Synaptic plasticity in distal dendrites can optimize probabilistic top-down predictions in pyramidal cells

Arjun Rao[*,1], Robert Legenstein[*,1], Anand Subramoney[1,2], and Wolfgang Maass[1,+]

[*]First authors

[1]Institute for Theoretical Computer Science, Graz University of Technology, 8010, Austria
[2]Institute for Neural Computation, Ruhr University Bochum, Germany
[+]Corresponding author: maass@igi.tugraz.at

March 4, 2021

Abstract

Predictive coding has been identified as a major driver of computation and learning in cortical microcircuits. But it has remained unknown which synaptic plasticity processes install and maintain predictive coding capability. Predictions are inherently uncertain, and learning rules that aim at discriminating linearly separable classes of inputs – such as the perceptron learning rule – are not suitable for learning to predict. We show that experimental data on synaptic plasticity in distal dendrites of pyramidal cells support a functionally more powerful learning rule that produces a spike-based approximation to logistic regression, i.e., to probabilistic prediction. We also show that experimentally found interactions between different dendritic branches support learning of predictions for more complex probability distributions. The resulting learning theory for top-down inputs to pyramidal cells provides a normative framework for evaluating experimental data, and suggests further experiments for tracking the emergence of predictive coding through synaptic plasticity in distal dendrites.

1 Introduction

A hallmark of brain intelligence is the capability of the brain to use internal models of the world, predictions, and internal goals that are maintained in higher brain areas to shape the processing of sensory input streams in lower brain areas. This ubiquituous integration of context into computations on sensory inputs distinguishes brain computations from most other computational paradigms, and is well supported by stereotypical laminar cortical microcircuits: Inputs from higher brain areas typically arrive in pyramidal cells of lower brain areas through synaptic connections in distal dendrites that are located in the upper layers of the laminar microcircuit. Hence in order to understand how context information from higher brain areas shapes processing of sensory stimuli by these pyramidal cells in lower brain areas, we need to understand the plasticity mechanisms in distal dendrites of these pyramidal cells that regulate the impact which neurons in higher brain areas have on their activity. One intriguing hypothesis is that pyramidal cells learn to predict bottom-up sensory inputs on the basis of information from other brain areas, including motor areas [Leinweber et al., 2017], that reaches these neurons through synapses on their apical dendrites

[Larkum, 2013] [Doron et al., 2020]. This hypothesis implies that plasticity of synapses at distal dendrites plays a key role for the emergence of such predictive coding. Interestingly, a number of experimental data suggest that plasticity processes in distal dendrites differ in essential aspects from those that implement correlative Hebbian or spike-timing-dependent plasticity (STDP)-like rules for synaptic plasticity at other synaptic connections to pyramidal cells; see for example the recent review on higher order thalamus [Aru et al., 2020], especially their Fig. 1c and 3. In particular two differences have been emphasized: Firstly, the firing of the postsynaptic neurons, i.e., its somatic spikes, do not appear to play an essential role in the induction of this synaptic plasticity [Kampa et al., 2006, Kampa et al., 2007, Gambino et al., 2014]. Secondly, instead of the somatic spikes, plasticity of these synapses is gated by instructive signals [Magee and Grienberger, 2020], such as Ca^{2+}-spikes [Kampa et al., 2006, Kampa et al., 2007], NMDAR spikes [Gambino et al., 2014], more general forms of plateau potentials within the dendritic arbor of pyramidal cells in area CA1 [Bittner et al., 2015, Bittner et al., 2017, Magee and Grienberger, 2020], or synaptic input from perirhinal cortex to its apical dendrites [Doron et al., 2020].

In the case of Ca^{2+}-spikes as instructive signals for synaptic plasticity, these Ca^{2+}-spikes are typically generated if synaptic input to apical dendrites roughly coincides in time with somatic spikes of the pyramidal cell in layer 5. Here, somatic spikes may be caused by bottom-up sensory inputs to synapses to the neuron on layers below the dendritic tuft [Larkum, 2013]. Hence in this case, prediction of a Ca^{2+}-spike becomes equivalent to predicting coincidence of a depolarization of a distal dendrite with depolarization caused by sensory inputs. These gating signals for synaptic plasticity have the common property that they are longer lasting events than postsynaptic spikes, extending over at least 50 ms. Hence we will refer to them as instructive signals in line with the terminology of [Magee and Grienberger, 2020].

However, a theoretical framework for these types of plasticity rules is missing, and the goal of this article is to establish such a framework. In fact, we also propose a normative model for synaptic plasticity at distal dendrites. This normative model is based on the commonly conjectured learning goal for this synaptic plasticity, which differs from that for plasticity of other synaptic connections. If one considers for example the prediction of sensory inputs as one prominent learning goal for synaptic plasticity in distal dendrites, one needs to take into account that learning to predict is inherently different from learning to classify sensory inputs. In particular, prediction is inherently uncertain and cannot be learned with learning rules such as the perceptron learning rule that are geared for learning classification of linearly separable classes. Furthermore, predictions are most useful in a computational context if they are accompanied with estimates of their uncertainty. For that one has to resort to models for learning probabilistic prediction, rather than classification.

We will first present a theoretical framework for learning probabilistic predictions through synaptic plasticity in a single dendritic branch in the dendritic tuft of a pyramidal cell, ignoring its interaction with other dendritic branches. We then extend the framework to the case that there are several dendritic branch in the apical tuft. We consider the consequences of synaptic plasticity for two types of interactions between dendritic branches that have been reported, a saturating summation [Naud et al., 2014, Poirazi et al., 2003], as well as a soft XOR for pyramidal cells on layer 2/3 of the human cortex [Gidon et al., 2020]. We include illustrative examples for computational tasks that can be learned with the proposed model for synaptic plasticity in distal dendrites, both for the case of a single branch, and for the two types of interaction between several branches. For simplicity, we will refer to those dendritic branches in the dendritic tuft of pyramidal cells in layer 2/3 or 5 to which our learning theory applies just as apical dendrites.

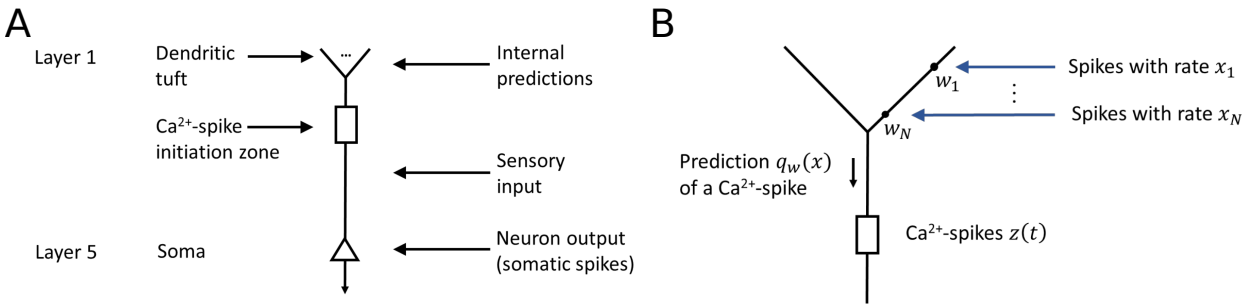

Figure 1: Scheme of a pyramidal cell with apical dendrites in its dendritic tuft, and learning model for synaptic input to this dendritic tuft. A) Generic model of a pyramidal cell in layer 5 of the stereotypical laminar cortical microcircuit, integrating two different inputs streams, see [Larkum, 2013] for details. A Ca^{2+}-spike, which also propagates backwards into the dendritic tuft, is generated at the Ca^{2+}-spike initiation zone when sufficiently strong depolarization arises from both input streams. This Ca^{2+}-spike serves as instructive signal in the sense of the plasticity model of [Magee and Grienberger, 2020] **B)** Illustration of the notation in our learning model (simplest case of a single dendritic branch). Synaptic input from neurons in higher brain areas with firing rates $x_1, \ldots, x_N$ is integrated with synaptic weights $w_1, \ldots . w_N$ in the dendritic branch. We present a theoretical framework for the plasticity of these weights, that leads to a depolarization $q_{\mathbf{w}}(\mathbf{x})$ that is viewed as output of this dendritic branch.

2 Results

2.1 A theoretically tractable model for prediction learning through synaptic plasticity in apical dendrites

We are considering a model for synaptic plasticity in apical dendrites that aims at being consistent with the main body of experimental data, and simultaneously at being as simple as possible so that it can be analyzed theoretically. Our model is closely related to the simple model for computations in apical dendrites of layer 5 pyramidal cells (L5Ps) proposed in [Larkum, 2013], see Fig. 1A. We have added to that a simple model for synaptic plasticity in these dendrites, see Fig. 1B.

The vector $\mathbf{x} = [x_1, x_2, \ldots, x_N]^T$ represents the current firing rates of N neurons in higher brain areas that provide synaptic input to this dendritic branch. We write $\mathbf{w} = [w_1, w_2, \ldots, w_N]^T$ for the synaptic weights of these synapses. We use the variable z to indicate the presence of a gating signal for synaptic plasticity, with $z(t) = 1$ if synaptic plasticity is enabled at time t, else $z(t) = 0$. This gating signal is referred to as instructive signal in [Magee and Grienberger, 2020].

