Sequence anticipation and spike-time-dependent-plasticity emerge from a predictive learning rule

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

[a]*Ernst-Strüngmann Institute for Neuroscience in Cooperation with Max Planck Society, Frankfurt Am Main, Germany*
[b]*IMPRS for Neural Circuits, Max-Planck Institute for Brain Research, Frankfurt Am Main, Germany*
[c]*Donders Centre for Neuroscience, Department of Neuroinformatics, Radboud University Nijmegen, 6525 Nijmegen, the Netherlands*
[d]*Correspondence: matteo.saponati@esi-frankfurt.de, martin.vinck@esi-frankfurt.de*

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 signalling and recall in a recurrent network. 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 planning [3]. Predicting future states entails that a system can anticipate and signal events ahead of time. Indeed, there is evidence for anticipatory neural activity in various brain systems [4, 5, 6, 7, 8, 9, 10]. Furthermore, the predictability of sensory events can evoke different neuronal signals, in particular enhanced firing rates for surprising inputs, which may guide the update of model predictions in other brain areas [11, 12, 10, 13]. Yet, the associations among sensory events and their predictability should not only result in specific patterns of neural activity, but should also have specific consequences for synaptic plasticity and neuronal outputs [14].

In particular, one would expect that synaptic inputs that carry much information about the future receive high credit, whereas those synaptic inputs that are redundant and predicted by other inputs are downregulated. We conceptualize this credit assignment as a form of predictive plasticity. As a consequence of credit assignment to predictive synaptic inputs, neurons might learn to anticipate and signal future events that are predicted, which can then lead to adaptive behavior of the organism. Importantly, predictive relationships between events must eventually 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 [15, 16, 17], spike-timing-dependent plasticity (STDP) [18, 19, 20] and homeostatic plasticity [21, 22]. These experimental studies have shown 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. These plasticity mechanisms greatly enrich the computational capabilities of neurons [23] and may underlie the biological substrate for the association between events across long temporal sequences. They may further account for the observation that repeated sequential activity is associated with subsequent recall or replay of sequences at compressed time scales. Yet, a computational understanding of how these plasticity processes may contribute to the prediction of the future has to be reached.

We hypothesized that predictive plasticity may account for the existence of learning processes inside individual neurons, allowing neurons to learn temporal sequences and anticipate future events. We recapitulate this predictive mechanism as a spiking neuron model, where the cell anticipates future inputs by learning a low-dimensional model of its high-dimensional synaptic inputs. Based on this principle, we derive a local predictive learning rule. We show how single neurons can learn to anticipate and recall sequences over long timescales, and that the described learning rule gives rise to several experimentally observed STDP mechanism.

Results

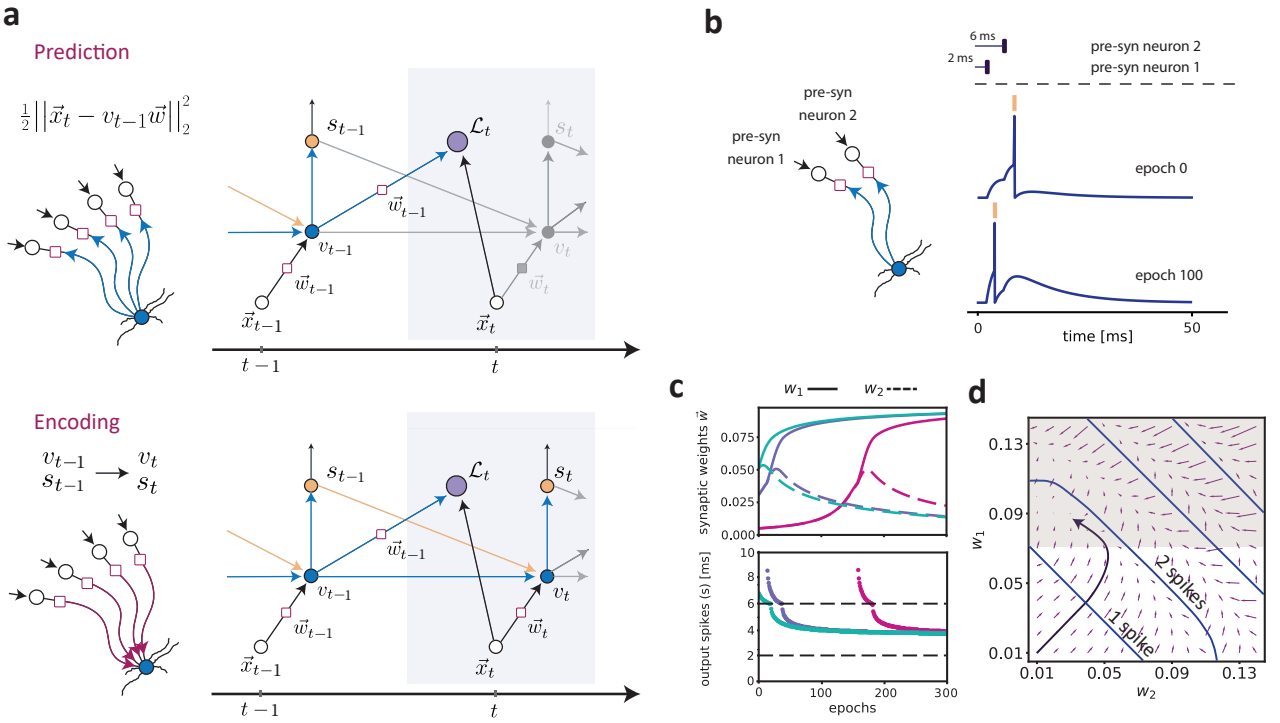

Figure 1: Description of the predictive plasticity rule. **a**) Illustration of the model and the computational graph corresponding to the learning algorithm. Top: at time step t, the neuron computes a prediction of the new input $\vec{x}_t$ from the previous membrane potential v_{t-1} and synaptic weight vector $\vec{w}_{t-1}$ (see Equation (2)). The mismatch between the present input and its prediction is computed locally at every synapse. The prediction error is used to drive synaptic plasticity and update the synaptic weight vector $\vec{w}_{t-1}$ (see Equation (3)). Bottom: the neuron updates its membrane potential by encoding the actual input $\vec{x}_t$ via the learned weight vector $\vec{w}_t$ and its previous internal state v_{t-1} (see Equation (1)). If the voltage exceeds the threshold, an output spike is emitted (shown in yellow) and this spiking event reduces the membrane potential by a constant value at the next time step. Otherwise, the value of the membrane potential v_t is kept and passed to the next time step. **b**) In the simulation illustrated here, we considered a pattern of two pre-synaptic spikes from two different pre-synaptic neurons with a relative delay of 4 ms. Shown are the dynamics of the membrane potential at the first training epoch and after 100 iterations. The neuron learns to fire ahead of the input that arrives at 6 ms (i.e. pre-syn neuron 2). **c**) Top: Dynamics of the weights for different initial conditions (i.e. the weights at epoch 0). The unbroken and dashed lines correspond, respectively, to the pre-synaptic inputs arriving at 2 ms (w_1, pre-synaptic neuron 1) and 6 ms (w_2, pre-synaptic neuron 2). Bottom: evolution of the output spike times across epochs. The bottom and top plot have the same color code. **d**) The flow field in the parameter space was obtained by computing the difference between the weight vector (w_1, w_2) in the first epoch and after 10 epochs. The blue lines represent the partition given by the number of spikes that are fired. Note that when the synaptic weights are larger, the neuron fires more spikes. 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 shaded region shows the section of the parameter space where the neuron fires ahead of the input at 6 ms from neuron 2.

Model of prediction at the single neuron level

We formalized the proposed predictive process in the following single-neuron model: In this model, at each moment in time t, the neuron 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 (see the Methods section for a detailed account of the model and analytical derivations). We first 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 $t-1$ (Equation (1)). The membrane potential at a given time is the result of the temporal summation of previous synaptic inputs and the membrane potential thereby encodes a compression of the high-dimensional input dynamics in time. This is described by the system of equations

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

Here, the temporal integration of the inputs $\vec{x}_t$ is weighted by a synaptic weight vector $\vec{w}_t$, which gives different credit to different synapses. Together with the recurrent dynamics of the membrane voltage, we set a spiking threshold in Equation (1). Accordingly, if the membrane potential reaches the threshold at time step $t-1$, the cell fires a post-synaptic spike s_t and the voltage is decreased by v_{th} at the next timestep.

The objective of the neuron is to recursively compute a local prediction of its own inputs by using the temporal relations in the input spike trains. The prediction of the incoming pre-synaptic input at time step t is given by the weight of the associated synapse and the previous state of the membrane potential, i.e.

