all synapses that were coactive with it (i.e., the input pattern), while depressing all other synapses. If we consider a bistable synaptic weight that can either be in a “potentiated” (1, i.e., w_{max}) or “depressed” (0, i.e., w_{min}) position, this supervisory plateau signal in BTSP is effectively a “write” command which tells the neuron to store the state of its inputs, i.e., active (1) or inactive (0) as its weights. If this ‘stored pattern’ contains at least k active inputs, and if $w_{max} * k$ is larger than the spike threshold b , then the neuron will spike whenever it sees that input pattern again. We call this “one-shot flip-flop” learning, because the supervising signal overwrites all of the neuron's weights to new binary values in a single timestep, ensuring that the neuron only fires in response to the newly stored pattern (or to a pattern that overlaps with the new pattern by j active synapses such that $w_{max} * j > b$).

To implement this in the calcitron, we again assume binary inputs, and we also enforce binary synaptic weights by setting $\eta(C) = 1$ in both the depressive and potentiative regions of calcium, so a synapse will immediately be set to w_{min} whenever the $[Ca^{2+}]$ in the depressive region and to w_{max} whenever there is a potentiative $[Ca^{2+}]$ value. As a simple demonstration of this instantaneous behavior, we show that a series of calcium pulses of different strengths leads to immediate weight changes of maximal magnitude (Fig 6A1–6A3). For the 1SFF rule, we will consider the presynaptic activity and the supervisory signal while ignoring heterosynaptic plasticity and the postsynaptic spike, so we will only adjust the coefficients for C_{local}^i (α) and C_{sprv} (δ); β and γ will be set to 0. Assuming that the supervisory signal $Z = 1$ when active, one-shot flip-flop learning can be implemented in the calcitron by enforcing $\alpha < \theta_d$ (presynaptic input alone doesn't cause plasticity) $\theta_d < \delta < \theta_p$ (the supervisory signal depresses all inactive synapses) and $\delta + \alpha > \theta_p$ (the supervisory signal potentiates all active synapses) (Fig 6B).

To illustrate the effect of one-shot flip-flop learning, we present the calcitron with a repeated sequence of 4 binary input patterns, labeled with the numbers 0–3, to simulate a mouse running on a circular track (Fig 6C1). These 4 input patterns can be thought of as locations on a track traversed by the mouse. We then provide a supervisory signal at random intervals (Fig 6C2). Every time the supervisory signal is presented, the binary state of the weights at the following time step are set to the binary state of the inputs that were active in tandem with the supervisory signal (Fig 6C3–6C5). The next time the neuron encounters the input pattern (location) at which it previously received a supervisory signal, the neuron fires. In other words, the supervisory signal turned the neuron to a place cell for that location. When another supervisory signal comes at a different location, the previous place field is overwritten with a new place field at the new location (Fig 6C6).




Fig 6. “One-shot flip-flop” (1SFF) plasticity. (A1) Fixed points ($F(Ca(t))$, black line, left y-axis) and learning rates ($\eta(Ca(t))$, pink line, right y-axis) for the different regions of $[Ca^{2+}]$. For 1SFF learning, the learning rate is set to 1 in the depressive and potentiative regions of $[Ca^{2+}]$ for immediate switch-like plasticity. (A2) Exemplar stimulus illustrating plasticity dynamics. An instantaneous $[Ca^{2+}]$ pulse is generated at three timesteps over the course of the experiment. (A3) Synaptic weights over time in response to stimulus presented in A2. (B) 1SFF plasticity rule. Local input alone does not reach the depression threshold, a plateau potential alone induces a depressive $[Ca^{2+}]$, but local input combined with a plateau potential induces potentiation. (C1) A repeated sequence of input patterns (0,1,2,3) corresponding to locations on a circular track that a mouse traverses at each timestep as it runs multiple laps. (C2) Externally generated supervisory signal (plateau potential), Z , presented at different locations over the course of the experiment. (C3–C5) $[Ca^{2+}]$, plasticity, and weights, respectively for each time step. Synaptic inputs at the time of the supervisory signal are “written” to the synaptic weights at the following respective time step. (C6) Calcitron output at each time step. Red circles indicate spikes. Note that the neuron spikes at the location at which the supervisory signal occurred in the previous lap.

<https://doi.org/10.1371/journal.pcbi.1012754.g006>

Homeostatic plasticity with both internal and circuit mechanisms

Another important form of plasticity does not involve storing new information per se, but rather maintaining a regular average firing rate. This form of plasticity is known as homeostatic plasticity. Homeostatic plasticity can be important for the health of the neuron (i.e.,

too much firing can deplete neuronal resources, potentially leading to cell death) as well as maintaining stability and regularity within neural circuits [44–47]. While some models of homeostatic plasticity involve synaptic competition for resources [48,49], we focus here on calcium-based mechanisms for homeostatic plasticity.

Before we demonstrate how homeostatic plasticity can be implemented in the calcitron, we first note that our solution for the calcium-dependent mechanism of homeostatic plasticity differs from what has been observed experimentally, which involves a different calcium signaling pathway and takes place over much longer time scales (many hours, instead of seconds or minutes) [46]. One form of experimentally-observed homeostatic plasticity seems to depend on *somatic* calcium concentrations – when a neuron is firing too slowly, the low calcium levels set off a signaling pathway to globally increase synaptic strengths, and when a neuron fires too much, the high calcium levels initiate a signaling pathway to decrease synaptic strengths. This form of homeostatic plasticity thus involves a different relationship between calcium and the direction of plasticity – low levels of somatic calcium induce potentiation, medium levels induce no change, and high levels induce depression [46]. For the purposes of our work, however, we propose an alternative strategy that depends on the calcium at the spine with the standard plasticity thresholds: low levels of calcium induce no change, medium levels induce depression, and high levels induce potentiation. As such, the calcitron version of homeostatic plasticity is a speculative exploration of how the brain could potentially implement homeostatic plasticity using the standard synaptic plasticity mechanisms and calcium thresholds.

In a rate model version of the calcitron (e.g., where the activation function g is linear), we can think of homeostatic plasticity as the problem of trying to keep the output firing rate, $\hat{y}$, within a target range, between $\hat{y}_{min}^*$ and $\hat{y}_{max}^*$. We will also assume that due to the postsynaptic refractory period, neurons have a maximum rate that it is physically possible to fire at, $\hat{y}_{max}$, and that neurons also have a minimum physically possible firing rate $\hat{y}_{min}$ (trivially, neural firing rates can't be negative, so we will always have $\hat{y}_{min} \geq 0$). In general, then, we have that $\hat{y}_{min} \leq \hat{y}_{min}^* \leq \hat{y}_{max}^* \leq \hat{y}_{max}$.

For simplicity, if we assume that the synaptic inputs to the calcitron are binary (even though the output is not), there are broadly two strategies we can take with a homeostatic plasticity rule. One strategy, which we term “global homeostasis”, is that whenever the calcitron's output $\hat{y}$ is too large (or too small), we can depress (potentiate) all the calcitron's synapses irrespective of their input. This will eventually result in the calcitron's output being within the target range on average.