Since the analysis of synaptic plasticity in apical dendrites of these neurons [Kampa et al., 2006, Kampa et al., 2007] focused on the Ca^{2+}-spike as a gating factor for synaptic plasticity in L5Ps, and since L5Ps are used as a running example for our analysis, we will refer in the following to this gating factor as a Ca^{2+}-spike. But our analysis also holds if synaptic plasticity in apical dendrites is gated by a different signal, such as input from perirhinal cortex according to [Doron et al., 2020], or a plateau potential in pyramidal cells of area CA1 that may emerge stochastically or triggered by attention or salience signals [Bittner et al., 2015, Bittner et al., 2017, Magee and Grienberger, 2020]. A common property of all these experimentally reported gating signals is that they last for 50 ms or longer.

We assume that the goal of synaptic plasticity of the weights $\mathbf{w}$ of synapses on the apical dendrite is the prediction of the gating signal z. Thus if this gating signal is triggered for example in pyramidal cells in layer 5 when the synaptic input from higher brain areas to the apical dendrites coincides with

Gating signal z	weight update $\Delta\mathbf{w}$
0	$-\ q_{\mathbf{w}}(\mathbf{x})\ x_j$
1	$\mathbf{x}\ -\ q_{\mathbf{w}}(\mathbf{x})\ x_j$

Table 1: Updates for synaptic weight w_j according to DLR: The firing rates of presynaptic neurons in higher brain areas are denoted by $\mathbf{x}$, and z is the gating signal. $q_{\mathbf{w}}(\mathbf{x})$ is a function of the local dendritic membrane potential defined in equation 2.

sensory inputs, this learning goal amounts to predicting coinciding sensory input to the neuron, see [Larkum, 2013]. Furthermore we assume that the apical dendrites produce this prediction in the form of the depolarization $q_{\mathbf{w}}(\mathbf{x})$ of the membrane potential that results from the current synaptic inputs $\mathbf{x}$ with weights $\mathbf{w}$ onto these dendrites. If one denotes the joint distribution of the input to the considered dendritic branch by $\mathbf{x}$ and the Ca^{2+}-spike z by $p_{\text{target}}(\mathbf{x}, z)$, then the optimal value of the model estimate $q_{\mathbf{w}}(\mathbf{x})$ is the conditional probability $p_{\text{target}}(z|\mathbf{x})$. An important computational role of the resulting contribution $q_{\mathbf{w}}(\mathbf{x})$ of distal dendrites to the processing of sensory information by a pyramidal cell in layer 5 was elucidated in [Takahashi et al., 2016] and [Takahashi et al., 2020]: it shifts the perceptual threshold.

Since Ca^{2+}-spikes and the other gating signals that we are considering have a duration of 50 ms or more, we are looking for a synaptic plasticity rule which is consistent with this time scale, i.e., with a time scale that is slower than that of single spikes and EPSPs. In particular, we focus on a realization of the prediction $q_{\mathbf{w}}(\mathbf{x})$ on this slower time scale: through the average membrane potential $\bar{u}$ that results from the firing rates $\mathbf{x} = [x_1, x_2, \ldots, x_N]^T$ of neurons in higher brain areas that have formed synaptic connections on the considered apical dendrite. We approximate the resulting average membrane potential in this dendrite by:

$$\bar{u} = \sum_{j=1}^{N} x_j w_j = \mathbf{x}^T \mathbf{w}. \tag{1}$$

It turns out that the resulting plasticity rule assumes a particularly simple form if one uses the following functional form for $q_{\mathbf{w}}(\mathbf{x})$, which also takes into account a saturation effect:

$$q_{\mathbf{w}}(\mathbf{x}) = \sigma\left(\beta\left(\bar{u} - u_0\right)\right), \tag{2}$$

where $\sigma(x) = \frac{1}{1+e^{-x}}$ is the standard sigmoid function, u_0 is a constant threshold, and β is a constant scaling factor.

In order to optimize $q_{\mathbf{w}}(\mathbf{x})$ as approximation of $p_{\text{target}}(z|\mathbf{x})$, we use the cross entropy loss as a standard measure of the distance between the two distributions. A useful property of this loss function is that it is convex. Gradient descent for this loss function yields the following normative rule for synaptic plasticity of the weights $\mathbf{w}$ of synapses onto the apical dendrite:

$$\Delta w_j \ = \ \eta\left(z - q_{\mathbf{w}}(\mathbf{x})\right) x_j \ , \tag{3}$$

where Δw_j is the change of the weight of synapse j and $\eta > 0$ is a constant scaling factor. We show in Section 3.1, that this weight update performs stochastic gradient descent for the previously described cross entropy loss. Because of the convexity of this loss function, this update is guaranteed to converge to its global optimum. Hence this plasticity rule is in a well-defined sense optimal, and

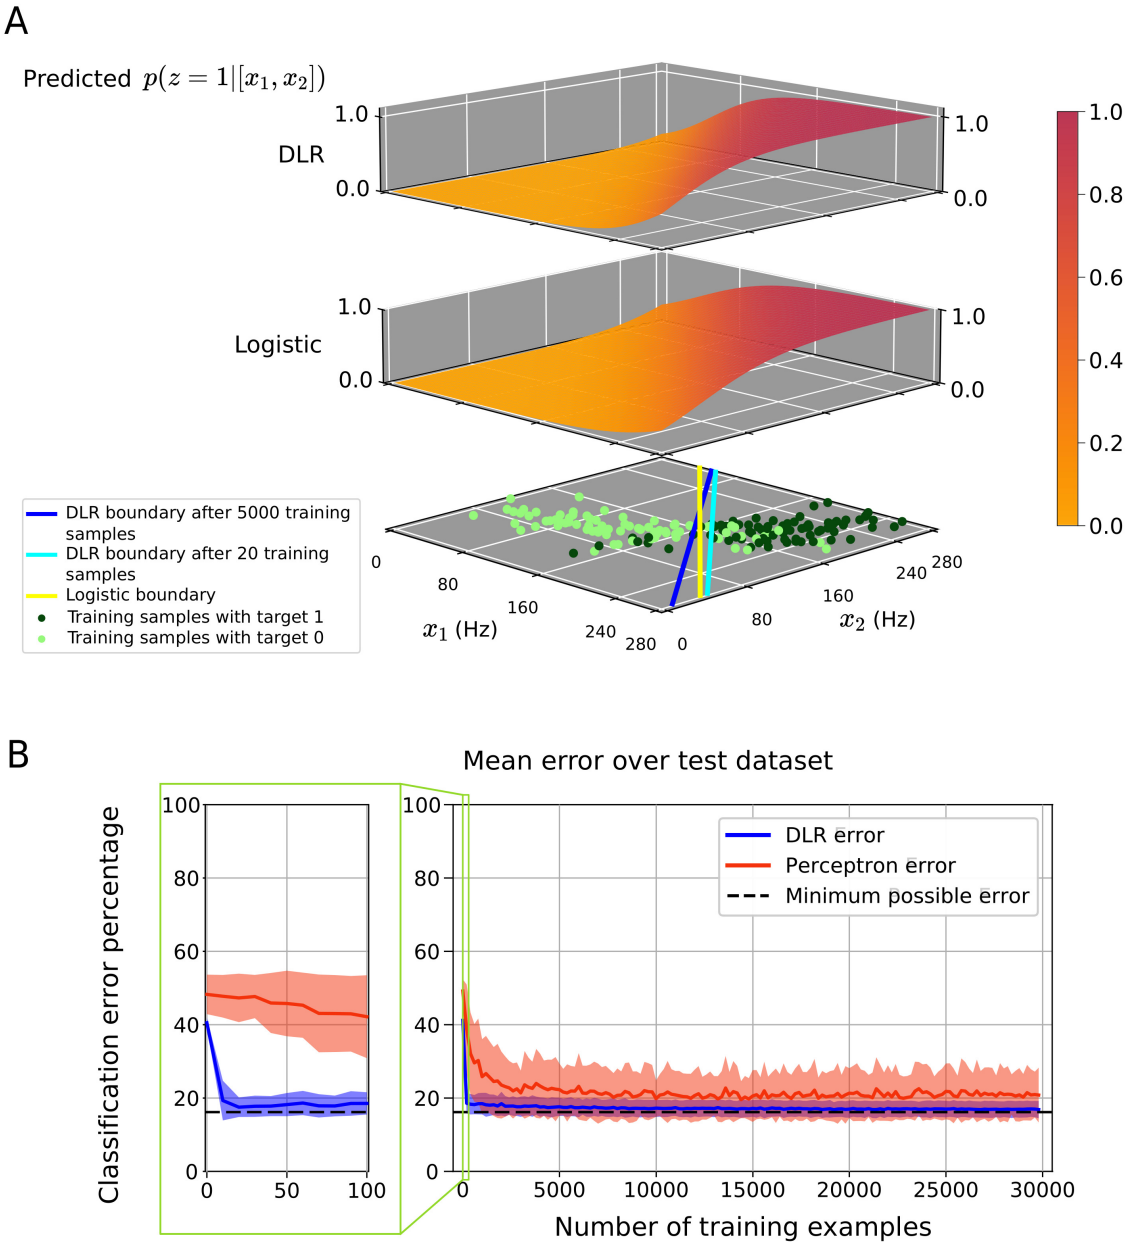

Figure 2: Illustration of emergent prediction capabilities of apical dendrites through DLR. A) Plot of probabilistic predictions for different firing rate patterns $\mathbf{x}$ of neurons in higher brain area that project to the considered apical dendrites, both as learned through DLR (at the top) and theoretically optimal predictions (below). The mean difference between the predicted probabilities is just 0.036. Also the resulting hyperplanes for resulting binary predictions for DLR and theoretically optimal logistic regression are very similar (bottom part). We also see a very rapid convergence of the DLR separation boundary after training with only 20 samples. **B)** Speed and accuracy of learning to guess (classify) from which cluster a presynaptic firing pattern $\mathbf{x}$ was drawn for the DLR learning rule and the perceptron learning rule, as considered in [Moldwin and Segev, 2020]. One sees that DLR enables much faster learning, and that the perceptron learning rule never converges to an optimal prediction performance. Mean and variance of the error are shown for 100 learning episodes with independently drawn initial weights and data points.

this analysis provides the foundation for our normative model. This plasticity rule can equivalently be described in the form of Table 1. Note that it contains a weight decay term that is activated when the apical dendrite is depolarized, i.e., $q_{\mathbf{w}}(\mathbf{x})$) is large, but there is no gating signal, i.e., $z(t) = 0$. Both the LTP and the LTD component of the plasticity rule are proportional to the pre-synaptic firing rate x_j.

