

---

# Dendritic cortical microcircuits approximate the backpropagation algorithm

---

**João Sacramento\***

Department of Physiology  
University of Bern, Switzerland  
sacramento@pyl.unibe.ch

**Rui Ponte Costa†**

Department of Physiology  
University of Bern, Switzerland  
costa@pyl.unibe.ch

**Yoshua Bengio‡**

Mila and Université de Montréal, Canada  
yoshua.bengio@mila.quebec

**Walter Senn**

Department of Physiology  
University of Bern, Switzerland  
senn@pyl.unibe.ch

## Abstract

Deep learning has seen remarkable developments over the last years, many of them inspired by neuroscience. However, the main learning mechanism behind these advances – error backpropagation – appears to be at odds with neurobiology. Here, we introduce a multilayer neuronal network model with simplified dendritic compartments in which error-driven synaptic plasticity adapts the network towards a global desired output. In contrast to previous work our model does not require separate phases and synaptic learning is driven by local dendritic prediction errors continuously in time. Such errors originate at apical dendrites and occur due to a mismatch between predictive input from lateral interneurons and activity from actual top-down feedback. Through the use of simple dendritic compartments and different cell-types our model can represent both error and normal activity within a pyramidal neuron. We demonstrate the learning capabilities of the model in regression and classification tasks, and show analytically that it approximates the error backpropagation algorithm. Moreover, our framework is consistent with recent observations of learning between brain areas and the architecture of cortical microcircuits. Overall, we introduce a novel view of learning on dendritic cortical circuits and on how the brain may solve the long-standing synaptic credit assignment problem.

## 1 Introduction

Machine learning is going through remarkable developments powered by deep neural networks (LeCun et al., 2015). Interestingly, the workhorse of deep learning is still the classical backpropagation of errors algorithm (backprop; Rumelhart et al., 1986), which has been long dismissed in neuroscience on the grounds of biologically implausibility (Grossberg, 1987; Crick, 1989). Irrespective of such concerns, growing evidence demonstrates that deep neural networks outperform alternative frameworks in accurately reproducing activity patterns observed in the cortex (Lillicrap and Scott, 2013; Yamins et al., 2014; Khaligh-Razavi and Kriegeskorte, 2014; Yamins and DiCarlo, 2016; Kell et al., 2018). Although recent developments have started to bridge the gap between neuroscience

---

\*Present address: Institute of Neuroinformatics, University of Zürich and ETH Zürich, Zürich, Switzerland

†Present address: Computational Neuroscience Unit, Department of Computer Science, SCEEM, Faculty of Engineering, University of Bristol, United Kingdom

‡CIFAR Senior Fellow

and artificial intelligence (Marblestone et al., 2016; Lillicrap et al., 2016; Scellier and Bengio, 2017; Costa et al., 2017; Guerguiev et al., 2017), how the brain could implement a backprop-like algorithm remains an open question.

In neuroscience, understanding how the brain learns to associate different areas (e.g., visual and motor cortices) to successfully drive behaviour is of fundamental importance (Petreanu et al., 2012; Manita et al., 2015; Makino and Komiyama, 2015; Poort et al., 2015; Fu et al., 2015; Pakan et al., 2016; Zmarz and Keller, 2016; Attinger et al., 2017). However, how to correctly modify synapses to achieve this has puzzled neuroscientists for decades. This is often referred to as the synaptic credit assignment problem (Rumelhart et al., 1986; Sutton and Barto, 1998; Roelfsema and van Ooyen, 2005; Friedrich et al., 2011; Bengio, 2014; Lee et al., 2015; Roelfsema and Holtmaat, 2018), for which the backprop algorithm provides an elegant solution.

Here we propose that the prediction errors that drive learning in backprop are encoded at distal dendrites of pyramidal neurons, which receive top-down input from downstream brain areas (we interpret a brain area as being equivalent to a layer in machine learning) (Petreanu et al., 2009; Larkum, 2013). In our model, these errors arise from the inability to exactly match via lateral input from local interneurons (e.g. somatostatin-expressing; SST) the top-down feedback from downstream cortical areas. Learning of bottom-up connections (i.e., feedforward weights) is driven by such error signals through local synaptic plasticity. Therefore, in contrast to previous approaches (Marblestone et al., 2016), in our framework a given neuron is used simultaneously for activity propagation (at the somatic level), error encoding (at distal dendrites) and error propagation to the soma without the need for separate phases.

We first illustrate the different components of the model. Then, we show analytically that under certain conditions learning in our network approximates backpropagation. Finally, we empirically evaluate the performance of the model on nonlinear regression and recognition tasks.

## 2 Error-encoding dendritic cortical microcircuits

### 2.1 Neuron and network model

Building upon previous work (Urbanczik and Senn, 2014), we adopt a simplified multicompartment neuron and describe pyramidal neurons as three-compartment units (schematically depicted in Fig. 1A). These compartments represent the somatic, basal and apical integration zones that characteristically define neocortical pyramidal cells (Spruston, 2008; Larkum, 2013). The dendritic structure of the model is exploited by having bottom-up and top-down synapses converging onto separate dendritic compartments (basal and distal dendrites, respectively), a first approximation in line with experimental observations (Spruston, 2008) and reflecting the preferred connectivity patterns of cortico-cortical projections (Larkum, 2013).

Consistent with the connectivity of SST interneurons (Urban-Ciecko and Barth, 2016), we also introduce a second population of cells within each hidden layer with both lateral and cross-layer connectivity, whose role is to *cancel the top-down input* so as to leave only the backpropagated errors as apical dendrite activity. Modelled as two-compartment units (depicted in red, Fig. 1A), such interneurons are predominantly driven by pyramidal cells within the same layer through weights  $\mathbf{W}_{k,k}^{\text{IP}}$ , and they project back to the apical dendrites of the same-layer pyramidal cells through weights  $\mathbf{W}_{k,k}^{\text{PI}}$  (Fig. 1A). Additionally, cross-layer feedback onto SST cells originating at the next upper layer  $k+1$  provide a weak nudging signal for these interneurons, modelled after Urbanczik and Senn (2014) as a conductance-based somatic input current. We modelled this weak top-down nudging on a one-to-one basis: each interneuron is nudged towards the potential of a corresponding upper-layer pyramidal cell. Although the one-to-one connectivity imposes a restriction in the model architecture, this is to a certain degree in accordance with recent monosynaptic input mapping experiments show that SST cells in fact receive top-down projections (Leinweber et al., 2017), that according to our proposal may encode the weak interneuron ‘teaching’ signals from higher to lower brain areas.