Although global homeostasis may be effective if the calcitron's output is consistently too low or too high irrespective of the input, globally depressing or potentiating all synapses is a drastic measure that can destroy previously stored information (unless scaling is multiplicative, see [44,45]). Instead, we can use a more fine-grained approach, “targeted homeostasis”, which only modifies the synapses that were active at the time when the output firing rate was out of range, thus only correcting “errant” synapses. (See also [50,51] regarding the consequences of the global vs. local strategies of homeostatic plasticity when considering a realistic dendritic tree.) For all the homeostatic plasticity rules presented here, we do not require C_{het} , so we set $\beta = 0$. The use of other calcium sources will depend on the particular plasticity rule.

For the global homeostasis strategy, we will always set $\alpha = 0$ so we can ignore presynaptic calcium, and for the targeted homeostasis strategy we will always set $0 < \alpha < \theta_d$ so the presynaptic calcium can break the symmetry between active and inactive synapses but isn't so large that it would be sufficient on its own to induce plasticity.

We first consider whether either of these homeostatic plasticity strategies can be implemented using exclusively internal mechanisms (i.e., without an external supervisory signal). If our neuron is firing too strongly, i.e. $\hat{y} > \hat{y}_{max}^*$, we would like the resultant calcium

to be in the depressive $[Ca^{2+}]$ region, but we also want to ensure that if the neuron is firing very strongly (i.e., $\hat{y} = \hat{y}_{max}$) that the $[Ca^{2+}]$ doesn't cross over into the potentiative region. For the global homeostasis strategy, this can be implemented by setting $\gamma > \frac{\theta_D}{\hat{y}_{max}^*}$ and $\theta_p > \gamma \hat{y}_{max}$, and in the targeted homeostasis strategy by setting $\gamma > \frac{\theta_D - \alpha}{\hat{y}_{max}^*}$ and $\theta_p > \gamma \hat{y}_{max} + \alpha$.

Unfortunately, setting the parameters in this manner creates a problem: because $\theta_p > \gamma \hat{y}_{max}$ in both strategies, there is no longer any way to potentiate synapses when the firing rate is too low, as even the maximum physically possible firing rate will not produce enough $[Ca^{2+}]$ to induce potentiation, so certainly firing rates that are too low will not be able to induce potentiation. We therefore must instead rely on an external supervisory mechanism to potentiate synapses when the calcitron's output is too low.

To construct a potentiation supervisor, we consider a simple disinhibitory circuit. The calcitron forms a synapse onto an inhibitory neuron, which inhibits a supervisory neuron that supervises the calcitron. When the calcitron's output rises above $\hat{y}_{min}^*$, the inhibitory neuron is active, thus preventing the supervisory neuron from sending the calcitron an supervisory signal. When the calcitron's output falls below $\hat{y}_{min}^*$, however, the inhibitory neuron becomes inactive, thus permitting the supervisory neuron to send a potentiative supervisory signal $Z_p = 1$ to the calcitron (Fig 7A and 7B). For the global homeostasis strategy, by setting $\delta > \theta_p$, the supervisor induces potentiation at all synapses whenever $\hat{y} < \hat{y}_{min}^*$. For the targeted homeostasis strategy, we set $\delta > \theta_p - \alpha$ to ensure that active synapses are potentiated, and we apply an additional constraint that $\delta < \theta_D$, to ensure that non-active synapses aren't affected by the supervisory signal (Fig 7C).

It is also possible to implement both homeostatic potentiation and depression using external supervisors, instead of using an internal supervisor for depression. To do this, we set $\gamma = 0$, so the postsynaptic spike itself doesn't induce calcium influx, and we use the same disinhibitory circuit mechanism for depression as described above. To implement an external depression supervisor, we consider an additional circuit mechanism where the calcitron also synapses directly onto a new "depression supervisor" neuron, which synapses back onto the calcitron with a supervising synapse. This depression supervisor will be active whenever the calcitron's firing rate $\hat{y}$ exceeds $\hat{y}_{max}^*$ (Fig 7D).

Importantly, the depression supervisor and potentiation supervisor give supervisory signals of different strengths, Z_D and Z_p , respectively. Without loss of generality, we can set $\delta = 1$, so it is only necessary to set the magnitude of the supervising signals. For the global homeostasis strategy, we set $Z_D > \theta_D$ and $Z_p > \theta_p$, which provides the calcitron with a global potentiative signal when $\hat{y} < \hat{y}_{min}^*$ and a global depressive signal when $\hat{y} > \hat{y}_{max}^*$ (Fig 7E). For targeted homeostasis, we set $Z_D > \theta_D - \alpha$ and $Z_p > \theta_p - \alpha$, and we enforce that $Z_D < Z_p < \theta_D$ so that the supervising signals only modify active synapses (Fig 7F).

To compare the two different supervisory circuits (external potentiation and internal depression vs. external potentiation and depression) and the targeted vs. global strategies, we created a calcitron whose even-numbered synapses were initialized to small weights and whose odd-numbered synapses were initialized to large weights. We then randomly present input patterns that either only activate the even-numbered synapses ("even patterns") or only activate the odd-numbered synapses ("odd patterns"). Initially, the even patterns produce an output which is too low (i.e., below $\hat{y}_{min}^* = 1$) and the odd patterns produce an output which is too high (i.e., above $\hat{y}_{max}^* = 3$). Over the course of presenting the patterns multiple times, the homeostatic mechanisms succeeded to increase the weights for the even synapses and decrease the weights for the odd synapses such that eventually the neuron fired within the target output range ($\hat{y}_{min}^* < \hat{y} < \hat{y}_{max}^*$) in response to both even and odd patterns. All four combinations of supervisory circuit structure and plasticity strategy (global vs. targeted) succeeded in this task,




Fig 7. Different mechanisms for homeostatic plasticity. (A) Supervision circuit for homeostatic plasticity using only a potentiation supervisor. If the calcitron ("C") fires above the target minimum rate (i.e., $\hat{y} > \hat{y}_{min}^*$) it activates an inhibitory population ("I"), which prevents the potentiation supervisor

(S_p) from producing a supervisory signal. If the calcitron's output falls below its target range (i.e., $\hat{y} < \hat{y}_{min}^*$) the supervisor is disinhibited, sending a potentiative calcium signal (Z_p) to the calcitron. (B) Plasticity rule for global homeostatic plasticity using an internal mechanism for depression and a circuit mechanism for potentiation, as in (A). Here $\hat{y}_{min}^* = 1$ and $\hat{y}_{max}^* = 3$. (B1) Overly strong outputs ($\hat{y} > 3$) produce sufficient calcium to depress all synapses; Overly weak outputs ($\hat{y} < 1$) result in the activation of S_p , setting $Z_p = 1$, potentiating all synapses. Note that because this is a “global” strategy, the presynaptic input x_i does not affect the plasticity outcomes. (B2) Two input patterns (only even synapses or only odd synapses) are presented to the neuron in random order. (B3) Z_p occurs whenever $\hat{y} < 1$. (B4–B6) $[Ca^{2+}]$, plasticity, and weights over the course of the simulation. Even-numbered synapses are initialized to low weights; odd-numbered synapses are initialized to large weights. (B7) Neural rate output. Blue ‘x’ indicates output below $\hat{y}_{min}^*$ (lower dashed line), red ‘x’ indicates output above $\hat{y}_{max}^*$ (upper dashed line), green circles indicate output in the acceptable range. (C) Plasticity rule for targeted homeostatic plasticity using the postsynaptic firing rate and the circuit from (A) as well as local $[Ca^{2+}]$. Here, only synapses that are active when the firing rate is too low (even synapses) are potentiated, and only synapses that are activated when the firing rate is too high (odd synapses) are depressed, as the C_{local} is necessary to bring the C_{total} above the plasticity thresholds. (D) Supervision circuit for homeostatic plasticity using both a potentiation (S_p) and depression (S_D) supervisor. In addition to the disinhibitory circuit for the control of S_p as in (A), when the calcitron's output is above the target output range ($\hat{y} > \hat{y}_{max}^*$), a depression supervisor (S_D) is activated, sending a depressive signal (Z_D) to the calcitron. (E) Plasticity rule for global homeostatic plasticity using both S_p and S_D . Here, the postsynaptic spike-dependent calcium C_{BAP} is not used; only the supervisory calcium signals C_{SPRV} are necessary. Note the different strengths of S_p and S_D in E1 and E3. (F) Plasticity rule for targeted homeostatic plasticity using S_p and S_D in combination with C_{local}^i at active synapses, so even and odd synapses will be differentially potentiated and depressed.