Both the loss function that we minimize and the resulting type of plasticity rule are closely related to a well known optimization method in statistics and machine learning called Logistic Regression. Hence we refer to the resulting synaptic plasticity rule as Dendritic Logistic Regression (DLR).

2.1.1 A simple illustration of the capability of DLR to enable probabilistic predictions

Assume that there are two clusters of firing patterns $\mathbf{x}$ of neurons in higher brain areas that project to an apical dendrite of a pyramidal cell in layer 5, and that one can model each of these clusters by a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$, with parameters $(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ and $(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$. Furthermore assume that firing patterns $\mathbf{x}$ from this distribution for k = 1 predict that a subsequent sensory input will arrive at the target neuron (thereby causing a Ca^{2+}-spike in the neuron, i.e., a gating signal z = 1), and firing patterns from the distribution for k = 0 predict that this sensory input will not arrive (and that z = 0). Now assume that these two distributions overlap – as indicated by the clouds of dark green and light green dots in Fig. 2A. This occurs if many firing patterns $\mathbf{x}$ have significant probability under both of these distributions. In order to facilitate visualization, we consider here just two-dimensional firing patterns $\mathbf{x}$, but the same principle applies however for input patterns of arbitrary dimension. These two point clouds are not linearly separable, and hence the perceptron learning rule (as investigated in [Moldwin and Segev, 2020]) cannot do a good job in separating them, as indicated by the red trace in Fig. 2B. However synaptic weights in the apical dendrite can learn this inherently uncertain prediction task very fast through DLR, see the blowup of results for the first few trials in Fig. 2B. Furthermore the apical dendrites learn through DLR to predict the probability that a spike input pattern $\mathbf{x}$ predicts a subsequent sensory input through the resulting value of $q_{\mathbf{w}}(\mathbf{x})$ not only for those samples $\mathbf{x}$ that occurred during training, but they could generalize this probabilistic predictions to all potential values of presynaptic firing rates $\mathbf{x}$. Furthermore the resulting predicting capability was on the same level as the theoretically optimal logistic regression learning rule, as shown in Fig. 2A. The resulting prediction capability through DLR can be measured through the negative log likelihood (NLL) in nats (lower is better) for the training data. Across a 100 independent training runs, it achieves an average of 0.378 nats compared with 0.361 nats achieved through logistic regression.

2.1.2 Predicting future sensory input caused by a moving object

In order to illustrate predictive coding capabilities that emerge through DLR we consider the case that neurons in higher brain areas provide through their firing rates $\mathbf{x}$ an estimate of the initial position, speed, and movement direction of a moving object, see Fig. 3A. This information from neurons in higher brain areas is assumed to be projected to a sheet, modelled as a 30×30 grid, of pyramidal cells in layer 5 of a lower area through axons that form synapses on the apical dendrites of these pyramidal cells (we assume here for simplicity that there are no other synaptic connections between the neurons involved). We assume that each of these neurons receives also bottom-up visual input that informs it when a moving object currently covers its receptive field. This simple architecture supports already learning of predictions of future positions of a moving object through self-supervised learning via DLR: We assume that during training the gating signal z assumes value

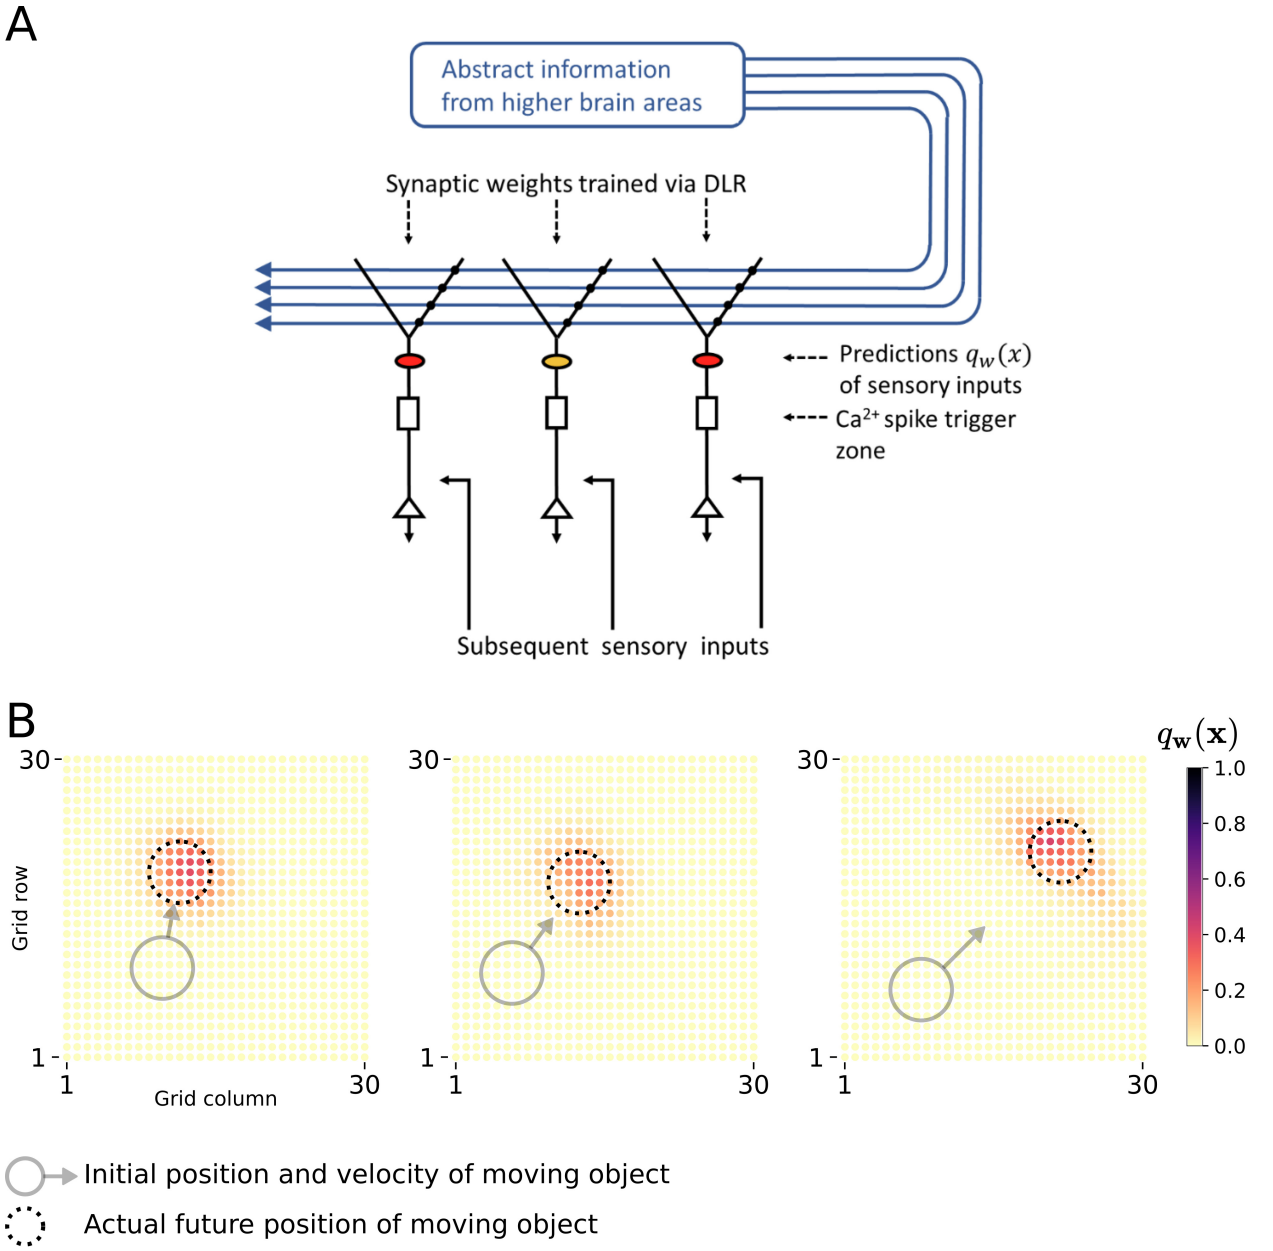

Figure 3: Performance of DLR for self-supervised learning to predict whether the receptive field of a neuron will be covered by a moving object. A) Scheme of the simple network architecture that we consider. We assume that initial position, speed, and movement direction of a moving object are computed in higher brain areas, and transmitted to primary visual cortex through synaptic connections in apical dendrites of a sheet of pyramidal cells whose receptive field forms a 30×30 grid in the visual field. **B)** Shown are in color-coding for three different new combinations of movement parameters the probabilistic predictions $q_{\mathbf{w}}(\mathbf{x})$ of each of pyramidal cells in a 30×30 array. These predictions emerged through self-supervised learning with DLR. One can compare the predicted probabilities of each of the neurons having its receptive field covered by the moving object at target time T to it's actual position at time T

1 when its receptive field is actually covered by the moving object at the target time point T. For simplicity we assume in Fig. 3B that the receptive field of each of these pyramidal cells consists of one grid point in a 30×30 grid. After self-supervised learning through DLR each pyramidal cell in the 30×30 array produces in its apical dendrites a prediction $q_{\mathbf{w}}(\mathbf{x})$, indicated through color coding in Fig. 3B, whether its receptive field will be covered by the moving object for given initial position, speed, and movement direction of the moving object. This is shown in Fig. 3B for three trials with combinations of initial position, speed, and movement direction that did not occur during training.