The somatic membrane potentials of pyramidal neurons and interneurons evolve in time according to

$$\frac{d}{dt} \mathbf{u}_k^{\text{P}}(t) = -g_{\text{lk}} \mathbf{u}_k^{\text{P}}(t) + g_{\text{B}} (\mathbf{v}_{\text{B},k}^{\text{P}}(t) - \mathbf{u}_k^{\text{P}}(t)) + g_{\text{A}} (\mathbf{v}_{\text{A},k}^{\text{P}}(t) - \mathbf{u}_k^{\text{P}}(t)) + \sigma \boldsymbol{\xi}(t) \quad (1)$$

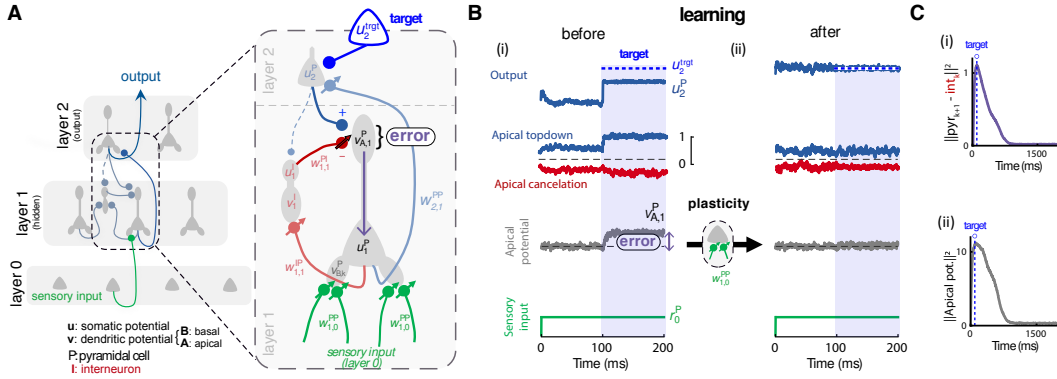
$$\frac{d}{dt} \mathbf{u}_k^{\text{I}}(t) = -g_{\text{lk}} \mathbf{u}_k^{\text{I}}(t) + g_{\text{D}} (\mathbf{v}_k^{\text{I}}(t) - \mathbf{u}_k^{\text{I}}(t)) + \mathbf{i}_k^{\text{I}}(t) + \sigma \boldsymbol{\xi}(t), \quad (2)$$

with one such pair of dynamical equations for every hidden layer  $0 < k < N$ ; input layer neurons are indexed by  $k = 0$ ,  $g$ 's are fixed conductances,  $\sigma$  controls the amount of injected noise. Basal and apical dendritic compartments of pyramidal cells are coupled to the soma with effective transfer conductances  $g_B$  and  $g_A$ , respectively. Subscript lk is for leak, A is for apical, B for basal, D for dendritic, superscript I for inhibitory and P for pyramidal neuron. Eqs. 1 and 2 describe standard conductance-based voltage integration dynamics, having set membrane capacitance to unity and resting potential to zero for clarity. Background activity is modelled as a Gaussian white noise input,  $\xi$  in the equations above. To keep the exposition brief we use matrix notation, and denote by  $\mathbf{u}_k^P$  and  $\mathbf{u}_k^I$  the vectors of pyramidal and interneuron somatic voltages, respectively. Both matrices and vectors, assumed column vectors by default, are typed in boldface here and throughout. Dendritic compartmental potentials are denoted by  $\mathbf{v}$  and are given in instantaneous form by

$$\mathbf{v}_{B,k}^P(t) = \mathbf{W}_{k,k-1}^{PP} \phi(\mathbf{u}_{k-1}^P(t)) \quad (3)$$

$$\mathbf{v}_{A,k}^P(t) = \mathbf{W}_{k,k+1}^{PP} \phi(\mathbf{u}_{k+1}^P(t)) + \mathbf{W}_{k,k}^{PI} \phi(\mathbf{u}_k^I(t)), \quad (4)$$

where  $\phi(\mathbf{u})$  is the neuronal transfer function, which acts componentwise on  $\mathbf{u}$ .



**Figure 1: Learning in error-encoding dendritic microcircuit network.** (A) Schematic of network with pyramidal cells and lateral inhibitory interneurons. Starting from a self-predicting state – see main text and supplementary material (SM) – when a novel teaching (or associative) signal is presented at the output layer ( $\mathbf{u}_2^{\text{trgt}}$ ), a prediction error in the apical compartments of pyramidal neurons in the upstream layer (layer 1, ‘error’) is generated. This error appears as an apical voltage deflection that propagates down to the soma (purple arrow) where it modulates the somatic firing rate, which in turn leads to plasticity at bottom-up synapses (bottom, green). (B) Activity traces in the microcircuit before and after a new teaching signal is learned. (i) Before learning: a new teaching signal is presented ( $\mathbf{u}_2^{\text{trgt}}$ ), which triggers a mismatch between the top-down feedback (grey blue) and the cancellation given by the lateral interneurons (red). (ii) After learning (with plasticity at the bottom-up synapses ( $\mathbf{W}_{1,0}^{PP}$ )), the network successfully predicts the new teaching signal, reflected on no distal ‘error’ (top-down and lateral interneuron input cancel each other). (C) Interneurons learn to predict the backpropagated activity (i), while simultaneously silencing the apical compartment (ii), even though the pyramidal neurons remain active (not shown).

