Sequence anticipation and STDP emerge from a voltage-based predictive learning rule

Matteo Saponati[a,b,d], Martin Vinck[a,b,d,c]

[a]Ernst-Strüngmann Institute for Neuroscience in Cooperation with Max Planck Society, Frankfurt Am Main, Germany
[b]Donders Centre for Neuroscience, Department of Neuroinformatics, Radboud University Nijmegen, 6525 Nijmegen, the Netherlands
[c]Lead contact
[d]Correspondence: matteo.saponati@esi-frankfurt.de (M.S.), martin.vinck@esi-frankfurt.de (M.V.)

Abstract

Intelligent behavior depends on the brain's ability to anticipate future events. However, the learning rules that enable neurons to predict and fire ahead of sensory inputs remain largely unknown. We propose a plasticity rule based on predictive processing, where the neuron learns a low-rank model of the synaptic input dynamics in its membrane potential. Neurons thereby amplify those synapses that maximally predict other synaptic inputs based on their temporal relations, which provide a solution to an optimization problem that can be implemented at the single-neuron level using only local information. Consequently, neurons learn sequences over long timescales and shift their spikes towards the first inputs in a sequence. We show that this mechanism can explain the development of anticipatory motion signaling and recall in the visual system. Furthermore, we demonstrate that the learning rule gives rise to several experimentally observed STDP (spike-timing-dependent plasticity) mechanisms. These findings suggest prediction as a guiding principle to orchestrate learning and synaptic plasticity in single neurons.

Introduction

Predicting the future is pivotal in guiding interactions with the world, for example in reward learning [1, 2] and in action planing [3]. For a system to predict the future, it is critical to give credit to those inputs that are informative about future events, and discard those inputs that are expected and redundant. That is, predicting the future entails that a system can anticipate and signal events ahead of time. These requirements are computationally challenging, as associations must be learned between sequences of events that are temporally distant, and these associations must give rise to specific consequences for synaptic plasticity and neuronal outputs. Consistent with Predictive Coding models, it is known that the predictability of sensory events can evoke different neuronal signals, in particular enhanced firing rates for surprising information, which then leads to signal transmission and updates of predictions in other brain areas [4, 5, 6, 7]. However, according to the predictability of sensory events, there also need to be specific patterns of synapse formation in order to associate events that predict each other. In particular, one would expect that those synaptic inputs that are carrying information about the future receive high credit, whereas those synaptic inputs that are redundant and predicted are downregulated, which we conceptualize as a form of predictive plasticity. Credit to the synaptic inputs that convey predictive information about the future should lead to a shift of neuronal firing ahead in time. As a result, neurons would be able to anticipate and signal future events that are predicted, which can then lead to adaptive behavior of the organism. Eventually, predictive relationships between events must lead to plasticity formation at the level of a single neuron, which receives a limited set of inputs. However, it remains unclear how synaptic plasticity formation in individual neurons relates to predictive processing.

Experimental evidences suggest numerous and complex synaptic plasticity mechanisms for single neurons, e.g. heterosynaptic plasticity [8, 9, 10], spike-timing-dependent plasticity (STDP) [11, 12, 13] and homeostatic plasticity [14, 15]. These mechanisms indicate that synaptic adjustment is sensitive to the relative firing times of pre-synaptic inputs, the temporal relation between pre- and post-synaptic firing, and that neurons can simultaneously orchestrate plasticity at multiple synapses. This greatly enriches the computational capabilities of neurons [16] and may underlie the biological substrate for association of events across long temporal sequences. Yet, a computational understanding of how these plasticity processes may contribute to prediction of the future has yet to be reached.

We hypothesized that predictive plasticity may account for the existence of learning processes inside individual neurons, allowing to learn temporal sequences and anticipate future events. We recapitulate this predictive mechanism as a spiking neuron model where the cell anticipates future input by learning a low-dimensional model of its high-dimensional synaptic inputs. On the basis of this principle we derive a local learning rule and we show how single neurons can learn to anticipate and recall sequences over long timescales. Furthermore, we show that different synaptic plasticity rules observed experimentally can be interpreted as mechanisms for the evaluation of prediction errors in time, as they emerge naturally from the proposed learning rule.

Results

Model of prediction at the single neuron level

We developed a single neuron model which, at each moment in time t, integrates the present pre-synaptic inputs in the current state of the membrane potential and extracts from its dynamics a prediction of the future input states (Fig 1a). We formalized the proposed predictive process as follows:

First, we defined the membrane potential v_t as a linear filter, such that the neuron updates its membrane potential recursively by encoding the actual input at time t and the previous value of the membrane potential at time $t-1$ (Fig 1a, left). In fact, the membrane potential at a given time is the result of the temporal summation of previous synaptic inputs and, as such, it encodes a low-dimensional compression of the high-dimensional input dynamics. The temporal integration of the pre-synaptic inputs is weighted by a weight vector $\vec{w}$ giving different credit to different synapses. Besides the integration of synaptic inputs on the timescale of the membrane constant, the neuron can produce action potentials, which are nonlinear operations enriching its computational capabilities [17]. We accounted for this by setting a spiking threshold as in Integrate-and-Fire (I&F) models [18]: if the membrane potential reaches the threshold at the timestep $t-1$ the cell fires a post-synaptic spike and voltage gets reset to its resting value at the next timestep (Fig 1a, left).

Next, the neuron computes a local prediction at the same timestep t: the prediction of the future pre-synaptic input is given by the weight of the associated synapse and the previous state of the membrane potential (see Methods) which controls the sign and amplitude of plasticity as observed experimentally [19, 20, 21] (Fig 1a, right). We derived a predictive learning rule analytically by minimizing the mismatch between the actual input and the prediction. This mismatch can be interpreted as a prediction error which can be computed with local synaptic states and in real-time with the dynamics of the inputs (see Methods). Consequently, synaptic weights undergo potentiation/depotentiation proportionally to the unpredictability/predictability of their inputs. Thus, a synapse gets potentiated if the associated input anticipates other pre-synaptic inputs and it gets suppressed if the related input can be predicted by the past.

On a long timescale, the neuron learns a specific set of synaptic strengths by adjusting the synaptic weight continuously as it collects evidence in its membrane potential, in real-time with the input dynamics (Fig 1a). In the following section, we will show that the neuron fires for the most predictive inputs which are indeed maximally potentiated, thereby anticipating the occurrence of the predictable part of a temporal input pattern. We refer to the Methods section for a detailed account of the model and for the analytical derivation of the learning rule.

Dynamics of anticipation and predictive plasticity

The emission of a post-synaptic spike and the consequent reset of the membrane potential gives sensitivity to the timing relations among pre-synaptic inputs and also relative to the post-synaptic spikes. Let us consider a case in which a subset of pre-synaptic inputs in an incoming sequence triggers a post-synaptic spike before the end of the sequence. The output spike influences the evaluation of future predictions as the dynamics encoded in the membrane potential gets reset to the resting potential (Fig 1a). The model suggests the hypothesis that post-synaptic spikes are timed decisions to anticipate predictable synaptic inputs and thereby lead to predictive processes happening at the single neuron level. To test this mechanism, we exposed the model to a sequence of two input spikes coming from by two different pre-synaptic neurons having a relative delay of 4 ms (Fig 1b). In this simple scenario, the first pre-synaptic input is predictive of the following pre-synaptic input and should thus be potentiated, driving the neuron to fire ahead of the EPSP caused by the second input spike.

We trained the model by repeating the input pattern for 300 epochs of duration $T = 500$ ms. During the training period, the model learns to adjust its output spike time and to eventually fire ahead of the input coming at 6 ms (Fig 1b). The model learns this anticipatory solution by selective adjustment of the synaptic weights (Fig 1c, left). In particular, the neuron gives credit to the input coming at 2 ms (solid line) and suppresses the input coming at 6 ms (dashed line). The symmetry of the weights dynamics is broken when the membrane potential reaches the threshold and an output spike is elicited. Accordingly, this follows in the anticipation of the predictable input spike (Fig 1c, right). In fact, the anticipatory solution is a consequence of the constrained latent space given by the threshold. The membrane potential v is constrained to a specific set of values and its reset influences the dynamics of the weight vector. The parameter space given by (w_1, w_2) is thus partitioned in different regions depending on the amount of output spikes elicited (Fig 1d). This partition gives rise to symmetry breaking: the model converges to the fixed point representing high credit to the input with predictive power and suppression of the predicted part. The learning dynamics are qualitatively the same when the initial conditions lie in regions of multiple output spikes (Supplementary Figure S1).

Anticipation of spikes for simple input patterns has been shown in neural networks with specific architecture and endowed with a Temporal Difference (TD) learning algorithm [22]. Our results show that even a single neuron can anticipate the predictable spikes by learning the temporal association between events and selectively potentiating or suppressing synaptic inputs. The model can anticipate the predicted inputs independently from specific initial values of the synaptic weights. Anticipatory firing is obtained as the weight vector evolves following the proposed learning rule (see Fig 1d and Methods).

Anticipation of spike patterns in single neurons

The simple case of two pre-synaptic inputs firing one spike in a deterministic temporal patterns suggests that in the predictive plasticity model, the neuron can predict future inputs and adjust anticipatory signals. However, neurons in the brain receive many hundreds of synaptic inputs where temporal pattern are hidden in high-dimensional synaptic inputs and embedded in background, stochastic firing [23]. Spatio-temporal spike patterns can precisely repeat over time [24], but their detection and eventual anticipation *in-vivo* is a computationally

