



## OPEN ACCESS

## EDITED BY

Petia D. Koprinkova-Hristova,  
Institute of Information and Communication  
Technologies (BAS), Bulgaria

## REVIEWED BY

Shuangming Yang,  
Tianjin University, China  
Rufin VanRullen,  
Centre National de la Recherche Scientifique  
(CNRS), France

## \*CORRESPONDENCE

Pau Vilimelis Aceituno

✉ pvilimelis@ethz.ch

Matilde Tristany Farinha

✉ mtristany@ethz.ch

Benjamin F. Grewe

✉ bgrewe@ethz.ch

†These authors have contributed equally to this work

RECEIVED 02 January 2023

ACCEPTED 25 April 2023

PUBLISHED 24 May 2023

## CITATION

Aceituno PV, Farinha MT, Loidl R and Grewe BF  
(2023) Learning cortical hierarchies with  
temporal Hebbian updates.  
*Front. Comput. Neurosci.* 17:1136010.  
doi: 10.3389/fncom.2023.1136010

## COPYRIGHT

© 2023 Aceituno, Farinha, Loidl and Grewe.  
This is an open-access article distributed under  
the terms of the [Creative Commons Attribution  
License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or  
reproduction in other forums is permitted,  
provided the original author(s) and the  
copyright owner(s) are credited and that the  
original publication in this journal is cited, in  
accordance with accepted academic practice.  
No use, distribution or reproduction is  
permitted which does not comply with these  
terms.

# Learning cortical hierarchies with temporal Hebbian updates

Pau Vilimelis Aceituno<sup>1,2\*†</sup>, Matilde Tristany Farinha<sup>1\*†</sup>,  
Reinhard Loidl<sup>1</sup> and Benjamin F. Grewe<sup>1,2\*†</sup>

<sup>1</sup>Institute of Neuroinformatics, University of Zurich and ETH Zurich, Zurich, Switzerland, <sup>2</sup>ETH AI Center, ETH Zurich, Zurich, Switzerland

A key driver of mammalian intelligence is the ability to represent incoming sensory information across multiple abstraction levels. For example, in the visual ventral stream, incoming signals are first represented as low-level edge filters and then transformed into high-level object representations. Similar hierarchical structures routinely emerge in artificial neural networks (ANNs) trained for object recognition tasks, suggesting that similar structures may underlie biological neural networks. However, the classical ANN training algorithm, backpropagation, is considered biologically implausible, and thus alternative biologically plausible training methods have been developed such as Equilibrium Propagation, Deep Feedback Control, Supervised Predictive Coding, and Dendritic Error Backpropagation. Several of those models propose that local errors are calculated for each neuron by comparing apical and somatic activities. Notwithstanding, from a neuroscience perspective, it is not clear how a neuron could compare compartmental signals. Here, we propose a solution to this problem in that we let the apical feedback signal change the postsynaptic firing rate and combine this with a differential Hebbian update, a rate-based version of classical spiking time-dependent plasticity (STDP). We prove that weight updates of this form minimize two alternative loss functions that we prove to be equivalent to the error-based losses used in machine learning: the inference latency and the amount of top-down feedback necessary. Moreover, we show that the use of differential Hebbian updates works similarly well in other feedback-based deep learning frameworks such as Predictive Coding or Equilibrium Propagation. Finally, our work removes a key requirement of biologically plausible models for deep learning and proposes a learning mechanism that would explain how temporal Hebbian learning rules can implement supervised hierarchical learning.

## KEYWORDS

cortical hierarchies, deep learning, credit assignment, synaptic plasticity, backpropagation, spiking time-dependent plasticity, target propagation, differential Hebbian learning

## 1. Introduction

To survive in complex natural environments, humans and animals transform sensory input into neuronal signals which in turn generate and modulate behavior. Learning of such transformations often amounts to a non-trivial problem, since sensory inputs can be very high-dimensional and complex. The complexity of sensory inputs requires hierarchical information processing, which relies on multilayer networks. To form hierarchies, cortical networks need to process these sensory signals and convey plasticity signals down to every neuron in the hierarchy so that the output of the network (e.g., the motor output or behavior) improves during learning. In deep learning, this is known as the credit assignment (CA)

problem and it is commonly addressed by the error backpropagation (BP) method. During BP learning, neurons in the lower hierarchies change their afferent synapses by integrating a backpropagated error signal. A neuron's afferent weight update is then calculated as the product of the presynaptic activity and its non-local output error. However, several key aspects of BP are still at odds with learning in biological neural networks (Crick, 1989; Lillicrap et al., 2020). For example, ANNs separate the processing or encoding of neuronal activity signals from the weight update signals, they utilize distinct phases and they implement an exact weight symmetry of forward and feedback pathways. Moreover, plasticity in biological synapses is local in space and time and tightly coupled to the timing of the pre- and post-synaptic activity (Bi and Poo, 1998).

Attempting to address some of these implausibilities, recent cortical-inspired ANN models leverage network dynamics to directly couple changes in neuronal activity to weight updates (Whittington and Bogacz, 2017; Sacramento et al., 2018). Those models postulate multi-compartment pyramidal neurons with a highly specialized dendritic morphology that use their apical dendrite to integrate a feedback signal that modulates feedforward plasticity (Figure 1A, left schematic). Although multi-compartment models agree with some biological constraints, such as the spatial locality of learning rules and the fact that feedback not only generates plasticity but also affects neuronal activity (Gilbert and Li, 2013), the apical “dendritic-error” learning approach still requires tightly coordinated and highly specific error signaling circuits (Whittington and Bogacz, 2017; Sacramento et al., 2018). To avoid these highly specific error circuits, we recently developed a novel class of cortical-inspired ANNs that utilizes the same dendritic-error learning rule, but does not require highly specific error circuits and is capable of online learning of all weights without requiring separate forward and backward passes (Meulemans et al., 2021a,b, 2022b). In this model, known as “Deep Feedback Control” (DFC), we dynamically tailor the feedback to each hidden neuron until the network output reaches the desired target. The weight update of the feedforward pathway is then calculated upon convergence as the difference in neural activities when the effect of top-down apical feedback is fully taken into account or not. Still, this model relies on the same dendritic-error learning rule as its predecessors (Whittington and Bogacz, 2017; Sacramento et al., 2018), and it is unclear how a neuron would be able to compare the activities of its basal and apical compartments (Figure 1A, left scheme). In this work, we argue that dendritic learning rules can be substituted by experimentally validated temporal Hebbian learning rules (e.g., STDP) and we use the DFC framework as an example of how a deep network can learn with this mechanism. We argue that single-compartment neurons, whose firing rate is strongly affected by apical input, can use the difference between consecutive instances of their activity as a learning signal (Figure 1B), as opposed to comparing the changes in two different compartments. Based on this dynamic change in the postsynaptic activity we can thus encode the learning signal while being consistent with experimentally observed learning rules such as STDP (Figure 1C).

### Spike-Timing Dependent Plasticity (STDP)

When using the term STDP, we here refer to the well-established observation that the precise timing of pre- and post-synaptic spikes significantly determines the sign and magnitude of synaptic plasticity (Markram et al., 1997; Bi and Poo, 1998). In cortical pyramidal neurons, a presynaptic spike that precedes a postsynaptic spike within a narrow time window induces long-term potentiation (LTP) (Markram et al., 1997; Bi and Poo, 1998; Nishiyama et al., 2000; Sjöström et al., 2001; Wittenberg and Wang, 2006; Feldman, 2012); if the order is reversed it leads to long-term depression (LTD). Using this classical STDP profile (Figure 1A), multiple theoretical models were able to predict biological plasticity by assuming a simple superposition of spike pairs (Gerstner et al., 1996; Kempter et al., 1999; Abbott and Nelson, 2000; Song et al., 2000; van Rossum et al., 2000; Izhikevich and Desai, 2003; Güttig, 2016).