As expected, uncertainty of the position of the moving object at the target time T, which is compounded in this experiment through Gaussian noise on the top-down input to apical dendrites, is larger for higher movement speeds. For more quantitative measures of performance, refer to Methods section 3.3.2

2.2 Extensions to the case of multiple dendritic branches

When there are several dendritic branches, the question arises how the contributions from different branches are combined for the prediction $q_{\mathbf{w}}(\mathbf{x})$. Here, we denote the weights and inputs into each dendritic branch B by $\mathbf{w}^B = \left[w_1^B, w_2^B, \ldots, w_N^B\right]^T, \mathbf{x}^B = \left[x_1^B, x_2^B, \ldots, x_N^B\right]^T$ respectively, and the weights and inputs across all branches as $\mathbf{w}, \mathbf{x}$ respectively. We denote the predicted distribution by $q_{\mathbf{w}}(\mathbf{x})$ (similarly to equation 2), which estimates the target conditional distribution $p_{\text{target}}(z = 1|\mathbf{x})$.

It was shown in [Naud et al., 2014] that a neuron model with a somatic and dendritic compartment can accurately predict spike responses of cortical pyramidal neurons. Similar two-layer models have been proposed for hippocampal CA1 pyramidal cells [Poirazi et al., 2003] and even for cortical interneurons [Tzilivaki et al., 2019]. In general, it is believed that two-layer models provide a viable model for neocortical and hippocampal pyramidal cells [Jadi et al., 2014]. In this model the activity of each dendritic branch saturates according to a sigmoidal nonlinearity. Dendritic outputs are then summed and another non-linearity is applied to model the neuron output activity. In line with this model, we assume that the contributions from each dendritic branch (f^B) are a function of the mean branch local membrane potential $\bar{u}^B = \sum_i x_i^B w_i^B$ (equation 1) and consider a model of $q_{\mathbf{w}}(\mathbf{x})$ that accumulates contributions from the dendritic branches:

$$q_{\mathbf{w}}(\mathbf{x}) = \Psi\Big(\sum_B f^B(\bar{u}^B)\Big) \triangleq \Psi\left(f_{\text{total}}\right). \tag{4}$$

Consistent with the single-branch model (equation 2), the contribution of each branch is given by a sigmoidal function of its weighted input rate, i.e.

$$f^B = \sigma(\beta\left(\bar{u}^B - u_0\right)). \tag{5}$$

For the function Ψ we consider two possibilities. According to the classical two-layer model, the neuron output activity is given by a saturating nonlinearity applied to the summed dendritic contributions. Hence, as a first possibility, we model Ψ as a sigmoidal function in Subsection 2.2.1 below. More recently, a more complex relationship was found for single human neurons [Gidon et al., 2020] which will be considered in Subsection 2.2.2.

2.2.1 Saturating model

In the first possibility, taking into consideration the thresholding and saturation dynamics of pyramidal neurons [Naud et al., 2014, Jadi et al., 2014], we chose the sigmoidal shape for Ψ to get

$$\Psi\left(f_{\text{total}}\right) = \sigma\left(\alpha\left(f_{\text{total}} - \vartheta\right)\right), \tag{6}$$

where α and θ are suitable constants. Note that by choice of α and ϑ, one can model different types of (soft-) logic operations. A low ϑ lesser than 1 implements a soft OR operation (a minimum of one branch needs to be active to get a high probability output). A ϑ close to the number of branches implements a soft AND, and in-between we obtain a general threshold condition on the number of active branches. See Fig. 4B (left and middle) for an example with two branches.

The goal of the learning rule in this case is to get $q_{\mathbf{w}}(\mathbf{x})$ to predict the probability $p_{\text{target}}(z = 1|\mathbf{x})$ of Ca^{2+}-spikes that occurs during training. Hence a normative model for synaptic plasticity in these dendritic branches should minimize the distance between these two terms, formulated again as cross-entropy error.

The above choices for f^B, Ψ leads to the following weight update rule for the weights incoming to a particular dendritic branch B:

$$\Delta\, w_j^B = \eta\, \left(z - \Psi\left(f_{\text{total}}\right)\right)\, \dot{f}^B(u^B)\, x_j^B, \tag{7}$$

with learning rate η. This update is similar to the single-dendrite update, with two differences: (1) the term $\dot{f}^B(u^B)$ means that there is a nonlinear dependence on the local branch potential, and (2) the update at a coincidence event (LTP) depends on $\Psi\left(f_{\text{total}}\right)$. This quantity is not local to the dendritic branch. In order to minimize the information about $\Psi\left(f_{\text{total}}\right)$ at the branch, one can consider a gating of updates based on thresholds on $\Psi\left(f_{\text{total}}\right)$ which approximates the term $\left(z - \Psi\left(f_{\text{total}}\right)\right)$ in the update rule 7. This leads to the following updates of synaptic efficacies:

$$\begin{aligned}
&\text{For } z = 0: \\
&\qquad \Delta\, w_j^B = \begin{cases} -\dot{f}^B(u^B) x_j^B & , \quad \text{if } \Psi\left(f_{\text{total}}\right) \geq q^- \\ 0 & , \quad \text{otherwise} \end{cases} \\
&\text{For } z = 1: \\
&\qquad \Delta\, w_j^B = \begin{cases} \dot{f}^B(u^B) x_j^B & , \quad \text{if } \Psi\left(f_{\text{total}}\right) \leq q^+ \\ 0 & , \quad \text{otherwise} \end{cases}
\end{aligned} \tag{8}$$

for plasticity thresholds q^- and q^+. To see the correspondence with rule 7, first consider the case without Ca^{2+}-spike ($z = 0$). Here, the term $\left(z - \Psi\left(f_{\text{total}}\right)\right)$ in equation 7 is between -1 and 0. It becomes very small when $\Psi\left(f_{\text{total}}\right)$ is close to zero. This is approximated by a gating mechanism which blocks plasticity if $\Psi\left(f_{\text{total}}\right)$ is below some threshold q^- close to 0. When there is a Ca^{2+}-spike ($z = 1$), the term $\left(z - \Psi\left(f_{\text{total}}\right)\right)$ is between 0 and 1. It becomes very small when $\Psi\left(f_{\text{total}}\right)$ is close to one. Hence, plasticity is blocked if $\Psi\left(f_{\text{total}}\right)$ is above some threshold q^+ close to 1 in the approximation 8.

2.2.2 Soft-XOR model

This second choice for Ψ is motivated by recent findings from human L2/3 pyramidal cells. According to [Gidon et al., 2020], Ca^{2+}-spikes are initiated in these cells at a threshold activation ϑ of dendrites, but if dendrites are even more strongly activated, these spikes are not elicited. We can model such behavior by a nonlinearity $\Psi(f_{\text{total}})$ that is close to 1 in a certain activation region $(\vartheta, \vartheta + \Delta\vartheta)$, see Fig. 4B (right) and Methods.

A plasticity rule for this model is derived in Methods, where the plasticity at the synapse depends on f_{total}. Again, we can consider a coarse influence of the f_{total} on the local plasticity, where it switches between LTP, LTD, and no update based on three thresholds f^- (low f_{total}), f^{mid} (medium f_{total}), and f^+ (high f_{total}), see Fig. 4B (right). This leads to the following updates of synaptic efficacies:

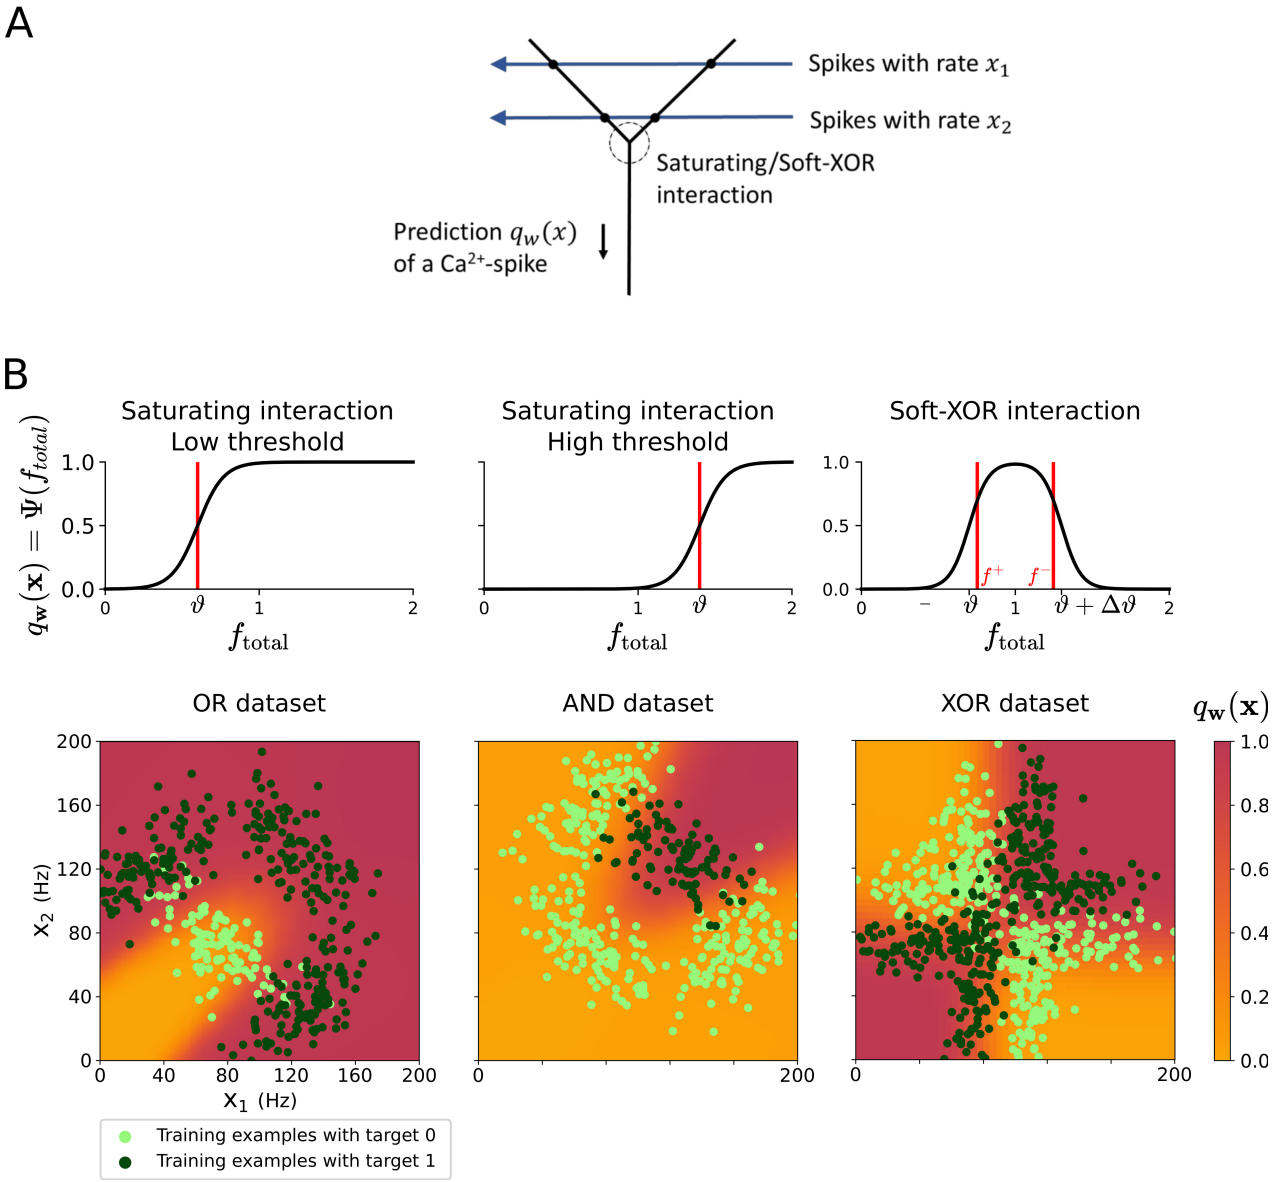

Figure 4: Illustration of emergent prediction capabilities through DLR of apical dendrites in the case of multiple branches in the dendritic tuft. A) We consider here for simplicity the case of two branches of the dendritic tuft that both receive synaptic input from the same two neurons in a higher brain area, for two experimentally reported types of interaction of their contributions: saturating and soft-XOR. **B)** The resulting probabilistic predictions $\Psi(f_{\text{total}})$ of the apical tuft are shown after training on three different datasets. Note that for neither one of these datasets a single-branch probabilistic prediction as shown in Fig. 2 suffices. But both experimentally observed saturating and XOR-like interactions between different dendritic branches enable DLR to produce suitable distributions of probabilities, indicated again through color coding, that separate the two point clouds. Note that like in the case of Fig. 2, a number of these training examples could have been generated by either one of the two distributions that generate dark green and light green datapoints. Hence probabilistic predictions are needed, like in the case of Fig. 2. But with DLR running in parallel in several apical dendrites, also complex distributions of datapoints can be handled (Table 2)

$$
\begin{aligned}
&\text{For } z = 0: \\
&\Delta\, w_j^B = \begin{cases} -\dot{f}^B(u^B)x_j^B & , \quad \text{if } f^- \leq f_{\text{total}} < f^{\text{mid}} \\ \dot{f}^B(u^B)x_j^B & , \quad \text{if } f^{\text{mid}} \leq f_{\text{total}} \leq f^+ \\ 0 & , \quad \text{otherwise} \end{cases} \\
&\text{For } z = 1: \\
&\Delta\, w_j^B = \begin{cases} \dot{f}^B(u^B)x_j^B & , \quad \text{if } f_{\text{total}} < f^- \\ -\dot{f}^B(u^B)x_j^B & , \quad \text{if } f_{\text{total}} > f^+ \\ 0 & , \quad \text{otherwise} \end{cases}
\end{aligned}
\tag{9}
$$

We can interpret these plasticity dynamics as follows. Assume first the case of a Ca^{2+}-spike $z = 1$. For low f_{total} ($< f^-$), the activation has to increase in order to increase the probability of the model. Hence active synapses are potentiated. For medium f_{total} ($f^- \leq f_{\text{total}} \leq f^+$), we have already a high probability for a Ca^{2+}-spike. Hence, no update is necessary. For high f_{total} ($> f^+$), the activation is too high and depression is needed. Consider now the case of no Ca^{2+}-spike $z = 0$. If the neuron is at medium f_{total}, the probability can be decreased either by decreasing the activation (LTD) or increasing it (LTP). In the upper half of the "high probability bump" ($f^{\text{mid}} < f_{\text{total}} \leq f^+$), LTP is chosen, and in the lower half ($f^- \leq f_{\text{total}} < f^{\text{mid}}$) LTD is chosen, which shift f_{total} in the direction leading to a decreased probability. Outside of this, i.e. for $f_{\text{total}} > f^+$ and $f_{\text{total}} < f^-$, we already predict a low probability of Ca^{2+}-spike and thus no updates are necessary.

2.2.3 Demonstration of the enhanced probabilistic prediction capabilities that emerge through DLR in the case of several apical dendrites

Similarly as in Fig. 2 we assume here that data points with labels $z = 0$ or 1 are drawn from some overlapping probability distributions. But in contrast to Fig. 2 we assume here that these distributions are not just Gaussians, that can be separated by an optimal linear decision boundary. Examples for the resulting clouds of data points are shown for three examples of more complex distributions in Fig. 4B.

Using a model with two dendritic branches, we consider three different kinds of interactions between these two branches – a saturating interaction with a low threshold, high threshold, and a soft-XOR interaction. One sees in Fig. 4B that each of these three types of interactions is useful for learning probabilistic predictions for three different types of distributions of data points. Furthermore one sees that training via DLR enables the dendritic branches to provide also meaningful predictions for presynaptic firing rates $\mathbf{x}$ that never occurred during training. Quantitative results are given in Methods.

Table 2 shows how well the predictions of various models fit the three datasets in terms of the negative log-likelihood. When using two dendritic branches, the cases of saturating interaction with low and high value of the strong spike threshold perform very well on the OR and AND tasks respectively. The Soft-XOR interaction performs much better than the saturating interaction in modelling the XOR dataset. This validates the intuitive understanding of the different interactions and how they correspond to different boolean operations. Adding more branches with a saturating interaction with an intermediate threshold value ϑ, we see that these configurations perform very well

across all the three datasets thus showing that increasingly complex distributions can be modelled by a single configuration given enough branches.

3 Methods

3.1 Derivation of DLR for a single branch

We detail here the model for a single dendritic branch, the associated regression formulation, and derive the DLR rule for this case.

In general, each dendritic branch indexed by B receives top-down input $\mathbf{x}^B$ weighted using input weights $\mathbf{w}^B$. Here, we consider the case of a single dendritic branch. Hence, to simplify notation, we will in the following suppress the superscripts and simply write $\mathbf{x}$ for the inputs to this branch and $\mathbf{w}$ for the corresponding weights. In our model the input $\mathbf{x}$ represents the vector of firing rates of the neurons pre-synaptic to the dendritic branch. We assume that a specific input $\mathbf{x} = [x_1, \ldots, x_N]^T$ is presented to the neuron for a fixed time duration as input spike trains $s_1(t), \ldots, s_N(t)$ where $s_n(t)$ is a Poisson spike train of rate equal to x_n for each n.