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

We then 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 based on the dynamics of the inputs (see Methods). By letting the synaptic weights $\vec{w}_t$ evolve in real-time with the dynamics of the input, we obtained our predictive plasticity rule

$$\vec{w}_t = \vec{w}_{t-1} + \eta \Big[\vec{\epsilon}_t \, v_{t-1} + (\vec{\epsilon}_t^\top \vec{w}_{t-1}) \, \vec{p}_{t-1} \Big] . \tag{3}$$

Here, η defines the timescale of plasticity, $\vec{p}_{t-1}$ is an input-specific eligibility trace (see Methods), and $\vec{\epsilon}_t$ is the prediction error

$$\vec{\epsilon}_t \equiv \vec{x}_t - v_{t-1} \vec{w} \,, \tag{4}$$

which defines the sign and amplitude of plasticity. Consequently, synaptic weights undergo potentiation or depotentiation depending on the predictability of the inputs. Thus, a synapse gets respectively potentiated or suppressed if the associated input anticipates or is anticipated by other pre-synaptic inputs.

The computational steps of the predictive neuron model are as follows: (1) At each time step t, the objective function $\mathcal{L}_t$ is evaluated as the neuron learns to predict the current input (Figure 1a, top); (2) the prediction error ϵ_t is used to drive plasticity and update the synaptic weights via Equation (3), and the current input $\vec{x}_t$ is encoded by updating the state variables of the neuron (Figure 1a, bottom). The rule is composed of three terms: (1) A first-order correlation term $\vec{x}_t v_{t-1}$; (2) a normalization term $-v_{t-1}^2 \vec{w}_{t-1}$, which stabilizes learning [24, 25, 26] as has been observed experimentally [27, 28, 29]; (3) an heterosynaptic plasticity term $(\vec{\epsilon}_t^\top \vec{w}_{t-1}) \, \vec{p}_{t-1}$ [16, 17]. Accordingly, the prediction of future inputs can be computed at the synaptic level in the point-neuron approximation based only on locally available information. 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.

To illustrate the development of anticipatory firing for a simple example, we exposed the neuron to a sequence of two input spikes coming from two different pre-synaptic neurons that fire with a relative delay of 4 ms (Figure 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 (excitatory post-synaptic potential) 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 neuron learns to adjust its output spike time and to eventually fire ahead of the pre-synaptic input 2, which arrives at 6 ms (Figure 1b). The neuron converges onto an anticipatory "solution" by a selective adjustment of the synaptic weights (Figure 1c, top). In particular, the neuron assigns credit to the pre-synaptic input 1, which arrives at 2 ms, and depotentiates the strength of the input arriving at 6 ms. Accordingly, this leads to the anticipation of the predictable input (Figure 1c, bottom). We further observed that the parameter space given by (w_1, w_2) is partitioned in different regions depending on the amount of spikes fired by the post-synaptic neuron (Figure 1d). The symmetry of the weight dynamics is broken when the membrane potential reaches the threshold and an output spike is fired (Figure 1d). The learning dynamics are qualitatively the same when the initial conditions lie in regions of multiple output spikes (Supplementary Figure S1).

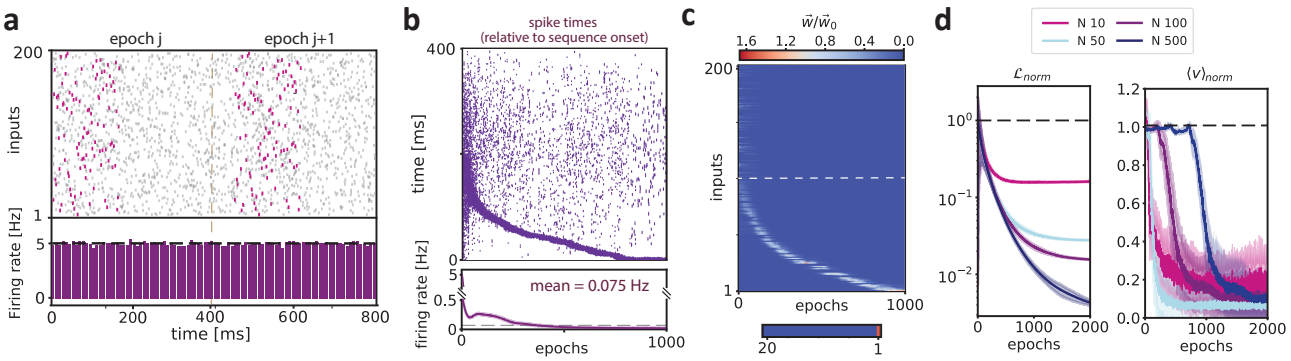

Figure 2: Anticipation of spiking sequences. **a**) Top: Example spike sequence during different training epochs. A spiking sequence is defined by the correlated activity of a subset ($N = 100$) of pre-synaptic neurons. These N pre-synaptic neurons fire sequentially with relative delays of 2 ms, resulting in a total sequence length of 200 ms (pink spike pattern). In each epoch, there are three different sources of noise: (1) jitter of the spike times (random jitter between -2 and 2 ms); (2) random background firing following an homogeneous Poisson process with rate λ distributed between 0 and 10 Hz (see Methods); (3) another subset of 100 pre-synaptic neurons that fired randomly according to an homogeneous Poisson process with randomly distributed rates between 0 and 10 Hz. For each training epoch, the onset of the spike sequences is drawn from an uniform distribution with values between 0 and 200 ms. The bottom plots shows the population firing rate over 10 ms time bins (neuron membrane time constant). **b**) Dynamics of the post-synaptic spiking activity during learning. The spike times are defined relative to the actual onset of the sequence in each respective epoch. The bottom plot shows the neuron's output firing rate within each training epoch. This firing rate was computed across 100 independent simulations (shown are mean and standard deviation). **c**) Top: Dynamics of the normalized synaptic weights $\vec{w}/\vec{w}_0$ as a function of the training epochs. Here $\vec{w}_0$ is the weight vector in epoch 0. Above the dashed white line are the 100 background pre-synaptic neurons that do not participate in the sequence. The synaptic weights are ordered along the y-axis from 1 to 100 following the temporal order of the sequence. Bottom: normalized weights of the first 20 inputs at epoch 1000, showing only the first input has been assigned credit. **d**) Left: Normalized objective function $\mathcal{L}_{norm}$ (left plot) as a function of the training epochs. Different colors correspond to a different number of neurons participating in the sequence. Right: normalized cumulative membrane potential $\langle v \rangle$. The cumulative membrane potential was computed as the sum of the v_t at each time step in the simulations. The panels shows the mean and standard deviation computed over 100 different simulations.

Prediction of temporal structures in the input spike trains

The simple case of two pre-synaptic inputs in a fixed sequence shown in Figure 1b-d suggests that in the predictive plasticity model, the neuron can predict future inputs and generate anticipatory signals. However, in the brain, neurons receive many hundreds of synaptic inputs, yielding high-dimensional sequences that may be embedded in background, stochastic firing [30]. To investigate a relatively complex scenario, we considered a temporal sequence determined by N pre-synaptic neurons that fire sequentially with fixed delays (Figure 2a). To produce stochastic firing patterns, we used two types of noise sources, 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 according to a homogeneous Poisson process. For each epoch, the population firing rate of the pre-synaptic inputs was constant across time, indicating that the sequence could not be detected based on the population firing rate alone (Figure 2a). In addition, the onset of the input sequence was random during each training epoch (Figure 2a). Because of these sources of noise and jitter, the post-synaptic neuron received different realizations of the input pattern for each training epoch. We numerically solved the learning dynamics and studied the output spike pattern during learning.

We observed that during the first presentation of the stimulus, the neuron fired in a random fashion for the entire duration of the epoch. Subsequently, the predictive learning mechanism led to structured output spike trains, and the neuron started to group its activity earlier in time, such that it eventually learned to fire for the first inputs in the sequence (Figure 2b). During learning, the neuron kept a low output firing rate that reflects its selectivity (Figure 2b, bottom plot). The anticipation of the pre-synaptic pattern is driven by the update of the synaptic weights (Figure 2c). Initially, the neuron assign uniform credit to all of the pre-synaptic inputs, while firing randomly across the entire sequence. Subsequently, the neuron potentiates the inputs that anticipate the ones that are driving post-synaptic spikes, eventually assigning most credit to the first inputs in the sequence. Because the learning dynamics follow the direction of reducing the overall prediction error, the objective function $\mathcal{L}_{norm}$ decreases across epochs (Figure 2d, left). Furthermore, during learning, the total amount of depolarization across one stimulus presentation is reduced (Figure 2d, right).

Further analyses demonstrate that the neuron model was able to predict and anticipate input sequences for a substantial range of model parameters (Figure S2a-b). First, we find that the main results do not depend on the initial weight vector (Figure S3a). Second, we show that anticipatory firing emerges even for longer sequences (Figure S3b) and increased noise amplitude and number of distractors (Figure S3c). Third, the noise source do not qualitatively affect the behavior of the model across training epochs (Figure S4a-d). Finally, we considered the case where the input pattern is composed by different sub-sequences which were spaced in time and belonged to independent subsets of pre-synaptic neurons. We show that the neuron exhibits anticipatory firing also in case of multiple sub-sequences (Figure S5).

Together, these results show that an Integrate-and-Fire-like neuron with a predictive learning rule can learn to anticipate high-dimensional input sequences over short and long timescales. The neuron effectively uses the timing of each input spike and its temporal context across the spike pattern in a self-supervised manner. A synapse gets potentiated if, on average, the corresponding pre-synaptic input anticipates successive inputs that initially trigger post synaptic spikes.