<https://doi.org/10.1371/journal.pcbi.1012754.g007>

demonstrating that the same desired functional result can emerge from different underlying mechanisms (Fig 7).

Perceptron learning algorithm with calcium-based plasticity

We now show that it is possible to implement the perceptron learning algorithm with calcium-based plasticity. The perceptron learning algorithm (Rosenblatt, 1958) is the procedure by which a single threshold-linear neuron (see Eq. (1)) can learn to solve classification tasks, such as distinguishing between images of cats and dogs, by modifying its synaptic weights.

The perceptron learning rule is a supervised learning rule. Namely, each input pattern comes with an associated target outcome – that the neuron should either spike or not spike. Formally, we have a set of P input patterns $\{\mathbf{x}^1 \dots \mathbf{x}^\mu \dots \mathbf{x}^P\}$ and associated labels $\{y^1 \dots y^\mu \dots y^P\}$, where the bolded $\mathbf{x}^\mu$ is an N -dimensional vector of activity of the μ th input pattern and $y^\mu \in \{0, 1\}$ is the associated target label, or class. The goal of the perceptron learning algorithm is to ensure that the neuron's output on each pattern ($\hat{y}^\mu$) matches the target label, i.e., $\hat{y}^\mu = y^\mu$.

The perceptron rule states that if the neuron makes an error on an input pattern $\hat{y}^\mu \neq y^\mu$, we increase or decrease the perceptron's synaptic weights in a manner proportional to the input vector. Formally, if $\hat{y}^\mu = 0$ and $y^\mu = 1$ (a false negative; the neuron should have spiked but didn't) we update each weight w_i according to the rule $w_i \leftarrow w_i + \eta x_i$, where η is the learning rate. If $\hat{y}^\mu = 1$ and $y^\mu = 0$ (a false positive; the neuron spiked when it wasn't supposed to) we update each weight w_i according to the rule $w_i \leftarrow w_i - \eta x_i$. If the neuron produced the correct output, $\hat{y}^\mu = y^\mu$, we don't modify the weights. The change in each synaptic weight at each time step, $\Delta w_i(t) = w_i(t+1) - w_i(t)$, for the perceptron learning rule can be described by Table 1.

We first note that it is only possible to exactly replicate the perceptron learning rule with the linear calcium-based rule (Eq. (9)). However, because we would like to demonstrate perceptron-like dynamics with the more biologically realistic FPLR rule (Eq. (11)), we propose an “asymptotic perceptron learning rule” which functions similarly to the original perceptron rule, except that weights increase or decrease asymptotically towards w_{min} or w_{max} instead of being able to increase or decrease indefinitely, as in the standard perceptron rule.

Formally, if $\hat{y}^\mu = 0$ and $y^\mu = 1$ (false negative), we update w_i according to the rule $w_i \leftarrow w_i + \eta(w_{max} - w_i)x_i$. If $\hat{y}^\mu = 1$ and $y^\mu = 0$ (false positive) we update each weight w_i according to the rule $w_i \leftarrow w_i + \eta(w_{min} - w_i)x_i$ (Table 2). Because it is always the case that

Table 1. Weight update Δw in the standard perceptron learning rule.

	$y = 0$	$y = 1$
$\hat{y} = 0$	0	ηx
$\hat{y} = 1$	$-\eta x$	0

<https://doi.org/10.1371/journal.pcbi.1012754.t001>

Table 2. Weight update Δw in the asymptotic perceptron learning rule.

	$y = 0$	$y = 1$
$\hat{y} = 0$	0	$\eta_p (w_{\max} - w) x$
$\hat{y} = 1$	$\eta_D (w_{\min} - w) x$	0

<https://doi.org/10.1371/journal.pcbi.1012754.t002>

$w_{\min} < w_i < w_{\max}$, the sign of the weight update in the asymptotic perceptron rule for each case is consistent with the sign of the update in the original perceptron rule; the rules only differ in whether the weights change linearly or asymptotically as a function of the current weight.

To implement the perceptron learning rule in the calcitron, we again stipulate that the synaptic inputs are binary, so we only have to worry about the direction of synaptic change, not the magnitude. Because the perceptron is a supervised learning algorithm, we will also have a supervisory signal. In our first attempt at implementing the perceptron, the supervisory signal (the “label supervisor”) will simply indicate the value of label (Fig 8A). In other words, for each input pattern $\mathbf{x}^\mu$, we have $Z^\mu = y^\mu$. The challenge here is to set the calcium thresholds and coefficients such that each quadrant of Table 2 is satisfied. (For the perceptron rule, we again ignore C_{het} , so we set $\beta = 0$.)

To satisfy the upper left quadrant ($\hat{y}^\mu = 0$ and $y^\mu = 0$), we stipulate that $\alpha < \theta_D$, so that in the absence of a postsynaptic spike and a supervising signal the presynaptic $[Ca^{2+}]$ alone is too low to induce plasticity. For the lower left quadrant, ($\hat{y}^\mu = 1$ and $y^\mu = 0$, false positives) we require that a postsynaptic spike induces depression at an active synapse but does nothing to inactive synapses. To accomplish this, we enforce $\gamma < \theta_D$ and $\theta_D < \gamma + \alpha < \theta_p$. For the upper right quadrant, ($\hat{y}^\mu = 0$ and $y^\mu = 1$, false negatives), we want the supervisor in the absence of a postsynaptic spike to induce potentiation at active synapses but do nothing at inactive synapses. For this we require that $\delta < \theta_D$ and $\delta + \alpha > \theta_p$.

A problem arises, however, when we get to the lower right quadrant ($\hat{y}^\mu = 1$ and $y^\mu = 1$, true positives). When there is both a postsynaptic spike and a supervising signal, we require that none of the synapses will be updated. But we already enforced that $\delta + \alpha > \theta_p$, so now when we have both a postsynaptic spike and a supervisory signal, we have a $[Ca^{2+}]$ of $\delta + \alpha + \gamma$ at active synapses, which is certainly greater than θ_p ! We solve this by adding a third region of $[Ca^{2+}]$, the post-potentiative neutral zone, or PPNZ ($\theta_{PPNZ} > \theta_p$), wherein the $[Ca^{2+}]$ is so high that it ceases to potentiate synapses (there is some evidence for this, see [52], Fig 3). We then enforce that $\delta + \gamma > \theta_{PPNZ}$, which gives us our fourth quadrant and allows us to reproduce the perceptron learning rule in its entirety (Fig 8B).