For simplicity, we reduce pyramidal output neurons to two-compartment cells: the apical compartment is absent ( $g_A = 0$  in Eq. 1) and basal voltages are as defined in Eq. 3. Although the design can be extended to more complex morphologies, in the framework of dendritic predictive plasticity two compartments suffice to compare desired target with actual prediction. Synapses proximal to the soma of output neurons provide direct external teaching input, incorporated as an additional source of current  $\mathbf{i}_N^P$ . In practice, one can simply set  $\mathbf{i}_N^P = g_{\text{som}}(\mathbf{u}_N^{\text{trgt}} - \mathbf{u}_N^P)$ , with some fixed somatic nudging conductance  $g_{\text{som}}$ . This can be modelled closer to biology by explicitly setting the somatic excitatory and inhibitory conductance-based inputs (Urbanczik and Senn, 2014). For a given output neuron,  $i_N^P(t) = g_{\text{exc},N}^P(t) (E_{\text{exc}} - u_N^P(t)) + g_{\text{inh},N}^P(t) (E_{\text{inh}} - u_N^P(t))$ , where  $E_{\text{exc}}$  and  $E_{\text{inh}}$  are excitatory and inhibitory synaptic reversal potentials, respectively, where the inputs are balanced according to

$g_{\text{exc},N}^{\text{P}} = g_{\text{som}} \frac{u_N^{\text{tgt}} - E_{\text{inh}}}{E_{\text{exc}} - E_{\text{inh}}}$ ,  $g_{\text{inh},N}^{\text{P}} = -g_{\text{som}} \frac{u_N^{\text{tgt}} - E_{\text{exc}}}{E_{\text{exc}} - E_{\text{inh}}}$ . The point at which no current flows,  $i_N^{\text{P}} = 0$ , defines the target teaching voltage  $u_N^{\text{tgt}}$  towards which the neuron is nudged<sup>4</sup>.

Interneurons are similarly modelled as two-compartment cells, cf. Eq. 2. Lateral dendritic projections from neighboring pyramidal neurons provide the main source of input as

$$\mathbf{v}_k^{\text{I}}(t) = \mathbf{W}_{k,k}^{\text{IP}} \phi(\mathbf{u}_k^{\text{P}}(t)), \quad (5)$$

whereas cross-layer, top-down synapses define the teaching current  $i_k^{\text{I}}$ . This means that an interneuron at layer  $k$  permanently (i.e., when learning or performing a task) receives balanced somatic teaching excitatory and inhibitory input from a pyramidal neuron at layer  $k+1$  on a one-to-one basis (as above, but with  $\mathbf{u}_{k+1}^{\text{P}}$  as target). With this setting, the interneuron is nudged to follow the corresponding next layer pyramidal neuron. See SM for detailed parameters.

## 2.2 Synaptic learning rules

The synaptic learning rules we use belong to the class of dendritic predictive plasticity rules (Urbanczik and Senn, 2014; Spicher et al., 2018) that can be expressed in its general form as

$$\frac{d}{dt}w = \eta (\phi(u) - \phi(v)) r, \quad (6)$$

where  $w$  is an individual synaptic weight,  $\eta$  is a learning rate,  $u$  and  $v$  denote distinct compartmental potentials,  $\phi$  is a rate function, and  $r$  is the presynaptic input. Eq. 6 was originally derived in the light of reducing the prediction error of somatic spiking, when  $u$  represents the somatic potential and  $v$  is a function of the postsynaptic dendritic potential.

In our model the plasticity rules for the various connection types are:

$$\frac{d}{dt} \mathbf{W}_{k,k-1}^{\text{PP}} = \eta_{k,k-1}^{\text{PP}} (\phi(\mathbf{u}_k^{\text{P}}) - \phi(\hat{\mathbf{v}}_{\text{B},k}^{\text{P}})) (\mathbf{r}_{k-1}^{\text{P}})^T, \quad (7)$$

$$\frac{d}{dt} \mathbf{W}_{k,k}^{\text{IP}} = \eta_{k,k}^{\text{IP}} (\phi(\mathbf{u}_k^{\text{I}}) - \phi(\hat{\mathbf{v}}_k^{\text{I}})) (\mathbf{r}_k^{\text{P}})^T, \quad (8)$$

$$\frac{d}{dt} \mathbf{W}_{k,k}^{\text{PI}} = \eta_{k,k}^{\text{PI}} (\mathbf{v}_{\text{rest}} - \mathbf{v}_{\text{A},k}^{\text{P}}) (\mathbf{r}_k^{\text{I}})^T, \quad (9)$$

where  $(\cdot)^T$  denotes vector transpose and  $\mathbf{r}_k \equiv \phi(\mathbf{u}_k)$  the layer  $k$  firing rates. The synaptic weights evolve according to the product of dendritic prediction error and presynaptic rate, and can undergo both potentiation or depression depending on the sign of the first factor (i.e., the prediction error).

For basal synapses, such prediction error factor amounts to a difference between postsynaptic rate and a local dendritic estimate which depends on the branch potential. In Eqs. 7 and 8,  $\hat{\mathbf{v}}_{\text{B},k}^{\text{P}} = \frac{g_{\text{B}}}{g_{\text{ik}} + g_{\text{B}} + g_{\text{A}}} \mathbf{v}_{\text{B},k}^{\text{P}}$  and  $\hat{\mathbf{v}}_k^{\text{I}} = \frac{g_{\text{D}}}{g_{\text{ik}} + g_{\text{D}}} \mathbf{v}_k^{\text{I}}$  take into account dendritic attenuation factors of the different compartments. On the other hand, the plasticity rule (9) of lateral interneuron-to-pyramidal synapses aims to silence (i.e., set to resting potential  $\mathbf{v}_{\text{rest}} = \mathbf{0}$ , here and throughout zero for simplicity) the apical compartment; this introduces an attractive state for learning where the contribution from interneurons balances (or cancels out) top-down dendritic input. This learning rule of apical-targeting interneuron synapses can be thought of as a dendritic variant of the homeostatic inhibitory plasticity proposed by Vogels et al. (2011); Luz and Shamir (2012).

In experiments where the top-down connections are plastic, the weights evolve according to

$$\frac{d}{dt} \mathbf{W}_{k,k+1}^{\text{PP}} = \eta_{k,k+1}^{\text{PP}} (\phi(\mathbf{u}_k^{\text{P}}) - \phi(\hat{\mathbf{v}}_{\text{TD},k}^{\text{P}})) (\mathbf{r}_{k+1}^{\text{P}})^T, \quad (10)$$

with  $\hat{\mathbf{v}}_{\text{TD},k}^{\text{P}} = \mathbf{W}_{k,k+1} \mathbf{r}_{k+1}^{\text{P}}$ . An implementation of this rule requires a subdivision of the apical compartment into a distal part receiving the top-down input (with voltage  $\hat{\mathbf{v}}_{\text{TD},k}^{\text{P}}$ ) and another distal compartment receiving the lateral input from the interneurons (with voltage  $\mathbf{v}_{\text{A},k}^{\text{P}}$ ).

<sup>4</sup>Note that in biology a target may be represented by an associative signal from the motor cortex to a sensory cortex (Attinger et al., 2017).