## 2. Results

### 2.1. Single neuron supervised learning with STDP

We first demonstrate how an STDP learning rule can be used to train a single neuron on a linear classification task (Figure 2A). We use a neuron with the sigmoid activation function, which gets both feedforward basal inputs from two other neurons (A and B) and a feedback apical input, resulting in the rate-based dynamics

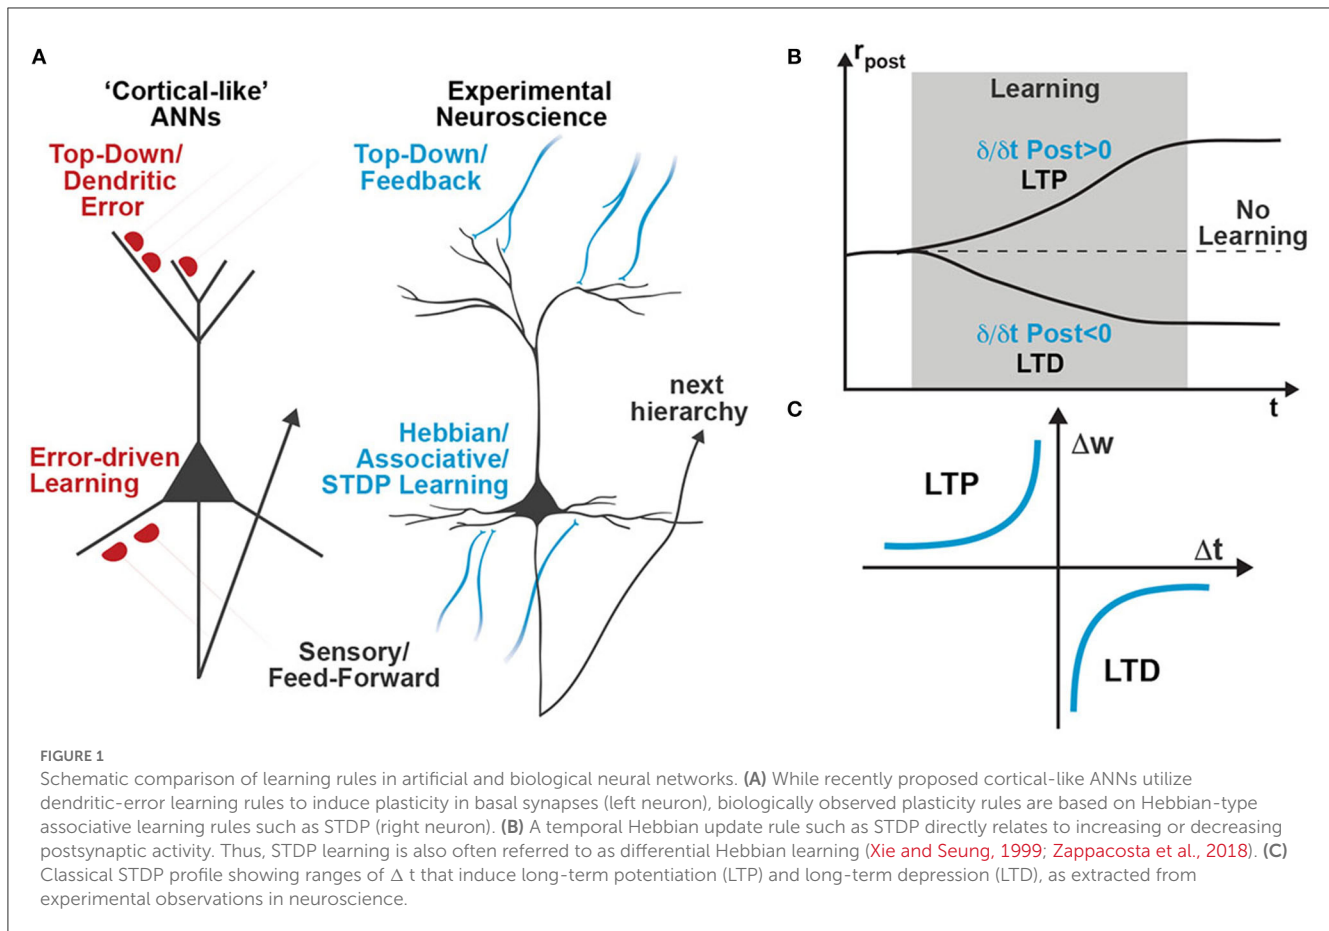
$$\begin{aligned} \dot{v}_{\text{post}}(t) &= -v_{\text{post}}(t) + w_A r_A + w_B r_B + c(t) \\ r_{\text{post}}(t) &= \phi(v_{\text{post}}(t)) \end{aligned} \quad (1)$$

where  $w_A, w_B$  are the synaptic strengths of the connection from the input neurons to the output neuron,  $r_A, r_B$  are the firing rates of neurons A, B, respectively,  $r_{\text{post}}(t)$  and  $v_{\text{post}}(t)$  are the output firing rate and membrane potential at time  $t$ , and  $c(t)$  is the apical feedback given to the output neuron. In our simple example, the neuron can get two incoming stimuli, from neurons A and B, and the apical feedback  $c(t)$  changes the output firing rate to be high when B is presented and low when A is presented.

The firing rate variables are converted into spike trains with an inhomogeneous Poisson process, where at every time step the probability of spiking in each neuron is given by  $r_{\text{post}}(t)$ ,  $r_A, r_B$ , respectively. These spike trains are then used to induce synaptic weight changes by STDP (see Figure 2B).

We observe in Figure 2B that, as learning progresses, the weights evolve to the expected values (high for  $w_B$ , low for  $w_A$ ), and that this changes the dynamics of  $r_{\text{post}}(t)$ , causing  $r_{\text{post}}$  to start closer to its target value and, thus, shortening the time and the feedback required to produce the desired output (Figures 2C, D). Such changes can be understood in terms of the following set of equivalent loss functions that are minimized:

- The initial distance to the target activity can be computed as the Mean Squared Error (MSE), denoted by  $\mathcal{L}$ . This loss is commonly used in the machine learning literature as a standard performance measure.
- The feedback required to maintain or reach the target activity, denoted by  $\mathcal{H}$ . This loss is equivalent to the one presented in previous works on using feedback to train neural networks



(Gilra and Gerstner, 2017; Meulemans et al., 2020) and relates to the intuition from Predictive Coding that a trained ANN minimizes the feedback needed to correctly process the input (Rao and Ballard, 1999).

- The time delay to reach the target is denoted by  $\mathcal{T}$ . This loss function represents the amount of time a neuron takes to reach its target value. This idea appears in previous works based on STDP models (Masquelier et al., 2009; Vilimelis Aceituno et al., 2020) and is also implicitly used in models for learning in deep networks (Luczak et al., 2022).

To relate the three losses to temporal Hebbian learning, we re-express the STDP update through its rate-based form, known as the differential Hebbian (DH) learning rule (Xie and Seung, 1999; Saudargiene et al., 2004; Bengio et al., 2017),