If we don’t wish to add a post-potentiative neutral zone, we can still implement the perceptron learning rule with the normal calcium thresholds if we use a “smart” supervisory signal. Instead of just telling the calcitron what the correct label is, we can have a “critic supervisor” which compares the calcitron’s output $\hat{y}^\mu$ to the target output y^μ , and gives supervisory signals of different strengths depending on whether the input pattern resulted in a false positive




Fig 8. Perceptron learning with the calcitron. (A) Supervision circuit for perceptron learning using a “target” supervisor. Whenever the target label is 1, the supervisor sends a potentiative supervisory signal Z_p to the calcitron. (B1) Plasticity rule for the perceptron with a target supervisor. Note that

an additional calcium threshold (dashed green line) for a post-potentiative neutral zone (PPNZ) where no plastic change occurs has been added to the plasticity rule. **(B2)** Six patterns, half of which are arbitrarily assigned to the positive class and half to the negative class ($y = 1$ or $y = 0$, respectively, see tick labels on x-axis) are repeatedly presented to the calcitron in random order over several epochs. **(B3)** Supervisory signal. Appears whenever the target label $y = 1$. **(B4–B6)** $[Ca^{2+}]$, plasticity, and weights over the course of the simulation. **(B7)** Calcitron output. Red circle: true positive, red 'x': false positive, black circle: true negative, black 'x': false negative. **(C)** Supervision circuit for perceptron learning using a “critic” supervisor. The supervisor compares the target label y to the calcitron output $\hat{y}$. If the trial was a false negative ($y = 1, \hat{y} = 0$), the supervisor sends a potentiative supervisory signal Z_p to the calcitron. If the trial was a false positive ($y = 0, \hat{y} = 1$), the supervisor sends a depressive supervisory signal Z_d to the calcitron. **(D1–D7)** Perceptron learning with the “critic” supervisory circuit. Note the different magnitudes of the supervisory signal in (D3) – the large signal corresponds to Z_p and the small signal corresponds to Z_d .

<https://doi.org/10.1371/journal.pcbi.1012754.g008>

or false negative. Importantly, the supervisor will no longer give any supervisory signal when the pattern is classified correctly, circumventing the problem we had with true positives using the label supervisor. Here, we will not explicitly describe the circuit necessary to construct such a supervisor, we will simply assume that some other brain circuit can perform the operation of comparing the calcitron’s output to the target output and produce the appropriate supervisory signals (Fig 8C).

Once we have a critic supervisor, it is no longer necessary to use the calcium from the postsynaptic spike, so we set $\gamma = 0$. Now the strategy for implementing the perceptron is straightforward. For true positives ($\hat{y} = 1$ and $y = 1$) and true negatives ($\hat{y} = 0$ and $y = 0$), there is no supervisory calcium, so we just have to ensure that an active synapse by itself doesn’t induce plasticity by setting $\alpha < \theta_D$. For false negatives ($\hat{y} = 0$ and $y = 1$) we need the presynaptic calcium combined with the active synapse to induce potentiation, while the supervisor alone doesn’t induce any plasticity at inactive synapses. To do this we set $\delta = 1$, and we construct a potentiation supervisor for false negatives, Z_p , such that $Z_p < \theta_D$ and $Z_p + \alpha > \theta_p$. For false positives ($\hat{y} = 1$ and $y = 0$), we similarly construct a supervisor Z_d , such that $Z_p < \theta_D$ and $\theta_D < Z_d + \alpha < \theta_p$. These constraints satisfy the rules specified in Table 2 (and Table 1) in a much simpler manner.

To illustrate that the calcitron can indeed implement the perceptron learning rule just by setting the calcium coefficients and thresholds, we performed a simple classification experiment using a calcitron with the constraints and critic supervisor described above. We generated six binary patterns with 24 synapses each. Half of the patterns were arbitrarily assigned to the positive class and half to the negative class ($y = 1$ or $y = 0$, respectively). We repetitively presented these patterns in random order to both to the “critic” and “target” calcitron supervisory circuits. Initially, the calcitron makes mistakes on some of the patterns, of both the “false positive” and “false negative” variety. After a sufficient number of presentations of each input pattern, however, both supervisory circuits succeed to ensure that the calcitron correctly classifies all patterns (Fig 8).

We have thus demonstrated that the perceptron learning algorithm can be implemented with the calcitron, either with a “label supervisor” – which requires an additional plasticity threshold, or with a “critic supervisor” which can use the standard plasticity thresholds.

Discussion

Summary

We have shown that the calcitron, a simple model neuron that uses four different sources of calcium to modify its synaptic weights, can implement a wide variety of learning and plasticity rules. We have demonstrated that merely by appropriately setting the amount of calcium influx synapses receive from each calcium source and the calcium thresholds for plasticity, it is possible to reproduce classic learning and rules, such as various forms of Hebbian and anti-Hebbian learning (Figs 2 and 3), frequency-dependent plasticity (Fig 4), and

unsupervised recognition of frequently repeating input patterns (Fig 5). Moreover, by devising simple neural circuits to provide supervisory signals, we show how the calcitron can implement BTSP-inspired one-shot learning (Fig 6), homeostatic plasticity (Fig 7), and perceptron learning (Fig 8).

The calcium control hypothesis was originally developed as an explanation for Hebbian and anti-Hebbian plasticity [18]. Subsequent mathematical formulations of the calcium control hypothesis were designed to reproduce frequency-dependent and spike timing-dependent plasticity [14,16,18,30]. With the calcitron, however, we have expanded these earlier results into a generalized, simple model which can explain and predict a much wider set of plasticity and learning results from first principles. The calcitron can thus be helpful in providing an intuition for the range of possible calcium-based mechanisms underlying experimentally-observed plasticity, including forms of plasticity that are as-yet undiscovered.

The mathematical formalism of the calcitron also makes it easier to understand the limitations of calcium-based plasticity. The calcitron equations impose constraints on the potential learning rules that can emerge from the calcium control hypothesis, which can help us determine whether it is necessary to posit additional mechanisms, such as the supervisory circuit we proposed for homeostatic plasticity (Fig 7). Importantly, we do not claim that every possible learning rule we propose here is implemented in biology as we have described it. Rather, the calcitron is intended to serve as a framework for exploring different predictions of the calcium control hypothesis.

Future directions and additional biological considerations

The calcitron model was formulated to be as simple as possible in order to provide straightforward mathematical intuitions about the calcium basis of synaptic plasticity. This simplicity comes with certain drawbacks. For example, the choice of using a perceptron-like neuron model that does not temporally integrate information means that the calcitron, as formulated here, cannot implement spike timing-dependent plasticity (STDP). Prior work has shown that a leaky integrator model with calcium dynamics can indeed generate STDP ([14,16,30]). As such, future work can consider leaky integrator versions of the calcitron that can implement STDP and other temporally-sensitive plasticity rules. However, this temporal sensitivity introduces the need to finely tune calcium decay time constants in order to produce different plasticity rules, which introduces a level of complexity beyond the scope of the current work. It is also possible to explicitly model the kinetics of the relevant ion channels (i.e., VGCCs and NMDA receptors), adding even more biological detail at the expense of greater complexity.

The calcitron is also linear, both in terms of its input-output function (excluding the activation function) and in terms of how calcium from different sources combines to produce C_{total}^i . Linear point neurons are commonly used to model neural phenomena, although experimental and theoretical work indicate that real neurons may integrate information in a nonlinear fashion [53–64]. Of particular relevance is the superlinear activation function of the NMDA receptor, whose conductance exhibits sigmoidal sensitivity to local voltage at the synapse location [25,26]. A more realistic model for local calcium influx could thus include this voltage-dependent nonlinearity, in line with experimental work showing that NMDA spikes induce plasticity [65].