## 2.3 Comparison to previous work

It has been suggested that error backpropagation could be approximated by an algorithm that requires alternating between two learning phases, known as contrastive Hebbian learning (Ackley et al., 1985). This link between the two algorithms was first established for an unsupervised learning task (Hinton and McClelland, 1988) and later analyzed (Xie and Seung, 2003) and generalized to broader classes of models (O’Reilly, 1996; Scellier and Bengio, 2017).

The concept of apical dendrites as distinct integration zones, and the suggestion that this could simplify the implementation of backprop has been previously made (Körding and König, 2000, 2001). Our microcircuit design builds upon this view, offering a concrete mechanism that enables apical error encoding. In a similar spirit, two-phase learning recently reappeared in a study that exploits dendrites for deep learning with biological neurons (Guerguiev et al., 2017). In this more recent work, the temporal difference between the activity of the apical dendrite in the presence and in the absence of the teaching input represents the error that induces plasticity at the forward synapses. This difference is used directly for learning the bottom-up synapses without influencing the somatic activity of the pyramidal cell. In contrast, we postulate that the apical dendrite has an explicit error representation by simultaneously integrating top-down excitation and lateral inhibition. As a consequence, we do not need to postulate separate temporal phases, and our network operates continuously while plasticity at all synapses is always turned on.

Error minimization is an integral part of brain function according to predictive coding theories (Rao and Ballard, 1999; Friston, 2005). Interestingly, recent work has shown that backprop can be mapped onto a predictive coding network architecture (Whittington and Bogacz, 2017), related to the general framework introduced by LeCun (1988). A possible network implementation is suggested by Whittington and Bogacz (2017) that requires intricate circuitry with appropriately tuned error-representing neurons. According to this work, the only plastic synapses are those that connect prediction and error neurons. By contrast, in our model, lateral, bottom-up and top-down connections are all plastic, and errors are directly encoded in dendritic compartments.

## 3 Results

### 3.1 Learning in dendritic error networks approximates backprop

In our model, neurons implicitly carry and transmit errors across the network. In the supplementary material, we formally show such propagation of errors for networks in a particular regime, which we term *self-predicting*. Self-predicting nets are such that when no external target is provided to output layer neurons, the lateral input from interneurons cancels the internally generated top-down feedback and renders apical dendrites silent. In this case, the output becomes a feedforward function of the input, which can in theory be optimized by conventional backprop. We demonstrate that synaptic plasticity in self-predicting nets approximates the weight changes prescribed by backprop.

We summarize below the main points of the full analysis (see SM). First, we show that somatic membrane potentials at hidden layer  $k$  integrate feedforward predictions (encoded in basal dendritic potentials) with backpropagated errors (encoded in apical dendritic potentials):

$$\mathbf{u}_k^P = \mathbf{u}_k^- + \lambda^{N-k+1} \mathbf{W}_{k,k+1}^{PP} \left( \prod_{l=k+1}^{N-1} \mathbf{D}_l^- \mathbf{W}_{l,l+1}^{PP} \right) \mathbf{D}_N^- (\mathbf{u}_N^{\text{tgt}} - \mathbf{u}_N^-) + \mathcal{O}(\lambda^{N-k+2}).$$

Parameter  $\lambda \ll 1$  sets the strength of feedback and teaching versus bottom-up inputs and is assumed to be small to simplify the analysis. The first term is the basal contribution and corresponds to  $\mathbf{u}_k^-$ , the activation computed by a purely feedforward network that is obtained by removing lateral and top-down weights from the model (here and below, we use superscript ‘-’ to refer to the feedforward model). The second term (of order  $\lambda^{N-k+1}$ ) is an error that is backpropagated from the output layer down to  $k$ -th layer hidden neurons; matrix  $\mathbf{D}_k$  is a diagonal matrix with  $i$ -th entry containing the derivative of the neuronal transfer function evaluated at  $u_{k,i}^-$ .

Second, we compare model synaptic weight updates for the bottom-up connections to those prescribed by backprop. Output layer updates are exactly equal by construction. For hidden neuron synapses,

we obtain

$$\Delta \mathbf{W}_{k,k-1}^{\text{PP}} = \eta_{k,k-1}^{\text{PP}} \lambda^{N-k+1} \left( \prod_{l=k}^{N-1} \mathbf{D}_l^- \mathbf{W}_{l,l+1}^{\text{PP}} \right) \mathbf{D}_N^- (\mathbf{u}_N^{\text{trgt}} - \mathbf{u}_N^-) (\mathbf{r}_{k-1}^-)^T + \mathcal{O}(\lambda^{N-k+2}).$$

Up to a factor which can be absorbed in the learning rate, this plasticity rule becomes equal to the backprop weight change in the weak feedback limit  $\lambda \rightarrow 0$ , provided that the top-down weights are set to the transpose of the corresponding feedforward weights.

In our simulations, top-down weights are either set at random and kept fixed, in which case the equation above shows that the plasticity model optimizes the predictions according to an approximation of backprop known as feedback alignment (Lillicrap et al., 2016); or learned so as to minimize an inverse reconstruction loss, in which case the network implements a form of target propagation (Bengio, 2014; Lee et al., 2015).

### 3.2 Deviations from self-predictions encode backpropagated errors

To illustrate learning in the model and to confirm our analytical insights we first study a very simple task: memorizing a single input-output pattern association with only one hidden layer; the task naturally generalizes to multiple memories.

Given a self-predicting network (established by microcircuit plasticity, Fig. S1, see SM for more details), we focus on how prediction errors get propagated backwards when a novel teaching signal is provided to the output layer, modeled via the activation of additional somatic conductances in output pyramidal neurons. Here we consider a network model with an input, a hidden and an output layer (layers 0, 1 and 2, respectively; Fig. 1A).