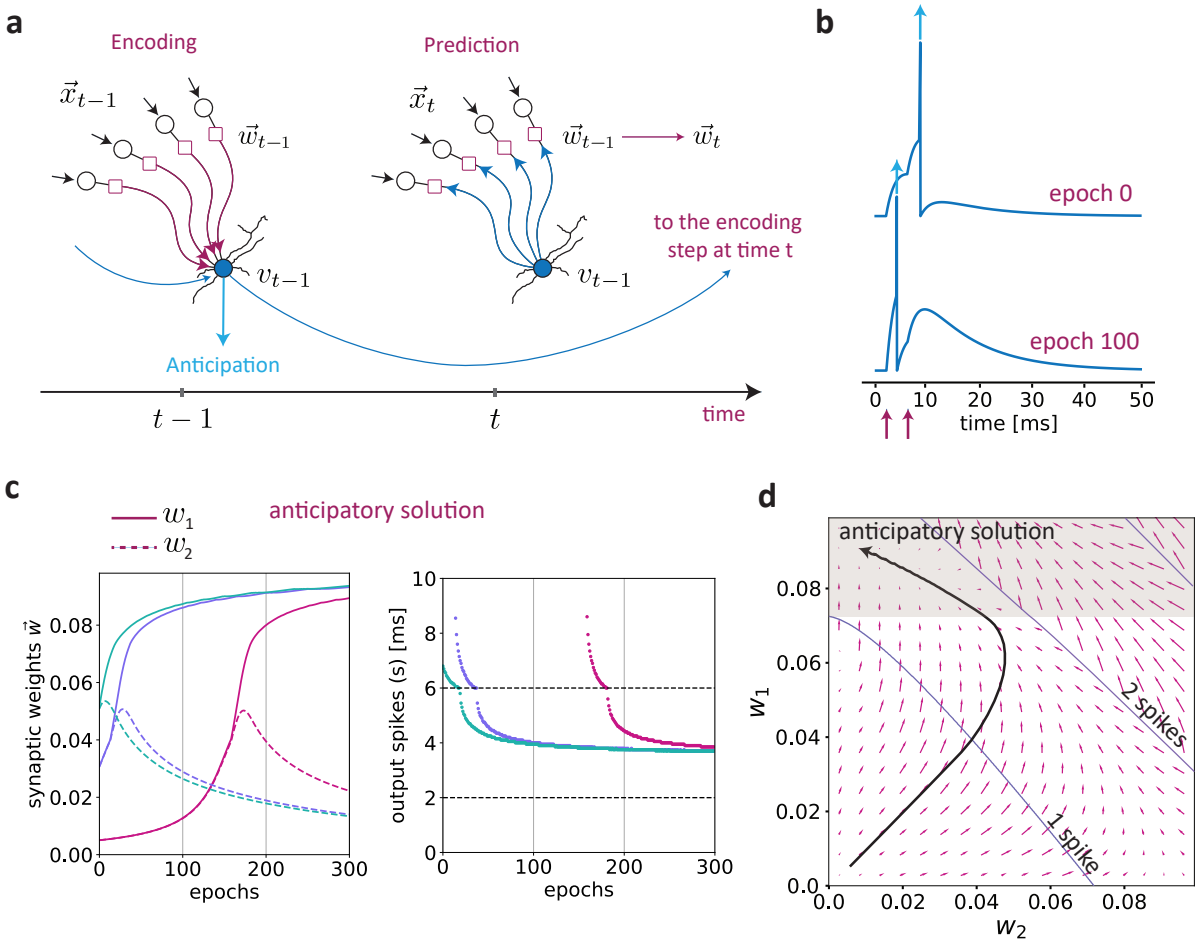

Figure 1: Description of the predictive plasticity rule. **a**) Illustration of the model and of the relative computational graph. At the timestep $t-1$ the neuron updates the its membrane potential by encoding the actual input $\vec{x}_{t-1}$ through the synaptic weight vector $\vec{w}_{t-1}$ and its previous internal state. If the voltage exceeds the threshold, an output spike is emitted (light blue arrow) and this event reset the membrane voltage to the resting potential at the next timestep (see Methods). Otherwise, the value of the membrane potential is kept and passed to the next iteration. Next, at the timestep t, the neuron decodes a prediction of the new input $\vec{x}_t$ from the previous membrane potential v_{t-1} and synaptic weight vector $\vec{w}_{t-1}$. The mismatch between the present input and its prediction is computed locally at every synapse. The prediction error is then used to drive synaptic plasticity and update the synaptic weight vector $\vec{w}_{t-1}$. Afterwards, the neuron updates again its membrane potential by encoding the actual input $\vec{x}_t$ via the learned weight vector $\vec{w}_t$ and its previous internal state to update the prediction for the next timestep. See Methods and for a detailed description of the model and the related computational flow graph. **b**) In the numerical simulations we considered a pattern of two pre-synaptic spike from two different pre-synaptic neurons with a relative delay of 4 ms. Specifically, the two pre-synaptic spike times are 2 ms and 6 ms. Dynamics of the membrane potential v at the first training epoch and after 100 iterations: the neuron learns to fire ahead of the input coming at 6 ms. **c**) Right: dynamics of the weights for difference initial conditions - the straight (dashed) line correspond to the input at 2 ms (6 ms). Left: evolution of the output spike times s across epochs. The color code corresponds to the dynamics of the weights in the right plot. **d**) Flow field in the parameter space (w_1, w_2) obtained by computing the difference between the weight vector at the first epoch and after 10 epochs. The blue lines represent the partition given by the spiking threshold. The black arrow shows the trajectory of the weights obtained by training the model for 500 epochs with initial conditions $\vec{w}_0 = (0.005, 0.005)$. The shared region shows the section of the parameter space where the neuron fires ahead of the second input ($s < 6$ ms).

expensive problem. We studied if the single neuron model can achieve this hard computational task by adjusting its weight for predictive plasticity.

We considered a temporal sequence determined by N pre-synaptic neurons firing sequentially with fixed delays (Fig 2a). To reproduce firing patterns similar to *in vivo* conditions, we used two types of noise source, namely jitter of spike times in the sequence and random firing of the pre-synaptic neurons. The sequences were also embedded into a higher dimensional input pattern where N additional pre-synaptic neurons fired randomly following an homogeneous Poisson process. For each epoch the population spike density of the pre-synaptic afferents was constant during the sequence presentation as well as during spontaneous activity, indicating that nothing characterized the pattern in terms of firing rate (Fig 2a). In addition, the onset of the input sequence was random during each training epoch, that is there was no implicit time reference in the total input pattern (Fig 2a). The post-synaptic neuron received different realizations of such input pattern for each training epoch. We numerically solved the dynamics and studied the output spike pattern during learning (Fig 2b). During the first presentations of the stimulus, the neuron fires in a random fash-

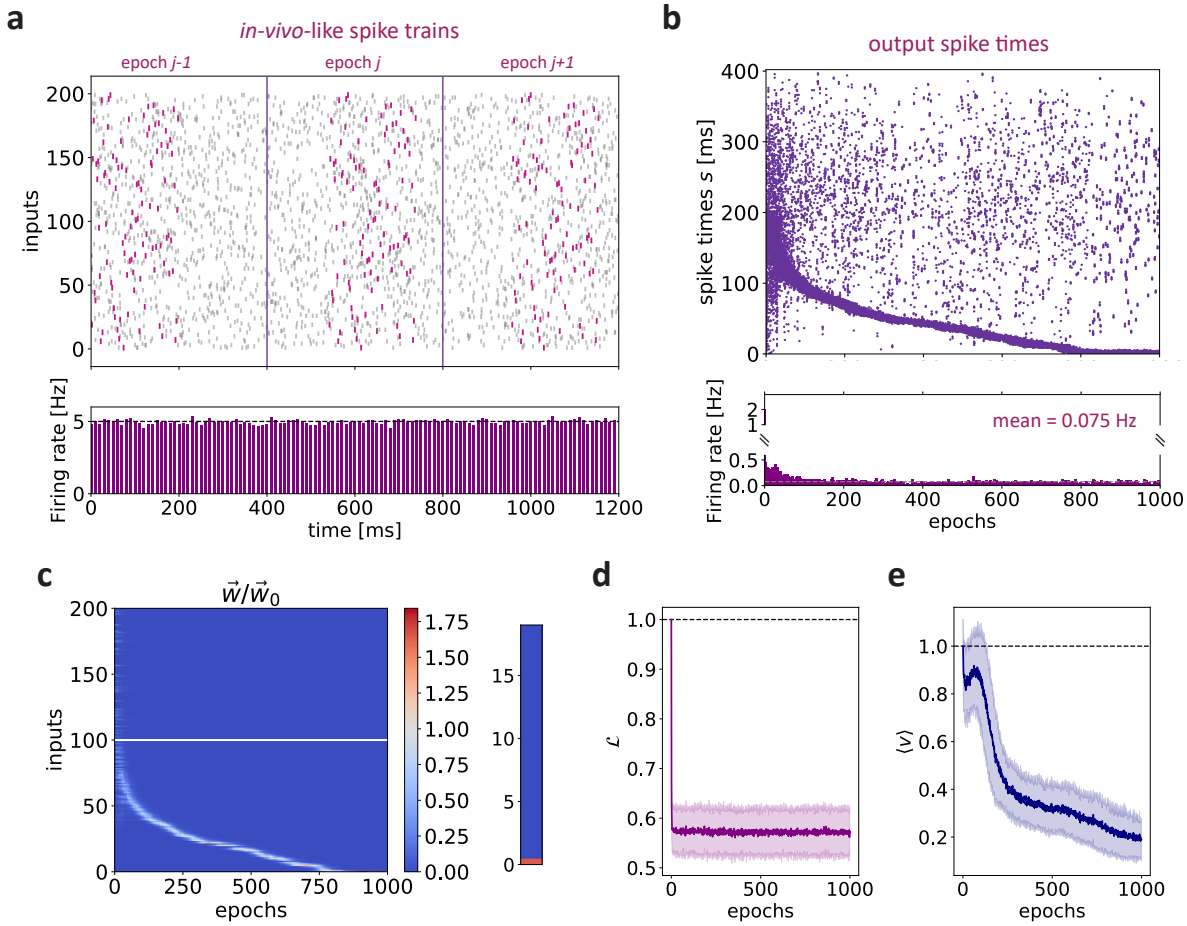

Figure 2: Anticipation of sequences embedded in realistic spike trains. **a**) Example of the input spike pattern during different training epochs. The spike pattern contains a sequence given by the correlated activity of a subset of pre-synaptic neurons. Here we considered a set of 100 pre-synaptic neurons firing one spike per neuron sequentially every 2 ms, resulting in a total sequence length of 200 ms (purple spike pattern). Each example contains two different sources of noise: jitter of the spike times (maximal Δt = 2 ms) and random firing following an homogeneous Poisson process with constant rate λ uniformly distributed between 0 and 10 Hz (see Methods). We also included 100 more pre-synaptic neurons firing randomly according to an homogeneous Poisson processes with randomly distributed rate as before. During each training epoch of duration T = 400 ms the sequence is presented once at random times. The sequence onset times are homogeneously distributed between 0 and 200 ms and they are different for each training epoch. The bottom plots shows the population firing rate averaged over 10 ms time bins (neuron membrane time constant). **b**) Dynamics of the output spike pattern s during learning. Bottom plot shows neuron output firing rate averaged over each training epoch and its mean value across epochs. **c**) Dynamics of normalized synaptic weights $\vec{w}/\vec{w}_0$ as a function of the training epochs. Inputs above the white line are the background pre-synaptic neurons which do not participate in the sequence. The synaptic weights are ordered from 1 to 100 following the temporal order of the sequence. Right side of the panel shows the normalized weights of the first 20 inputs at epoch 1000. At the end of training, the first input is the only one getting maximally amplified. **d**) Normalized objective function $\mathcal{L}$ as a function of the training epochs. The panel shows the mean and standard deviation computed over 100 different simulations. **e**) Normalized cumulative membrane potential $\langle v \rangle$ as a function of epochs. The cumulative membrane potential is computed as the sum of the v_t at each time step in the simulations. The panel shows the mean and standard deviation computed over 100 different simulations.

ion for the entire duration of the sequence. Subsequently, the predictive learning mechanism leads to structured output spike trains and the neuron starts grouping its activity earlier in time, such that it eventually learns to fire for the first inputs in the sequence. During learning, the neuron keeps a low output spike density that represents its selectivity (Fig 2b, bottom plot). As in the case of two pre-synaptic inputs, the anticipation of the pre-synaptic sequence follows from the update of the synaptic weights (Fig 2c). Initially, the neuron gives uniform credit to all the pre-synaptic inputs while firing randomly across the whole sequence. Subsequently, this symmetry is broken as the neuron potentiates the inputs that anticipate the ones driving post-synaptic spikes, eventually favouring the first inputs in the sequence. These results do not depend on the initial weight vector as the anticipatory solution is a convergence point of learning (Fig S2a). Since the learning dynamics follows the direction of reducing prediction error, the objective function $\mathcal{L}$ decreases across epochs (Fig 2d). Further to this, the total amount of depolarization across one stimulus presentation is reduced during learning (Fig 2e). The learned solution of anticipating the input sequence thus decreases the amount of spikes used and the energy consumed by the neuron, consistent with efficient and pre-

dictive coding theories [25, 26]. The predictive neuron model is able to predict and anticipate sequence of different lengths, demonstrating good performances even for longer sequences (Fig S2b) and increased noise amplitude (Fig S2c).