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. The learned solution of anticipating the input sequence thus decreases the number of fired spikes and the energy consumed by the neuron, which can be understood as a form of efficient coding [31, 32].

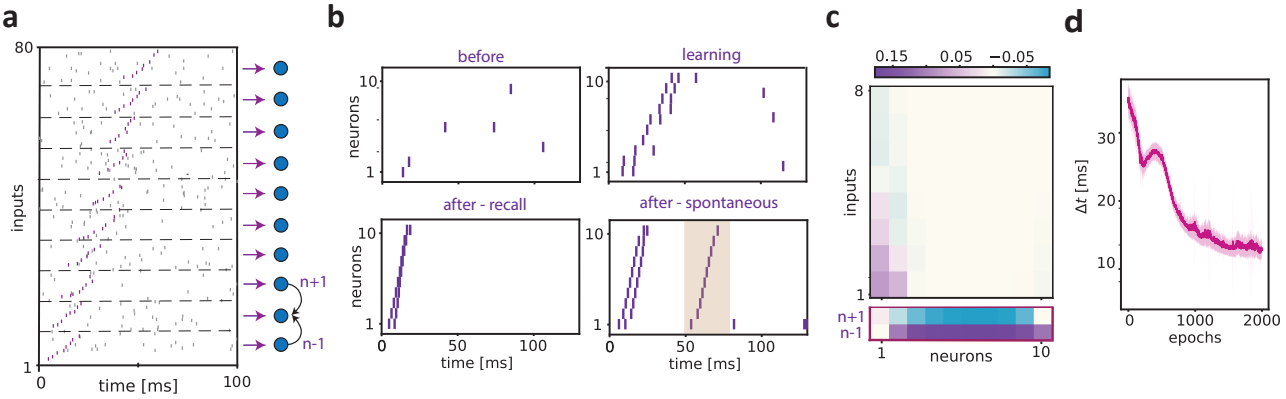

Figure 3: Sequence anticipation and recall in a network with recurrent connectivity. **a**) In this example, we simulated a network of 10 neurons with nearest-neighbour recurrent connectivity, that is each neuron n in the network received inputs from the $n-1$-th and $n+1$-th adiacent neurons. The first and the last neuron only received inputs from the second and second last neuron in the network, respectively. Shown are the connections to the second neuron. Each neurons in the network received inputs from 8 pre-synaptic neurons that fire sequentially with relative delays of 2 ms, resulting in a total sequence length of 16 ms (pink spike pattern). The sequence onset of pre-synaptic inputs for the $n+1$-th neuron started 4 ms after the sequence onset for the n-th neuron in the network, etc. Each epoch contained two different sources of noise: (1) random jitter of the spikes in the sequence (between -2 and 2 ms); random background firing of the pre-synaptic neurons according to an homogeneous Poisson process with rate $\lambda = 10$ Hz. Both the connections from the pre-synaptic neurons to the neurons in the network, as well as the connections between the neurons in the network were plastic and modified according to the predictive learning rule described in the main text. **b**) Raster plot of the network's activity during different epochs of training: (1) The "before" case, where only the pre-synaptic neurons corresponding to the first neuron in the network exhibited sequential firing. In this case the background stochastic firing was still present in all of the $8 \times 10 = 80$ pre-synaptic neurons. (2) The "learning" or conditioning case, where we presented the entire sequence (which was repeated 2000 times). (3) The "after" or "recall" condition, which was the same as the before condition (now after learning). (4) Same as (3), but an example where spontaneous recall occurs due to the background stochastic firing. The neurons are ordered as in panel **a**. **c**) The synaptic weights matrix obtained at the end of training (epoch 1000). Top: The i-th column corresponds to the synaptic weights learned by the i-th neuron in the network, where the 8 entries correspond to the synaptic weights for the pre-synaptic inputs. Bottom: the nearest-neighbour connections in the network towards the i-th neurons. Note that the first and last neuron do not receive inputs from the $n-1$-th and $n+1$-th neuron, respectively. **e**) Evolution of the duration of network activity across epochs. We computed the temporal difference between the last spike of the last neuron and the first spike of the first neuron to estimate the total duration of the network's activity. We computed the average duration and the standard deviation from 100 simulations with different stochastic background firing and random jitter of the spike times.

Sequence anticipation and recall in a network with recurrent connectivity

In the previous section, we studied the emergence of anticipatory firing in a single neuron receiving many pre-synaptic inputs. However, in cortical networks, each neuron may receive a large set of pre-synaptic inputs from other areas, as well as recurrent inputs from neurons in the same local network. We therefore investigated a more complex scenario of a network of recurrently coupled neurons that were each endowed with a predictive

learning rule. Our simulations were inspired by experimental observations of recall and spontaneous replay after learning. For example, a previous study in rat V1, has shown that the repeated presentation of a sequence of flashes (at different retinotopic locations) gradually leads to a reorganization of spiking activity in the order of the presented sequence [7]. The same study also showed that the presentation of only the first stimulus in the sequence leads to a compressed recall of the entire sequence [7]. Likewise, the sequential activation of neurons in prefrontal cortex and hippocampus is known to lead to subsequent replay at a compressed time-scale [33, 34, 35, 36, 37]. These findings have been interpreted in terms of a local reorganization of synaptic weight distributions as a result of repeated activation with an input sequence [7, 33, 34, 37, 36]. We wondered if a network of recurrently connected neurons with the predictive learning rule described above can develop sequence anticipation as well as (stimulus-evoked) sequence recall and spontaneous replay. We explored the dynamics of a network model where 10 neurons received an input sequence distributed across 80 external units. Each neuron in the network received a unique set inputs from 8 pre-synaptic neurons, which fired sequentially and exhibited stochastic background firing (purple and black in Figure 3a, respectively). The neurons in the network were activated sequentially, such that the inputs into the first neuron arrived earliest, the inputs into the second neuron arrived slightly later, etc. (Figure 3a). The network had a recurrent, nearest-neighbour connectivity scheme, such that each n-th neuron was connected to the neighbouring $n-1$-th and $n+1$-th neuron (Figure 3a). Thus, each neuron received a set of "afferent" pre-synaptic inputs together with the inputs from the neighbouring neurons (Figure 3a). Both the synaptic connections from the afferent inputs and the recurrent inputs were adjusted by plasticity according to the predictive learning rule described above.

We show the activity of the network for three cases in Figure 3b: (1) The "before" case, where only the pre-synaptic neurons corresponding to the first neuron in the network exhibited sequential firing. In this case the background stochastic firing was still present in all of the 80 pre-synaptic neurons. (2) The "learning" or conditioning case, where we presented the entire sequence (which is repeated 2000 times). (3) The "after" or "recall" condition, which was the same as the before condition, but after learning. We observed that in the before condition, the network activity was relatively unstructured, with firing occurring in the period after the sequence due to the background stochastic firing. During learning, the neurons were active during a relatively long part of the sequence and showed a sequential activation pattern. After learning, the network showed sequential firing upon the presentation of the inputs to the first neuron in the network in the order of the sequence. This sequential firing took place at a compressed time scale. We found that the recall effect was due to the potentiation of the inputs from the $n-1$-th neuron to the n-th neuron, as well as potentiation of the first pre-synaptic inputs to the first neurons (Figure 3c). Finally, we observed that sequential firing could also be triggered spontaneously due to the background stochastic activity of the pre-synaptic neurons (Figure 3c; after-spontaneous). Thus, the network exhibited a form of activity that resembles spontaneous replay of sequences.

We further characterized the evolution of the network's output during learning, and the reconfiguration of synaptic weights. We found that after several hundreds of epochs, the network converged onto a stable, sequential output that was time-compressed (Figure 3d). We furthermore quantified the number of neurons that needed to be activated in order for the network to recall the full sequence. We found that this required number of neurons decreased gradually across epochs, indicating a gradual reorganization of the synaptic weight distribution during learning (Figure S6).

To generalize these findings, we also studied a network with all-to-all connectivity, i.e. each neuron was recurrently connected to all of the other neurons in the network. In this case, the network also learned to recall the full sequence on a relatively fast timescale (Figure S7). The output of the network with all-to-all connectivity however differed from the example with recurrent connectivity between neighbours: After prolonged learning, the other neurons in the all-to-all network all fired shortly after the first neuron was activated (Figure S7).

Together, these results show that a recurrently connected network of neurons each endowed with a predictive learning rule can spontaneously organize to fire preferentially at the beginning of a sequence, and recall (or replay) sequences at a compressed time scale.

Emergence of spike-timing-dependent-plasticity rules

The results shown above clearly demonstrate that the potentiation of synaptic weights depends on the timing relationships between inputs. This suggests that there may be a connection between the predictive learning rule described here and the experimentally observed spike-time-dependent-plasticity (STDP) rule [19].