$$\Delta w \propto \int r_{pre}(t) \dot{r}_{post}(t) dt, \tag{2}$$

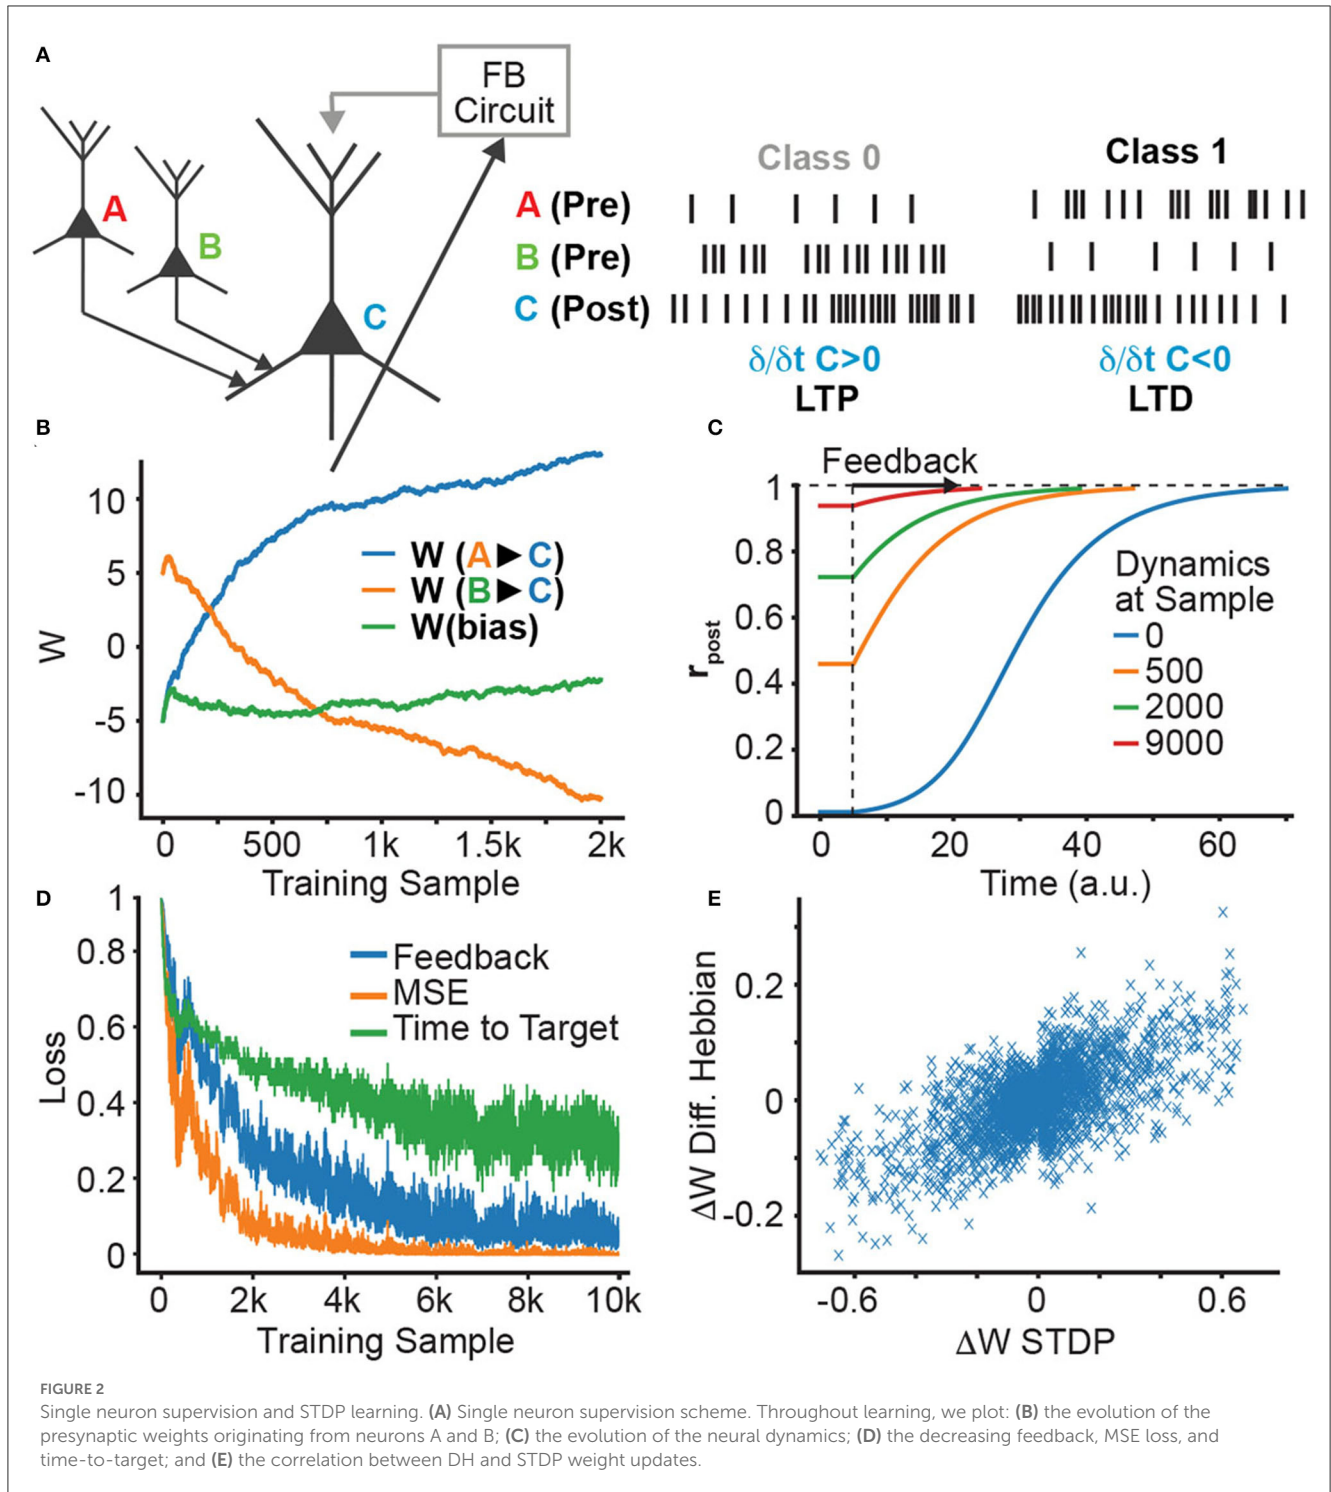
where  $\Delta w$  is the change in feedforward synaptic strength,  $r_{pre}(t)$  is the presynaptic activity and  $\dot{r}_{post}(t)$  is the derivative of the postsynaptic activity, which corresponds to the change in firing probability. As we see in Figure 2E, the DH learning rule is indeed similar to STDP, albeit with noise induced by the inherent stochasticity of the Poisson neuron.

To understand how this rule relates to the three loss functions mentioned above, we note that in the single neuron example, the presynaptic firing rate is fixed, which simplifies the previous rule to

$$\Delta w \propto r_{pre} \int \dot{r}_{post}(t) dt = r_{pre} [r_{post}(T) - r_{post}(0)], \tag{3}$$

where  $r_{post}(T)$  is the postsynaptic activity after reaching the target state. This corresponds to the dendritic-error learning rule (Gilra and Gerstner, 2017; Sacramento et al., 2018; Meulemans et al., 2020). The correlation of the weight updates for this rule and STDP is shown in Figure 2E. In this single neuron setting, it is clear that both the STDP and DH learning mechanisms decrease the three loss functions: having an initial activity that is closer to the target activity implies that the MSE loss is lower at the beginning, and also that the change in activity is smaller. Hence, it needs less feedback and the target can be reached much faster (see Appendix, Section 2 for a detailed explanation).

The next key question is whether we can use STDP and DH learning in a similar manner for hierarchical credit assignment, i.e., for training multilayer neural networks.

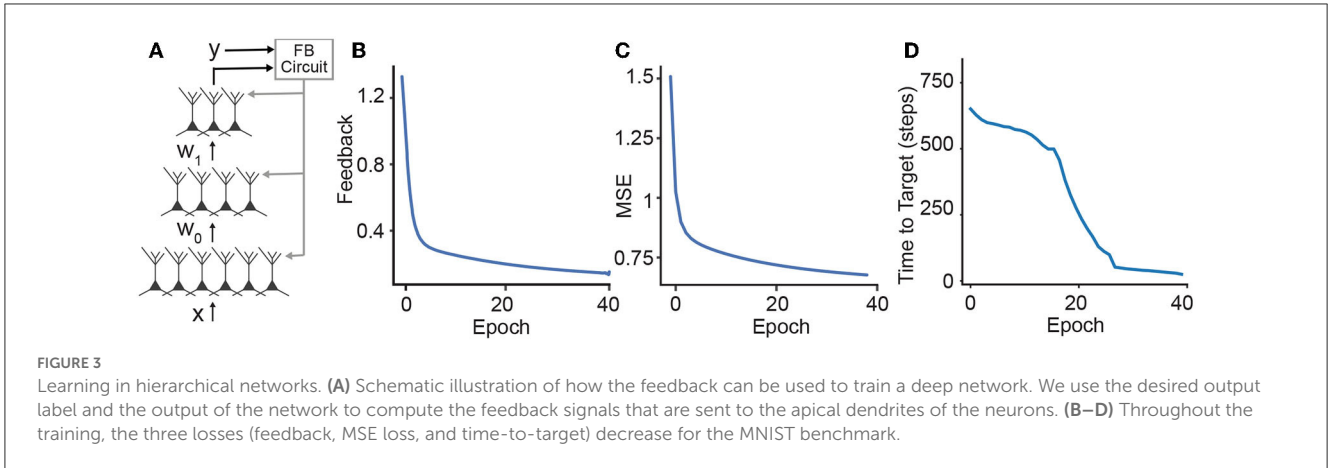


## 2.2. Differential Hebbian can train multilayer networks

To extend our results to multilayer neural networks, the feedback must be received by the neurons in all layers. To compute the appropriate feedback signals, we use the framework of deep feedback control (DFC) from our previous work (Meulemans et al., 2021a), which we detail here for completeness.