Together, these results show that an Integrate-and-Fire neuron following a predictive principle can learn to anticipate high-dimensional input sequences over short and long timescales. The neuron exploits the timing of each input spike and its temporal context across the spike pattern in a self-supervised manner. The plasticity of a pre-synaptic input is defined by its relative timing and it gets potentiated if, on average, it anticipates successive inputs triggering post-synaptic spikes. The output spiking threshold lets the neuron adjust the potentiation and depression scheme dynamically. This is in agreement with experimental results on the role of the first input spikes for synaptic plasticity and learning of complex spike patterns [27, 28]. The predictive plasticity mechanism relies solely on the temporal relation between inputs, it does not depend on initial conditions and it is robust to several pattern disruptions.

Network prediction signals and recall of learned sequences

Classical theories of predictive coding assume that the generative model belongs to higher order areas along the cortical hierarchy and, by signalling predictions, it acts on the activity of the lower areas [4, 29, 30, 31]. Nonetheless, experimental evidence in the visual cortex of rodents suggests that the anticipation of a repeated sequence occurs already in area V1, suggesting that it could therefore depend on a local reorganization of synaptic weight distributions [32]. Specifically, experimental data suggests that V1 populations learn to reproduce sequences in the visual stimuli depending on the relative timing between inputs, and that the neuronal populations can recall the sequence when parts of the sequence are omitted [32]. Furthermore, it has been shown that specific neurons adapt to anticipate the activity of bottom-up inputs [33], and that the timescale of the recalled sequence is shorter than the actual input length [34]. The output of a neuron anticipating a given sequence can be decisive as other neurons could credit this internally generated signal and fire ahead of similar incoming inputs. This self-organized mechanism can lead to recall of input pattern and, more importantly, to the anticipation of sequences as observed experimentally. We asked if a network following our predictive plasticity protocol can show this type of behaviour without the need of global, top-down predictions.

We explored the dynamics of a network model where N neurons received an input sequence distributed across $2N$ external units and embedded into background random firing (see Methods). The network has a recurrent, all-to-all connectivity scheme and is arranged along the time dimension of the sequence, such that each neuron receives part of the input in a subsequent manner (Fig 3a). Accordingly, the pattern that each neuron tries to predict is composed by the external pre-synaptic input and by the internally generated patterns $\vec{s}$ (Fig 3a). We simulated the protocol in [34] and we numerically solved the model when only part of the sequence was shown, before and after training. Before learning, only the first neurons in the network can respond to the first part of the input sequence (Fig 3b, top). During the first training epochs, the network shows a symmetrical output pattern spanning the whole sequence (Fig 3b, center). Yet, after learning, the network can reproduce the recall of the sequence as observed in V1 [34] when only the first part of the input is shown (Fig 3b, bottom). The network can also perform a pattern completion when part of the input is omitted (Fig 3b, bottom) as observed experimentally [32]. Interestingly, the network structured its state space such that background activity can lead to recall of the anticipatory sequence even when the stimulus is not present. The neurons thereby learn to give credit to the previous units along the network since their output firing carries predictive information about the future (Fig 3c). Indeed, only the first neuron potentiated mostly the first of its own external inputs. The internally generated sequence is faster as the neuron firing represents a compressed replay of the input sequence, thereby learning to anticipate the input pattern (Fig 3d).

Thus, these results show that a network of neurons endowed with a predictive learning rule can spontaneously organize to learn and recall sequences as observed experimentally in rodent V1. In [34] the authors reported almost 4-times faster recall of sequences by V1 neurons, in qualitative agreement with our simulations (Fig 3b). Our results suggest that the underlying mechanism of recalling can be understood from the perspective of prediction within the same network, where each neuron aims to anticipate its own input sequence and this process results in the generation of a organized population spike pattern. This population signal lies on a faster manifold than the input sequence and the network can retrieve the sequence also during spontaneous activity, similarly to what has been observed in the hippocampus [35].

Emergence of spike-timing-dependent plasticity

Lastly, we wondered how our predictive learning rule relates to experimentally observed plasticity rules, in particular the spike-time-dependent-plasticity (STDP) rule [12]. In Figure 1, we showed that the predictive learning scheme leads to asymmetric synaptic weights and that credit assignment critically depends on the pre- and post-synaptic spike times. During training, different regions of the parameter space (w_1, w_2) converge to the anticipatory solution at different times (Fig 1c). Accordingly, for a fixed number of training epochs, the final degree of asymmetry between the weights will be different. In the following, we show that the learning rule naturally gives rise to STDP and related experimental phenomena.

To quantify the asymmetry between synaptic weights we defined an asymmetry index and compared its values in the parameter space at different stages of training on the sequence of two deterministic input spikes (Fig 4a). We trained the model by repeating a sequence of two input spikes coming from by two different pre-synaptic neurons having a relative delay of 4 ms for several epochs of duration $T = 500$ ms. After 100 epochs some regions of the parameter space still shows symmetrical weights while others converged to the emission of an anticipatory spike as the region of post-synaptic firing expands. After 300 epochs every initial condition led to the asymmetric solution and thus to a positive value of the asymmetry index.

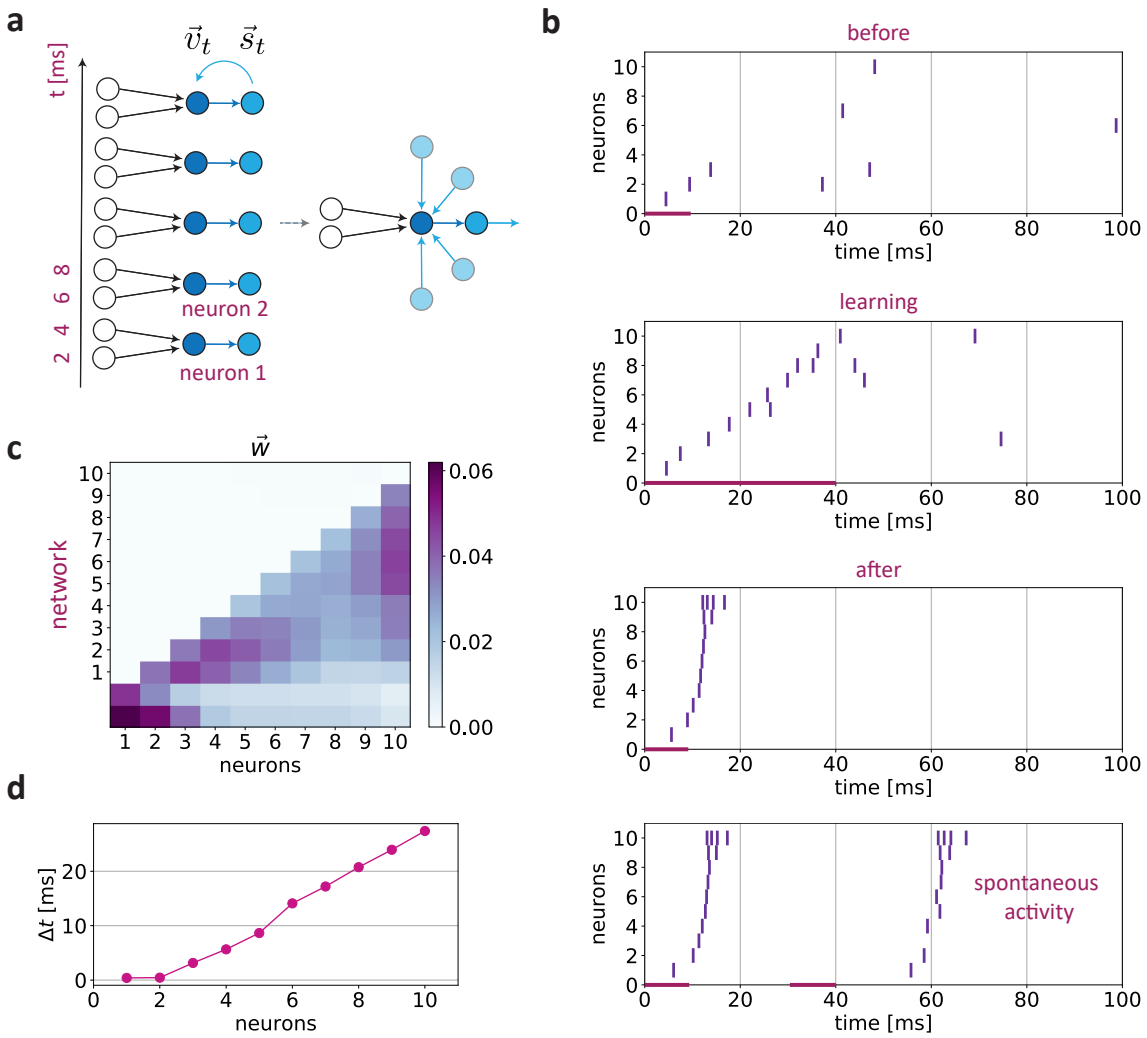

Figure 3: Emergence of anticipatory signals and self-generation of sequences at the network level. **a**) Schematic of the network model. We considered a sequence of 20 input spikes distributed across 20 pre-synaptic inputs and a network of 10 neurons with recurrent, all-to-all connectivity. For simplicity, we only represented 5 neurons in the schematic of the network. The input spikes are allocated evenly along the sequence with a relative spacing of 2 ms and each neuron is connected to 2 pre-synaptic units. We included background activity as each pre-synaptic neuron fires randomly following a Poisson process with constant rate λ drawn uniformly between 0 and 10 Hz and jitter of the spike times (maximal $\Delta t = 2$ ms). For clarity, we sort the neurons following the temporal order of their relative inputs: neuron 1 receives inputs at 2 and 4 ms, neuron 2 receives inputs at 6 and 8 ms and so on. The total input $\vec{x}_i$ to the i-th neuron is given by the pre-synaptic spikes and by the recurrent inputs $\{s_j\}_{j\neq i}$ which are internally generated by the network. **b**) *Top*: output spikes $\vec{s}$ when we presented only part of the input sequence as represented by the red line. The stimulus was composed by the first 6 input spikes to the first 3 neurons in the network. *Middle*: output spikes $\vec{s}$ during the first training epoch. The network output follows the input pattern during the whole duration of the sequence (red line). *Bottom*: the first plot shows the output spikes $\vec{s}$ after 900 training epochs as we presented only the first 6 input spikes, as before. The second plot shows the output spikes $\vec{s}$ as we presented only the first 6 and the the last 4 input spikes (red line). Note that the random background firing can also elicit spontaneous recall of the sequence. **c**) Matrix of synaptic weights where each column represent the weight vector learned by the respective neuron. The last two elements are the synaptic weights of the relative external inputs while the others are the recurrent connections listed by the relative position of the pre-synaptic neurons. **d**) Difference between the output spike times of the network at the first epoch of training and after training.