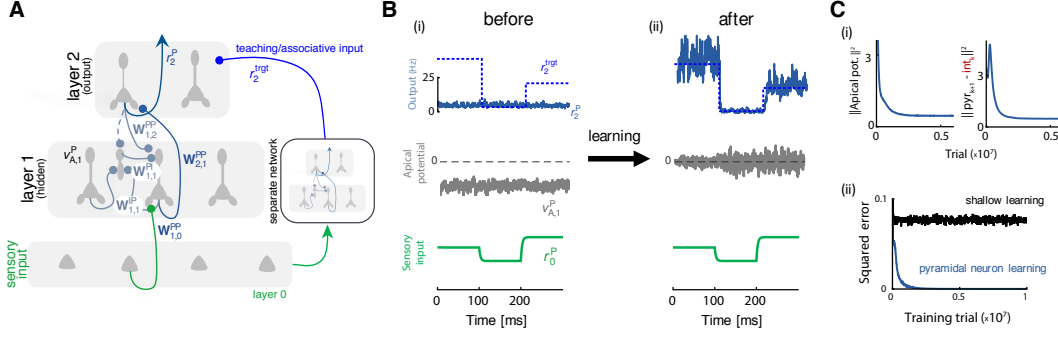
When the pyramidal cell activity in the output layer is nudged towards some desired target (Fig. 1B (i)), the bottom-up synapses  $\mathbf{W}_{2,1}^{\text{PP}}$  from the lower layer neurons to the basal dendrites are adapted, again according to the plasticity rule that implements the dendritic prediction of somatic spiking (see Eq. 7). What these synapses cannot explain away encodes a dendritic error in the pyramidal neurons of the lower layer 1. In fact, the self-predicting microcircuit can only cancel the feedback that is produced by the lower layer activity.

The somatic integration of apical activity induces plasticity at the bottom-up synapses  $\mathbf{W}_{1,0}^{\text{PP}}$  (Eq. 7). As the apical error changes the somatic activity, plasticity of the  $\mathbf{W}_{1,0}^{\text{PP}}$  weights tries to further reduce the error in the output layer. Importantly, the plasticity rule depends only on local information available at the synaptic level: postsynaptic firing and dendritic branch voltage, as well as the presynaptic activity, in par with phenomenological models of synaptic plasticity (Sjöström et al., 2001; Clopath et al., 2010; Bono and Clopath, 2017). This learning occur concurrently with modifications of lateral interneuron weights which track changes in the output layer. Through the course of learning the network comes to a point where the novel top-down input is successfully predicted (Fig. 1B,C).

### 3.3 Network learns to solve a nonlinear regression task

We now test the learning capabilities of the model on a nonlinear regression task, where the goal is to associate sensory input with the output of a separate multilayer network that transforms the same sensory input (Fig. 2A). More precisely, a pyramidal neuron network of dimensions 30-50-10 (and 10 hidden layer interneurons) learns to approximate a random nonlinear function implemented by a held-aside feedforward network of dimensions 30-20-10. One teaching example consists of a randomly drawn input pattern  $\mathbf{r}_0^{\text{P}}$  assigned to corresponding target  $\mathbf{r}_2^{\text{trgt}} = \phi(k_{2,1} \mathbf{W}_{2,1}^{\text{trgt}} \phi(k_{1,0} \mathbf{W}_{1,0}^{\text{trgt}} \mathbf{r}_0^{\text{P}}))$ , with scale factors  $k_{2,1} = 10$  and  $k_{1,0} = 2$ . Teacher network weights and input pattern entries are sampled from a uniform distribution  $U(-1, 1)$ . We used a soft rectifying nonlinearity as the neuronal transfer function,  $\phi(u) = \gamma \log(1 + \exp(\beta(u - \theta)))$ , with  $\gamma = 0.1$ ,  $\beta = 1$  and  $\theta = 3$ . This parameter setting led to neuronal activity in the nonlinear, sparse firing regime.

The network is initialized to a random initial synaptic weight configuration, with both pyramidal-pyramidal  $\mathbf{W}_{1,0}^{\text{PP}}$ ,  $\mathbf{W}_{2,1}^{\text{PP}}$ ,  $\mathbf{W}_{1,2}^{\text{PP}}$  and pyramidal-interneuron weights  $\mathbf{W}_{1,1}^{\text{IP}}$ ,  $\mathbf{W}_{1,1}^{\text{PI}}$  independently drawn from a uniform distribution. Top-down weight matrix  $\mathbf{W}_{1,2}^{\text{PP}}$  is kept fixed throughout, in the spirit of feedback alignment (Lillicrap et al., 2016). Output layer teaching currents  $\mathbf{i}_2^{\text{P}}$  are set so as to nudge  $\mathbf{u}_2^{\text{P}}$  towards the teacher-generated  $\mathbf{u}_2^{\text{trgt}}$ . Learning rates were manually chosen to yield best



**Figure 2: Dendritic error microcircuit learns to solve a nonlinear regression task online and without phases.** (A-C) Starting from a random initial weight configuration, a 30-50-10 fully-connected network learns to approximate a nonlinear function (‘separate network’) from input-output pattern pairs. (B) Example firing rates for a randomly chosen output neuron ( $r_2^P$ , blue noisy trace) and its desired target imposed by the associative input ( $r_2^{\text{trgt}}$ , blue dashed line), together with the voltage in the apical compartment of a hidden neuron ( $v_{A,1}^P$ , grey noisy trace) and the input rate from the sensory neuron ( $r_0^P$ , green). Traces are shown before (i) and after learning (ii). (C) Error curves for the full model and a shallow model for comparison.

performance. Some learning rate tuning was required to ensure the microcircuit could track the changes in the bottom-up pyramidal-pyramidal weights, but we did not observe high sensitivity once the correct parameter regime was identified. Error curves are exponential moving averages of the sum of squared errors loss  $\|\mathbf{r}_2^P - \mathbf{r}_2^{\text{trgt}}\|^2$  computed after every example on unseen input patterns. Test error performance is measured in a noise-free setting ( $\sigma = 0$ ). Plasticity induction terms given by Eqs. 7-9 are low-pass filtered with time constant  $\tau_w$  before being definitely consolidated, to dampen fluctuations; synaptic plasticity is kept on throughout. Plasticity and neuron model parameters are as defined above.

We let learning occur in continuous time without pauses or alternations in plasticity as input patterns are sequentially presented. This is in contrast to previous learning models that rely on computing activity differences over distinct phases, requiring temporally nonlocal computation, or globally coordinated plasticity rule switches (Hinton and McClelland, 1988; O’Reilly, 1996; Xie and Seung, 2003; Scellier and Bengio, 2017; Guerguiev et al., 2017). Furthermore, we relaxed the bottom-up vs. top-down weight symmetry imposed by backprop and kept the top-down weights  $\mathbf{W}_{1,2}^{\text{PP}}$  fixed.

Forward  $\mathbf{W}_{1,2}^{\text{PP}}$  weights quickly aligned to  $\sim 45^\circ$  of the feedback weights ( $\mathbf{W}_{2,1}^{\text{PP}})^T$  (see Fig. S1), in line with the recently discovered feedback alignment phenomenon (Lillicrap et al., 2016). This simplifies the architecture, because top-down and interneuron-to-pyramidal synapses need not be changed. We set the scale of the top-down weights, apical and somatic conductances such that feedback and teaching inputs were strong, to test the model outside the weak feedback regime ( $\lambda \rightarrow 0$ ) for which our SM theory was developed. Finally, to test robustness, we injected a weak noise current to every neuron.