Another important feature of biological neurons not incorporated into the calcitron is the spatial distribution of synapses on a neuron's dendrites. This is especially relevant for heterosynaptic plasticity, which depends on the location of synapses relative to each other as well as their absolute location on the dendritic tree [66–70]. The calcitron can be augmented

to include location-dependent heterosynaptic plasticity on a single dendrite following the schemes of the clusteron [56] or the G-clusteron [57] or to explicitly include a branching dendrite to account for the branch-dependent hierarchical heterosynaptic plasticity effect we posited in previous work [27]. The spatial structure of the dendrite can also influence homeostatic plasticity [50,51] as well as how inhibitory inputs affect plasticity at excitatory synapses [71].

A crucial assumption we made in formulating the calcitron is that local calcium is always a function of presynaptic activity. If we had allowed for synapse-specific supervisory signals, it would be possible to implement arbitrary learning rules without any concern for constraints, as we could simply engineer a supervisor to potentiate or depress each synapse independently. In biology, however, it is possible that synapse-specific plasticity supervisors exist. One candidate for such a mechanism are the internal calcium stores in the endoplasmic reticulum (ER) [72,73]. The ER can be localized to individual spines [74], making it a candidate for targeted calcium release at selected synapses. It is usually assumed that calcium released from the ER is mediated by calcium influx via a process called calcium-induced calcium release, or CICR. If the initial calcium needed to induce CICR at the ER in individual spines comes from extracellular sources, CICR is still broadly consistent with our model; the internal calcium sources may simply be amplifying the calcium which enters via external sources. If, however, there are other endogenous mechanisms that can differentially induce calcium release from the ER at different synapses, the ER may be able to play the role of a synapse-specific supervisor. An additional candidate for synapse-specific supervisors are astrocytes, which target individual dendritic spines [75] and can influence plasticity via a variety of signaling pathways [76]. Either of these mechanisms can enable a less constrained, more diverse palette of neural learning rules when combined with the more classical calcium sources we suggested here.

With respect to the supervisory signal, we have focused on plasticity supervisors that operate via calcium mechanisms, such as are observed in Purkinje neurons [35] and possibly in the CA1 region of the hippocampus during BTSP [30,31,32]. However, there are other forms of plasticity that are more dependent on neuromodulation. For example, associative learning is often thought to require dopamine as a supervisory signal [77,78]. Future work can explore the interplay between calcium-based and neuromodulator-dependent mechanisms of plasticity.

There are a variety of ways to build on the calcitron. In this work, we have focused on building neurons and circuits that implement one learning rule at a time. It may be desirable, however, for a neuron to simultaneously implement, for example, Hebbian plasticity and homeostatic plasticity. The calcitron equations constrain which kind of learning rules can be simultaneously implemented in this fashion. It would be worthwhile to explore which learning rule combinations are mathematically valid in the calcitron framework.

Another natural extension of the calcitron would be to incorporate multiple calcitron nodes into a network to see how calcium-based plasticity affect network activity. It would be interesting to see how networks containing neurons with different calcium coefficients and thresholds (and thus different learning rules) would operate. Alternatively, it is possible to build large networks containing subnetworks with different calcium parameters, to explore how brain regions with different plasticity rules might interact with each other. The calcitron framework thus provides a theoretical tool to explore the network-level consequences of biologically realistic plasticity rules.

Methods

All simulations were performed on a Windows PC using Python and plotted using the Matplotlib library [79]. Parameters for each simulation are given in the table below (Table 3):

Table 3. Simulation parameters.

Fig	Panel	θ_D	F_D	η_D	θ_P	F_P	η_P	α	β	γ	δ	b	g	N	Z	Z_P	Z_D	η_{PPNZ}	F_PPNZ
Fig 1	E-G	0.5		-0.1	1		0.2												
	H-L	0.5	0	0.1	1	1	0.1												
Fig 2	A1-F1	0.5	0	0.1	0.8	1	0.1	0.4	0	0.45	0	-1.8	step	10					
	A2-F2	0.5	0	0.1	0.8	1	0.1	0.55	0	0.7	0	-1.8	step	10					
	A3-F3	0.5	0	0.1	0.8	1	0.1	0.4	0	0.3	0	-1.8	step	10					
	A4-F4	0.8	0	0.1	0.5	1	0.1	0.55	0	0.6	0	-1.8	step	10					
Fig 4	A	0.5	0	0.1	0.8	1	0.1	0.3	0	0.3	0								
	B1	0.5	0	0.1	0.8	1	0.1	0.3	0	0	0								
	B2	0.5	0	0.1	0.8	1	0.1	0.3	0	0.15	0								
	B3	0.5	0	0.1	0.8	1	0.1	0.3	0	0.3	0								
	B4	0.5	0	0.1	0.8	1	0.1	0.15	0	0.3	0								
	B5	0.5	0	0.1	0.8	1	0.1	0	0	0.3	0								
Fig 5	All	0.7	1	0.025	2.6	3	0.025	2	0.2	0	0	-8	step	12					
Fig 6	All	0.5	0	1	0.7	1	1	0.2	0	0	0.6	-6	step	14	1				
Fig 7	B	0.5	0	0.02	1	0.5	0.02	0	0	0.2	1	0	linear	20		1.05			
	C	0.5	0	0.02	0.7	0.5	0.02	0.35	0	0.06	1	0	linear	20		0.35			
	E	0.5	0	0.02	1	0.5	0.02	0	0	0	1	0	linear	20		1.05	0.55		
	F	0.5	0	0.02	0.7	0.5	0.02	0.3	0	0	1	0	linear	20		0.45	0.25		
Fig 8	B	0.5	0	0.2	0.6	1	0.2	0.45	0	0.1	1	-3	step	24		0.3		0	0.25
	D	0.6	0	0.2	0.9	1	0.2	0.45	0	0	1	-3	step	24		0.5	0.2		

<https://doi.org/10.1371/journal.pcbi.1012754.t003>

Supplementary information

S1 Fig. Possible plasticity rules for presynaptic input- and spike-dependent- calcium for reversed calcium thresholds. (A) First scenario for reversed calcium thresholds (i.e., $\theta_p < \theta_D$), where the potentiative region is larger than the pre-potentiative region, i.e., $\theta_D - \theta_p > \theta_p$. Asterisk indicates rule (PPP) that can't be implemented under the second threshold scenario from (C). (B1–B13) Each of the 13 regions from panel (A) represented as a bar plot. (C) Second calcium-threshold scenario, where the pre-potentiative region is larger than the potentiative region, i.e., $\theta_p > \theta_D - \theta_p$. Asterisk indicates rule (NND) that can't be implemented under the first threshold scenario from (A). (D) Bar plot for the NND rule from panel (C). (TIFF)

Author contributions

Conceptualization: Toviah Moldwin, Li Shay Azran.

Formal analysis: Toviah Moldwin, Li Shay Azran.

Funding acquisition: Idan Segev.

Investigation: Toviah Moldwin, Li Shay Azran.

Methodology: Toviah Moldwin.

Project administration: Toviah Moldwin.

Supervision: Toviah Moldwin, Idan Segev.

Validation: Idan Segev.

Visualization: Toviah Moldwin, Li Shay Azran.