For the duration of the input presentation, the membrane potential u is given by the weighted sum of the spike trains convolved temporally with a linear post-synaptic potential kernel k_{PSP}

$$u(t) = \sum_{n=1}^{N} w_n \left(k_{PSP} * s_n\right)(t). \tag{10}$$

For any PSP shape, the linear summation of incoming input spikes means that the expected mean membrane voltage for this interval $\bar{u}$ is proportional to $\mathbf{x}^T\mathbf{w}$. Without loss of generality we choose a PSP kernel for which the constant of proportionality is 1, giving

$$\bar{u} = \mathbf{x}^T\mathbf{w}. \tag{11}$$

The Ca^{2+}-spikes are denoted by z where $z = 1$ denotes the presence of the calcium spike, and $z = 0$ denotes the absence of a Ca^{2+}-spike for the duration of the input. We denote the joint distribution of the input $\mathbf{x}$ and the calcium spikes z as $p_{\text{target}}(\mathbf{x}, z)$. Our regression formulation then seeks to provide an estimator parametrized by $\mathbf{w}$ ($q_{\mathbf{w}}(\mathbf{x})$) which is optimized to approximate $p_{\text{target}}(z = 1|\mathbf{x})$. Namely, we aim to train the weights $\mathbf{w}$ so that $q_{\mathbf{w}}(\mathbf{x})$ accurately predicts the probability of the occurence of a Ca^{2+}-spike given a particular top-down input $\mathbf{x}$.

The cross entropy loss function that measures the distance between $q_{\mathbf{w}}(\mathbf{x})$ and the data distribution p_{target} is defined by

$$\mathcal{LL}(\mathbf{w}) = -E_{\mathbf{x},z}\left[z \log q_{\mathbf{w}}(\mathbf{x}) + (1 - z) \log\left(1 - q_{\mathbf{w}}(\mathbf{x})\right)\right], \tag{12}$$

where $\mathbf{x}, z$ are sampled from the data distribution p_{target}.

The corresponding gradient with respect to the parameters $\mathbf{w}$ is given by

$$\nabla_{\mathbf{w}}\, \mathcal{LL}(\mathbf{w}) = -E_{\mathbf{x},z}\left[\frac{\nabla_{\mathbf{w}}\, q_{\mathbf{w}}(\mathbf{x})}{q_{\mathbf{w}}(\mathbf{x})\left(1 - q_{\mathbf{w}}(\mathbf{x})\right)}\left(z - q_{\mathbf{w}}(\mathbf{x})\right)\right]. \tag{13}$$

For the special case where $q_{\mathbf{w}}(\mathbf{x})$ takes the simple form $q(\mathbf{w}^T\mathbf{x})$ for some nonlinearity q, one gets

$$\nabla_{\mathbf{w}}\, \mathcal{LL}(\mathbf{w}) = -E_{\mathbf{x},z}\left[\frac{\dot{q}(\mathbf{w}^T\mathbf{x})}{q(\mathbf{w}^T\mathbf{x})\left(1 - q(\mathbf{w}^T\mathbf{x})\right)}\left(z - q(\mathbf{w}^T\mathbf{x})\right)\mathbf{x}\right]. \tag{14}$$

A further simplification arises if one uses a logistic sigmoid function as nonlinearity q such that:

$$q_{\mathbf{w}}(\mathbf{x}) = \sigma\left(\beta(\mathbf{w}^T\mathbf{x} - u_0)\right) = \sigma\left(\beta(\bar{u} - u_0)\right), \tag{15}$$

where $\sigma(a) = \frac{1}{1+e^{-a}}$ and β and u_0 are model constants. Hence, in this case, the predictive depolarization that results from synaptic currents at the dendrite is a saturated version of $\bar{u}$ as given in equation 2. Then, using the fact that $\sigma'(a) = \sigma(a)\,(1 - \sigma(a))$, one obtains

$$\nabla_{\mathbf{w}}\, \mathcal{LL}(\mathbf{w}) = -E_{\mathbf{x},z}\left[(z - q_{\mathbf{w}}(\mathbf{x}))\ \mathbf{x}\right]. \tag{16}$$

When we plug in the definition of $q_{\mathbf{w}}(\mathbf{x})$ from equation 15 into the loss function $\mathcal{LL}(\mathbf{w})$, we find that this loss function corresponds to a well known linear regression scheme called Logistic Regression. It is known that this loss function is a convex loss function.

Thus, with $\mathbf{x}$ encoded as the mean input spike rate, and $\bar{u} = \mathbf{w}^T\mathbf{x}$ (equation 11), we see that the weight update rule described in equation 3 performs stochastic descent along the gradient of the convex logistic cross entropy loss function. This guarantees the convergence of the weight update to the global optimum of the loss function.

3.2 Derivation of DLR for the case of multiple branches

Here, $\mathbf{x}^B, \mathbf{w}^B$ represent the weights and input to dendritic branch B, and $\mathbf{x}, \mathbf{w}$ represent the full input and weights. The corresponding cross entropy loss with $q_{\mathbf{w}}(\mathbf{x})$ defined as in equation 4 is given by

$$\begin{aligned} L(\mathbf{w}) &= -E_{\mathbf{x},z}\left[z \log q_{\mathbf{w}}(\mathbf{x}) + (1-z)\log\left(1 - q_{\mathbf{w}}(\mathbf{x})\right)\right] \\ &= -E_{\mathbf{x},z}\left[z \log \Psi\left(f_{\text{total}}\right) + (1-z)\log\left(1 - \Psi\left(f_{\text{total}}\right)\right)\right]. \end{aligned} \tag{17}$$

The derivative of the above loss with respect to the weights coming to a particular branch $\mathbf{w}^B$ evaluates to

$$\nabla_{\mathbf{w}^B} L(\mathbf{w}) = -E_{\mathbf{x},z}\left[\frac{\dot{\Psi}\left(f_{\text{total}}\right)}{\Psi\left(f_{\text{total}}\right)\left(1 - \Psi\left(f_{\text{total}}\right)\right)}\left(z - \Psi\left(f_{\text{total}}\right)\right)\dot{f}^B(\mathbf{w}^{B^T}\mathbf{x}^B)\ \mathbf{x}^B\right]. \tag{18}$$

3.2.1 Derivation of DLR for the saturating model

In this model, a sigmoidal shape for Ψ is chosen

$$\Psi\left(f_{\text{total}}\right) = \sigma\left(\alpha\left(f_{\text{total}} - \vartheta\right)\right), \tag{19}$$

where α and θ are suitable constants. Using this in equation 18 we obtain the gradient for the saturating model:

$$\nabla_{\mathbf{w}^B} L(\mathbf{w}) = -\,\alpha\ \left(z - \Psi\left(f_{\text{total}}\right)\right)\ \dot{f}^B(u^B)\,\mathbf{x}^B, \tag{20}$$

which leads to the weight update rule in equation 7.

In order to minimize the information about $\Psi\left(f_{\text{total}}\right)$ at the branch, one can consider a gating of updates based on thresholds on $\Psi\left(f_{\text{total}}\right)$ which approximates the term $\left(z - \Psi\left(f_{\text{total}}\right)\right)$ in the update rule 7. We consider the two cases $z = 0$ and $z = 1$ separately. For $z = 0$, the term $\left(z - \Psi\left(f_{\text{total}}\right)\right)$ is between -1 and 0. It becomes very small when $\Psi\left(f_{\text{total}}\right)$ is close to zero. A gating mechanism which blocks plasticity if $\Psi\left(f_{\text{total}}\right)$ is below some threshold q^- close to 0 approximates this term. For $z = 1$, the term $\left(z - \Psi\left(f_{\text{total}}\right)\right)$ is between 0 and 1. It becomes very small when $\Psi\left(f_{\text{total}}\right)$ is close to one. A gating mechanism which blocks plasticity if $\Psi\left(f_{\text{total}}\right)$ is above some threshold q^+ close to 1 approximates this term. Hence, we can formulate the approximate update 8.

3.2.2 Derivation of DLR for the soft-XOR model

In the soft-XOR model, Ψ is given by

$$\Psi(f_{\text{total}}) = \sigma\left(\alpha\left(f_{\text{total}} - \vartheta\right)\right)\sigma\left(\alpha\left(-f_{\text{total}} + \vartheta + \Delta\vartheta\right)\right) \tag{21}$$

$$\equiv \sigma_1(f_{\text{total}})\sigma_2(f_{\text{total}}). \tag{22}$$

Suitable parameters α, θ, and $\Delta\theta$ lead to an XOR-like behavior (high probability in a certain activation region), see Fig. 4D. Using this nonlinearity in equation 18, we obtain

$$\nabla_{\mathbf{w}^B} L(\mathbf{w}) = -E_{\mathbf{x},z}\left[\frac{\sigma_2(f_{\text{total}}) - \sigma_1(f_{\text{total}})}{(1 - \Psi\left(f_{\text{total}}\right))}\left(z - \Psi\left(f_{\text{total}}\right)\right)\dot{f}^B(\mathbf{w}^{B^T}\mathbf{x}^B)\,\mathbf{x}^B\right]. \tag{23}$$

Considering a coarse influence of the f_{total} on the local plasticity, we can formulate an approximate rule which switches between LTP, LTD, and no update based on three thresholds f^- (close to ϑ), f^{mid} $(\vartheta + \Delta\vartheta/2)$, and f^+ (close to $\vartheta + \Delta\vartheta$), see Fig. 4B. This leads to the updates given in equation 9.