Our network was able to learn this harder task (Fig. 2B), performing considerably better than a shallow learner where only hidden-to-output weights were adjusted (Fig. 2C). Useful changes were thus made to hidden layer bottom-up weights. The self-predicting network state emerged throughout learning from a random initial configuration (see SM; Fig. S1).

### 3.4 Microcircuit network learns to classify handwritten digits

Next, we turn to the problem of classifying MNIST handwritten digits. We wondered how our model would fare in this benchmark, in particular whether the prediction errors computed by the interneuron microcircuit would allow learning the weights of a hierarchical nonlinear network with multiple hidden layers. To that end, we trained a deeper, larger 4-layer network (with 784-500-500-10 pyramidal neurons, Fig. 3A) by pairing digit images with teaching inputs that nudged the 10 output neurons towards the correct class pattern. We initialized the network to a random but self-predicting

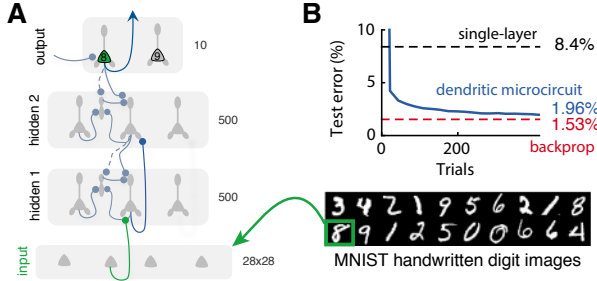


Figure 3: **Dendritic error networks learn to classify handwritten digits.** (A) A network with two hidden layers learns to classify handwritten digits from the MNIST data set. (B) Classification error achieved on the MNIST testing set (blue; cf. shallow learner (black) and standard backprop<sup>6</sup>(red)).

configuration where interneurons cancelled top-down inputs, rendering the apical compartments silent before training started. Top-down and interneuron-to-pyramidal weights were kept fixed.

Here for computational efficiency we used a simplified network dynamics where the compartmental potentials are updated only in two steps before applying synaptic changes. In particular, for each presented MNIST image, both pyramidal and interneurons are first initialized to their bottom-up prediction state (3),  $\mathbf{u}_k = \mathbf{v}_{B,k}$ , starting from layer 1 up to the top layer  $N$ . Output layer neurons are then nudged towards their desired target  $\mathbf{u}_N^{\text{trgt}}$ , yielding updated somatic potentials  $\mathbf{u}_N^{\text{P}} = (1 - \lambda_N) \mathbf{v}_{B,N} + \lambda_N \mathbf{u}_N^{\text{trgt}}$ . To obtain the remaining final compartmental potentials, the network is visited in reverse order, proceeding from layer  $k = N - 1$  down to  $k = 1$ . For each  $k$ , interneurons are first updated to include top-down teaching signals,  $\mathbf{u}_k^{\text{I}} = (1 - \lambda_I) \mathbf{v}_k^{\text{I}} + \lambda_I \mathbf{u}_{k+1}^{\text{P}}$ ; this yields apical compartment potentials according to (4), after which we update hidden layer somatic potentials as a convex combination with mixing factor  $\lambda_k$ . The convex combination factors introduced above are directly related to neuron model parameters as conductance ratios. Synaptic weights are then updated according to Eqs. 7-10. Such simplified dynamics approximates the full recurrent network relaxation in the deterministic setting  $\sigma \rightarrow 0$ , with the approximation improving as the top-down dendritic coupling is decreased,  $g_A \rightarrow 0$ .

We train the models on the standard MNIST handwritten image database, further splitting the training set into 55000 training and 5000 validation examples. The reported test error curves are computed on the 10000 held-aside test images. The four-layer network shown in Fig. 3 is initialized in a self-predicting state with appropriately scaled initial weight matrices. For our MNIST networks, we used relatively weak feedback weights, apical and somatic conductances (see SM) to justify our simplified approximate dynamics described above, although we found that performance did not appreciably degrade with larger values. To speed-up training we use a mini-batch strategy on every learning rule, whereby weight changes are averaged across 10 images before being applied. We take the neuronal transfer function  $\phi$  to be a logistic function,  $\phi(u) = 1/(1 + \exp(-u))$  and include a learnable threshold on each neuron, modelled as an additional input fixed at unity with a plastic weight. Desired target class vectors are 1-hot coded, with  $r_N^{\text{trgt}} \in \{0.1, 0.8\}$ . During testing, the output is determined by picking the class label corresponding to the neuron with highest firing rate. We found the model to be relatively robust to learning rate tuning on the MNIST task, except for the rescaling by the inverse mixing factor to compensate for teaching signal dilution (see SM for the exact parameters).

The network was able to achieve a test error of 1.96%, Fig. 3B, a figure not overly far from the reference mark of non-convolutional artificial neural networks optimized with backprop (1.53%) and comparable to recently published results that lie within the range 1.6-2.4% (Lee et al., 2015; Lillicrap et al., 2016; Nøkland, 2016). The performance of our model also compares favorably to the 3.2% test error reported by Guerguiev et al. (2017) for a two-hidden-layer network. This was possible despite the asymmetry of forward and top-down weights and at odds with exact backprop, thanks to a feedback alignment dynamics. Apical compartment voltages remained approximately silent when output nudging was turned off (data not shown), reflecting the maintenance of a self-predicting state throughout learning, which enabled the propagation of errors through the network. To further demonstrate that the microcircuit was able to propagate errors to deeper hidden layers, and that the task was not being solved by making useful changes only to the weights onto the topmost hidden layer, we re-ran the experiment while keeping fixed the pyramidal-pyramidal weights connecting the two hidden layers. The network still learned the dataset and achieved a test error of 2.11%.



As top-down weights are likely plastic in cortex, we also trained a one-hidden-layer (784-1000-10) network where top-down weights were learned on a slow time-scale according to learning rule (10). This inverse learning scheme is closely related to target propagation (Bengio, 2014; Lee et al., 2015). Such learning could play a role in perceptual denoising, pattern completion and disambiguation, and boost alignment beyond that achieved by pure feedback alignment (Bengio, 2014). Starting from random initial conditions and keeping all weights plastic (bottom-up, lateral and top-down) throughout, our network achieved a test classification performance of 2.48% on MNIST. Once more, useful changes were made to hidden synapses, even though the microcircuit had to track changes in both the bottom-up and the top-down pathways.