During convergence the neuron potentiates/suppresses input spikes arriving before/after a post-synaptic action potential, for different initial conditions. These mechanisms are at the basis of STDP where the precise temporal relation between pre-synaptic inputs and post-synaptic activity is key for synaptic plasticity [37, 38]. Our results suggest that the STDP observed *in-vitro* could be interpreted as the cell learning relations between pre-synaptic inputs and current injection and adjusting the synaptic weights according to a predictive principle.

To test this hypothesis, we replicated the standard STDP protocol with numerical simulations (see Methods). We trained the predictive plasticity model with a sequence of two input spikes from two different pre-synaptic inputs x_1 and x_2 and a relative delay Δt. The initial conditions were chosen such that x_1 induced a sub-threshold depolarization while the input x_2 triggered a post-synaptic spike, and vice versa. The temporal

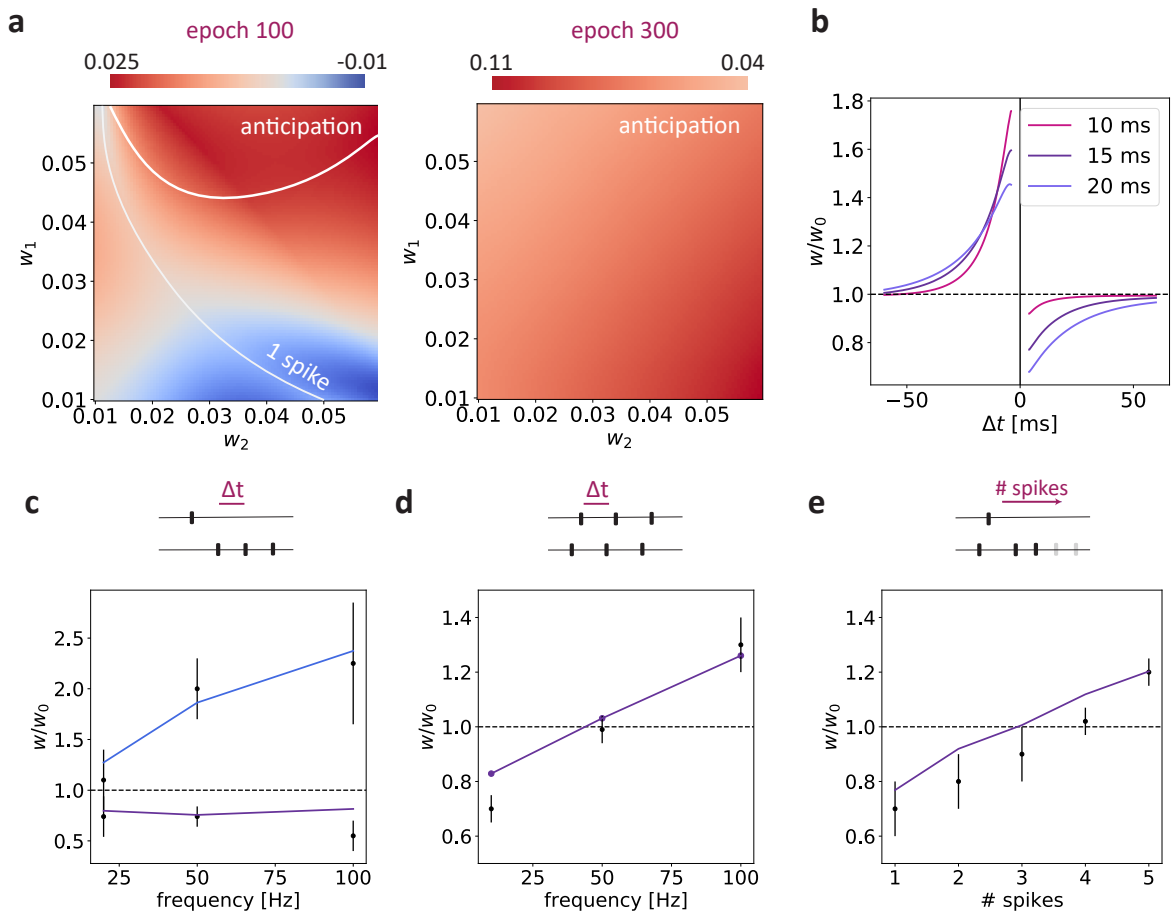

Figure 4: Spike-timing-dependent-plasticity features as an emergent phenomenon of predictive plasticity. **a**) Asymmetry index computed as the difference between the initial weight vector (w_1^0, w_2^0) and the final vector after k epochs: $(w_1^k - w_1^0) - (w_2^k - w_2^0)$. Left: asymmetry index computed with $k = 100$ epochs. The white lines represents the expansion of the one output spike region and the sub-region where the neuron fires ahead of the second input. Right: asymmetry index computed with $k = 300$ epochs. The predictive model converged to the anticipatory solution for each initial condition. **b**) Weight change (in percentage) of the sub-threshold input as a function of the delay between the two input spikes (see Methods). We repeated the same simulation for different values of the membrane time constant τ as reported in the legend. **c**) Weights change (in percentage) as a function of the frequency of post-synaptic bursts (3 post-synaptic spikes per burst). The blue and purple lines refer to pre-post and post-pre pairing, respectively. Data redrawn from Nevian and Sakmann [36], RMS error: 0.146 for pre-post pairing and 0.157 for post-pre pairing. **d**) Weights change induced by increasing the frequency pairing in a post-pre protocol (10 ms inter-spike-interval). Data redrawn from Froemke et al [27], RMS error: 1.017. **e**) Weights change as a function of the number of post-synaptic spike. Data redrawn from Froemke et al [27], RMS error: 0.089. We refer to the Methods for the implementation details of the STDP protocols.

order and the relative delay Δt of the inputs completely defined the patterns. We quantified the total update of the sub-threshold weight as a function of the delay and the order of input spikes (Figure 4b). The model displays an anti-symmetric learning window which depends on the membrane time constant τ. The model exhibits such learning window even though we did not explicitly assume any spike-timing-dependent LTP and LTD rule or any specific learning kernel. In fact, the model can also reproduce a symmetric STDP window observed experimentally [39] (Figure S3a) as well as a LTD-based symmetric STDP window [40] (Figure S3b) just by varying the model parameters and the initial conditions for the weight vector $\vec{w}$. This would be difficult to achieve with phenomenological learning rules which are specifically designed to reproduce asymmetric Hebbian [37, 41, 42]. The specific shape of the learning window can be Hebbian or anti-Hebbian depending on model parameters, without losing its predictive property.

In addition to the classical dependence on the relative delay between inputs, many other pre- and post-synaptic factors can influence the sign and amplitude of synaptic plasticity [28, 43, 44]. The nonlinear and history-dependent interactions in STDP are particularly relevant when neurons receive complex input spike patterns as in *in vivo* scenarios. We consider some of those complex STDP protocols to test if we can reproduce these nonlinear effects by means of the same learning rule:

(1) Experimental evidence indicates that the frequency of post-synaptic bursts after a pre-synaptic input can boost LTP while LTD remains unchanged [36]. We quantified the total weight change after training the model on a burst-dependent protocol equivalent to the one used in [36]. The same effect of post-synaptic bursting on synaptic weight changes is quantitatively reproduced by the model (Figure 4c).

(2) The frequency of pre- and post-synaptic spike pairings can influence the plasticity and convert LTD to LTP for high frequency bursts [20, 27, 28]. The predictive plasticity model quantitatively reproduces the dependence on the frequency pairing as observed in [27] (Figure 4d).

(3) Finally, Froemke et al showed that adding more post-synaptic spikes after a post-pre pairing can convert LTD into LTP (see Figure 6 in [27]). We tested the model on such multi-spike protocol and observed that it can reproduce quantitatively the transition measured in the experiments (Figure 4e).

Altogether, we showed that the predictive learning model can reproduce several linear and nonlinear STDP features both qualitatively and quantitatively. We stress that the model is not of a phenomenological type, i.e. it is not directly designed to reproduce experimental data. The mechanistic explanation for our results relies on the constrains caused by the spiking threshold (Fig 1). The neuron potentiates the inputs which anticipate the one triggering an action potential while the inputs arriving after a post-synaptic spike are suppressed, in agreement with STDP. The amplitude of plasticity depends on how well the membrane potential can capture the relation between the inputs, leading to the exponential shape of the learning window as a function of the time delay between inputs.

Discussion

The process of anticipation requires a system to learn the temporal relations between events and to act ahead of predicable inputs as the temporal sequence unfolds in time. The task of labeling temporal events as predictive or predictable increases in complexity with their relative delays. Our model supports the interpretation of synaptic plasticity as a means to associate temporal events by controlling the future effects of pre-synaptic inputs on the dynamics of the membrane potential. The sign and magnitude of synaptic plasticity is determined by the timing of the pre-synaptic inputs in the context of the current state of the cell: credit assignment is employed as potentiating the synapses which are predictive of future input states, i.e. which give temporal cues about the upcoming inputs (Fig 1). The proposed predictive plasticity mechanism is complementary to the theory that neurons assign credit by means of changes in firing rates as postulated by Predictive Coding theories [4, 6, 45]. Predictive coding is a signal processing strategy based on principles of expectation and prediction, where neurons transmit surprising (predicted) inputs by enhanced (decreased) firing rates [5]. The objective of the proposed predictive learning scheme is not to minimize surprise, but rather to learn a compressed model of the input dynamics, thereby associating events in time and anticipating predictable outcomes. Hence, prediction errors are expressed as intracellular signals (membrane potential fluctuations) and act directly at the synaptic level rather than in the changes of the firing rates of neurons. Thus, cells have direct access to the the result of matching inputs with the predictions provided by fluctuations of the membrane potential. This translates in assigning credit to the synaptic input with most predictive power. The neuron then learns to fire post-synaptic spikes at the beginning of a temporal sequence of pre-synaptic inputs, where the beginning of the temporal sequence represents the non-predictable part of the sequence (Fig 2). The proposed predictive learning scheme is in line with the notion of "coding for prediction", in which neurons primarily transmit information about sensory inputs that carry predictive information about the future, as observed in the retinal neural circuits [46].