In DFC, each neuron receives a feedforward basal input and an apical feedback signal that is computed by a controller whose goal is to achieve a target output response (Figure 3A). The neuronal dynamics is described by

$$\dot{v}_{\text{post}}(t) = -v_{\text{post}}(t) + \underbrace{W_{\{\text{pre},\text{post}\}}r_{\text{pre}}(t)}_{\text{feedforward}} + \underbrace{Q_{\text{post}}c(t)}_{\text{feedback}}, \quad (4)$$



where the postsynaptic membrane potential  $v_{\text{post}}(t)$  at each neuron is given by the presynaptic firing rate  $r_{\text{pre}}(t)$ , originating from neurons in the previous layer, multiplied by the feedforward synaptic weights  $W_{\text{[pre,post]}}$ . In order to compute the target signal for every neuron, DFC uses a global PI controller that affects all the neurons in the network denoted by  $c(t)$ ,

$$c(t) = c^{\text{int}}(t) + ke(t), \quad \tau_u \dot{c}^{\text{int}}(t) = e(t)\alpha - c^{\text{int}}(t), \quad (5)$$

where  $k$  is the proportional control constant,  $e(t)$  is the difference between the target output activity and the network's current output, and  $\alpha$  is the leak constant. The  $Q$  matrix contains the top-down feedback weights that map the controller signal into each hidden neuron and is pre-trained using local anti-Hebbian learning rules, as done in Meulemans et al. (2021a), but then kept fixed throughout the learning of the feedforward weights. We start with a random  $Q$  weight matrix and add independent zero-mean noise into the network,

$$\dot{v}_{\text{post}}(t) = -v_{\text{post}}(t) + W_{\text{[pre,post]}}r_{\text{pre}}(t) + Q_{\text{post}}c(t) + \epsilon, \quad (6)$$

where the fluctuations ( $\epsilon$ ) on every neuron propagate through the feedforward network and affect the output layer, which in turn creates fluctuations in  $c(t)$  that the controller then acts to eliminate them. We then use an anti-Hebbian learning rule of the form

$$\dot{Q}_{\text{post}}(t) = -v_{\text{post}}(t)c(t)^T - \beta Q_{\text{post}}, \quad (7)$$

where parameter  $\beta$  controls the strength of the feedback weights. As proven in Meulemans et al. (2021a), learning  $Q$  with this rule ensures that the model does principled optimization, meaning it converges in learning.

We test the DH learning with the DFC setting on MNIST (LeCun, 1998), a widely accepted standard computer vision benchmark that aims to classify  $28 \times 28$  pixel grayscale images of handwritten digits between 0 and 9. We show that DH with feedback computed through the DFC framework can train a three-hidden layer network ( $256 \times 256 \times 256$ ) to match state-of-the-art performances and compare our framework with BP as well as the original DFC framework based on dendritic-error learning (Table 1). We find that the testing classification error rates of BP, DFC, and DH-DFC are on par with this benchmark.

To complement our analysis, we investigate the training loss in DH-DFC. We note that the amount of feedback required to reach the target decreases throughout the training (Figure 3B), implying that DH-DFC also decreases the required feedback. The MSE loss also decreases (Figure 3C), hence DH-DFC also learns by minimizing an implicit error. Finally, we show that the latency to reach the target is also reduced (see Figure 3D), implying that the latency-reduction nature of temporally asymmetric learning rules (Masquelier et al., 2009; Vilimelis Aceituno et al., 2020) is reflected in our framework.

In addition, we experimentally calculate the similarity of the weight updates arising from different learning rules (Figure 4). We find that both the DFC and the BP updates are strongly positively correlated with DH-DFC, with coefficients of determination of 0.804 and 0.966, respectively.

Finally, we compare the DH-DFC weight updates with STDP updates evaluated on spike trains, which have a positive correlation, with a coefficient of determination of 0.008, but a very noisy alignment due to the randomness induced by using Poisson neurons (see Appendix, Section 5). This randomness can be reduced by computing several parallel conversions of firing rates to Poisson spike trains and averaging the resulting STDP updates, although here we find that limitations in computer memory prevent us from reaching state-of-the-art accuracies (see Appendix, Section 5).

It is worth noting that the DFC framework we use as a baseline is not the only model that uses feedback to train deep neural networks. In the next section, we argue that the use of temporal Hebbian rules is not restricted to the DFC framework, being instead applicable to other feedback-based learning models.

### 2.3. Differential Hebbian learning applies to other feedback-based networks

We extend our results on the single neuron framework and DFC multilayer model and prove that DH learning works in a general framework where some feedback is given to each neuron in the network so that the neuron reaches its target state. In contrast to the single neuron set-up, the DH learning rule is not equivalent to a simple delta rule. Since the presynaptic firing rate of most synapses

TABLE 1 MNIST classification error for BP, DFC, DH-DFC, PC, and DH-PC.

	MNIST (%)
BP	1.74 $\pm$ 0.10
DFC	1.98 $\pm$ 0.05
DH-DFC	1.89 $\pm$ 0.15
PC*	2.71 $\pm$ 0.2
DH-PC*	2.69 $\pm$ 0.15

We compare the testing accuracies reached by our training procedure (DH-DFC) with those achieved by BP, DFC (reported in Meulemans et al., 2021b), PC, and PC with DH (DH-PC) in the MNIST dataset. The reported values correspond to the average testing classification error obtained from five different random seeds. \*The PC and DH-PC are based on the code from Tschantz (2020) without any fine-tuning.

changes in time, the DH learning rule can then be expressed as

$$\Delta w \propto \int r_{\text{pre}}(t)\dot{r}_{\text{post}}(t)dt = r_{\text{pre}}(T) \int \dot{r}_{\text{post}}(t)dt - \int \tilde{r}_{\text{pre}}(t)\dot{r}_{\text{post}}(t)dt, \quad (8)$$

where the extra term includes the difference between the presynaptic firing rate at time  $t$  and its target,  $\tilde{r}_{\text{pre}}(t) = r_{\text{pre}}(T) - r_{\text{pre}}(t)$ .

To understand why DH learning works despite being different from the classical dendritic-error learning rule, it is useful to note that  $\tilde{r}_{\text{pre}}(t) \rightarrow 0$  as learning proceeds, and, thus, this term disappears around the convergence point of the weights. By using an inductive argument, this can be extended to other layers (see Appendix, Section 3.1 for a detailed derivation).

A critical point of our convergence proof is that it does not depend on how feedback is computed. In fact, the only requirements are that feedback somehow pushes the neurons toward their target states. This suggests that the logic of DH learning should also work with other feedback-based learning models such as the original DFC (Meulemans et al., 2022a), but also models relying on Predictive Coding (PC) (Whittington and Bogacz, 2017; Rosenbaum, 2022), or Equilibrium Propagation (Scellier and Bengio, 2017). For DFC, we already saw that the weight updates align with DFC-DH and the performances are equally comparable (Figure 4 and Table 1); we further complement this by analytically showing that the learning rules converge to the same network configurations after learning (see Appendix, Section 4.1). For PC, we find that using the prediction error (implemented through error neurons) as implicit feedback leads to the same convergence proof as in the DH-DFC (see Appendix, Section 4.2). Moreover, in simulations, we find that the performance of PC using DH (DH-PC) is similar to that of PC and DH-DFC (see Table 1). For Equilibrium Propagation (Scellier and Bengio, 2017), we note that our rule is analytically equivalent to a modified version of DH that accounts for the specific architectural constraints as noted in the original work (see Appendix, Section 4.2).