Writing – original draft: Toviah Moldwin, Li Shay Azran.

Writing – review & editing: Toviah Moldwin, Li Shay Azran, Idan Segev.

References

1. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015;521(7553):436–44. <https://doi.org/10.1038/nature14539> PMID: [26017442](https://pubmed.ncbi.nlm.nih.gov/26017442/)
2. McCulloch WS, Pitts W. A logical calculus of the ideas immanent in nervous activity. *Bull Math Biophys*. 1943;5:115–33.
3. Shao F, Shen Z. How can artificial neural networks approximate the brain? *Front Psychol*. 2023;13:970214. <https://doi.org/10.3389/fpsyg.2022.970214> PMID: [36698593](https://pubmed.ncbi.nlm.nih.gov/36698593/)
4. Oja E. A simplified neuron model as a principal component analyzer. *J Math Biol*. 1982;15(3):267–73. <https://doi.org/10.1007/BF00275687> PMID: [7153672](https://pubmed.ncbi.nlm.nih.gov/7153672/)
5. Rosenblatt F. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol Rev*. 1958;65(6):386–408. <https://doi.org/10.1037/h0042519> PMID: [13602029](https://pubmed.ncbi.nlm.nih.gov/13602029/)
6. Rumelhart D, McClelland J. Parallel distributed processing: explorations in the microstructure of cognition. 1986 [cited 2021 Mar 27]. Available from: <https://cds.cern.ch/record/111912>
7. Rumelhart DE, Hinton GE, Williams RJ. Learning representations by back-propagating errors. *Nature*. 1986;323(6088):533–6. <https://doi.org/10.1038/323533a0>
8. Lillicrap TP, Santoro A, Marris L, Akerman CJ, Hinton G. Backpropagation and the brain. *Nat Rev Neurosci*. 2020;21(6):335–46. <https://doi.org/10.1038/s41583-020-0277-3> PMID: [32303713](https://pubmed.ncbi.nlm.nih.gov/32303713/)
9. Marr D. A theory of cerebellar cortex. *J Physiol*. 1969;202(2):437–70. <https://doi.org/10.1113/jphysiol.1969.sp008820> PMID: [5784296](https://pubmed.ncbi.nlm.nih.gov/5784296/)
10. Sievers M, Motta A, Schmidt M, Yener Y, Loomba S, Song K, et al. Connectomic reconstruction of a cortical column. *bioRxiv*. 2024;2024.03.22.586254. <https://doi.org/10.1101/2024.03.22.586254>
11. Millidge B, Tschantz A, Buckley CL. Predictive coding approximates backprop along arbitrary computation graphs. *Neural Comput*. 2022;34(6):1329–68. https://doi.org/10.1162/neco_a_01497 PMID: [35534010](https://pubmed.ncbi.nlm.nih.gov/35534010/)
12. Hinton G. The forward-forward algorithm: some preliminary investigations. 2022 [cited 2023 Sep 11]. Available from: <https://arxiv.org/abs/2212.13345v1>
13. Lillicrap TP, Cownden D, Tweed DB, Akerman CJ. Random synaptic feedback weights support error backpropagation for deep learning. *Nat Commun*. 2016;7:13276. <https://doi.org/10.1038/ncomms13276> PMID: [27824044](https://pubmed.ncbi.nlm.nih.gov/27824044/)
14. Graupner M, Brunel N. Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. *Proc Natl Acad Sci USA*. 2012;109(10):3991–6. <https://doi.org/10.1073/pnas.1109359109>
15. Lisman J, Goldring M. Evaluation of a model of long-term memory based on the properties of the Ca^{2+} /calmodulin-dependent protein kinase. *J Physiol (Paris)*. 1988;83(3):187–97. PMID: [2856110](https://pubmed.ncbi.nlm.nih.gov/2856110/)
16. Shouval HZ, Bear MF, Cooper LN. A unified model of NMDA receptor-dependent bidirectional synaptic plasticity. *Proc Natl Acad Sci U S A*. 2002;99(16):10831–6. <https://doi.org/10.1073/pnas.152343099> PMID: [12136127](https://pubmed.ncbi.nlm.nih.gov/12136127/)
17. Shouval HZ, Wang SS-H, Wittenberg GM. Spike timing dependent plasticity: a consequence of more fundamental learning rules. *Front Comput Neurosci*. 2010;4:19. <https://doi.org/10.3389/fncom.2010.00019> PMID: [20725599](https://pubmed.ncbi.nlm.nih.gov/20725599/)
18. Lisman J. A mechanism for the Hebb and the anti-Hebb processes underlying learning and memory. *Proc Natl Acad Sci U S A*. 1989;86(23):9574–8. <https://doi.org/10.1073/pnas.86.23.9574> PMID: [2556718](https://pubmed.ncbi.nlm.nih.gov/2556718/)
19. Lisman JE. Three Ca^{2+} levels affect plasticity differently: the LTP zone, the LTD zone and no man's land. *J Physiol*. 2001;532(Pt 2):285. <https://doi.org/10.1111/j.1469-7793.2001.0285f.x> PMID: [11306649](https://pubmed.ncbi.nlm.nih.gov/11306649/)