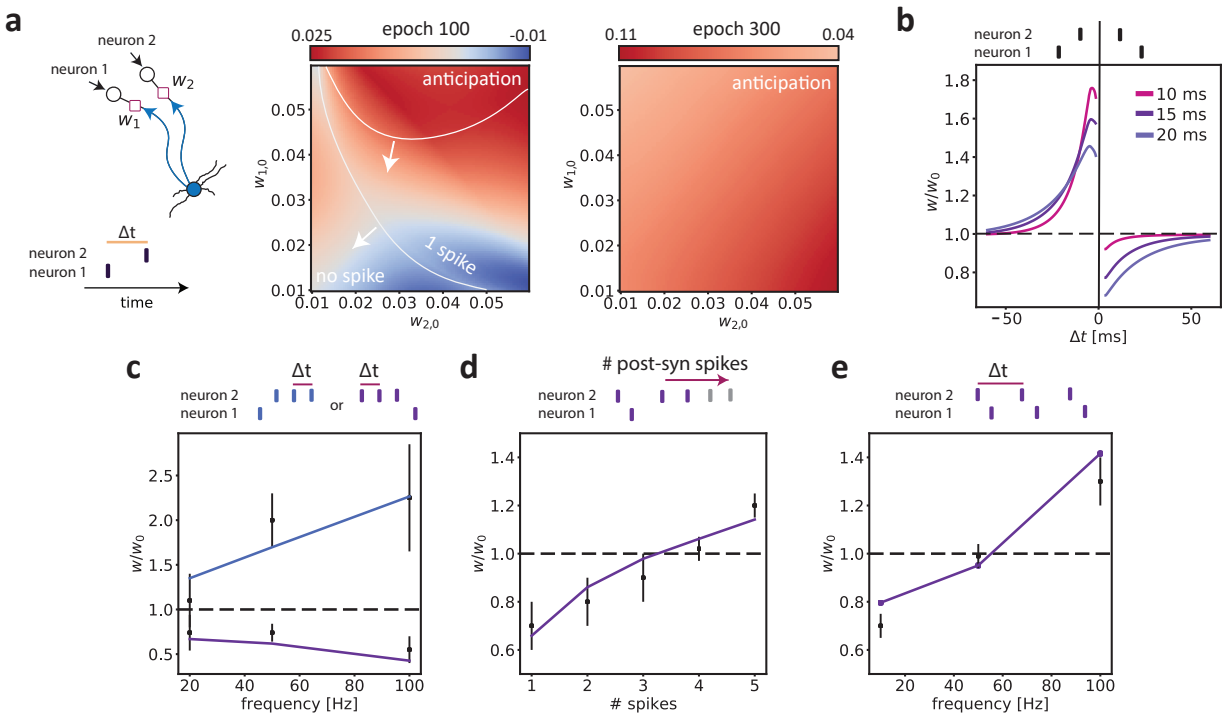

Figure 4: Predictive learning rule gives rise to spike-timing-dependent-plasticity mechanisms. **a**) Left: Illustration of the protocol, where a neuron receives inputs from two pre-synaptic neurons (associated with weights w_1 and w_2) with a delay of $\Delta t = 4$ ms. These inputs were repeated across epochs. Middle and right: Asymmetry index computed as the difference between the initial weight vector $(w_{1,0}, w_{2,0})$ and the final vector after j epochs: $(w_{1,j} - w_{1,0}) - (w_{2,j} - w_{2,0})$. Positive values of the asymmetry index thus indicate that w_1 increases relative to w_2. Shown are the asymmetry index after 100 and 300 epochs, as a function of the initial weights. The white lines divide three regions: (1) No spike; (2) A single spike; (3) A single spike before the second input (i.e. anticipation). The right panel shows, that for all initial weight conditions, the weight of the first input showed a relative increase as compared to the second input. **b**) In order to model classic STDP protocols with current injection, one of the two inputs (pre-synaptic neuron 2) had a strong initial weight, and the other input (pre-synaptic neuron 1) was sub-threshold (i.e. did not evoke a spike). In this simulation, the weights for both inputs could be adjusted via the predictive learning rule (see Figure S10 when the second input has a fixed weight). The y-axis shows the weight change (in percentage relative to the initial weight) of the sub-threshold input (i.e. input 1) as a function of the delay Δt between the two input spikes (see Methods). Negative and positive values of Δt indicate that input 1 preceded or lagged input 2, respectively. Shown are the weight changes for different membrane time constants after 60 epochs. **c**) In this simulation, the second input contained a burst of 3 spikes, which arrived after the first input, and each triggered a spike in the post-synaptic neuron. The input from pre-synaptic neuron 1 only had a sub-threshold effect. Shown is the weights change (as in **b**) versus the firing frequency, i.e. $1/\Delta t$, within the burst (total of 3 spikes per burst). The blue and purple lines refer to the case that input 1 preceded input 2 or lagged input 2, respectively. Experimental data were redrawn from Nevian and Sakmann [38], RMS error: 0.868 for pre-post pairing and 0.206 for post-pre pairing. **d**) Weights change (as in **b**) as a function of the number of spikes in the second input. The inputs from pre-synaptic neuron 2 each triggered a spike in the post-synaptic neuron. The input from pre-synaptic neuron 1 only had a sub-threshold effect. Experimental data were redrawn from Froemke et al [39], RMS error: 0.089. **e**) Weights change (as in **b**) induced by increasing the frequency pairing. Here, the inputs from pre-synaptic neuron 2 always triggered a spike in the post-synaptic neuron, whereas the input from pre-synaptic neuron 1 only had a sub-threshold effect. The inputs from neuron 2 arrived 6 ms before the inputs from neuron 1. Data redrawn from Froemke et al [39], RMS error: 0.057.

To systematically investigate the dependence of potentiation and depotentiation on the timing relationships between pre-synaptic inputs, we considered the simplified case of two inputs (as in Figure 1). In Figure 1, we had shown that the predictive learning scheme leads to asymmetric synaptic weights when two pre-synaptic inputs have different arrival times. To quantify how the asymmetry between the synaptic weights of the first and second input evolved with learning, we defined an asymmetry index (Figure 4a). The asymmetry index was defined as $d_j - d_0$, where d_j was the difference in weights at the j-th epoch $d_j = w_{1,j} - w_{2,j}$, and d_0 the initial difference in weights. Thus, positive values of the asymmetry index indicated that the synaptic weight for the first input became relatively large compared to the synaptic weight for the second input. To illustrate the behavior of the asymmetry index, we trained the model by repeating a sequence of two input spikes coming from two different pre-synaptic neurons having a relative delay of 4 ms for several epochs of duration $T = 500$ ms (as in Figure 1). The parameter space in this case was defined by the initial weights $(w_{1,0}, w_{2,0})$. After 100 epochs, a small region of the parameter space still showed asymmetry indices around zero, while for most

of the parameter space there were positive asymmetry indices, indicating a convergence towards anticipatory activity. After 300 epochs every initial condition led to the asymmetric solution and thus to a positive value of the asymmetry index. Thus, the observation that the first input was potentiated was generally observed for different initial states of the synaptic weights (Figure 4a).

To directly investigate the relation between STDP and the predictive learning rule described here, we investigated the the dependence of plasticity on the relative timing between inputs. To this end, we performed a simulation that resembled the standard STDP protocol (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 arriving at a relative delay Δt. To approximate the STDP protocol with a current injection that triggers a post-synaptic spike, the initial conditions were chosen such that x_2 triggered a post-synaptic spike, and x_1 was a sub-threshold input. A negative and positive delay Δt indicated that x_1 arrived before or after x_2, respectively. We found that the potentiation of the first input was determined by the relative delay Δt (Figure 4b). Specifically, we observed an anti-symmetric learning window with a similar time dependence as has been experimentally observed for STDP [18, 19] (Figure 4b).

This window expanded as a function of the membrane time constant τ_m (Figure 4b). The model exhibited such anti-symmetric learning window even though we did not explicitly implement any spike-timing-dependent LTP and LTD rule with a specific learning kernel.

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 [40, 41, 42]. The nonlinear and history-dependent interactions associated with STDP are especially relevant when neurons receive complex input spiking patterns. We considered some of those complex STDP protocols to test if we can reproduce the nonlinear effects by means of the predictive 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 [38]. To emulate this, we simulated a case where the inputs from pre-synaptic neuron 2 arrived in a burst, with each spike in the burst triggering a post-synaptic spike. We quantified the total weight change after training the model on this protocol. The effect of the intra-burst frequency on the synaptic weight change, as experimentally observed [38], was reproduced by our model (Figure 4c). (2) 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 [39]). We tested the model on such multi-spike protocol and observed that our model can reproduce the transition measured in the experimental data (Figure 4d). (3) The frequency of pre- and post-synaptic spike pairings can influence plasticity and convert LTD to LTP for high-frequency bursts [43, 39, 40]. The predictive plasticity model reproduced the dependence on the frequency pairing as observed in [39] (Figure 4e).