As we conclude that there are multiple feedback-based learning models to which a DH learning rule generalizes, it is natural to inquire whether the specific combination of DH-DFC has any advantage over its predecessor. In the original studies of DFC, Predictive Coding, and Equilibrium Propagation, the learning rules are applied after the neural dynamics have converged to equilibrium. Learning is then based on an error-like component

that corresponds to the difference between the activities (or membrane potentials) before and after feedback has shaped them. In DFC, this error is obtained by having two-compartment neurons, while in Predictive Coding errors are accumulated (usually as error neurons); in both cases, this raises the number of variables from  $N$  feedforward neurons to  $2N$ . In contrast, both DH-DFC and Equilibrium Propagation rely on the same  $N$  neurons for the feedforward pass and the learning updates. However, Equilibrium Propagation requires a symmetry of the weights and weight updates, imposing a specific feedback architecture and an *ad-hoc* learning rule. In summary, we note that the DH-DFC is more parsimonious in the sense that it makes very simple assumptions on the feedback and requires less complex model architectures.

### 3. Discussion


Building upon previous studies, our work represents another leap forward to understanding the different aspects of hierarchical learning in biological networks. In the following sections, we go through the relationship between our work and previous works on computational and experimental neuroscience as well as limitations and future directions.

#### 3.1. How does our work fit into the existing literature

A key contribution of our work is the connection between experimentally observed learning rules and computational models that can train deep networks. In this section, we discuss how this work fits with (1) the electrophysiology literature on learning rules, (2) temporal Hebbian learning rules both at the neuron and network level, (3) Predictive Coding and the combination of bottom-up inputs and top-down feedback, and (4) other bioplausible deep learning models.

##### 3.1.1. Electrophysiological observations that agree with our model

In biological neural networks, LTP and LTD are one of the most prevalent forms of synaptic plasticity, and various studies have shown that LTP is induced when presynaptic spikes precede postsynaptic ones. In the case of multiple spike pairs, this is consistent with our model in that an increase in postsynaptic activity would lead to LTP and a decrease in LTD. Interestingly, recent work suggests that classical STDP-inducing protocols might fail under physiological extracellular calcium concentrations, suggesting that additional mechanisms might be required to act on the intracellular calcium levels (Larkum et al., 1999; Inglebert et al., 2020). In pyramidal neurons, intracellular calcium levels can be modulated by backpropagating action potential-evoked calcium (BAC) spikes that arise when apical inputs arrive shortly after basal inputs, resulting in action potential bursts (Larkum et al., 1999). Our model is consistent with this notion that delayed feedback into the apical dendrite drives plasticity while basal feedforward input does not. Future neuroscience experiments should explore if high




**FIGURE 4** Comparison of learning algorithms. We calculate the synaptic weight update in our deep network using different algorithms and synaptic plasticity rules. We compare our DH weight updates (x-axis) to the updates given by other algorithms. Weight update correlations between DH-DFC and: (A) BP; (B) DFC; (C) and STDP. We observe a clear correlation between DH-DFC and both BP and DFC, and a significant but weaker correlation with STDP.

calcium concentrations resulting from BAC spikes and bursts are indeed suitable to restore LTP and LTD induction when using a classical STDP protocol (Inglebert et al., 2020).

Finally, our model requires feedback that is specific to every neuron. Therefore, the synaptic weights to the apical dendrite have very specific values that must be computed by some biological mechanism. In previous work, we showed that these weights can be learned in a bioplausible manner by an anti-Hebbian learning rule (Meulemans et al., 2020). In biology, anti-Hebbian learning rules appear in disinhibitory GABAergic synapses (Lamsa et al., 2007), suggesting that the target used for learning in our model would be fed back into excitatory neurons through disinhibitory circuits. This nicely relates our work to the role of coupled apical and basal inputs in learning and the regulation of this coupling by disinhibitory circuits (Zhang et al., 2014; Avital et al., 2019; Williams and Holtmaat, 2019), and therefore use connectivity that matches the requirements of our feedback-based target propagation framework. Future theoretical investigations should continue this line of work by looking beyond Hebbian-like learning rules and integrating the knowledge of BAC-firing dynamics, the effects of calcium on plasticity, and the role of disinhibitory circuits in bioplausible models of deep learning.

### 3.1.2. Learning with temporal Hebbian learning rules

Temporal Hebbian learning rules such as STDP or DH rules have been mostly used for unsupervised learning (Gerstner et al., 1996; Toyozumi et al., 2005; Lazar et al., 2009; Sjöström and Gerstner, 2010) or as an enhancement of supervised learning in shallow networks (Diehl and Cook, 2015). In order to use these rules in a supervised setting, they require a teaching signal, which can be implemented either through a neuromodulator or a third-factor learning rule (Frémaux and Gerstner, 2016). However, such approaches do not go beyond shallow networks (Illing et al., 2019) and, although it has been suggested that STDP or DH could be adopted for error-driven hierarchical learning (Xie and Seung, 1999; Hinton, 2007; Bengio et al., 2017), a suitable network architecture and dynamics to combine time-dependent Hebbian learning rules with deep networks has not been proposed yet (Bengio et al., 2015). Our work fills this gap by presenting an approach that is able to train deep hierarchies with a learning



**FIGURE 5** Surprise triggers a large feedback signal that alters neuronal activities. Across learning, the change in the post-synaptic activity driven by the apical feedback reduces as neurons reach their target rates. However, when the labels are randomly swapped (indicated by the arrow in epoch 20), the apical feedback is notably increased. Note that the label switch did not set the network to its baseline state, because the required feedback decreased to the pre-shift level much faster than on the first epochs.

rule that retains the time-based principles of STDP. This in turn connects deep network optimization to latency reduction, a well-known effect of STDP where neurons fire earlier in time every time that an input sequence is presented (Masquelier et al., 2009; Vilimelis Aceituno et al., 2020; Saponati and Vinck, 2021). This had been studied only at the level of neurons but we now turned it into a systems-level optimization process.

### 3.1.3. Predictive Coding and top-down feedback

Due to the close relation of our model to Predictive Coding (PC), we next compare our approach to PC. In the PC literature, learning decreases the amount of top-down feedback. This process intrinsically generates expedited neuronal responses after stimulus presentation, which are often interpreted as predictions (Friston and Kiebel, 2009; Whittington and Bogacz, 2017; Keller and Mrsic-Flogel, 2018). The PC framework goes beyond explanations of these

activities by proposing neural circuits that could implement this behavior (Rao and Ballard, 1999; Bastos et al., 2012).

However, PC as a mechanistic theory for neural circuits requires explicit error encoding (Koch and Poggio, 1999; Rao and Ballard, 1999; Bastos et al., 2012), a requirement which is problematic for making valid testable predictions (Kogo and Trengove, 2015). In contrast, our framework can exhibit a similar reduction of top-down feedback and anticipated neuronal responses. Still, since it is based on the target activities of neurons rather than on errors, it does not require explicit errors to be encoded. This shows that it is actually possible to design neural circuits that can reproduce the relevant PC features while representing errors implicitly with the temporal neuronal dynamics. To illustrate this effect, in Figure 5 we plot the feedback that modulates a deep network during training. The feedback decreases as the model learns but, when we randomly shuffle the labels—which can be considered a surprising response—the feedback signal increases substantially, thereby changing the neuronal activity in accordance with experimental observations (Keller and Mrcic-Flogel, 2018).