20. Artola A, Bröcher S, Singer W. Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature*. 1990;347(6288):69–72. <https://doi.org/10.1038/347069a0> PMID: 1975639
21. Piochon C, Titley HK, Simmons DH, Grasselli G, Elgersma Y, Hansel C. Calcium threshold shift enables frequency-independent control of plasticity by an instructive signal. *Proc Natl Acad Sci U S A*. 2016;113(46):13221–6. <https://doi.org/10.1073/pnas.1613897113> PMID: 27799554
22. Citri A, Malenka RC. Synaptic plasticity: multiple forms, functions, and mechanisms. *Neuropsychopharmacology*. 2008;33(1):18–41. <https://doi.org/10.1038/sj.npp.1301559> PMID: 17728696
23. Sanes JR, Lichtman JW. Can molecules explain long-term potentiation? *Nat Neurosci*. 1999;2(7):597–604. <https://doi.org/10.1038/10154> PMID: 10404178
24. Rodrigues YE, Tigaret CM, Marie H, O'Donnell C, Veltz R. A stochastic model of hippocampal synaptic plasticity with geometrical readout of enzyme dynamics. *Elife*. 2023;12:e80152. <https://doi.org/10.7554/eLife.80152> PMID: 37589251
25. Jahr CE, Stevens CF. Voltage dependence of NMDA-activated macroscopic conductances predicted by single-channel kinetics. *J Neurosci*. 1990;10(9):3178–82. <https://doi.org/10.1523/JNEUROSCI.10-09-03178.1990> PMID: 1697902
26. Jahr CE, Stevens CF. A quantitative description of NMDA receptor-channel kinetic behavior. *J Neurosci*. 1990;10(6):1830–7. <https://doi.org/10.1523/JNEUROSCI.10-06-01830.1990> PMID: 1693952
27. Moldwin T, Kalmenson M, Segev I. Asymmetric voltage attenuation in dendrites can enable hierarchical heterosynaptic plasticity. *eNeuro*. 2023;10(7):ENEURO.0014-23.2023. <https://doi.org/10.1523/ENEURO.0014-23.2023> PMID: 37414554
28. Stuart GJ, Sakmann B. Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature*. 1994;367(6458):69–72. <https://doi.org/10.1038/367069a0> PMID: 8107777
29. Bi GQ, Poo MM. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci*. 1998;18(24):10464–72. <https://doi.org/10.1523/JNEUROSCI.18-24-10464.1998> PMID: 9852584
30. Moldwin T, Azran LS, Segev I. A Generalized framework for the calcium control hypothesis describes weight-dependent synaptic plasticity. *bioRxiv*. 2023;2023.07.13.548837. <https://doi.org/10.1101/2023.07.13.548837>
31. Bittner KC, Grienberger C, Vaidya SP, Milstein AD, Macklin JJ, Suh J, et al. Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. *Nat Neurosci*. 2015;18(8):1133–42. <https://doi.org/10.1038/nn.4062> PMID: 26167906
32. Bittner KC, Milstein AD, Grienberger C, Romani S, Magee JC. Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science*. 2017;357(6355):1033–6. <https://doi.org/10.1126/science.aan3846> PMID: 28883072
33. Grienberger C, Magee JC. Entorhinal cortex directs learning-related changes in CA1 representations. *Nature*. 2022;611(7936):554–62. <https://doi.org/10.1038/s41586-022-05378-6> PMID: 36323779
34. Milstein AD, Li Y, Bittner KC, Grienberger C, Soltesz I, Magee JC, et al. Bidirectional synaptic plasticity rapidly modifies hippocampal representations. *Elife*. 2021;10:e73046. <https://doi.org/10.7554/eLife.73046> PMID: 34882093
35. Konnerth A, Dreessen J, Augustine GJ. Brief dendritic calcium signals initiate long-lasting synaptic depression in cerebellar Purkinje cells. *Proc Natl Acad Sci U S A*. 1992;89(15):7051–5. <https://doi.org/10.1073/pnas.89.15.7051> PMID: 1323125
36. Bienenstock EL, Cooper LN, Munro PW. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci*. 1982;2(1):32–48. <https://doi.org/10.1523/JNEUROSCI.02-01-00032.1982> PMID: 7054394
37. Evans RC, Blackwell KT. Calcium: amplitude, duration, or location? *Biol Bull*. 2015;228(1):75–83. <https://doi.org/10.1086/BBLv228n1p75> PMID: 25745102
38. O'Connor DH, Wittenberg GM, Wang SS-H. Dissection of bidirectional synaptic plasticity into saturable unidirectional processes. *J Neurophysiol*. 2005;94(2):1565–73. <https://doi.org/10.1152/jn.00047.2005> PMID: 15800079
39. Bliss TV, Lomo T. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J Physiol*. 1973;232(2):331–56. <https://doi.org/10.1113/jphysiol.1973.sp010273> PMID: 4727084
40. Dudek SM, Bear MF. Homosynaptic long-term depression in area CA1 of hippocampus and effects of N-methyl-D-aspartate receptor blockade. *Proc Natl Acad Sci U S A*. 1992;89(10):4363–7. <https://doi.org/10.1073/pnas.89.10.4363> PMID: 1350090

41. Clopath C, Büsing L, Vasilaki E, Gerstner W. Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nat Neurosci*. 2010;13(3):344–52. <https://doi.org/10.1038/nn.2479> PMID: [20098420](https://pubmed.ncbi.nlm.nih.gov/20098420/)
42. Yeung LC, Shouval HZ, Blais BS, Cooper LN. Synaptic homeostasis and input selectivity follow from a calcium-dependent plasticity model. *Proc Natl Acad Sci U S A*. 2004;101(41):14943–8. <https://doi.org/10.1073/pnas.0405555101> PMID: [15466713](https://pubmed.ncbi.nlm.nih.gov/15466713/)
43. Bittner KC, Milstein AD, Grienberger C, Romani S, Magee JC. Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science*. 2017;357(6355):1033–6. <https://doi.org/10.1126/science.aan3846> PMID: [28883072](https://pubmed.ncbi.nlm.nih.gov/28883072/)
44. Turrigiano GG, Nelson SB. Homeostatic plasticity in the developing nervous system. *Nat Rev Neurosci*. 2004;5(2):97–107. <https://doi.org/10.1038/nrn1327> PMID: [14735113](https://pubmed.ncbi.nlm.nih.gov/14735113/)
45. Turrigiano GG. The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell*. 2008;135(3):422–35. <https://doi.org/10.1016/j.cell.2008.10.008> PMID: [18984155](https://pubmed.ncbi.nlm.nih.gov/18984155/)
46. Turrigiano G. Homeostatic synaptic plasticity: local and global mechanisms for stabilizing neuronal function. *Cold Spring Harb Perspect Biol*. 2012;4(1):a005736. <https://doi.org/10.1101/cshperspect.a005736> PMID: [22086977](https://pubmed.ncbi.nlm.nih.gov/22086977/)
47. Turrigiano GG, Leslie KR, Desai NS, Rutherford LC, Nelson SB. Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*. 1998;391(6670):892–6. <https://doi.org/10.1038/36103> PMID: [9495341](https://pubmed.ncbi.nlm.nih.gov/9495341/)
48. Peters A, Schweiger U, Pellerin L, Hubold C, Oltmanns KM, Conrad M, et al. The selfish brain: competition for energy resources. *Neurosci Biobehav Rev*. 2004;28(2):143–80. <https://doi.org/10.1016/j.neubiorev.2004.03.002> PMID: [15172762](https://pubmed.ncbi.nlm.nih.gov/15172762/)
49. Miller KD. Synaptic economics: competition and cooperation in synaptic plasticity. *Neuron*. 1996;17(3):371–4. [https://doi.org/10.1016/s0896-6273\(00\)80169-5](https://doi.org/10.1016/s0896-6273(00)80169-5) PMID: [8816700](https://pubmed.ncbi.nlm.nih.gov/8816700/)
50. Rabinowitch I, Segev I. The interplay between homeostatic synaptic plasticity and functional dendritic compartments. *J Neurophysiol*. 2006;96(1):276–83. <https://doi.org/10.1152/jn.00074.2006> PMID: [16554518](https://pubmed.ncbi.nlm.nih.gov/16554518/)
51. Rabinowitch I, Segev I. The endurance and selectivity of spatial patterns of long-term potentiation/depression in dendrites under homeostatic synaptic plasticity. *J Neurosci*. 2006;26(52):13474–84. <https://doi.org/10.1523/JNEUROSCI.4333-06.2006> PMID: [17192430](https://pubmed.ncbi.nlm.nih.gov/17192430/)
52. Tigaret CM, Olivo V, Sadowski JHLP, Ashby MC, Mellor JR. Coordinated activation of distinct Ca^{2+} sources and metabotropic glutamate receptors encodes Hebbian synaptic plasticity. *Nat Commun*. 2016;7:10289. <https://doi.org/10.1038/ncomms10289> PMID: [26758963](https://pubmed.ncbi.nlm.nih.gov/26758963/)
53. Moldwin T, Segev I. Perceptron learning and classification in a modeled cortical pyramidal cell. *Front Comput Neurosci*. 2020;14:1–33. <https://doi.org/10.3389/fncom.2020.00033> PMID: [32390819](https://pubmed.ncbi.nlm.nih.gov/32390819/)
54. Beniaguev D, Segev I, London M. Single cortical neurons as deep artificial neural networks. *Neuron*. 2021;109(17):2727–2739.e3. <https://doi.org/10.1016/j.neuron.2021.07.002> PMID: [34380016](https://pubmed.ncbi.nlm.nih.gov/34380016/)
55. Gordon U, Polsky A, Schiller J. Plasticity compartments in basal dendrites of neocortical pyramidal neurons. *J Neurosci*. 2006;26(49):12717–26. <https://doi.org/10.1523/JNEUROSCI.3502-06.2006> PMID: [17151275](https://pubmed.ncbi.nlm.nih.gov/17151275/)
56. Mel BW. The clusteron: toward a simple abstraction for a complex neuron. *NIPS*. 1991:35–42.
57. Moldwin T, Kalmenson M, Segev I. The gradient clusteron: a model neuron that learns to solve classification tasks via dendritic nonlinearities, structural plasticity, and gradient descent. *PLoS Comput Biol*. 2021;17(5):e1009015. <https://doi.org/10.1371/journal.pcbi.1009015> PMID: [34029309](https://pubmed.ncbi.nlm.nih.gov/34029309/)
58. Pagkalos M, Chavlis S, Poirazi P. Introducing the Dendripy framework for incorporating dendrites to spiking neural networks. *Nat Commun*. 2023;14(1):131. <https://doi.org/10.1038/s41467-022-35747-8> PMID: [36627284](https://pubmed.ncbi.nlm.nih.gov/36627284/)
59. Pagkalos M, Makarov R, Poirazi P. Leveraging dendritic properties to advance machine learning and neuro-inspired computing. 2023 [cited 2023 Sep 11]. Available from: <https://arxiv.org/abs/2306.08007v1>
60. Poirazi P, Mel BW. Impact of active dendrites and structural plasticity on the memory capacity of neural tissue. *Neuron*. 2001;29(3):779–96. [https://doi.org/10.1016/s0896-6273\(01\)00252-5](https://doi.org/10.1016/s0896-6273(01)00252-5) PMID: [11301036](https://pubmed.ncbi.nlm.nih.gov/11301036/)
61. Polsky A, Mel BW, Schiller J. Computational subunits in thin dendrites of pyramidal cells. *Nat Neurosci*. 2004;7(6):621–7. <https://doi.org/10.1038/nn1253> PMID: [15156147](https://pubmed.ncbi.nlm.nih.gov/15156147/)
62. Schiller J, Major G, Koester HJ, Schiller Y. NMDA spikes in basal dendrites of cortical pyramidal neurons. *Nature*. 2000;404(6775):285–9. <https://doi.org/10.1038/35005094> PMID: [10749211](https://pubmed.ncbi.nlm.nih.gov/10749211/)