3.3 Simulation details

3.3.1 Classification of 2D data (Fig. 2)

The training examples for this task consist of points generated from two 2D Gaussian clusters with means and covariances $\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k$ for $k \in \{0, 1\}$. All points generated from $\mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ have the target value $z = k$. The prior probabilities for the classes 0 and 1 were chosen to be equal (0.5). Thus to generate a training sample $(\mathbf{x}_i, z_i)$, we first pick the Gaussian component $k \in \{0, 1\}$ with equal probability i.e. $k \sim \text{bernoulli}(p = 0.5)$. We then pick $\mathbf{x}_i \sim \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ and assign the target $z_i = k$. This gives the following conditional probability distribution $p(z = 1|\mathbf{x})$

$$p(z = 1|\mathbf{x}) = \frac{p(\mathbf{x}|k = 1)}{p(\mathbf{x}|k = 0) + p(\mathbf{x}|k = 1)},$$

where $p(\mathbf{x}|k)$ is the Gaussian probability density function with mean $\boldsymbol{\mu}_k$ and covariance $\boldsymbol{\Sigma}_k$. It is the aim of the algorithm to learn this conditional distribution in the quantity $q_{\mathbf{w}}(\mathbf{x})$ in the dendrite (equation 2).

For this demonstrative simulation we consider the neuron to be receiving the input on a single dendritic branch, with the DLR being applied to the incoming weights to that branch. Each data point $\mathbf{x}_i$, is padded with an additional dimension containing the value 1 (in order to model intercept fitting). Each component of this data point is then given as input to the network encoded as a Poisson spike train with mean rate proportional to the value of the component. This spike train is provided to the neuron for a duration of 400 ms. Targets are fed to the network by enabling or disabling Ca^{2+}-spike firing corresponding to $z = 1$ and $z = 0$ respectively. Fig. 1B illustrates this neuron and learning setup.

3.3.2 Predicting the future sensory input from a moving object (Fig. 3)

In our simulation the moving object is a 2D circle which moves within a 30×30 grid and has a diameter of 5 Grid Units (GU). We consider here the task of predicting the sensory input corresponding to its position after a fixed delay of T. The top-down input corresponding to an instance of a moving object is specified by a four-tuple (r_x, r_y, v, θ) consisting of x and y coordinates r_x, r_y of the initial

position, the speed v, and the direction angle θ. This four-tuple is encoded in a population coded manner described below. This set of population coded inputs is connected to a 30×30 grid of L5P neurons (one neuron per grid position). The speed and direction are perturbed by a Gaussian noise to model inaccuracies in top-down measurement. This grid is trained to reproduce the image of the ball at its final position which is determined by the initial position and the true (unperturbed) velocity in the external environment. The incoming weights to each L5P neuron are trained using DLR with the target being the corresponding pixel value in this image.

The top-down inputs specifying the initial position and velocity of the ball are encoded as follows. The initial position (r_x, r_y) can take any value such that both r_x, r_y take values in the range of $S_r = [3, 27]$ GU. The speed v is defined in terms of the distance (in GU) that the moving object covers in the fixed delay T between the top-down and the bottom up inputs. This distance can take any value in the range $S_v = [4, 20]$ GU. The direction is specified by an angle θ to the horizontal which can take any value in the range $[0, 2\pi)$.

The encoding of these values into neuron populations is described below. We consider the sets M_r, M_v and M_θ, which consist of N_r, N_v, and N_θ equally spaced values in the ranges S_r, S_v, and S_θ respectively. Then, corresponding to each four-tuple $\big(\mu_x, \mu_y, \mu_v, \mu_\theta\big) \in M_r \times M_r \times M_v \times M_\theta$, we have a single input neuron that fires according to a Gaussian tuning curve centered at $\big(\mu_x, \mu_y, \mu_v, \mu_\theta\big)$. The neuron encodes the input $\big(r_x, r_y, v, \theta\big)$ by firing at a rate ρ given by

$$\rho = \frac{\rho_{\text{ total}}}{Z} e^{-\frac{(r_x-\mu_x)^2 (r_y-\mu_y)^2 (v-\mu_v)^2 (v-\mu_\theta)^2}{8{\sigma_r}^2 {\sigma_r}^2 {\sigma_v}^2 {\sigma_\theta}^2}}, \tag{24}$$

where Z is a normalization factor calculated so that the total spike rate of all the input neurons is $\rho_{\text{ total}} = 250$ Hz. Here the standard deviation of the response curve corresponding to position σ_r, velocity σ_v, and direction σ_θ are constant across all neurons.

For the simulation whose results are shown in Fig. 3, we use $N_r = 10$, $N_v = 10$, $N_\theta = 24$, $\sigma_r = 2.67$, $\sigma_v = 1.78$, $\sigma_\theta = 0.2618 = 15.0°$ This leads to a total input dimension of ${N_r}^2 N_v N_\theta = 24000$. The standard deviations of the Gaussian noise perturbations applied to the speed and direction components of the top-down input are equal to the width of the respective receptive fields σ_v, σ_θ. In order to train the network we make use of 168000 data points in the training dataset.

As a demonstration that the network predicts the position of the ball consistently, we measure the mean predicted $q_{\mathbf{w}}(\mathbf{x})$ for each neuron for which the moving object covers its receptive field after the fixed delay T. This corresponds to grid positions within a circle of diameter 5GU centered at the final position of the moving object. We then consider two more concentric slices of diameters $5 - 9$ GU and $9 - 13$ GU and compare the mean $q_{\mathbf{w}}(\mathbf{x})$ of all the corresponding neurons.

1. For pixels covered by the object at time T, mean $q_{\mathbf{w}}(\mathbf{x}) = 0.2374$
2. For the concentric slices of diameters $5 - 9$ GU, mean $q_{\mathbf{w}}(\mathbf{x}) = 0.1213$
3. For the concentric slices of diameters $9 - 13$ GU, mean $q_{\mathbf{w}}(\mathbf{x}) = 0.0433$

These results show that the predicted probability $q_{\mathbf{w}}(\mathbf{x})$ of a neuron receiving sensory input 1 at time T is highest for those corresponding to grid positions actually occupied by the moving object at T. This then decreases as we get further away from the object's final position (i.e. in the concentric slices). We also see that the uncertainty in velocity means that there is a significant non-zero probability assigned to grid positions that surround the true final position in the concentric slice $5 - 9$ GU in diameter.

The examples in Fig. 3 and the results above are calculated from combinations of initial position and velocity that were not seen when training.

		OR dataset	AND dataset	XOR dataset
Multi-branch DLR:				
Number of Branches	Branch Interaction			
2	Saturating (low ϑ)	0.21	0.60	0.61
2	Saturating (high ϑ)	0.65	0.18	0.60
2	Soft-XOR	0.19	0.33	0.36
4	Saturating	0.20	0.17	0.45
6	Saturating	0.18	0.16	0.36
Logistic Regression:		0.43	0.45	0.69

Table 2: Performance of DLR for the case of several interacting apical dendrites for probabilistic predictions in cases where there are no linear optimal decision boundaries: Mean negative log likelihood (NLL) in nats (lower is better) is given for the training sets of the three different data distributions considered in 4B, and for the three types of interactions among apical dendrites that we consider, as well as for different numbers of apical dendrites. One sees that using a saturating interaction with low threshold for two branches excels at the OR task, while the high threshold scheme excels at the AND task. The Soft-XOR interaction turns out to be substantially more effective for learning with two branches the probabilistic prediction for the case of the XOR data distribution. However, with more branches the saturating interaction supports learning of probabilistic predictions via DLR for all three tasks. Each of the mean NLL values above is calculated from the best 8 out of 14 independent training runs (to avoid outliers which do not converge).

3.4 Predicting complex distributions with multiple dendritic branches (Fig. 4)

In this experiment, the 2D data is generated from multiple gaussian clusters. The parameters and the class labels assigned to each cluster are chosen to create datasets corresponding to an OR, AND, and XOR operation shown in Fig. 4.

In order to solve this task, we consider a neuron with N_B (two or more) dendritic branches each with their respective incoming synapses with weights $\mathbf{w}^1, \ldots, \mathbf{w}^{N_B}$. All of these branches receive the same 2D input, i.e., $\mathbf{x}^B = [\mathbf{x}; 1] = [x_1, x_2, 1]$ for $B = 1, \ldots, N_B$. This input is provided as in the 2D regression task for a single branch, i.e., as a spike sequence with a spike rate proportional to the input value. The additional input component with value 1 is given in order to model the training of the branch threshold.

The various configurations in table 2 along with corresponding parameters are below.

1. 2 Branches, saturating interaction, low threshold $\vartheta = 0.6$: This implements the OR operation between the two branches.

2. 2 Branches, saturating interaction, high threshold $\vartheta = 1.2$: This implements the AND operation between the two branches.

3. 2 Branches, soft-XOR interaction, $\vartheta = 0.5$, $\Delta\vartheta = 1.0$. This implements the XOR operation between the two branches.