Interestingly, Figure 4a shows that in early training (i.e. repetition) epochs, certain regions of the parameter space can lead to an asymmetric index close to zero. This suggests that the STDP window might have different forms depending on the parameter space and the training epoch, even though the neuron eventually converges onto anticipatory firing. Consistent with this observation, experimental studies have also observed symmetric STDP windows that are either LTP-dominated [44] or LTD-dominated [45]. In Figure S8 we indeed show that the predictive learning rule can, for certain parameter settings, yield a symmetric STDP window that is either LTP- or LTD-dominated. Symmetric STDP windows can emerge even though the neuron eventually does converge onto an anticipatory solution. A key factor that determines the specific shape of the STDP window may be the initial strength of the synapse. In agreement, experimental work has shown that the amount of plasticity in a standard STDP protocol depends on the initial strength of the pre-synaptic weight [43]. To investigate this, we examined the potentiation of the first input depending on its initial synaptic weight. We found that there was a switch from potentiation to depotention as the initial synaptic strength increased (Figure S9), consistent with the experimental observations [43]. To further relate our findings to experimental observations in which a current injection protocol was used, we performed simulations in which we fixed the synaptic weight of the supra-threshold input (i.e., in this case it was not adjusted by plasticity). Also in this case, the model displayed an anti-symmetric STDP kernel (Figure S10).

Altogether, we showed that the predictive learning model can reproduce several linear and nonlinear STDP features.

Discussion

The anticipation of future events is a core feature of intelligence and critical for the survival of the organism. Here, we studied how individual neurons can learn to predict and fire ahead of sensory inputs. We propose a plasticity rule based on predictive processing, where an individual neuron learns a low-rank model of the synaptic input dynamics in its membrane potential. Accordingly, the sign and magnitude of synaptic plasticity are determined by the timing of the pre-synaptic inputs. That is, synapses are potentiated to the degree that they are predictive of future input states, which provides a solution to an optimization problem that can be implemented at the single-neuron level using only local information. We show that neurons endowed with such plasticity rule can learn sequences over long timescales and shift their spikes towards the first inputs in a sequence (i.e. anticipation). Furthermore, neurons represent their inputs in a more efficient manner (i.e. with reduced overall membrane potential). This anticipatory mechanism was able to explain the development of anticipatory signalling and recall in response to sequence stimuli. Finally, we demonstrated that the learning rule described here gives rise to several experimentally observed STDP mechanisms, including: asymmetric STDP kernels [19, 18], as well as symmetric ones [45, 44] given the initial conditions; the frequency-dependence of STDP [39]; the number of post-synaptic spikes in a burst or post-pre pairing [40]; the dependence of (de)potentiation on the initial synaptic strength [43]. Together, our results indicate that prediction may be a guiding principle that orchestrates learning and synaptic plasticity in single neurons, providing a novel interpretation of STDP phenomena.

We first discuss how our results relate to previous theories of coding in cortical networks that emphasize the importance of predictions. An influential theory of cortical function is hierarchical predictive coding (HPC). The basic understanding of HPC is that the brain maintains a model or representation of current and future states in the outside world, and updates this model as new information comes in. HPC posits that the inference process is implemented by the feedforward routing of surprising or unpredicted signals (i.e. prediction errors), and the routing of sensory predictions down the hierarchy via feedback (FB) projections [11, 46, 10]. The predictive plasticity mechanism that we described here differs from HPC models in several aspects, for example: (1) In HPC, prediction is the result of network interactions, in particular the cancellation of feedforward drive by inhibitory feedback. In our model, prediction results from plasticity at a single neuron level. (2) Different from HPC, in our model the neuron does not explicitly transmit (encode) prediction and error signals. (3) Both in HPC and our model, neurons may exhibit reduced firing for predicted as compared to unpredicted sensory inputs. Yet, in our model this is due to depotentiation of predictable inputs, whereas in HPC it is due to inhibitory feedback mediated by top-down projections. We note that our plasticity model is fully compatible with another flavor of predictive processing, namely "coding for prediction". According to this theory, neurons primarily transmit information about sensory inputs that carry predictive information about the future, as observed in the retinal neural circuits [47]. The findings here may also be relevant to understand the development of anticipatory firing in sensory systems [4, 7, 8], temporal difference learning [48, 49], as well as the compression of sequences during resting state based on prior experience [33, 50]. Finally, 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 [51]. However, prediction in the model of [51] does not happen in a unsupervised manner in time as their method relies on the combination of the single neuron prediction with a supervised teaching signal, a novel implementation of Contrastive Hebbian Learning [52].

Next, we discuss how our findings relate to STDP experiments and models, and the biological substrate of the learning rule described here. STDP is an established experimental phenomenon which has been widely observed *in-vitro* [18, 19]. There is evidence for a variety of STDP kernels [41], which dependent on several post-synaptic variables like backpropagating action potentials (bAP) [53], post-synaptic bursts [38] and the dendritic location of inputs [29]. These experimental findings are all obtained in *in-vitro* preparations. Thus it is unclear what the nature of STDP *in-vivo* is. The standard protocol for testing STDP has two major limitations that deviate from the normal physiological setting: (1) The protocol involves current injection in the post-synaptic neuron. The current injection itself is not subject to (physiological) plasticity and might therefore not be a good "proxy" for post-synaptic depolarization induced by natural pre-synaptic inputs *in-vivo*. (2) Several studies have pointed out that different post-synaptic signals (e.g. spike times, depolarization level, dendritic spikes) are relevant for STDP [53, 43, 54]. It is still a manner of debate what is the crucial post-synaptic variable for plasticity [55]. In principle, STDP models might apply both to cases with artificial currents as

well as physiological pre-synaptic inputs. An artificial depolarization caused by current injections can lead to plasticity in both our model and in STDP models. Yet it is an open experimental question what is the nature of learning rules when it comes to physiological synaptic inputs and their timing relationships. For example, it is known that to induce LTP (Long Term Poteniation), it is not necessary to evoke a post-synaptic spike [54, 56]. The learning rule proposed here predicts that, *in-vivo*, pre-synaptic inputs causing a post-synaptic spike will eventually become depotentiated if they are anticipated (i.e. predicted) by other pre-synaptic inputs in a sequence. Our model entails that the prediction of future inputs is driven by the interaction between different synapses, i.e. heterosynaptic plasticity, which has been observed experimentally [16, 17, 57] and proposed as a computationally powerful mechanism for learning in single neurons [58, 23]. NMDA receptors could naturally operate as voltage-gated units for prediction errors. This is for two reasons: (1) NMDARs are voltage-dependent; and (2) NMDARs allow for a comparisons between the internal state of the neuron with external inputs [59]. Finally, experimental evidence [60, 61, 62, 63] and theoretical studies [64] support the hypothesis that active ion-channels along dendritic compartments strongly enrich the dynamical repertoire of neurons and underpin higher computational capabilities.

Comparing the present results to previous work, we emphasize that we did not construct a learning rule to reproduce experimentally observed STDP phenomena. Indeed, several phenomenological (i.e. descriptive) STDP models have previously been proposed to fit experimental data [65, 66, 67, 68] providing mathematical tools to describe, model and predict the behavior of neurons. However, these phenomenological models might not fully explain the computational significance of STDP mechanisms, nor the algorithms from which these biological implementations can emerge. Our approach differs from these models in that we took an optimization problem based on prediction of future inputs as a starting point. From this optimization problem we derived a learning rule which gave rise to experimentally observed STDP mechanisms. Our results, together with previous studies [69, 70, 71], suggest that STDP is a consequence of a general learning rule given the particular state of the system, the stimulation protocol and the specific properties of the input. As a consequence, several STDP learning windows which are described by other phenomenological rules are predicted by our model, as well as the dependence on synaptic strength and depolarization level.

An example of an established phenomenological model of STDP is the one developed by Clopath et al [72]. The authors, guided by experimental evidence [73, 43, 29], modeled the role of the membrane voltage as the relevant post-synaptic variable for synaptic plasticity. The plasticity model described in [72] can accurately reproduce a wide range of experimental findings which, to our knowledge, is not possible with STDP learning rules that are only based on spike timing [68]. The Clopath rule is based on the two-threshold dynamics observed by Artola et al [73] and the authors assume an Adaptive-Exponential I&F (AeI&F) model for the voltage dynamics [74], together with additional variables for the spike after-potential and an adaptive threshold. In agreement with the results of [72], we were able to account for a wide range of phenomena with a simpler model of voltage dynamics, supporting the idea that the history-dependent effect of the membrane potential is pivotal to plasticity. Our model also predicts the experimental observation that the amount of LTP has an inverse dependence on the initial strength of the synaptic input [43]. To our knowledge, this finding is not described or predicted by the model of [72], because it does not include a dependence on the initial strength of the synapse. Another unique feature of the learning rule described here is that it can produce different several STDP kernels (e.g. asymmetrical, symmetrical) depending on the initial conditions.

The learning rule decribed here and phenomenological STDP models might also lead to similar behavior in terms of spiking output in response to sequences. In agreement with the present work, previous studies have shown that phase precession can lead to the learning of temporal sequences through an asymmetric learning window as in spike-timing-dependent plasticity (STDP) [5, 75, 76]. Modelling studies have shown that a post-synaptic neuron endowed with an LTD-dominated STDP model can exhibit potentiation of the first synaptic inputs in a temporal sequence, leading to a decrease in the latency of the post-synaptic response [77, 78]. A key difference with our work, however, is that the predictive plasticity rule described here does not produce asymmetric STDP under all conditions. In fact, the degree of potentiation and depotentiation in our model depends on the initial state. That is, there is no fixed STDP kernel in our model. Another difference is that our model can anticipate sequences independent from the initial conditions of synaptic weights (in contrast to [78]), for a wide range of sequence lengths, and pre-synaptic population size. In the model described here, we show that the anticipation of sequences is a convergence point during learning and thereby it is a general solution for a wide set of model parameters.