### 3.1.4. Alternative bioplausible deep learning models

Other bioplausible deep networks models such as Equilibrium Propagation, dendritic-error learning, or Burst Propagation require a learning signal to be computed either by using two separate phases (Scellier and Bengio, 2017), distinct dendritic and somatic compartments (Sacramento et al., 2018) or via multiplexing of feedback and feedforward signals as bursts and single spikes (Payeur et al., 2021), respectively. In contrast, our model encodes supervision signals as temporal changes in postsynaptic activities, which arrive at individual neurons via their apical dendrite with a short time delay. Table 2 provides a comprehensive comparison of our approach to the most recent alternative bioplausible deep learning methods and how they relate to experimental observations.

The relationship between temporal dynamics and bioplausible deep learning has been explored before. This was done through different methods, for instance: by making use of subsequent frames, usually in an unsupervised or self-supervised setting (Illing et al., 2020; Lotter et al., 2020); or having a combination of STDP and reward signals (Mozafari et al., 2019; Illing et al., 2020); or, more generally, with the so-called temporal error learning framework (Wittenberg and Wang, 2006). Our model applies a similar principle but with a supervised target and at the level of neuronal dynamics.

## 3.2. Limitations and future work

### 3.2.1. Limitations

On the experimental side, our framework requires a top-down controller to continuously compare the actual network output to the desired one, while sending feedback to the lower hierarchies. Although such a feedback controller can be easily realized as a neural circuit (Meulemans et al., 2021a), it is not clear yet if the brain employs any type of control circuit for learning. Future work

could look at whether the apical inputs going through disinhibitory circuits correspond to feedback inputs that drive neurons to a target activity that stabilizes the top-down feedback.

From a modeling perspective, weights from the same neuron can be positive and negative or even transition from negative to positive and vice-versa, which is in conflict with Dale's law. This is a common simplification of ANN models (Cornford et al., 2020). Violating Dale's law, however, can be corrected using a bias in the postsynaptic activity to turn negative weights into weak positive weights (Kriegeskorte and Golan, 2019). Moreover, recent studies showed that with certain network architectures and priors, Dale's law can be easily preserved while maintaining the same functional network properties (Cornford et al., 2020).

Another limitation of our work is that we use DH instead of STDP to train deep networks. This is due to the randomness induced by our implementation of spiking neurons using a Poisson model, which implicitly imposes noisy learning updates. Further work could use leaky integrate-and-fire neurons, which can reduce the effects of randomness. This would require computing feedback in an event-based network, which is a currently active area of research.

At the computational level, our method requires a long time to be simulated because the controller works by updating the neural activity in small incremental steps, requiring as many as a hundred forward passes for each sample, which is much more than off-the-shelf learning algorithms but in line with previous works using feedback mechanisms (Scellier and Bengio, 2017; Rosenbaum, 2022). Similarly to the previous point, the use of an event-based network would greatly reduce the computational costs of learning by reducing the control cost only to relevant events.

### 3.2.2. Future work

After learning, our model predicts an expedited onset of pyramidal neuron activity upon feedforward input (Figure 3) that is inversely correlated with the top-down feedback to alter neuronal activity. A related cortical micro-circuit hypothesis is that local inhibitory microcircuits projecting onto apical dendrites control the neuron's excitability and that their control strength reduces during learning. In an experimental setting, this temporal shift as well as the feedback strength attenuation could be tested using simultaneous *in vivo* 2-photon calcium imaging of excitatory and inhibitory populations (as in Han et al., 2019) combined with a plasticity-inducing whisker stimulation paradigm.

From a computational perspective, follow-up studies should go beyond modeling phenomenological learning rules such as STDP into hierarchical networks. For example, one direction could be to develop a more detailed mechanistic sub-cellular model that accounts for the coupling of intracellular voltage and calcium dynamics that are being differently modulated by inputs to apical and somatic synapses. Such sub-cellular mechanistic models might also include multiplicative effects of the apical input (Larkum et al., 2004) as well as apical-induced bursting (Segal, 2018) to further close the gap between the correlation-based models used in computational neuroscience and experimental observations showing, for example, the diverse intracellular effects of calcium on



TABLE 2 Comparison of diverse bioplausible hierarchical learning methods. For further details on these methods: Equilibrium Propagation (Scellier and Bengio, 2017), Predictive Coding (Whittington and Bogacz, 2017), Dendritic-error (Sacramento et al., 2018), Burst Propagation (Payeur et al., 2021).

Method	Implicit error		Explicit error		
	Temporal Hebbian	Equilibrium Propagation	Predictive Coding	Dendritic-error	Burst Propagation
Error type	Control	Contrastive	Prediction	Prediction	Multiplexed
Phases/Compartments	1/1	2/1	1/2	1/2	1/2
Network connectivity	Unconstrained	Unconstrained	Constrained	Constrained	Constrained
Weight symmetry	Not required	Required	Arises with learning	Not required	Arises with learning
Feedback magnitude	Strong feedback	Weak nudging	Weak error	Weak nudging	Weak error
Error encoding	Change in firing rate	Difference in neuronal activity	Activity of error neurons	Difference apical/somatic activity	Ratio bursts/single spikes
Update rule	Temporal Hebbian	Contrastive Hebbian	Hebbian	Apical/somatic error	Burst rate
Update timing	Continuous	Timed	Timed	Timed	Continuous
MNIST performance (%)	~1.9	~2–3	~1.7–1.8	~2.0	~1.1
Possible link to observations in neuroscience	Apical inputs affect postsynaptic activity, increased feedback activity upon new stimuli	Varying neural responses for different behaviors	Increased neural activity upon new stimuli	Apical inputs affect activity and plasticity	Apical inputs affect burst rate and plasticity
Bioplausibility	+++	+	++	++	+++

Note that the MNIST performance comparison does not take into account different model sizes. The +++ symbol corresponds to high biological plausibility, the ++ symbol corresponds to medium biological plausibility, and symbol corresponds the + to low biological plausibility.

learning and neuronal activity (Larkum et al., 1999, 2007; Larkum, 2013). Another logical future step would be to develop more explicit theoretical links between PC and our temporal Hebbian framework. This would require applying it to other problems, such as detecting deviations from learned time series (Garrido et al., 2009) or unsupervised image representations (Rao and Ballard, 1999) and comparing the reduction of feedback with the minimization of prediction errors or free energy (Friston and Kiebel, 2009). Showing such conceptual links would pave the way to design more cortical-like circuits that explain Predictive Coding features but avoid the problems emerging from explicit error neurons (Kogo and Trengove, 2015).

Our framework can be leveraged to build the theory in spiking neural networks, where the processing of time-centered losses is still in its infancy. For example, it would be interesting to see how the notion of control cost or latency to target response interplay with information theory metrics, which have been shown to be useful for continuous learning or few-shot learning (Yang et al., 2022b,c). Similarly, if using the multi-compartment neuron formulation of our model, one could include other relevant features such as working memory (Yang et al., 2022a).

Finally, the simplicity and locality of the model we propose makes it well-suited for on-chip event-based learning applications. This would require integrating a simple PI controller in a neuromorphic processor and further theoretical work on implementing our learning set-up with leaky integrate-and-fire neurons. Given that STDP can induce energy-efficient

representations (Vilimelis Aceituno et al., 2020), it is likely that training with STDP might even further improve the energy efficiency of neuromorphic devices. In addition, the fact that our framework can learn all weights in an online manner (Meulemans et al., 2021b) implies that a perfect model of the processor architecture is not required, which is a key problem when training neuromorphic devices off-line due to the so-called device mismatch (Pelgrom et al., 1989; Binas et al., 2016).

## 4. Conclusions

With this work, we present a new hierarchical learning framework in which the temporal order of neuronal signals is leveraged to encode top-down error signals. This reformulation of the error allows us to avoid unobserved learning rules while at the same time being consistent with classical ideas of Predictive Coding. Our work is a crucial step toward a more detailed understanding of how temporal Hebbian and STDP learning can be used for supervised learning in multilayer neural networks.

## Data availability statement

The code repository is publicly available in GitHub: <https://github.com/MatildeTristany/Learning-Cortical-Hierarchies-with-Temporal-Hebbian-Updates>.