There is increasing evidence for the existence of structured sequences in the brain, for example phase precession observed in the hippocampus [47] and phase coding in the visual cortex [48]. Furthermore, it is known that neurons can respond very precisely to repeated temporal input sequences [24]. Importantly, structured neuronal sequences are not only a consequence of regularities in the temporal evolution of sensory inputs, but also of the organized internal activity of the brain [47, 48]. Therefore, to learn associations between events is pivotal to classify, recall and anticipate sequences in a self-supervised manner. Indeed, a recent model [49] suggests that phase precession leads to the learning temporal sequence through an asymmetric learning window as in spike-timing-dependent plasticity (STDP).

The proposed model for synaptic plasticity is of an algorithmic nature. That is, 1) we formalized the core objective of prediction and we derived the resulting algorithm, 2) we found the biological substrate of the model as a local implementation at the single neuron level (see Methods) and 3) we found that the learning rule naturally leads to plasticity mechanisms which are observed experimentally, such as spike-timing-dependent plasticity (STDP) rules (Fig 4). This approach is different from the descriptive level of phenomenological STDP models, which are fundamental to understand underlying biophysical mechanisms, but might not necessarily explain their computational significance. Our results, together with previous studies [22, 50, 51, 52], suggest that STDP is a consequence of a general learning rule given the state of the system, the particular stimulation protocol and the specific properties of the input. We proposed a learning rule in which the sign and amplitude of plasticity is not static, and where the facilitation or suppression of synapses depends on the given state and the given starting point in the parameter space, but nonetheless follows the principle of anticipating predictable inputs (Fig 1). Our results agree with previous studies that have shown that a post-synaptic neuron endowed with standard STDP models can exhibit potentiation of the first synaptic inputs of a temporal pattern, leading to the decrease of latency of the post-synaptic response [53, 54]. Furthermore, our results suggests that this behavior is a consequence of a predictive plasticity rule. A key difference, however, is that the predictive plasticity rule does not produce asymmetric STDP under all conditions, in fact the degree of potentation and depotention in our model depends on the initial state (Fig 4 and Fig S3). That is, there is no fixed STDP kernel in our model. Furthermore, the prediction of future inputs is also driven by the interaction between different synapses, i.e. heterosynaptic plasticity, which has been observed experimentally [9, 10, 55] and proposed as a computationally powerful mechanism for learning in single neurons [56, 16].

The biological substrate of the proposed learning rule requires further experimental investigation. NMDA receptors

could naturally operate as voltage-gated units for prediction errors because: 1) their functions are voltage-dependent, 2) they can explicitly compare the internal state of the neuron and the external inputs [57], both mechanisms which AMPA receptors do not allow. Experimental evidence [58, 59, 60, 61] and theoretical studies [62] support the hypothesis that active ion-channels along dendritic compartments strongly enriches the dynamical repertoire of neurons and underpin higher computational capabilities.

Our work is closely related to a general model of STDP developed by Clopath et al [42] where the authors, guided by experimental evidence [19, 20, 21], studied the role of membrane voltage as the relevant post-synaptic variable for synaptic plasticity. The plasticity model proposed in [42] can accurately reproduce a wide range of experimental findings which, to our knowledge, is not possible with STDP learning rules based on spike timing only [63]. The Clopath rule is based on the two-threshold dynamics observed by Artola et al [19] and the authors assume an Adaptive-Exponential I&F (AeI&F) model for the voltage dynamics [64], together with additional variables for spike afterpotential and adaptive threshold. In agreement with the results of [42], we were able to account for a wide range of phenomena even with a much simpler model of voltage dynamics, suggesting that the history-dependent effect of the membrane potential is pivotal to plasticity (Fig 4).

As a final remark, we believe that a single neuron perspective on prediction and anticipatory mechanisms is important as the implementation of any plasticity rule is ultimately achieved at the neuronal level, thereby guiding behavior a the system level. Nonetheless, single neurons are embedded into networks and different means of communication can lead to more complex learning rules in which the proposed one might be a component. A recent work showed that neural activity in the auditory cortex can be predicted roughly 10-20 ms in advance and that these predictions can be exploited at the single neuron level to achieve high performance in classification tasks [65]. However, prediction in the model of [65] does not happen in a unsupervised manner in time as the proposed method relies on the combination of the single neuron prediction with a supervised teaching signal, a novel implementation of Contrastive Hebbian Learning [66]. Additionally, experimental evidence [67, 44] and theoretical findings [68] show that local learning rules implemented by neurons provide a substrate on which global feedback signals act [69]. This can give rise to complex learning strategies which can account as biological substrates of error backpropagation [70]. For example, [68] have shown that learning in spiking recurrent neural networks can be decomposed into two terms, a global loss, and an eligibility trace which depends on the local state of neurons and results in synaptic weight changes according to local Hebbian plasticity rules. Further investigations on single neurons anticipating local sets of inputs and on their interplay through network structure is key to understand how complex prediction strategies can emerge. Moreover, spatially segregated dendrites and spatio-temporal integration of events along the neuronal compartments could drastically increase the complexity of prediction obtainable at the single neuron level.

Methods

Neuron model

We used a leaky Integrate-and-Fire (LIF) model of the form

$$\tau \frac{dv(t)}{dt} = -v(t) + \vec{w}^\top \vec{x}(t) - \tau v_{th} \sum_j \delta(t - t_j) \tag{1}$$

where $v \in \mathbb{R}$ is the membrane potential, τ is the membrane time constant, $\vec{x}(t) \in \mathbb{R}^N$ is the pre-synaptic input, $\vec{w} \in \mathbb{R}^N$ is the weights vector and v_{th} is the spiking threshold. The sum in the last term runs over all the post-synaptic spike times t_j and $\delta(\cdot)$ is the Dirac delta function. Without loss of generality, we set the resting state of the membrane potential to zero. By defining a timestep h we discretize equation (1) to obtain a recurrence model of the form

$$\begin{cases} v_t = \alpha v_{t-1} - v_{th} s_{t-1} + \vec{w}^\top \vec{x}_t \\ s_t = H(v_t - v_{th}) \end{cases} \tag{2}$$

where $\alpha \equiv 1 - h/\tau$ and $H(\cdot)$ is the Heaviside function. The second term in the membrane potential equation represents the reset due to the spike emission. The variable $s_t \in \{0, 1\}$ takes binary values and indicates the presence or absence of an output spike at timestep t. We set $h = 0.05$ ms in all numerical simulations.

Derivation of the learning rule

The predictive plasticity model postulates that neurons predict future inputs by extracting information from the current state of the membrane potential. We formalized this as an optimization problem in time and we defined the objective function $\mathcal{L}$ as the cumulative error in a given time window T

$$\mathcal{L} \equiv \sum_{t=0}^{T} \mathcal{L}_t \equiv \sum_{t=0}^{T} \frac{1}{2} \left\| \vec{x}_t - v_{t-1} \vec{w} \right\|_2^2 \tag{3}$$

where $\|\cdot\|_2$ is the l_2-norm. The objective is to obtain the minimal difference between the input $\vec{x}_t$ and its prediction via v_{t-1} and $\vec{w}$. We assume that the mismatch is evaluated at each timestep t. The gradient of $\mathcal{L}$ w.r.t. to $\vec{w}$ can be computed in a recursive manner by unrolling the computation via backpropagation-through-time (BPTT) [71]. At each timestep t the exact gradient of $\mathcal{L}$ can be written as the contribution of two terms

$$\vec{\nabla}_{\vec{w}} \mathcal{L} = \sum_{t=0}^{T} \frac{1}{2} \Big(\vec{\nabla}_{\vec{w}} \mathcal{L}_t + \frac{\partial \mathcal{L}_t}{\partial v_{t-1}} \vec{\nabla}_{\vec{w}} v_{t-1} \Big) \tag{4}$$

The first term accounts for the direct effect of a weight change on $\mathcal{L}_t$ while the second for its indirect effect via the membrane potential v_{t-1}.

The first term of the gradient is given by

$$\vec{\nabla}_{\vec{w}} \mathcal{L}_t = -2(\vec{x}_t - v_{t-1} \vec{w}_t) v_{t-1} \tag{5}$$

which propagates the prediction error selectively to each input and scales it by the feedback signal determined by v_{t-1}. This term is equivalent to a time-shifted version of the Oja's rule

[72] where v_{t-1} plays the role of the linear output variable. The second term of equation (4) has a contribution given by the direct effect of v_{t-1} on the prediction

$$\frac{\partial \mathcal{L}_t}{\partial v_{t-1}} = -2(\vec{x}_t - v_{t-1}\vec{w})^\top \vec{w} \tag{6}$$

that is, the sum of the prediction errors weighted by their relative synaptic strengths. To obtain real-time learning we want to comprise the remaining term of equation (4) via dynamical updates. We obtain this by differentiating the dynamical rule of the membrane potential in (2)

$$\vec{\nabla}_{\vec{w}} v_t = \left(\alpha - v_{th}\frac{\partial s_{t-1}}{\partial v_{t-1}}\right)\vec{\nabla}_{\vec{w}} v_{t-1} + \vec{x}_t \tag{7}$$

The first term contains the Jacobian J_t of equation (2)

$$J_t = \frac{\partial v_t}{\partial v_{t-1}} = \alpha - v_{th}\frac{\partial s_{t-1}}{\partial v_{t-1}} \tag{8}$$

which holds the contribution of the linear recurrent term and of the threshold nonlinearity for the spiking output s_{t-1}. Similarly to the adjoint method [73], we define an influence vector $\vec{p}_t \equiv \vec{\nabla}_{\vec{w}} v_t$ such that it obeys the recursive equation

$$\vec{p}_t = J_t\, \vec{p}_{t-1} + \vec{x}_t \tag{9}$$

which is straightforward from equation (7) and gives a forward-pass dynamical update of the gradient [74]. Finally, we define the prediction error $\vec{\epsilon}_t$ at timestep t as

$$\vec{\epsilon}_t \equiv \vec{x}_t - v_{t-1}\vec{w} \tag{10}$$

which defines the sign and amplitude of plasticity. All together, the exact gradient of equation (4) can be written as

$$\vec{\nabla}_{\vec{w}} \mathcal{L} = -\sum_{t=0}^{T}\left[\vec{\epsilon}_t\, v_{t-1} + (\vec{\epsilon}_t^{\top}\vec{w})\,\vec{p}_{t-1}\right] \tag{11}$$

After the end of the period $[0, T]$ the exact gradient can be used to update the weight vector via gradient descent

$$\begin{aligned}\vec{w}_k &= \vec{w}_{k-1} - \eta\vec{\nabla}_{\vec{w}}\mathcal{L} \\ &= \vec{w}_{k-1} + \eta\sum_{t=0}^{T}\left[\vec{\epsilon}_t\, v_{t-1} + (\vec{\epsilon}_t^{\top}\vec{w}_{k-1})\,\vec{p}_{t-1}\right]\end{aligned} \tag{12}$$

where η is the learning rate parameter and the index k represents the k-th iteration during the training period.