Thus, we propose 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. Yet, it is obvious that single neurons are embedded into networks and different means of communication can lead to more complex learning rules in which the single-neuron learning rule described here might be one component. Indeed, it is possible that certain empirical phenomena like sequence recall additionally depend on network dynamics instead of single-neuron learning rules. For example, the faster recall of sequences in the visual system observed in [7] was reproduced in a recent work [79]. The authors developed a biologically realistic network model which differs from our implementation in several ways: (a) a network of both excitatory and inhibitory neurons, (b) a random gaussian connection probability, (c) a leaky I&F model with conductanced-based AMPA, GABA and NMDA synaptic currents, (d) several network hyperparameters such as synaptic delays, (e) a short-term depression model and a specific multiplicative, NN-STDP model. While the model in [79] gives a biological explanation based on the conductive properties of the ion-channels, the network implementation of our plasticity rule provides a principle approach to understand fast sequence recall as it is observed also in other brain areas, e.g different primary sensory areas [4, 6] or the hippocampus [50]. We qualitatively reproduced the faster recall of sequences with a much simpler model, supporting the pivotal role of spike times and excitability for the phenomenon.

Our work is further related to novel approximation algorithms for learning in neural networks such as e-prop [80] or surrogate-gradients techniques [81] as our model provides an online approximation for traning spiking neural networks (SNN). In [80] the authors showed 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. In our model the optimization problem is defined directly at the level of single neurons. Thus, all learning is local in space, i.e. there is no global loss at the network level. By definition, our model hence avoids the problem of propagating the gradient on recurrent connections [80]. An interesting question for future research is to combine these two approaches by obtaining a completely local learning rule for optimization at the single and network level simultaneously. Indeed, experimental evidence [82, 42] shows that local learning rules implemented by neurons provide a substrate on which global feedback signals act [83], which may provide a biological mechanism for error backpropagation [84]. As our loss function depends only on the membrane potential of the cell, our model avoids the problem of propagating the gradient through discrete spikes. It follows that we did not implement a surrogate-gradient approximation [81].

Looking forward, further investigation on single neurons anticipating local inputs and on their interplay through network interactions is key to understand how complex prediction strategies can emerge. Moreover, synapses can be located far from the soma along the dendritic arbor and might not be able to access the somatic membrane potential directly, with strong consequences on plasticity [29]. Other post-synaptic events such as NMDA spikes or plateau potentials can have an effect on plasticity rules based on membrane voltage, see e.g. [85]. Thereby, 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-like (LIF) model of the form

$$\tau_m \frac{dv(t)}{dt} = -v(t) + \vec{w}^\top \vec{x}(t) - v_{(\mathrm{th})} \sum_j \delta(t - t_j)\,. \tag{5}$$

Here, $v \in \mathbb{R}$ is the membrane potential, τ_m 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 weight vector and $v_{(\mathrm{th})}$ is the spiking threshold (the subscript $(_{\mathrm{th}})$: "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. We used a discrete time-step h to discretize equation (5), yielding a recurrence model of the form

$$\begin{cases} v_t = \alpha\, v_{t-1} - v_{(\mathrm{th})} s_{t-1} + \vec{w}^\top \vec{x}_t \\ s_t = H(v_t - v_{(\mathrm{th})})\,. \end{cases} \tag{6}$$

Here, $\alpha \equiv 1 - h/\tau_m$ and $H(\cdot)$ is the Heaviside function. The variable $s_t \in \{0, 1\}$ takes binary values and indicates the presence or absence of an output spike at timestep t. If the voltage exceeds the threshold, an output spike is emitted and this event reduces the membrane potential by a constant value $v_{(\mathrm{th})}$ at the next time step. This implementation of the membrane potential reset relates our model to the spike response model [86]. We set $h = 0.05$ ms in all numerical simulations.

Derivation of the learning rule

The predictive plasticity model entails 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

$$\mathcal{L} \equiv \sum_{t=0}^{T} \mathcal{L}_t \equiv \sum_{t=0}^{T} \frac{1}{2} \left\| \vec{x}_t - v_{t-1} \vec{w} \right\|_2^2 , \tag{7}$$

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) [87]. At each timestep t, the exact gradient of $\mathcal{L}$ can be written as the contribution of two terms

$$\vec{\nabla}_{\vec{w}} \mathcal{L} = \sum_{t=0}^{T} \frac{1}{2} \Big(\vec{\nabla}_{\vec{w}} \mathcal{L}_t + \frac{\partial \mathcal{L}_t}{\partial v_{t-1}} \vec{\nabla}_{\vec{w}} v_{t-1} \Big) . \tag{8}$$

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

The first term of the gradient is given by

$$\vec{\nabla}_{\vec{w}} \mathcal{L}_t = -2(\vec{x}_t - v_{t-1} \vec{w}_t) v_{t-1} , \tag{9}$$

which propagates the prediction error selectively to each input and scales it by the feedback signal determined by v_{t-1}. This term can be interpreted as a time-shifted version of the Oja's rule [25], where v_{t-1} plays the role of the linear output variable. In fact, Oja's rule can be derived as an approximated gradient descent method for dimensionality reduction problems [88]. The second term of equation (8) has a contribution given by the direct effect of v_{t-1} on the prediction

$$\frac{\partial \mathcal{L}_t}{\partial v_{t-1}} = -2(\vec{x}_t - v_{t-1} \vec{w})^\top \vec{w} , \tag{10}$$

i.e., 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 (8) via dynamical updates. We obtain this by differentiating the dynamical rule of the membrane potential in (1)

$$\vec{\nabla}_{\vec{w}} v_t = \Big(\alpha - v_{(\mathrm{th})} \frac{\partial s_{t-1}}{\partial v_{t-1}} \Big) \vec{\nabla}_{\vec{w}} v_{t-1} + \vec{x}_t . \tag{11}$$

The first term contains the Jacobian J_t of equation (1)

$$J_t = \frac{\partial v_t}{\partial v_{t-1}} = \alpha - v_{(\mathrm{th})} \frac{\partial s_{t-1}}{\partial v_{t-1}} , \tag{12}$$

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 [89], we define an influence vector $\vec{p}_t \equiv \vec{\nabla}_{\vec{w}} v_t$ such that it obeys the recursive equation

$$\vec{p}_t = J_t \, \vec{p}_{t-1} + \vec{x}_t , \tag{13}$$

which follows straightforward from equation (11) and gives a forward-pass dynamical update of the gradient [90]. Finally, we define the prediction error $\vec{\epsilon}_t$ at timestep t as

$$\vec{\epsilon}_t \equiv \vec{x}_t - v_{t-1} \vec{w} , \tag{14}$$

which defines the sign and amplitude of plasticity. All together, the exact gradient of equation (8) can be written as

$$\vec{\nabla}_{\vec{w}}\mathcal{L} = -\sum_{t=0}^{T}\left[\vec{\epsilon}_t\, v_{t-1} + (\vec{\epsilon}_t^{\top}\vec{w})\,\vec{p}_{t-1}\right]. \tag{15}$$

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

$$\begin{aligned}\vec{w}_k &= \vec{w}_{k-1} - \eta\vec{\nabla}_{\vec{w}}\mathcal{L} \\ &= \vec{w}_{k-1} + \eta\sum_{t=0}^{T}\left[\vec{\epsilon}_t\, v_{t-1} + (\vec{\epsilon}_t^{\top}\vec{w}_{k-1})\,\vec{p}_{t-1}\right].\end{aligned} \tag{16}$$

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

We are interested in an online learning rule where the weight update forms part of the dynamics of the model, and takes place in real-time with the prediction of the pre-synaptic inputs. This is a typical method in stochastic optimization theory for recursive objective functions [91] and online signal processing [88]. We approximate the learning equation (16) with the current estimate of the gradient

$$\vec{w}_t = \vec{w}_{t-1} - \eta\vec{\nabla}_{\vec{w}}\mathcal{L} \;\simeq\; \vec{w}_{t-1} - \eta\vec{\nabla}_{\vec{w}}\mathcal{L}_t\Big|_{\vec{w}=\vec{w}_{t-1}} = \vec{w}_{t-1} - \eta\Big[\vec{\epsilon}_t\, v_{t-1} + (\vec{\epsilon}_t^{\top}\vec{w}_{t-1})\,\vec{p}_{t-1}\Big]. \tag{17}$$

Different from other works [80], our loss function (7) is defined locally at the single neuron level and we thus directly avoid the problem of backpropagation of the gradient through the recurrent interactions between neurons. Therefore, our approximated learning rule is completely online as it only requires information available at time step t. Theoretical studies suggests that such stochastic approximation works when η is sufficiently small [88, 90]. In our case, the passive memory capacity of the membrane potential is given by its time constant τ_m. In the limit of $\eta\tau_m \ll 1$, the changes in the weights are slow compared to voltage changes and the following relation holds:

$$\vec{w}_{t'} \simeq \vec{w}_t \quad \forall\, t' \;:\; |t - t'| < \tau_m\,. \tag{18}$$

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

Jacobian and surrogate-gradient method

The first term in equation (12) 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. This latter term can be approximated following the surrogate-gradient method [81]

$$-v_{(\mathrm{th})}\frac{\partial s_{t-1}}{\partial v_{t-1}} \simeq \gamma f(v_{t-1}, v_{(\mathrm{th})})\,, \tag{19}$$

where f is a continuous function of v_t and γ is a scaling factor. In general, the backpropagation of the gradient through the reset mechanisms is neglected [92, 81, 80]. Here we defined the membrane potential v_t as the output variable of the system and the loss function and s_t as an hidden variable for the objective function. Therefore, our implementation directly avoids the problem of backpropagating through discrete output variables. Given these two arguments, we considered $\gamma = 0$ throughout the paper.