## Author contributions

PA, MF, and BG designed the project and the experiments and wrote the paper. PA developed the mathematical STDP framework and performed the single neuron and Predictive Coding simulations. MF performed the network simulations. RL provided biology insights about the project on both theoretical and experimental neuroscience and wrote part of the discussion section.

## Funding

This work was supported by the Swiss National Science Foundation (BG CRSII5-173721 and 315230189251), ETH project funding (BG ETH-20 19-01), and the Human Frontiers Science Program (RGY0072/2019). PA was supported by an ETH Zurich Postdoctoral fellowship Grant number ETH-1-007113-000. Open access funding by ETH Zurich.

## Acknowledgments

We thank Christoph van der Malsburg, Alexander Meulemans, and Matthew Cook for fruitful discussions. We thank Maria Cervera for initially exploring the Predictive Coding properties of

## References

- Abbott, L. F., and Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nat. Neurosci.* 3, 1178–1183. doi: 10.1038/81453
- Avital, A., Ruohe, Z., Myung, E. S., Ryohei, Y., and Gan, W.-B. (2019). Somatostatin-expressing interneurons enable and maintain learning-dependent sequential activation of pyramidal neurons. *Neuron* 102, 202–216. doi: 10.1016/j.neuron.2019.01.036
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., and Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron* 76, 695–711. doi: 10.1016/j.neuron.2012.10.038
- Bengio, Y., Lee, D.-H., Bornschein, J., Mesnard, T., and Lin, Z. (2015). Towards biologically plausible deep learning. *arXiv preprint arXiv:1502.04156*.
- Bengio, Y., Mesnard, T., Fischer, A., Zhang, S., and Wu, Y. (2017). Stdp-compatible approximation of back-propagation in an energy-based model. *Neural Comput.* 29, 555–577. doi: 10.1162/NECO\_a\_00934
- Bi, G.-Q., and Poo, M.-M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472. doi: 10.1523/JNEUROSCI.18-24-10464.1998
- Binas, J., Neil, D., Indiveri, G., Liu, S.-C., and Pfeiffer, M. (2016). Precise deep neural network computation on imprecise low-power analog hardware. *Comput. Sci. Neural Evol. Comput.* 1606.
- Cornford, J., Kalajdzievski, D., Leite, M., Lamarquette, A., Kullmann, D. M., and Richards, B. (2020). Learning to live with dale's principle: ANNs with separate excitatory and inhibitory units. *bioRxiv [Preprint]*. doi: 10.1101/2020.11.02.364968
- Crick, F. (1989). The recent excitement about neural networks. *Nature* 337, 129–132. doi: 10.1038/337129a0
- Diehl, P. U., and Cook, M. (2015). Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Front. Comput. Neurosci.* 9:99. doi: 10.3389/fncom.2015.00099
- Feldman, D. (2012). The spike-timing dependence of plasticity. *Neuron* 75, 556–571. doi: 10.1016/j.neuron.2012.08.001
- Frémaux, N., and Gerstner, W. (2016). Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. *Front. Neural Circuits* 9:85. doi: 10.3389/fncir.2015.00085
- Friston, K., and Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philos. Trans. R. Soc. B Biol. Sci.* 364, 1211–1221. doi: 10.1098/rstb.2008.0300
- Garrido, M. I., Kilner, J. M., Stephan, K. E., and Friston, K. J. (2009). The mismatch negativity: a review of underlying mechanisms. *Clin. Neurophysiol.* 120, 453–463. doi: 10.1016/j.clinph.2008.11.029
- Gerstner, W., Kempter, R., van Hemmen, J. L., and Wagner, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature* 383:76. doi: 10.1038/383076a0
- Gilbert, C. D., and Li, W. (2013). Top-down influences on visual processing. *Nat. Rev. Neurosci.* 14, 350–363. doi: 10.1038/nrn3476
- Gilra, A., and Gerstner, W. (2017). Predicting non-linear dynamics by stable local learning in a recurrent spiking neural network. *Elife* 6:e28295. doi: 10.7554/eLife.28295
- Gütig, R. (2016). Spiking neurons can discover predictive features by aggregate-label learning. *Science* 351:aab4113. doi: 10.1126/science.aab4113
- Han, S., Yang, W., and Yuste, R. (2019). Two-color volumetric imaging of neuronal activity of cortical columns. *Cell Rep.* 27, 2229–2240. doi: 10.1016/j.celrep.2019.04.075
- Hinton, G. (2007). "How to do backpropagation in a brain," in *Invited Talk at the NIPS'2007 Deep Learning Workshop, Vol. 656* (Vancouver, BC), 1–16.
- Illing, B., Gerstner, W., and Bellec, G. (2020). Towards truly local gradients with CLAPP: contrastive, local and predictive plasticity. *CoRR, abs/2010.08262*.
- Illing, B., Gerstner, W., and Brea, J. (2019). Biologically plausible deep learning-but how far can we go with shallow networks? *Neural Netw.* 118, 90–101. doi: 10.1016/j.neunet.2019.06.001
- Inglebert, Y., Aljadeff, J., Brunel, N., and Debanne, D. (2020). Synaptic plasticity rules with physiological calcium levels. *Proc. Natl. Acad. Sci. U.S.A.* 117, 33639–33648. doi: 10.1073/pnas.2013663117
- Izhikevich, E. M., and Desai, N. S. (2003). Relating STDP to BCM. *Neural Comput.* 15, 1511–1523. doi: 10.1162/089976603321891783
- Keller, G. B., and Mrsic-Flogel, T. D. (2018). Predictive processing: a canonical cortical computation. *Neuron* 100, 424–435. doi: 10.1016/j.neuron.2018.10.003
- Kempter, R., Gerstner, W., and van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. *Phys. Rev. E* 59, 4498–4514. doi: 10.1103/PhysRevE.59.4498
- Koch, C., and Poggio, T. (1999). Predicting the visual world: silence is golden. *Nat. Neurosci.* 2, 9–10. doi: 10.1038/4511
- Kogo, N., and Trengove, C. (2015). Is predictive coding theory articulated enough to be testable? *Front. Comput. Neurosci.* 9:111. doi: 10.3389/fncom.2015.00111

the standard DFC framework. We thank Sander de Haan for being part of the discussions and providing feedback on the paper.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

## Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fncom.2023.1136010/full#supplementary-material>