## 4 Conclusions

Our work makes several predictions across different levels of investigation. Here we briefly highlight some of these predictions and related experimental observations. The most fundamental feature of the model is that distal dendrites encode error signals that instruct learning of lateral and bottom-up connections. While monitoring such dendritic signals during learning is challenging, recent experimental evidence suggests that prediction errors in mouse visual cortex arise from a failure to locally inhibit motor feedback (Zmarz and Keller, 2016; Attinger et al., 2017), consistent with our model. Interestingly, the plasticity rule for apical dendritic inhibition, which is central to error encoding in the model, received support from another recent experimental study (Chiu et al., 2018).

A further implication of our model is that prediction errors occurring at a higher-order cortical area would imply also prediction errors co-occurring at earlier areas. Recent experimental observations in the macaque face-processing hierarchy support this (Schwiedrzik and Freiwald, 2017).

Here we have focused on the role of a specific interneuron type (SST) as a feedback-specific interneuron. There are many more interneuron types that we do not consider in our framework. One such type are the PV (parvalbumin-positive) cells, which have been postulated to mediate a somatic excitation-inhibition balance (Vogels et al., 2011; Froemke, 2015) and competition (Masquelier and Thorpe, 2007; Nessler et al., 2013). These functions could in principle be combined with our framework in that PV interneurons may be involved in representing another type of prediction error (e.g., generative errors).

Humans have the ability to perform fast (e.g., one-shot) learning, whereas neural networks trained by backpropagation of error (or approximations thereof, like ours) require iterating over many training examples to learn. This is an important open problem that stands in the way of understanding the neuronal basis of intelligence. One possibility where our model naturally fits is to consider multiple subsystems (for example, the neocortex and the hippocampus) that transfer knowledge to each other and learn at different rates (McClelland et al., 1995; Kumaran et al., 2016).

Overall, our work provides a new view on how the brain may solve the credit assignment problem for time-continuous input streams by approximating the backpropagation algorithm, and bringing together many puzzling features of cortical microcircuits.

## Acknowledgements

The authors would like to thank Timothy P. Lillicrap, Blake Richards, Benjamin Scellier and Mihai A. Petrovici for helpful discussions. WS thanks Matthew Larkum for many inspiring discussions on dendritic processing. JS thanks Elena Kreutzer, Pascal Leimer and Martin T. Wiechert for valuable feedback and critical reading of the manuscript.

This work has been supported by the Swiss National Science Foundation (grant 310030L-156863 of WS), the European Union’s Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 785907 (Human Brain Project), NSERC, CIFAR, and Canada Research Chairs.

## References

- Ackley, D. H., Hinton, G. E., and Sejnowski, T. J. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, 9(1):147–169.
- Attinger, A., Wang, B., and Keller, G. B. (2017). Visuomotor coupling shapes the functional development of mouse visual cortex. *Cell*, 169(7):1291–1302.e14.
- Bengio, Y. (2014). How auto-encoders could provide credit assignment in deep networks via target propagation. *arXiv:1407.7906*
- Bono, J. and Clopath, C. (2017). Modeling somatic and dendritic spike mediated plasticity at the single neuron and network level. *Nature Communications*, 8(1):706.
- Bottou, L. (1998). Online algorithms and stochastic approximations. In Saad, D., editor, *Online Learning and Neural Networks*. Cambridge University Press, Cambridge, UK.
- Chiu, C. Q., Martenson, J. S., Yamazaki, M., Natsume, R., Sakimura, K., Tomita, S., Tavalin, S. J., and Higley, M. J. (2018). Input-specific nmdar-dependent potentiation of dendritic gabaergic inhibition. *Neuron*, 97(2):368–377.
- Clopath, C., Büsing, L., Vasilaki, E., and Gerstner, W. (2010). Connectivity reflects coding: a model of voltage-based stdp with homeostasis. *Nature Neuroscience*, 13(3):344–352.
- Costa, R. P., Assael, Y. M., Shillingford, B., de Freitas, N., and Vogels, T. P. (2017). Cortical microcircuits as gated-recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 271–282.
- Crick, F. (1989). The recent excitement about neural networks. *Nature*, 337:129–132.
- Dornn, A. L., Yuan, K., Barker, A. J., Schreiner, C. E., and Froemke, R. C. (2010). Developmental sensory experience balances cortical excitation and inhibition. *Nature*, 465(7300):932–936.
- Friedrich, J., Urbanczik, R., and Senn, W. (2011). Spatio-temporal credit assignment in neuronal population learning. *PLOS Computational Biology*, 7(6):e1002092.
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 360(1456):815–836.
- Froemke, R. C. (2015). Plasticity of cortical excitatory-inhibitory balance. *Annual Review of Neuroscience*, 38(1):195–219.
- Fu, Y., Kaneko, M., Tang, Y., Alvarez-Buylla, A., and Stryker, M. P. (2015). A cortical disinhibitory circuit for enhancing adult plasticity. *eLife*, 4:e05558.
- Grossberg, S. (1987). Competitive learning: From interactive activation to adaptive resonance. *Cognitive Science*, 11(1):23–63.
- Guerguiev, J., Lillicrap, T. P., and Richards, B. A. (2017). Towards deep learning with segregated dendrites. *eLife*, 6:e22901.
- Hinton, G. E. and McClelland, J. L. (1988). Learning representations by recirculation. In Anderson, D. Z., editor, *Neural Information Processing Systems*, pages 358–366. American Institute of Physics.
- Kell, A. J., Yamins, D. L., Shook, E. N., Norman-Haignere, S. V., and McDermott, J. H. (2018). A task-optimized neural network replicates human auditory behavior, predicts brain responses, and reveals a cortical processing hierarchy. *Neuron*.
- Khaligh-Razavi, S.-M. and Kriegeskorte, N. (2014). Deep supervised, but not unsupervised, models may explain it cortical representation. *PLOS Computational Biology*, 10(11):1–29.
- Körding, K. P. and König, P. (2000). Learning with two sites of synaptic integration. *Network: Comput. Neural Syst.*, 11:1–15.
- Körding, K. P. and König, P. (2001). Supervised and unsupervised learning with two sites of synaptic integration. *Journal of Computational Neuroscience*, 11:207–215.
- Kumaran, D., Hassabis, D., and McClelland, J. L. (2016). What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in Cognitive Sciences*, 20(7):512 – 534.