63. Tran-Van-Minh A, Cazé RD, Abrahamsson T, Cathala L, Gutkin BS, DiGregorio DA. Contribution of sublinear and supralinear dendritic integration to neuronal computations. *Front Cell Neurosci*. 2015;9:67. <https://doi.org/10.3389/fncel.2015.00067> PMID: [25852470](#)
64. Beniaguev D, Shapira S, Segev I, London M. Dendro-plexing single input spikes by multiple synaptic contacts enriches the computational capabilities of cortical neurons and reduces axonal wiring. *bioRxiv*. 2022;2022.01.28.478132. <https://doi.org/10.1101/2022.01.28.478132>
65. Kumar A, Barkai E, Schiller J. Plasticity of olfactory bulb inputs mediated by dendritic NMDA-spikes in rodent piriform cortex. *Elife*. 2021;10:e70383. <https://doi.org/10.7554/eLife.70383> PMID: [34698637](#)
66. Chater TE, Goda Y. The role of AMPA receptors in postsynaptic mechanisms of synaptic plasticity. *Front Cell Neurosci*. 2014;8:401. <https://doi.org/10.3389/fncel.2014.00401> PMID: [25505875](#)
67. Chater TE, Goda Y. My Neighbour Hetero-deconstructing the mechanisms underlying heterosynaptic plasticity. *Curr Opin Neurobiol*. 2021;67:106–14. <https://doi.org/10.1016/j.conb.2020.10.007> PMID: [33160201](#)
68. Chistiakova M, Bannon NM, Bazhenov M, Volgushev M. Heterosynaptic plasticity: multiple mechanisms and multiple roles. *Neuroscientist*. 2014;20(5):483–98. <https://doi.org/10.1177/1073858414529829> PMID: [24727248](#)
69. Moldwin T, Kalmenson M, Segev I. Asymmetric voltage attenuation in dendrites can enable hierarchical heterosynaptic plasticity. *bioRxiv*. 2022;2022.07.07.499166. <https://doi.org/10.1101/2022.07.07.499166>
70. Tong R, Chater TE, Emptage NJ, Goda Y. Heterosynaptic cross-talk of pre- and postsynaptic strengths along segments of dendrites. *Cell Rep*. 2021;34(4):108693. <https://doi.org/10.1016/j.celrep.2021.108693> PMID: [33503435](#)
71. Bar-Ilan L, Gidon A, Segev I. The role of dendritic inhibition in shaping the plasticity of excitatory synapses. *Front Neural Circuits*. 2013;6:118. <https://doi.org/10.3389/fncir.2012.00118> PMID: [23565076](#)
72. Rose CR, Konnerth A. Stores not just for storage. intracellular calcium release and synaptic plasticity. *Neuron*. 2001;31(4):519–22. [https://doi.org/10.1016/s0896-6273\(01\)00402-0](https://doi.org/10.1016/s0896-6273(01)00402-0) PMID: [11545711](#)
73. Benedetti L, Fan R, Weigel AV, Moore AS, Houlihan PR, Kittisopikul M, et al. Periodic ER-plasma membrane junctions support long-range Ca^{2+} signal integration in dendrites. *Cell*. 2024;S0092-8674(24)01345-X. <https://doi.org/10.1016/j.cell.2024.11.029> PMID: [39708809](#)
74. Spacek J, Harris KM. Three-dimensional organization of smooth endoplasmic reticulum in hippocampal CA1 dendrites and dendritic spines of the immature and mature rat. *J Neurosci*. 1997;17(1):190–203. <https://doi.org/10.1523/JNEUROSCI.17-01-00190.1997> PMID: [8987748](#)
75. Haber M, Zhou L, Murai KK. Cooperative astrocyte and dendritic spine dynamics at hippocampal excitatory synapses. *J Neurosci*. 2006;26(35):8881–91. <https://doi.org/10.1523/JNEUROSCI.1302-06.2006> PMID: [16943543](#)
76. Barker AJ, Ullian EM. Astrocytes and synaptic plasticity. *Neuroscientist*. 2010;16(1):40–50. <https://doi.org/10.1177/1073858409339215> PMID: [20236948](#)
77. Puig MV, Antzoulatos EG, Miller EK. Prefrontal dopamine in associative learning and memory. *Neuroscience*. 2014;282:217–29. <https://doi.org/10.1016/j.neuroscience.2014.09.026> PMID: [25241063](#)
78. Lisman J, Grace AA, Duzel E. A neoHebbian framework for episodic memory; role of dopamine-dependent late LTP. *Trends Neurosci*. 2011;34(10):536–47. <https://doi.org/10.1016/j.tins.2011.07.006> PMID: [21851992](#)
79. Hunter JD. Matplotlib: a 2D graphics environment. *Comput Sci Eng*. 2007;9(3):90–5. <https://doi.org/10.1109/mcse.2007.55>
80. MICrONS Consortium. Functional connectomics spanning multiple areas of mouse visual cortex. *bioRxiv* 2021.07.28.454025. <https://doi.org/10.1101/2021.07.28.454025>