In fact, we are interested in an *online* learning rule where the weights update takes part to the dynamics of the model, in real-time with the prediction of the pre-synaptic inputs. This is a typical method in stochastic optimization theory for recursive objective functions [75] and online signal processing [76]. We approximate the learning equation (12) with the current estimate of the gradient

$$\vec{w}_t = \vec{w}_{t-1} - \eta\vec{\nabla}_{\vec{w}}\mathcal{L} \;\simeq\; \vec{w}_{t-1} - \eta\vec{\nabla}_{\vec{w}}\mathcal{L}_t\Big|_{\vec{w}=\vec{w}_{t-1}} \tag{13}$$

where

$$\vec{\nabla}_{\vec{w}}\mathcal{L}_t\Big|_{\vec{w}=\vec{w}_{t-1}} = \vec{\epsilon}_t\, v_{t-1} + (\vec{\epsilon}_t^{\top}\vec{w}_{t-1})\,\vec{p}_{t-1} \tag{14}$$

This approximation make the scheme completely online as it only requires information available at timestep t. Operationally, this approximation is equivalent to neglect the effects of a weight change on all future $\mathcal{L}_{t'}$ where $t' > t$ [68]. The computational steps of the predictive neuron model are as follow:

1) At each time step $t-1$, the current input $\vec{x}_{t-1}$ is encoded by updating the state variables of the neuron v_{t-1} and s_{t-1} following equation (2) (Fig S4a).

2) At the following time step t, the objective function $\mathcal{L}_t$ is evaluated and the prediction error ϵ_t is used to drive plasticity (Fig S4b). The plasticity mechanism updates the estimate $\vec{w}_t$ of the weight vector and this estimate is used to update the membrane voltage v_t and the output spike s_t for future predictions.

Theoretical studies suggests that such stochastic approximation works when η is sufficiently small [76, 74]. In our case, the passive memory capacity of the membrane potential is given by its time constant τ. In the limit of $\eta\tau \ll 1$ the changes in the weights are slow compared to voltage changes and the following relation holds

$$\vec{w}_{t'} \simeq \vec{w}_t \quad \forall\, t' \,:\, |t - t'| < \tau \tag{15}$$

Thus the weight change is negligible and the learning rule is approximately exact in the time window defined by the membrane time constant τ. The timescale separation discussed above holds true in biological neurons as the membrane time constant is in the order of $1 \div 10$ ms while synaptic plasticity happens in the order of $10^2 \div 10^3$ ms, requiring several repetitions of the same stimulation protocol.

We set a soft lower-bound to the weight dynamics, thereby setting the learning rate η to depend linearly on the weight vector $\vec{w}$ and obtaining the learning rule we used throughout the work

$$\vec{w}_t = \vec{w}_{t-1} + \eta\,\vec{w}_{t-1}\left[\vec{\epsilon}_t\, v_{t-1} + (\vec{\epsilon}_t^{\top}\vec{w}_{t-1})\,\vec{p}_{t-1}\right] \tag{16}$$

The rule is composed by a first order correlation term $\vec{x}_t v_{t-1}$, a normalization term $-v_{t-1}^2\vec{w}_{t-1}$ which is needed to stabilize learning [77, 72, 78] as observed experimentally [79, 80, 21] and by an heterosynaptic term $(\vec{\epsilon}_t^{\top}\vec{w}_{t-1})\,\vec{p}_{t-1}$ [9, 10]. Accordingly, prediction of future inputs can be computed at the synaptic level with locally available information.

Jacobian and surrogate-gradient method

The first term in equation (8) allows the gradient to flow at every time step t via the dynamics of the membrane potential v_t. The second term has a discontinuous effect in time - at the moment of the output spikes - and depends on the specific nonlinear function. In the paper we approximated this latter term following the surrogate-gradient method [81]

$$-v_{th}\frac{\partial s_{t-1}}{\partial v_{t-1}} \simeq \gamma\frac{1}{(1 + |v_{t-1} - v_{th}|^2)} \tag{17}$$

where γ is a scaling factor. Wdefined the membrane potential v_t as the output variable and s_t as an hidden variable for the objective function. We considered the case $\gamma = 0$ throughout the paper unless otherwise stated.

External inputs and training

Each input sequence was completely determined by the spike timings of each pre-synaptic inputs. Every pre-synaptic input contained a set of input spike times of amplitude $A_x = 1$ (a.u.). The inputs was convoluted with an exponential kernel with $\tau_x = 2$ ms to replicate the fast dynamics of post-synaptic currents. In Figure 2 and Figure 3 we considered two source of noise in the simulations: 1) The deterministic spikes are shifted by an amount Δt uniformly distributed between 0 and a maximal value (in ms). For each epoch we chose a different Δt for each pre-synaptic input. 2) Each input contained random firing following an homogeneous Poisson process with constant rate λ uniformly distributed between 0 and a maximal value (in Hz). For each epoch we used a different realization of the stochastic Poisson processes.

We trained the model by setting the training window T as twice the duration of the input sequence and by numerically solving the dynamics during each epoch. The sequence was presented at random times following an homogenenous distribution between 0 and the length of the sequence. The model is fully determined by 5 hyperparameters: the timestep h, the membrane time constant τ, the spiking threshold v_{th}, the learning rate η and the input time constant τ_x. To quantify the performance on the robustness tasks, we fixed the number of training epochs, we performed 100 numerical simulation for each condition and we labeled successful simulation based on two criteria: 1) The synaptic weight associated to the first spike in the sequence is bigger than all the other synaptic weights (input selectivity). 2) The output latency is smaller than 2 ms plus the maximal jitter considered (fast anticipation). As for the results in Figure 2, the maximal rate was 10 Hz and maximal jitter was 2 ms for the initial conditions and sequence length tests. For the rate and jitter tests, we neglected the other source of noise, that is spike jitter and background firing were respectively zero.

Spike-timing-dependent plasticity protocols

We used a 2-dimensional input $\vec{x}$ and we simulated classical pre-before-post and post-before-pre pairing as typically performed in STDP experiments [11, 12]. The initial conditions for the weight vector $\vec{w}$ were such that the input x_1 induced a sub-threshold depolarization while input x_2 induced a post-synaptic spike or viceversa, depending on the stimulation protocol (see Figure 4). We numerically solved the dynamics of the model by repeating the input pattern 60 times as usual in the experimental protocols. Those simulation were repeated as the temporal order and the relative delay of pre-synaptic spikes Δt was changed. We computed the weight changes of the sub-threshold input as as the ratio of the related synaptic weight before and after each simulation protocol.

We reproduced the same burst-dependent pairing protocol as in [36] by decreasing the delay Δt between post-synaptic spikes (Figure 4c). The bursts were triggered by three input spikes and the pairing was repeated 60 times with 1 s delay between repetitions as in the experiments.

We simulated the 1-n protocol as in [27] by pairing post-pre-post inputs with a 100 Hz burst frequency. The number of input spikes from the supra-threshold input was increased systematically as in the experimental protocol.

Acknowledgments

This was supported by an ERC Starting Grant (SPATEMP) and a BMF Grant (Bundesministerium fuer Bildung und Forschung, Computational Life Sciences, project BINDA, 031L0167). We thank Wolf Singer and Andreas Bahmer for helpful comments, edits, and insightful discussions.

Authorship contributions

Conceptualization: MS, MV. Mathematical analysis: MS, MV. Simulations: MS. Writing: MS, MV. Supervision: MV.

Declaration of Interests

The authors declare no competing interests.

References

[1] Richard S Sutton and Andrew G Barto. "Toward a modern theory of adaptive networks: expectation and prediction." In: *Psychological review* 88.2 (1981), p. 135.

[2] Richard S Sutton. "Learning to predict by the methods of temporal differences". In: *Machine learning* 3.1 (1988), pp. 9–44.

[3] David J Heeger. "Theory of cortical function". In: *Proceedings of the National Academy of Sciences* 114.8 (2017), pp. 1773–1782.

[4] Rajesh PN Rao and Dana H Ballard. "Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects". In: *Nature neuroscience* 2.1 (1999), pp. 79–87.

[5] Karl Friston. "The free-energy principle: a unified brain theory?" In: *Nature reviews neuroscience* 11.2 (2010), pp. 127–138.

[6] Andre M Bastos et al. "Canonical microcircuits for predictive coding". In: *Neuron* 76.4 (2012), pp. 695–711.

[7] Beren Millidge, Anil Seth, and Christopher L Buckley. "Predictive Coding: a Theoretical and Experimental Review". In: *arXiv preprint arXiv:2107.12979* (2021).

[8] Gary S Lynch, Thomas Dunwiddie, and Valentin Gribkoff. "Heterosynaptic depression: a postsynaptic correlate of long-term potentiation". In: *Nature* 266.5604 (1977), pp. 737–739.

[9] Sébastien Royer and Denis Paré. "Conservation of total synaptic weight through balanced synaptic depression and potentiation". In: *Nature* 422.6931 (2003), pp. 518–522.

[10] Won Chan Oh, Laxmi Kumar Parajuli, and Karen Zito. "Heterosynaptic structural plasticity on local dendritic segments of hippocampal CA1 neurons". In: *Cell reports* 10.2 (2015), pp. 162–169.

[11] Henry Markram et al. "Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs". In: *Science* 275.5297 (1997), pp. 213–215.

[12] Guo-qiang Bi and Mu-ming Poo. "Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type". In: *Journal of neuroscience* 18.24 (1998), pp. 10464–10472.

[13] Sen Song, Kenneth D Miller, and Larry F Abbott. "Competitive Hebbian learning through spike-timing-dependent synaptic plasticity". In: *Nature neuroscience* 3.9 (2000), pp. 919–926.

[14] Gina G Turrigiano and Sacha B Nelson. "Homeostatic plasticity in the developing nervous system". In: *Nature reviews neuroscience* 5.2 (2004), pp. 97–107.

[15] Tara Keck et al. "Synaptic scaling and homeostatic plasticity in the mouse visual cortex in vivo". In: *Neuron* 80.2 (2013), pp. 327–334.

[16] Friedemann Zenke, Everton J Agnes, and Wulfram Gerstner. "Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks". In: *Nature communications* 6.1 (2015), pp. 1–13.

[17] Wolfgang Maass. "Networks of spiking neurons: the third generation of neural network models". In: *Neural networks* 10.9 (1997), pp. 1659–1671.

[18] Wulfram Gerstner and Werner M Kistler. *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press, 2002.

[19] Alain Artola, S Bröcher, and Wolf Singer. "Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex". In: *Nature* 347.6288 (1990), pp. 69–72.

[20] Per Jesper Sjöström, Gina G Turrigiano, and Sacha B Nelson. "Rate, timing, and cooperativity jointly determine cortical synaptic plasticity". In: *Neuron* 32.6 (2001), pp. 1149–1164.

[21] Per Jesper Sjöström and Michael Häusser. "A cooperative switch determines the sign of synaptic plasticity in distal dendrites of neocortical pyramidal neurons". In: *Neuron* 51.2 (2006), pp. 227–238.

[22] Rajesh PN Rao and Terrence J Sejnowski. "Spike-timing-dependent Hebbian plasticity as temporal difference learning". In: *Neural computation* 13.10 (2001), pp. 2221–2237.