- Kriegeskorte, N., and Golan, T. (2019). Neural network models and deep learning. *Curr. Biol.* 29, R231–R236. doi: 10.1016/j.cub.2019.02.034
- Lamsa, K., Heeroma, J., Somogyi, P., Rusakov, D., and Kullmann, D. (2007). Anti-hebbian long-term potentiation in the hippocampal feedback inhibitory circuit. *Science* 315, 1262–1266. doi: 10.1126/science.1137450
- Larkum, M. (2013). A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. *Trends Neurosci.* 36, 141–151. doi: 10.1016/j.tins.2012.11.006
- Larkum, M. E., Senn, W., and Lüscher, H.-R. (2004). Top-down dendritic input increases the gain of layer 5 pyramidal neurons. *Cereb. Cortex* 14, 1059–1070. doi: 10.1093/cercor/bhh065
- Larkum, M. E., Waters, J., Sakmann, B., and Helmchen, F. (2007). Dendritic spikes in apical dendrites of neocortical layer 2/3 pyramidal neurons. *J. Neurosci.* 27, 8999–9008. doi: 10.1523/JNEUROSCI.1717-07.2007
- Larkum, M. E., Zhu, J. J., and Sakmann, B. (1999). A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* 398:338. doi: 10.1038/18686
- Lazar, A., Pipa, G., and Triesch, J. (2009). Sorn: a self-organizing recurrent neural network. *Front. Comput. Neurosci.* 3:23. doi: 10.3389/fncom.2009.10.023.2009
- LeCun, Y. (1998). *The mnist Database of Handwritten Digits*. Available online at: <http://yann.lecun.com/exdb/mnist/>
- Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J., and Hinton, G. (2020). Backpropagation and the brain. *Nat. Rev. Neurosci.* 21, 335–346. doi: 10.1038/s41583-020-0277-3
- Lotter, W., Kreiman, G., and Cox, D. (2020). A neural network trained for prediction mimics diverse features of biological neurons and perception. *Nat. Mach. Intell.* 2, 210–219. doi: 10.1038/s42256-020-0170-9
- Luczak, A., McNaughton, B. L., and Kubo, Y. (2022). Neurons learn by predicting future activity. *Nat. Mach. Intell.* 4, 62–72. doi: 10.1038/s42256-021-00430-y
- Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213–215. doi: 10.1126/science.275.5297.213
- Masquelier, T., Guyonneau, R., and Thorpe, S. J. (2009). Competitive STDP-based spike pattern learning. *Neural Comput.* 21, 1259–1276. doi: 10.1162/neco.2008.06.08.804
- Meulemans, A., Carzaniga, F., Suykens, J., Sacramento, J. a., and Grewe, B. F. (2020). A theoretical framework for target propagation. *Adv. Neural Inform. Process. Syst.* 33, 20024–20036.
- Meulemans, A., Farinha, M. T., Cervera, M. R., Sacramento, J., and Grewe, B. F. (2022a). “Minimizing control for credit assignment with strong feedback,” in *International Conference on Machine Learning* (Baltimore, MD), 15458–15483.
- Meulemans, A., Farinha, M. T., Ordóñez, J. G., Aceituno, P. V., Sacramento, J., and Grewe, B. F. (2021a). Credit assignment in neural networks through deep feedback control. *arXiv preprint arXiv:2106.07887*.
- Meulemans, A., Farinha, M. T., Ordóñez, J. G., Aceituno, P. V., Sacramento, J., and Grewe, B. F. (2021b). Credit assignment in neural networks through deep feedback control. *CoRR, abs/2106.07887*.
- Meulemans, A., Zucchet, N., Kobayashi, S., Von Oswald, J., and Sacramento, J. (2022b). The least-control principle for local learning at equilibrium. *Adv. Neural Inform. Process. Syst.* 35, 33603–33617
- Mozafari, M., Ganjtabesh, M., Nowzari-Dalini, A., and Masquelier, T. (2019). Spynetorch: efficient simulation of convolutional spiking neural networks with at most one spike per neuron. *Front. Neurosci.* 13:625. doi: 10.3389/fnins.2019.00625
- Nishiyama, M., Hong, K., Mikoshiba, K., Poo, M.-M., and Kato, K. (2000). Calcium stores regulate the polarity and input specificity of synaptic modification. *Nature* 408, 584–588. doi: 10.1038/35046067
- Payeur, A., Guerguiev, J., Zenke, F., Richards, B., and Naud, R. (2021). Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nat. Neurosci.* 24:1546. doi: 10.1038/s41593-021-00857-x
- Pelgrom, M. J., Duijnmaier, A. C., and Welbers, A. P. (1989). Matching properties of mos transistors. *IEEE J. Solid-State Circuits* 24, 1433–1439. doi: 10.1109/JSSC.1989.572629
- Rao, R. P., and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2, 79–87. doi: 10.1038/4580
- Rosenbaum, R. (2022). On the relationship between predictive coding and backpropagation. *PLoS ONE* 17:e0266102. doi: 10.1371/journal.pone.0266102
- Sacramento, J., Costa, R. P., Bengio, Y., and Senn, W. (2018). “Dendritic cortical microcircuits approximate the backpropagation algorithm,” in *Advances in Neural Information Processing Systems, Vol. 31*, 8721–8732.
- Saponati, M., and Vinck, M. (2021). Sequence anticipation and STDP emerge from a voltage-based predictive learning rule. *bioRxiv [Preprint]*. doi: 10.1101/2021.10.31.466667
- Saudargiene, A., Porr, B., and Wörgötter, F. (2004). How the shape of pre- and postsynaptic signals can influence STDP: a biophysical model. *Neural Comput.* 16, 595–625. doi: 10.1162/08997660472744929
- Scellier, B., and Bengio, Y. (2017). Equilibrium propagation: bridging the gap between energy-based models and backpropagation. *Front. Comput. Neurosci.* 11:24. doi: 10.3389/fncom.2017.00024
- Segal, M. (2018). Calcium stores regulate excitability in cultured rat hippocampal neurons. *J. Neurophysiol.* 120, 2694–2705. doi: 10.1152/jn.00447.2018
- Sjöström, J., and Gerstner, W. (2010). Spike-timing dependent plasticity. *Scholarpedia* 5, 1362. doi: 10.4249/scholarpedia.1362
- Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32, 1149–1164. doi: 10.1016/S0896-6273(01)00542-6
- Song, S., Miller, K. D., and Abbott, L. F. (2000). Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.* 3, 919–926. doi: 10.1038/78829
- Toyoizumi, T., Pfister, J.-P., Aihara, K., and Gerstner, W. (2005). Generalized bienstock-cooper-munro rule for spiking neurons that maximizes information transmission. *Proc. Natl. Acad. Sci. U.S.A.* 102, 5239–5244. doi: 10.1073/pnas.0500495102
- Tschantz, A. (2020). *A Python Implementation of An Approximation of the Error Backpropagation Algorithm in a Predictive Coding Network with Local Hebbian Synaptic Plasticity*. Available online at: <https://github.com/alec-tschantz/predcoding>
- van Rossum, M. C. W., Bi, G. Q., and Turrigiano, G. G. (2000). Stable hebbian learning from spike timing-dependent plasticity. *J. Neurosci.* 20, 8812–8821. doi: 10.1523/JNEUROSCI.20-23-08812.2000
- Vilimelis Aceituno, P., Ehsani, M., and Jost, J. (2020). Spiking time-dependent plasticity leads to efficient coding of predictions. *Biol. Cybern.* 114, 43–61. doi: 10.1007/s00422-019-00813-w
- Whittington, J. C., and Bogacz, R. (2017). An approximation of the error backpropagation algorithm in a predictive coding network with local hebbian synaptic plasticity. *Neural Comput.* 29, 1229–1262. doi: 10.1162/NECO\_a\_00949
- Williams, L. E., and Holtmaat, A. (2019). Higher-order thalamocortical inputs gate synaptic long-term potentiation via disinhibition. *Neuron* 101, 91–102. doi: 10.1016/j.neuron.2018.10.049
- Wittenberg, G. M., and Wang, S. S.-H. (2006). Malleability of spike-timing-dependent plasticity at the CA3-CA1 synapse. *J. Neurosci.* 26, 6610–6617. doi: 10.1523/JNEUROSCI.5388-05.2006
- Xie, X., and Seung, H. S. (1999). “Spike-based learning rules and stabilization of persistent neural activity,” in *Advances in Neural Information Processing Systems, Vol. 12*, eds S.olla, T. Leen, and K. Müller (Breckenridge, CO: MIT Press).
- Yang, S., Gao, T., Wang, J., Deng, B., Azghadi, M. R., Lei, T., et al. (2022a). Sam: a unified self-adaptive multicompartmental spiking neuron model for learning with working memory. *Front. Neurosci.* 16:850945. doi: 10.3389/fnins.2022.850945
- Yang, S., Linares-Barranco, B., and Chen, B. (2022b). Heterogeneous ensemble-based spike-driven few-shot online learning. *Front. Neurosci.* 16:850932. doi: 10.3389/fnins.2022.850932
- Yang, S., Tan, J., and Chen, B. (2022c). Robust spike-based continual meta-learning improved by restricted minimum error entropy criterion. *Entropy* 24:455. doi: 10.3390/e24040455
- Zappacosta, S., Mannella, F., Miroli, M., and Baldassarre, G. (2018). General differential Hebbian learning: capturing temporal relations between events in neural networks and the brain. *PLoS Comput. Biol.* 14:e1006227. doi: 10.1371/journal.pcbi.1006227
- Zhang, S., Xu, M., Kamigaki, T., Do, J. P. H., Chang, W.-C., Jenvay, S., et al. (2014). Long-range and local circuits for top-down modulation of visual cortex processing. *Science* 345, 660–665. doi: 10.1126/science.1254126