- Larkum, M. (2013). A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. *Trends in Neurosciences*, 36(3):141–151.
- LeCun, Y. (1988). A theoretical framework for back-propagation. In Touretzky, D., Hinton, G., and Sejnowski, T., editors, *Proceedings of the 1988 Connectionist Models Summer School*, pages 21–28. Morgan Kaufmann, Pittsburg, PA.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436–444.
- Lee, D.-H., Zhang, S., Fischer, A., and Bengio, Y. (2015). Difference target propagation. In *Machine Learning and Knowledge Discovery in Databases*, pages 498–515. Springer.
- Leinweber, M., Ward, D. R., Sobczak, J. M., Attinger, A., and Keller, G. B. (2017). A Sensorimotor Circuit in Mouse Cortex for Visual Flow Predictions. *Neuron*, 95(6):1420–1432.e5.
- Lillicrap, T. P., Cownden, D., Tweed, D. B., and Akerman, C. J. (2016). Random synaptic feedback weights support error backpropagation for deep learning. *Nature Communications*, 7:13276.
- Lillicrap, T. P. and Scott, S. H. (2013). Preference distributions of primary motor cortex neurons reflect control solutions optimized for limb biomechanics. *Neuron*, 77(1):168–179.
- Luz, Y. and Shamir, M. (2012). Balancing feed-forward excitation and inhibition via Hebbian inhibitory synaptic plasticity. *PLOS Computational Biology*, 8(1):e1002334.
- Makino, H. and Komiyama, T. (2015). Learning enhances the relative impact of top-down processing in the visual cortex. *Nature Neuroscience*, 18(8):1116–1122.
- Manita, S., Suzuki, T., Homma, C., Matsumoto, T., Odagawa, M., Yamada, K., Ota, K., Matsubara, C., Inutsuka, A., Sato, M., et al. (2015). A top-down cortical circuit for accurate sensory perception. *Neuron*, 86(5):1304–1316.
- Marblestone, A. H., Wayne, G., and Kording, K. P. (2016). Toward an integration of deep learning and neuroscience. *Frontiers in Computational Neuroscience*, 10:94.
- Masquelier, T. and Thorpe, S. (2007). Unsupervised learning of visual features through spike timing dependent plasticity. *PLOS Computational Biology*, 3.
- McClelland, J. L., McNaughton, B. L., and O’reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3):419.
- Nessler, B., Pfeiffer, M., Buesing, L., and Maass, W. (2013). Bayesian computation emerges in generic cortical microcircuits through spike-timing-dependent plasticity. *PLOS Computational Biology*, 9(4):e1003037.
- Nøkland, A. (2016). Direct feedback alignment provides learning in deep neural networks. In *Advances in Neural Information Processing Systems*, pages 1037–1045.
- O’Reilly, R. C. (1996). Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation*, 8(5):895–938.
- Pakan, J. M., Lowe, S. C., Dylida, E., Keemink, S. W., Currie, S. P., Coutts, C. A., Rochefort, N. L., and Mrsic-Flogel, T. D. (2016). Behavioral-state modulation of inhibition is context-dependent and cell type specific in mouse visual cortex. *eLife*, 5:e14985.
- Petreaanu, L., Gutnisky, D. A., Huber, D., Xu, N.-I., O’Connor, D. H., Tian, L., Looger, L., and Svoboda, K. (2012). Activity in motor-sensory projections reveals distributed coding in somatosensation. *Nature*, 489(7415):299–303.
- Petreaanu, L., Mao, T., Sternson, S. M., and Svoboda, K. (2009). The subcellular organization of neocortical excitatory connections. *Nature*, 457(7233):1142–1145.
- Poort, J., Khan, A. G., Pachitariu, M., Nemri, A., Orsolich, I., Krupic, J., Bauza, M., Sahani, M., Keller, G. B., Mrsic-Flogel, T. D., and Hofer, S. B. (2015). Learning enhances sensory and multiple non-sensory representations in primary visual cortex. *Neuron*, 86(6):1478–1490.
- Rao, R. P. and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87.
- Roelfsema, P. R. and Holtmaat, A. (2018). Control of synaptic plasticity in deep cortical networks. *Nature Reviews Neuroscience*, 19(3):166.

- Roelfsema, P. R. and van Ooyen, A. (2005). Attention-gated reinforcement learning of internal representations for classification. *Neural Computation*, 17(10):2176–2214.
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323:533–536.
- Scellier, B. and Bengio, Y. (2017). Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. *Frontiers in Computational Neuroscience*, 11:24.
- Schwiedrzik, C. M. and Freiwald, W. A. (2017). High-level prediction signals in a low-level area of the macaque face-processing hierarchy. *Neuron*, 96(1):89–97.e4.
- Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. (2001). Rate, Timing, and Cooperativity Jointly Determine Cortical Synaptic Plasticity. *Neuron*, 32(6):1149–1164.
- Spicher, D., Clopath, C., and Senn, W. (2018). Predictive plasticity in dendrites: from a computational principle to experimental data (in preparation).
- Spruston, N. (2008). Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews Neuroscience*, 9(3):206–221.
- Sutton, R. S. and Barto, A. G. (1998). *Reinforcement learning: An introduction*, volume 1. MIT Press, Cambridge, Mass.
- Urban-Ciecko, J. and Barth, A. L. (2016). Somatostatin-expressing neurons in cortical networks. *Nature Reviews Neuroscience*, 17(7):401–409.
- Urbanczik, R. and Senn, W. (2014). Learning by the dendritic prediction of somatic spiking. *Neuron*, 81(3):521–528.
- Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science*, 334(6062):1569–1573.
- Whittington, J. C. R. and Bogacz, R. (2017). An approximation of the error backpropagation algorithm in a predictive coding network with local Hebbian synaptic plasticity. *Neural Computation*, 29(5):1229–1262.
- Xie, X. and Seung, H. S. (2003). Equivalence of backpropagation and contrastive Hebbian learning in a layered network. *Neural Computation*, 15(2):441–454.
- Yamins, D. L. and DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19(3):356–365.
- Yamins, D. L., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., and DiCarlo, J. J. (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences*, 111(23):8619–8624.
- Zmarz, P. and Keller, G. B. (2016). Mismatch receptive fields in mouse visual cortex. *Neuron*, 92(4):766–772.