[23] Alfonso Renart et al. "The asynchronous state in cortical circuits". In: *science* 327.5965 (2010), pp. 587–590.

[24] Jean-Marc Fellous et al. "Discovering spike patterns in neuronal responses". In: *Journal of Neuroscience* 24.12 (2004), pp. 2989–3001.

[25] Horace B Barlow et al. "Possible principles underlying the transformation of sensory messages". In: *Sensory communication* 1.01 (1961).

[26] Eero P Simoncelli and Bruno A Olshausen. "Natural image statistics and neural representation". In: *Annual review of neuroscience* 24.1 (2001), pp. 1193–1216.

[27] Robert C Froemke et al. "Contribution of individual spikes in burst-induced long-term synaptic modification". In: *Journal of neurophysiology* (2006).

[28] Robert C Froemke, Dominique Debanne, and Guo-Qiang Bi. "Temporal modulation of spike-timing-dependent plasticity". In: *Frontiers in synaptic neuroscience* 2 (2010), p. 19.

[29] Peter Dayan et al. "The helmholtz machine". In: *Neural computation* 7.5 (1995), pp. 889–904.

[30] Karl Friston. "A theory of cortical responses". In: *Philosophical transactions of the Royal Society B: Biological sciences* 360.1456 (2005), pp. 815–836.

[31] Andy Clark. *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press, 2015.

[32] Jeffrey P Gavornik and Mark F Bear. "Learned spatiotemporal sequence recognition and prediction in primary visual cortex". In: *Nature neuroscience* 17.5 (2014), pp. 732–737.

[33] Aris Fiser et al. "Experience-dependent spatial expectations in mouse visual cortex". In: *Nature neuroscience* 19.12 (2016), pp. 1658–1664.

[34] Shengjin Xu et al. "Activity recall in a visual cortical ensemble". In: *Nature neuroscience* 15.3 (2012), pp. 449–455.

[35] Kamran Diba and György Buzsáki. "Forward and reverse hippocampal place-cell sequences during ripples". In: *Nature neuroscience* 10.10 (2007), pp. 1241–1242.

[36] Thomas Nevian and Bert Sakmann. "Spine Ca2+ signaling in spike-timing-dependent plasticity". In: *Journal of Neuroscience* 26.43 (2006), pp. 11001–11013.

[37] Wulfram Gerstner et al. "A neuronal learning rule for sub-millisecond temporal coding". In: *Nature* 383.6595 (1996), pp. 76–78.

[38] Henry Markram, Wulfram Gerstner, and Per Jesper Sjöström. "Spike-timing-dependent plasticity: a comprehensive overview". In: *Frontiers in synaptic neuroscience* 4 (2012), p. 2.

[39] Rajiv K Mishra et al. "Symmetric spike timing-dependent plasticity at CA3–CA3 synapses optimizes storage and recall in autoassociative networks". In: *Nature communications* 7.1 (2016), pp. 1–11.

[40] Jiang-teng Lu et al. "Spike-timing-dependent plasticity of neocortical excitatory synapses on inhibitory interneurons depends on target cell type". In: *Journal of Neuroscience* 27.36 (2007), pp. 9711–9720.

[41] Abigail Morrison, Markus Diesmann, and Wulfram Gerstner. "Phenomenological models of synaptic plasticity based on spike timing". In: *Biological cybernetics* 98.6 (2008), pp. 459–478.

[42] Claudia Clopath et al. "Connectivity reflects coding: a model of voltage-based STDP with homeostasis". In: *Nature neuroscience* 13.3 (2010), p. 344.

[43] Daniel E Feldman. "The spike-timing dependence of plasticity". In: *Neuron* 75.4 (2012), pp. 556–571.

[44] Wulfram Gerstner et al. "Eligibility traces and plasticity on behavioral time scales: experimental support of neohebbian three-factor learning rules". In: *Frontiers in neural circuits* 12 (2018), p. 53.

[45] Georg B Keller and Thomas D Mrsic-Flogel. "Predictive processing: a canonical cortical computation". In: *Neuron* 100.2 (2018), pp. 424–435.

[46] Stephanie E Palmer et al. "Predictive information in a sensory population". In: *Proceedings of the National Academy of Sciences* 112.22 (2015), pp. 6908–6913.

[47] William E Skaggs et al. "Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences". In: *Hippocampus* 6.2 (1996), pp. 149–172.

[48] Martin Vinck et al. "Gamma-phase shifting in awake monkey visual cortex". In: *Journal of neuroscience* 30.4 (2010), pp. 1250–1257.

[49] Eric Torsten Reifenstein, Ikhwan Bin Khalid, and Richard Kempter. "Synaptic learning rules for sequence learning". In: *Elife* 10 (2021), e67171.

[50] Geoffrey Hinton. "How to do backpropagation in a brain". In: *Invited talk at the NIPS'2007 deep learning workshop*. Vol. 656. 2007.

[51] Harel Z Shouval, Samuel S-H Wang, and Gayle M Wittenberg. "Spike timing dependent plasticity: a consequence of more fundamental learning rules". In: *Frontiers in computational neuroscience* 4 (2010), p. 19.

[52] Yoshua Bengio et al. "STDP as presynaptic activity times rate of change of postsynaptic activity". In: *arXiv preprint arXiv:1509.05936* (2015).

[53] Timothée Masquelier, Rudy Guyonneau, and Simon J Thorpe. "Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains". In: *PloS one* 3.1 (2008), e1377.

[54] Timothée Masquelier, Rudy Guyonneau, and Simon J Thorpe. "Competitive STDP-based spike pattern learning". In: *Neural computation* 21.5 (2009), pp. 1259–1276.

[55] Rachel E Field et al. "Heterosynaptic plasticity determines the set point for cortical excitatory-inhibitory balance". In: *Neuron* 106.5 (2020), pp. 842–854.

[56] Jen-Yung Chen et al. "Heterosynaptic plasticity prevents runaway synaptic dynamics". In: *Journal of Neuroscience* 33.40 (2013), pp. 15915–15929.

[57] Peter H Seeburg et al. "The NMDA receptor channel: molecular design of a coincidence detector". In: *Proceedings of the 1993 Laurentian Hormone Conference*. Elsevier. 1995, pp. 19–34.

[58] Panayiota Poirazi, Terrence Brannon, and Bartlett W Mel. "Pyramidal neuron as two-layer neural network". In: *Neuron* 37.6 (2003), pp. 989–999.

[59] Alon Polsky, Bartlett W Mel, and Jackie Schiller. "Computational subunits in thin dendrites of pyramidal cells". In: *Nature neuroscience* 7.6 (2004), pp. 621–627.

[60] Michael London and Michael Häusser. "Dendritic computation". In: *Annu. Rev. Neurosci.* 28 (2005), pp. 503–532.

[61] Albert Gidon et al. "Dendritic action potentials and computation in human layer 2/3 cortical neurons". In: *Science* 367.6473 (2020), pp. 83–87.

[62] Ilenna Simone Jones and Konrad Paul Kording. "Might a Single Neuron Solve Interesting Machine Learning Problems Through Successive Computations on Its Dendritic Tree?" In: *Neural Computation* 33.6 (2021), pp. 1554–1571.

[63] Claudia Clopath and Wulfram Gerstner. "Voltage and spike timing interact in STDP–a unified model". In: *Frontiers in synaptic neuroscience* 2 (2010), p. 25.

[64] Romain Brette and Wulfram Gerstner. "Adaptive exponential integrate-and-fire model as an effective description of neuronal activity". In: *Journal of neurophysiology* 94.5 (2005), pp. 3637–3642.

[65] Artur Luczak, Bruce L McNaughton, and Yoshimasa Kubo. "Neurons learn by predicting future activity." In: *bioRxiv* (2020).

[66] Fernando J Pineda. "Generalization of back-propagation to recurrent neural networks". In: *Physical review letters* 59.19 (1987), p. 2229.

[67] Verena Pawlak et al. "Timing is not everything: neuromodulation opens the STDP gate". In: *Frontiers in synaptic neuroscience* 2 (2010), p. 146.

[68] Guillaume Bellec et al. "A solution to the learning dilemma for recurrent networks of spiking neurons". In: *Nature communications* 11.1 (2020), pp. 1–15.

[69] J Yu Angela and Peter Dayan. "Uncertainty, neuromodulation, and attention". In: *Neuron* 46.4 (2005), pp. 681–692.

[70] Timothy P Lillicrap et al. "Backpropagation and the brain". In: *Nature Reviews Neuroscience* 21.6 (2020), pp. 335–346.

[71] Paul J Werbos. "Backpropagation through time: what it does and how to do it". In: *Proceedings of the IEEE* 78.10 (1990), pp. 1550–1560.

[72] Erkki Oja. "Simplified neuron model as a principal component analyzer". In: *Journal of mathematical biology* 15.3 (1982), pp. 267–273.

[73] Yann LeCun et al. "A theoretical framework for back-propagation". In: *Proceedings of the 1988 connectionist models summer school*. Vol. 1. 1988, pp. 21–28.

[74] Ronald J Williams and David Zipser. "A learning algorithm for continually running fully recurrent neural networks". In: *Neural computation* 1.2 (1989), pp. 270–280.

[75] Herbert Robbins and Sutton Monro. "A stochastic approximation method". In: *The annals of mathematical statistics* (1951), pp. 400–407.

[76] Bin Yang. "Projection approximation subspace tracking". In: *IEEE Transactions on Signal processing* 43.1 (1995), pp. 95–107.

[77] Elie L Bienenstock, Leon N Cooper, and Paul W Munro. "Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex". In: *Journal of Neuroscience* 2.1 (1982), pp. 32–48.

[78] Richard Kempter, Wulfram Gerstner, and J Leo Van Hemmen. "Hebbian learning and spiking neurons". In: *Physical Review E* 59.4 (1999), p. 4498.

[79] Gina G Turrigiano et al. "Activity-dependent scaling of quantal amplitude in neocortical neurons". In: *Nature* 391.6670 (1998), pp. 892–896.

[80] Gina G Turrigiano. "The self-tuning neuron: synaptic scaling of excitatory synapses". In: *Cell* 135.3 (2008), pp. 422–435.

[81] Emre O Neftci, Hesham Mostafa, and Friedemann Zenke. "Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks". In: *IEEE Signal Processing Magazine* 36.6 (2019), pp. 51–63.