4. 4 Branches, saturating interaction, $\vartheta = 1.4$

5. 6 Branches, saturating interaction, $\vartheta = 2.1$

For all the saturating interactions, we use a scale factor $\alpha = 12$. The parameters q^+, q^- used in the learning rule in 8 are set to $q^+ = 0.9$ and $q^- = 0.1$. For the Soft-XOR interaction, we use $\alpha = 8.0$ and set $f^- = \vartheta = 0.5$, $f^+ = \vartheta + \Delta\vartheta = 1.5$ and $f^{\text{mid}} = \vartheta + \Delta\vartheta/2 = 1.0$

3.5 Common Parameters for DLR:

Some of the common parameters and the values for the DLR rule for the three experiments are given in Table 3

Parameter	Single-branch DLR		Multi-branch DLR	
	2-D regression	Predicting input from moving object	Saturating interaction	Soft-XOR interaction
η_0	1.8	40.0	4.0	0.3
η_{final}	0.3	6.67	4.0	0.3
β	0.5	0.5	0.1	1.0
u_0	20.0	15.0	7.0	7.0

Table 3: Parameters related to DLR that are used for each experiment. η_0, η_{final} are the initial and final learning rate used. β and u_0 are the scale and threshold for the sigmoid function used to calculate $q_{\mathbf{w}}(\mathbf{x})$ and f^B in equations 2 and 5 respectively

4 Discussion

We have presented a mathematical framework and learning theory that enables a principled theoretical investigation of arguably the most important synaptic connection for the integration of top-down predictions from higher brain areas into computational processing of sensory input streams: Synapses that are formed by axons of neurons in higher brain areas in the dendritic tuft of pyramidal cells in lower brain areas. Experimental data show that plasticity at these synapses follows rules that differ from rules for correlated Hebbian learning, because they are gated by instructive signals [Magee and Grienberger, 2020]. Furthermore these instructive signals differ from commonly considered global third factors for synaptic plasticity, such as dopamine, that typically signal reward. These instructive signals for synaptic plasticity at distal dendrites either emerge locally within the neuron — such as Ca^{2+}-spikes [Kampa et al., 2006, Kampa et al., 2007], NMDAR spikes [Gambino et al., 2014], or other types of plateau potentials [Bittner et al., 2015, Bittner et al., 2017, Magee and Grienberger, 2020] – or arrive in the dendritic tuft through synaptic connections from brain areas such as perirhinal cortex [Doron et al., 2020] or higher order thalamus [Aru et al., 2020]. We have focused in our illustrative example in Fig. 1 and in the text on the case where these instructive signals, denoted by z in our mathematical models, take the form of Ca^{2+}-spikes. But our theoretical framework and normative learning theory also applies to the other cases. The only aspect on which this learning theory is based is that these instructive signals last for 50-100 ms or longer, so that the firing rates of presynaptic neurons become salient, rather than single spikes.

Apart from these neurophysiological differences, the impact of plasticity at these synapses on computational processing appears to differ from that of other synaptic connections. For example, the input from apical dendrites has been shown to shift the perceptual threshold in a graded manner [Takahashi et al., 2016]. In addition it carries predictions from higher brain areas that are inherently

uncertain, and therefore would be functionally most useful if it transmits probabilities, rather than categorical decisions. We have shown that this learning goal can be achieved through simple rules for LTP and LTD, gated by instructive signals z, at synapses onto distal dendrites. In fact, a suitable balancing of LTP and LTD can in principle result in a depolarization of the dendritic tuft that can be seen as logistic regression, i.e., a simple form of probabilistic prediction that is in a well-defined sense optimal. Hence we refer to this normative model for synaptic plasticity in the dendritic tuft as dendritic logistic regression (DLR). We have shown in illustrative examples that DLR enables the dendritic tuft to make meaningful probabilistic predictions in the case where the same prediction is sometimes correct and sometimes not (Fig. 2A). Furthermore these probability predictions emerge really fast (Fig. 2B). In particular, they provide substantially better performance than the perceptron learning rule [Moldwin and Segev, 2020]. The reason is that the latter is guaranteed to work well only if the training examples are linearly separable, in particular if they do not contain contradicting examples that commonly occur in the case of probabilistic predictions. We have also demonstrated in Fig. 3 that the DLR rule enables distal dendrites to predict future sensory inputs caused by a moving object.

A particular experimentally testable prediction of this normative model, i.e., DLR, is that after learning, the depolarization in a dendritic branch of the tuft is correlated with the probability that this same synaptic input coincides with an instructive signal, for example with a Ca^{2+}-spike. A more mechanistic prediction is that the current depolarization of the dendritic branch scales the amplitude of LTP and LTD during the induction of plasticity. In particular, if synaptic input to an apical dendrite coincides with a plateau potential, the expression of LTP is predicted to be weaker if the dendrite was already strongly depolarized at this time. Furthermore, synaptic input to an apical dendrite is predicted to cause LTD in the absence of a plateau potential, and the impact of LTD is stronger if the apical dendrite is more depolarized at the time of induction. An open question is whether LTP of synapses on distal dendrites of pyramidal cells in areas CA1 as studied in [Bittner et al., 2015, Bittner et al., 2017] is accompanied also by corresponding forms of LTD, and whether both processes combine to form mearningful probabilistic predictions.

We have extended our learning theory in the second part to the case of several interacting branches in the dendritic tuft. Two different types of interactions between depolarizations in different branches of the dendritic tuft have been found in experiments on different neurons: Summation of depolarization [Jadi et al., 2014, Poirazi et al., 2003], and competition between depolarized branches [Gidon et al., 2020]. We have characterized in Fig. 4 additional capabilities of probabilistic prediction that are enabled by DLR for these two types of interactions between dendritic branches. In both cases, probabilistic predictions emerge that can handle substantially more complex distributions or training examples as compared to the case of a single dendritic branch. In this sense, our learning theory also throws new light on likely computational contributions of the dendritic arbor.

Acknowledgments

This work was supported by the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreements No. 785907 (Human Brain Project) and No. 899265 (ADOPD).

References

[Aru et al., 2020] Aru, J., Suzuki, M., and Larkum, M. E. (2020). Cellular mechanisms of conscious processing. *Trends in Cognitive Sciences*.

[Bittner et al., 2015] Bittner, K. C., Grienberger, C., Vaidya, S. P., Milstein, A. D., Macklin, J. J., Suh, J., Tonegawa, S., and Magee, J. C. (2015). Conjunctive input processing drives feature selectivity in hippocampal ca1 neurons. *Nature neuroscience*, 18(8):1133–1142.

[Bittner et al., 2017] Bittner, K. C., Milstein, A. D., Grienberger, C., Romani, S., and Magee, J. C. (2017). Behavioral time scale synaptic plasticity underlies ca1 place fields. *Science*, 357(6355):1033–1036.

[Doron et al., 2020] Doron, G., Shin, J. N., Takahashi, N., Drüke, M., Bocklisch, C., Skenderi, S., de Mont, L., Toumazou, M., Ledderose, J., Brecht, M., et al. (2020). Perirhinal input to neocortical layer 1 controls learning. *Science*, 370(6523).

[Gambino et al., 2014] Gambino, F., Pagès, S., Kehayas, V., Baptista, D., Tatti, R., Carleton, A., and Holtmaat, A. (2014). Sensory-evoked ltp driven by dendritic plateau potentials in vivo. *Nature*, 515(7525):116–119.

[Gidon et al., 2020] Gidon, A., Zolnik, T. A., Fidzinski, P., Bolduan, F., Papoutsi, A., Poirazi, P., Holtkamp, M., Vida, I., and Larkum, M. E. (2020). Dendritic action potentials and computation in human layer 2/3 cortical neurons. *Science*, 367(6473):83–87.

[Jadi et al., 2014] Jadi, M. P., Behabadi, B. F., Poleg-Polsky, A., Schiller, J., and Mel, B. W. (2014). An augmented two-layer model captures nonlinear analog spatial integration effects in pyramidal neuron dendrites. *Proceedings of the IEEE*, 102(5):782–798.

[Kampa et al., 2006] Kampa, B. M., Letzkus, J. J., and Stuart, G. J. (2006). Requirement of dendritic calcium spikes for induction of spike-timing-dependent synaptic plasticity. *The Journal of physiology*, 574(1):283–290.

[Kampa et al., 2007] Kampa, B. M., Letzkus, J. J., and Stuart, G. J. (2007). Dendritic mechanisms controlling spike-timing-dependent synaptic plasticity. *Trends in Neurosciences*, 30(9):456–463.

[Larkum, 2013] Larkum, M. (2013). A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. *Trends in Neurosciences*, 36(3):141–151.

[Leinweber et al., 2017] 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.

[Magee and Grienberger, 2020] Magee, J. C. and Grienberger, C. (2020). Synaptic plasticity forms and functions. *Annual review of neuroscience*, 43:95–117.

[Moldwin and Segev, 2020] Moldwin, T. and Segev, I. (2020). Perceptron learning and classification in a modeled cortical pyramidal cell. *Frontiers in Computational Neuroscience*, 14.

[Naud et al., 2014] Naud, R., Bathellier, B., and Gerstner, W. (2014). Spike-timing prediction in cortical neurons with active dendrites. *Frontiers in computational neuroscience*, 8:90.

[Poirazi et al., 2003] Poirazi, P., Brannon, T., and Mel, B. W. (2003). Pyramidal neuron as two-layer neural network. *Neuron*, 37(6):989–999.

[Takahashi et al., 2020] Takahashi, N., Ebner, C., Sigl-Glöckner, J., Moberg, S., Nierwetberg, S., and Larkum, M. E. (2020). Active dendritic currents gate descending cortical outputs in perception. *Nature Neuroscience*, 23(10):1277–1285.

[Takahashi et al., 2016] Takahashi, N., Oertner, T. G., Hegemann, P., and Larkum, M. E. (2016). Active cortical dendrites modulate perception. *Science*, 354(6319):1587–1590.

[Tzilivaki et al., 2019] Tzilivaki, A., Kastellakis, G., and Poirazi, P. (2019). Challenging the point neuron dogma: Fs basket cells as 2-stage nonlinear integrators. *Nature communications*, 10(1):1–14.