Optimization and initialization scheme

Equation (3) defines a completely online optimization scheme which can be implemented locally by single neurons. In Figure 3, Figure S5, Figure S6, Figure S7 and Fig S9 we updated the synaptic weights following the online approximation of the gradient in Equation (17). For the results of Figure 1, Figure S1, Figure 2b-c, Figure 4, Figure S3, Figure S4, Figure S8 and Figure S10 we added a scaling term to the predictive plasticity rule as

$$\vec{w}_t = \vec{w}_{t-1} - \eta\,\vec{w}_{t-1}\Big[\vec{\epsilon}_t\, v_{t-1} + (\vec{\epsilon}_t^{\top}\vec{w}_{t-1})\,\vec{p}_{t-1}\Big]. \tag{20}$$

That is, we multiplied the learning rate with $\vec{w}_{t-1}$ to ensure non-negative values of the weights. Consequently, the weight update at time step t was proportional to the synaptic weight value at time step $t-1$. In Figure 2d and Figure S2 we did not use the online approximation of the gradient and we optimized the full model (Equation (16)) using the Adam optimizer [93].

For the example in Figure 1, for Figure 2b-c, Figure 3 and for the spike-timing dependent plasticity protocols in Figure 4, the initial weights were assigned to fixed values for the different cases considered. In Figure 2d, the initial weights were randomly drawn from a truncated normal distribution. We bonded the truncated normal distribution to obtain positive values of the initial weights. The variance of the normal distribution was scaled by the squared root of the total input size N.

Simulations

The spiking sequence were defined by a set of ordered spike times at which the pre-synaptic neurons were active. For all simulations, the inputs were convolved with an exponential kernel with $\tau_x = 2$ ms to replicate the fast dynamics of post-synaptic currents. We added two sources of noise to the simulations in Figure 2-3 and corresponding supplementary figures: 1) In each epoch, the spikes that were part of the sequence were randomly shifted by an amount Δt uniformly distributed between [-2, 2] (in ms). 2) Each pre-synaptic also exhibited stochastic background firing following an homogeneous Poisson process with constant rate λ uniformly distributed between 0 and 10 Hz. In addition, in Figure 2, the onset of the sequence relative to the time window was randomly chosen between 0 and 200 ms, and there were an additional 100 neurons that fired randomly. We trained the model by numerically solving the dynamics during each epoch. The model was fully determined by 5 hyperparameters: the timestep h, the membrane time constant τ_m, the spiking threshold v_{th}, the learning rate η and the input time constant τ_x. To quantify the performance (in terms of sequence anticipation) in Figure S3, we fixed the number of training epochs, we performed 100 numerical simulation for each condition and we labeled successful simulations 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 20 ms after the onset of the input sequence (fast anticipation). We computed the error as 1 minus the percentage of successful trials in the set of 100 simulations.

For the STDP simulations, 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 [18, 19]. To approximate the STDP protocol with a current injection that triggers a post-synaptic spike, the initial conditions were chosen such that x_2 triggered a post-synaptic spike, and x_1 was a sub-treshold input. For the results in Figure 4b-e, we changed the number and the timing of the pre-synaptic neuron spikes according to each experimental protocol. We numerically solved the dynamics of the model by repeating the input pattern a number of times as in the experimental protocols. We computed the weight change of the sub-threshold input as as the ratio of the synaptic weight before and after each simulation protocol.

We reproduced the same burst-dependent pairing protocol as in [38] by decreasing the delay Δt between post-synaptic spikes (Figure 4c). The bursts were triggered by three input spikes. We simulated the 1-n protocol as in [39] 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 (Figure 4d). We reproduced the frequency pairing protocol as in [39] by pairing 5 post-pre inputs with different intra-pairing frequencies as in the experimental protocol (Figure 4e).

Acknowledgments

This work 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] Michael J Berry et al. "Anticipation of moving stimuli by the retina". In: *Nature* 398.6725 (1999), pp. 334–338.

[5] Mayank R Mehta, Michael C Quirk, and Matthew A Wilson. "Experience-dependent asymmetric shape of hippocampal receptive fields". In: *Neuron* 25.3 (2000), pp. 707–715.

[6] Xiaofeng Lu and James Ashe. "Anticipatory activity in primary motor cortex codes memorized movement sequences". In: *Neuron* 45.6 (2005), pp. 967–973.

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

[8] 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.

[9] Peter E Keller, Giacomo Novembre, and Michael J Hove. "Rhythm in joint action: psychological and neurophysiological mechanisms for real-time interpersonal coordination". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 369.1658 (2014), p. 20130394.

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

[11] 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.

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

[13] Cem Uran et al. "Predictive coding of natural images by V1 firing rates and rhythmic synchronization". In: *Neuron* 110.7 (2022), pp. 1240–1257.

[14] Donald Olding Hebb. *The organization of behavior: A neuropsychological theory*. Psychology Press, 2005.

[15] 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.

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

[17] 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.

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

[19] 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.

[20] 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.

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

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

[23] 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.

[24] 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.

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

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

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

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

[29] 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.

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

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

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

[33] David R Euston, Masami Tatsuno, and Bruce L McNaughton. "Fast-forward playback of recent memory sequences in prefrontal cortex during sleep". In: *science* 318.5853 (2007), pp. 1147–1150.

[34] Thomas J Davidson, Fabian Kloosterman, and Matthew A Wilson. "Hippocampal replay of extended experience". In: *Neuron* 63.4 (2009), pp. 497–507.

[35] Albert K Lee and Matthew A Wilson. "Memory of sequential experience in the hippocampus during slow wave sleep". In: *Neuron* 36.6 (2002), pp. 1183–1194.

[36] Zoltán Nádasdy et al. "Replay and time compression of recurring spike sequences in the hippocampus". In: *Journal of Neuroscience* 19.21 (1999), pp. 9497–9507.

[37] William E Skaggs and Bruce L McNaughton. "Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience". In: *Science* 271.5257 (1996), pp. 1870–1873.

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

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

[40] 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.

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

[42] 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.

[43] 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.

[44] 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.

[45] 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.

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

[47] 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.

[48] Shogo Ohmae and Javier F Medina. "Climbing fibers encode a temporal-difference prediction error during cerebellar learning in mice". In: *Nature neuroscience* 18.12 (2015), pp. 1798–1803.

[49] 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.

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

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

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

[53] Jeffrey C Magee and Daniel Johnston. "A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons". In: *Science* 275.5297 (1997), pp. 209–213.

[54] Nace L Golding, Nathan P Staff, and Nelson Spruston. "Dendritic spikes as a mechanism for cooperative long-term potentiation". In: *Nature* 418.6895 (2002), pp. 326–331.

[55] John Lisman and Nelson Spruston. "Questions about STDP as a general model of synaptic plasticity". In: *Frontiers in synaptic neuroscience* 2 (2010), p. 140.

[56] Jason Hardie and Nelson Spruston. "Synaptic depolarization is more effective than back-propagating action potentials during induction of associative long-term potentiation in hippocampal pyramidal neurons". In: *Journal of Neuroscience* 29.10 (2009), pp. 3233–3241.

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

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

[59] 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.

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

[61] 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.

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

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

[64] 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.

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

[66] Jean-Pascal Pfister and Wulfram Gerstner. "Triplets of spikes in a model of spike timing-dependent plasticity". In: *Journal of Neuroscience* 26.38 (2006), pp. 9673–9682.

[67] Michael Graupner and Nicolas Brunel. "Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location". In: *Proceedings of the National Academy of Sciences* 109.10 (2012), pp. 3991–3996.

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

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

[70] 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.

[71] Manu Srinath Halvagal and Friedemann Zenke. "The combination of Hebbian and predictive plasticity learns invariant object representations in deep sensory networks". In: *bioRxiv* (2022).

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

[73] 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.

[74] 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.

[75] Jason J Moore et al. "Linking hippocampal multiplexed tuning, Hebbian plasticity and navigation". In: *Nature* 599.7885 (2021), pp. 442–448.