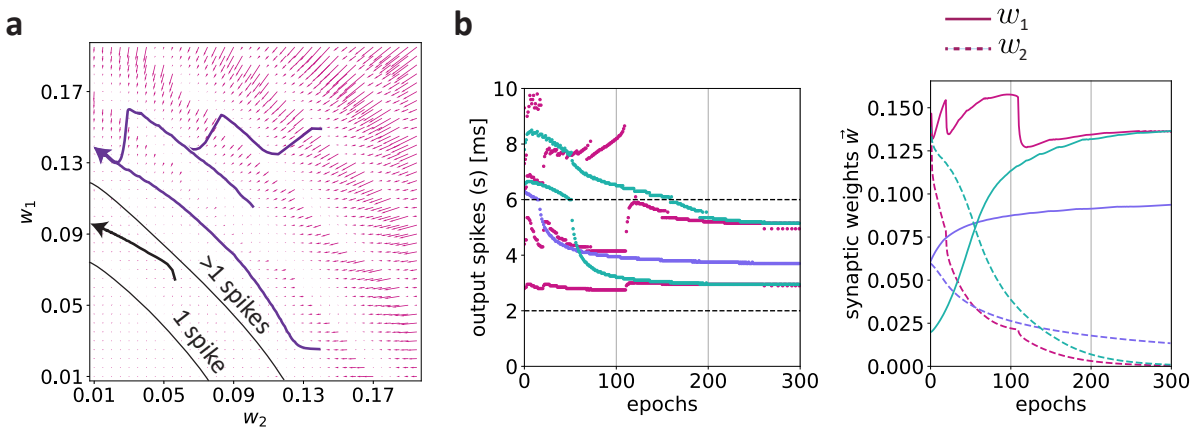

Figure S1: Dynamics in the parameter space. **a**) Flow field in the parameter space (w_1, w_2) obtained by computing the difference between the weight vector at the first epoch and after 10 epochs. The blue lines represent the partition given by the spiking threshold, in particular the region of 1 and more than 1 post-synaptic spikes. The black arrow shows the trajectory of the weights obtained by training the model with different initial conditions, in particular $\vec{w}_0$ = (0.06,0.06), $\vec{w}_0$ = (0.1,0.1), $\vec{w}_0$ = (0.15,0.02), $\vec{w}_0$ = (0.15,0.15). All the regions of parameter space with 2 or more spikes converge to a fixed point which maximally credits w_1 and depotentiates w_2. **b**) Right: dynamics of the weights for difference initial conditions - the straight (dashed) line correspond to the input at 2 ms (6 ms). The purple, red and green line corresponds to the first, second and third initial condition, respectively. Left: evolution of the output spike times s across epochs. The color code corresponds to the dynamics of the weights in the right plot.

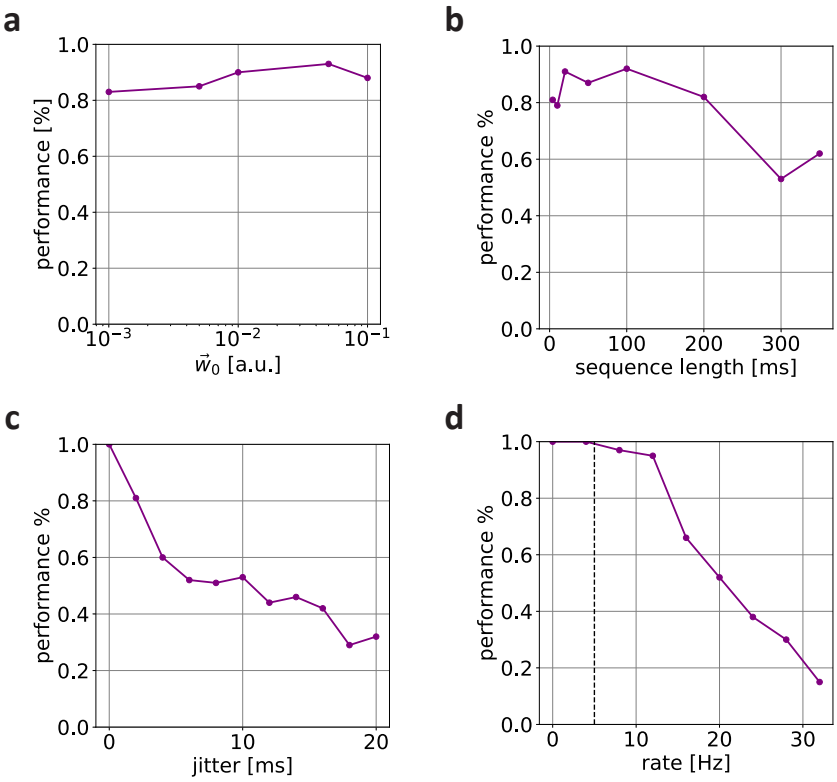

Figure S2: Performance metrics. To quantify the performance we fixed the number of training epochs (1000 epochs), we performed 100 numerical simulation for each condition and we labeled successful simulation based on the criteria of input selectivity and fast anticipation (see Methods). **a**) Performance as a function of the initial weight vector $\vec{w}_0$. The performance do not strongly depend on the initial conditions as the anticipatory solution is a fixed point in the parameter space. **b**) Performance as a function of the sequence length. We performed the simulations as follows: given a sequence length T we considered $T/2$ input neurons each one firing one spike with a relative delay of 2 ms and $T/2$ background neurons (see Methods). **c**) Performance as a function of the maximal spike time jitter. **d**) Performance as a function of the maximal rate of background firing. The dotted line represent the intrinsic firing rate of the sequence, where each neuron fires 1 spike in a 200 ms time window (rate = 5 Hz).

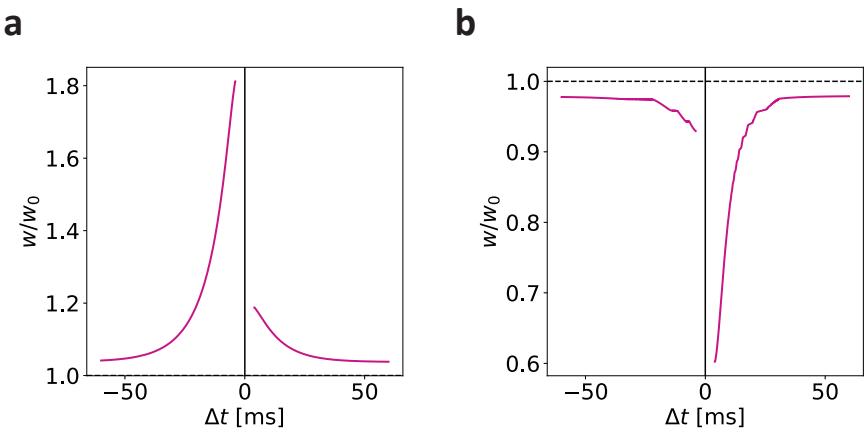

Figure S3: Different spike-timing-dependent plasticity windows. **a**) Weight change (in percentage) of the sub-threshold input as a function of the delay between the two input spikes (see Methods). We obtained the learning window by setting $\tau = 10$ ms, $v_{th} = 1.5$ and initial conditions $\vec{w}_0 = ((0.06, 0.01)$ for the pre-post pairings and $\vec{w}_0 = (0.01, 0.06)$ for the post-pre pairings. **b**) Weight change (in percentage) of the sub-threshold input as a function of the delay between the two input spikes (see Methods). We obtained the learning window by setting $\tau = 10$ ms, $v_{th} = 2$ and initial conditions $\vec{w}_0 = ((0.02, 0.1)$ for the pre-post pairings and $\vec{w}_0 = (0.1, 0.02)$ for the post-pre pairings.

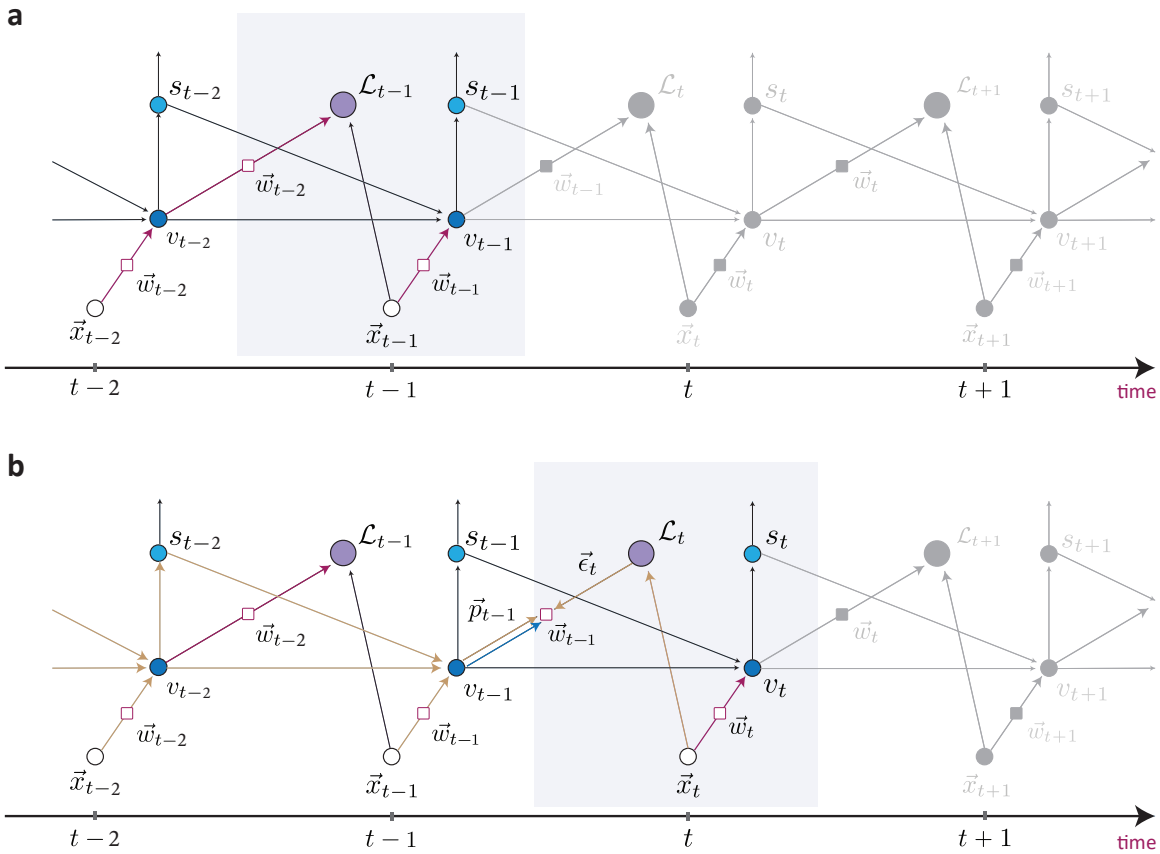

Figure S4: Computational flow graph. **a**) At timestep $t-1$ the neuron encodes the input $\vec{x}_{t-1}$ via the current estimate of the weight vector $\vec{w}_{t-1}$ and it updates its state variables following equation (2). **b**) At timestep t the neuron computes a prediction of the current input state $\vec{x}_t$ via the membrane potential v_{t-1}. The current objective function $\mathcal{L}_t$ is computed via the previous estimate of the weight vector $\vec{w}_{t-1}$. Next, the weight vector is updated by combining the prediction error $\vec{\epsilon}_t$ made at the current timestep with the influence vector $\vec{p}_{t-1}$ from the recursive part of the gradient (9) (golden arrows). The weight vector is updated in real-time with the update of the neuron state variables v_t and s_t.