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

[77] Rudy Guyonneau, Rufin VanRullen, and Simon J Thorpe. "Neurons tune to the earliest spikes through STDP". In: *Neural Computation* 17.4 (2005), pp. 859–879.

[78] 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.

[79] Xuhui Huang et al. "Different propagation speeds of recalled sequences in plastic spiking neural networks". In: *New Journal of Physics* 17.3 (2015), p. 035006.

[80] 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.

[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.

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

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

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

[85] Jacopo Bono, Katharina A Wilmes, and Claudia Clopath. "Modelling plasticity in dendrites: from single cells to networks". In: *Current opinion in neurobiology* 46 (2017), pp. 136–141.

[86] Wulfram Gerstner et al. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.

[87] 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.

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

[89] 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.

[90] 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.

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

[92] Guillaume Bellec et al. "Long short-term memory and learning-to-learn in networks of spiking neurons". In: *arXiv preprint arXiv:1803.09574* (2018).

[93] Diederik P Kingma and Jimmy Ba. "Adam: A method for stochastic optimization". In: *arXiv preprint arXiv:1412.6980* (2014).

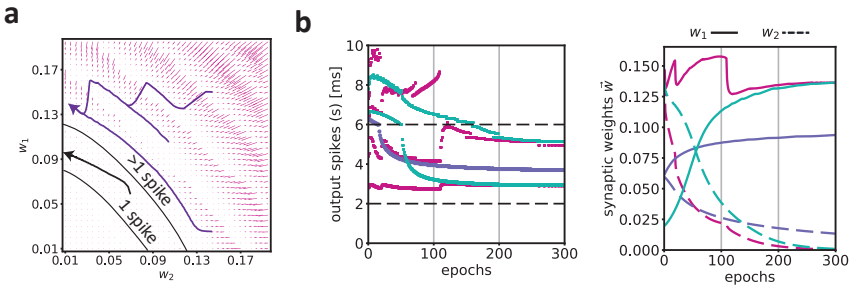

Figure S1: Dynamics in the parameter space. Relates to Figure 1 of the main text. a) We show here that the learning dynamics of the model shown in Figure 1c-d are qualitatively the same when the initial conditions lie in regions of multiple output spikes. Shown is the flow field as in Figure 1d. The black arrow shows the trajectory of the weights obtained by training the model with four 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). It can be seen, that all the regions of the parameter space with two or more post-synaptic spikes converge to a fixed point, at which there is maximal credit for the first input (i.e. $w_1 > w_2$). **b**) Right: same as in Figure 1c), for the four initial conditions shown in **a**.

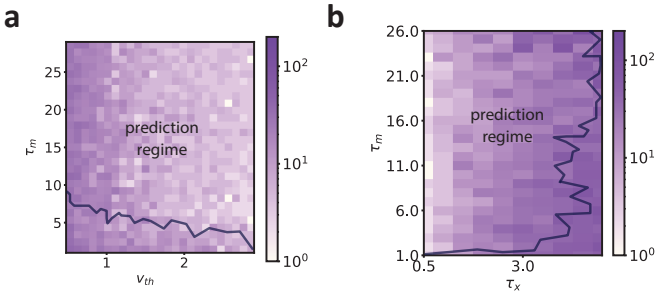

Figure S2: Effect of model parameters. Relates to Figure 2 of the main text. **a**) Simulations were performed as in Figure 2, but now with different values of model parameters τ_m and v_{th}. The color map corresponds to the total number of output spikes at the end of learning. We show regions for the parameter space where the total number of spikes was smaller than 200. The black line outlines the region of parameter space where the neuron (in epoch 1000) fires between 2 and 10 ms after the onset of the input sequence, which can be interpreted as the predictive or anticipatory solution. For the vast portion of the parameter space, the model converges to the anticipatory solution. **b**) Same as in a) for different values of model parameters τ_m and τ_x.

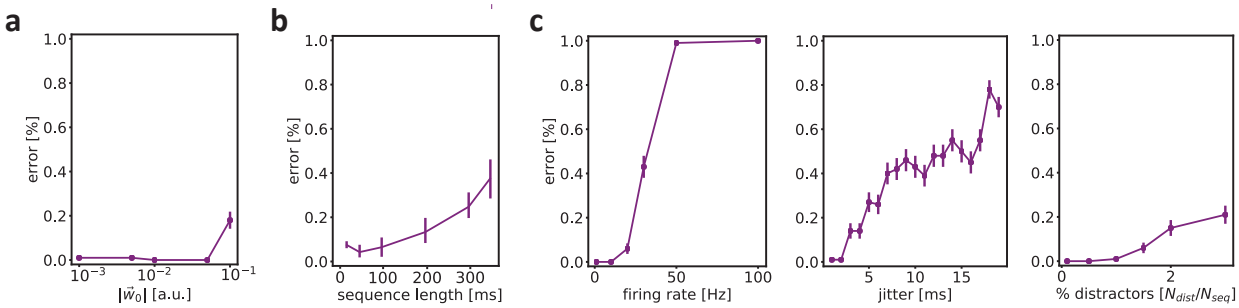

Figure S3: Dependence of performance, in terms of anticipation of sequences, depending on model parameters. Relates to Figure 2 of the main text. **a**) Simulations were performed as in Figure 2. To quantify the performance in terms of anticipating the sequence, we labeled successful simulation based on the criteria of input selectivity and fast anticipation (see Methods). Performance is then shown as the fraction of simulations in which the model did not reach an anticipatory solution. Performance is shown as a function of the initial weight vector $\vec{w}_0$, with the weight the same for each pre-synaptic input. The performance does not strongly depend on the initial conditions. For each weight, we performed 100 simulations (shown are the standard errors of the mean). For each simulation, we used 2000 epochs. **b**) As **a**, but now as a function of the sequence length. **c**) Left: Performance as a function of the maximal rate of background firing. Center: Performance as a function of the maximal spike time jitter. Right: Performance as a function of the proportion of pre-synaptic distractor neurons, that is pre-synaptic neurons that do not participate in the sequence.

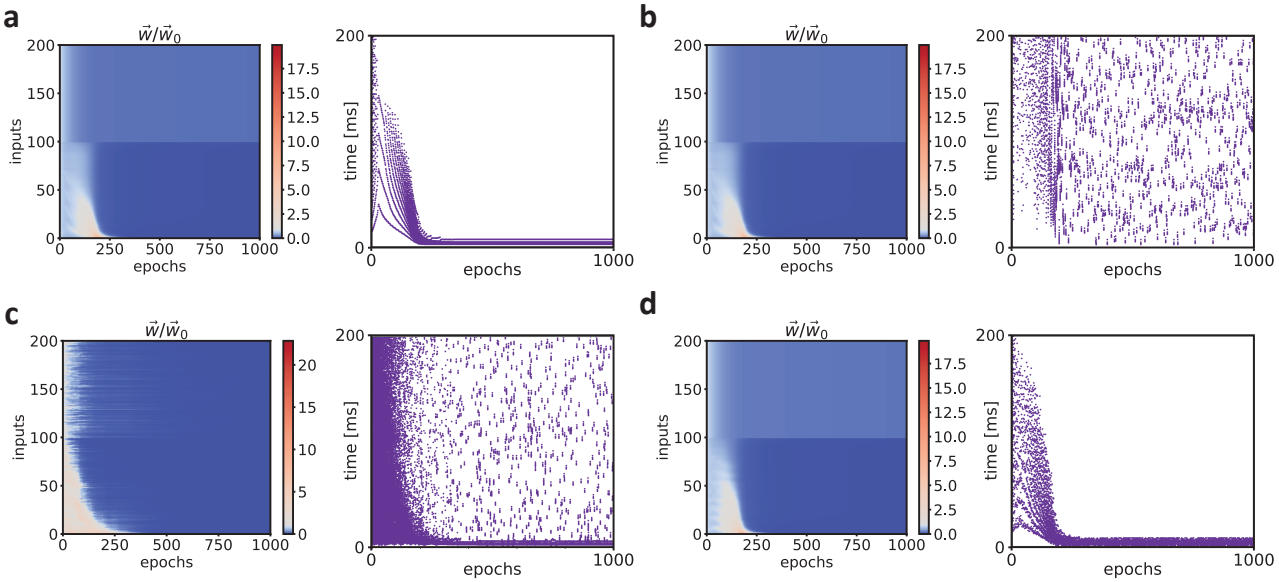

Figure S4: Effect of individual noise sources - Relates to Figure 2 of the main text. **a**) Simulations were performed as in Figure 2, however without any noise source (no spike time jitter, no distractor neurons, no background firing, no random onset of the sequence). Dynamics of the synaptic weights $\vec{w}$ (left) and of the post-synaptic spiking activity (right) across training, without any source of noise. **b**) Same as in **a**, but now with one noise source, in this case the sequence onset drawn from a uniform distribution with values between 0 and 100 ms. **c**) Same as in **a**, with now one noise source, namely random background firing following an homogeneous Poisson process with rate distributed between 0 and 10 Hz. **d**) Same as in **a**, with again one noise source, namely jitter of the input spike times (random jitter between -2 and 2 ms).

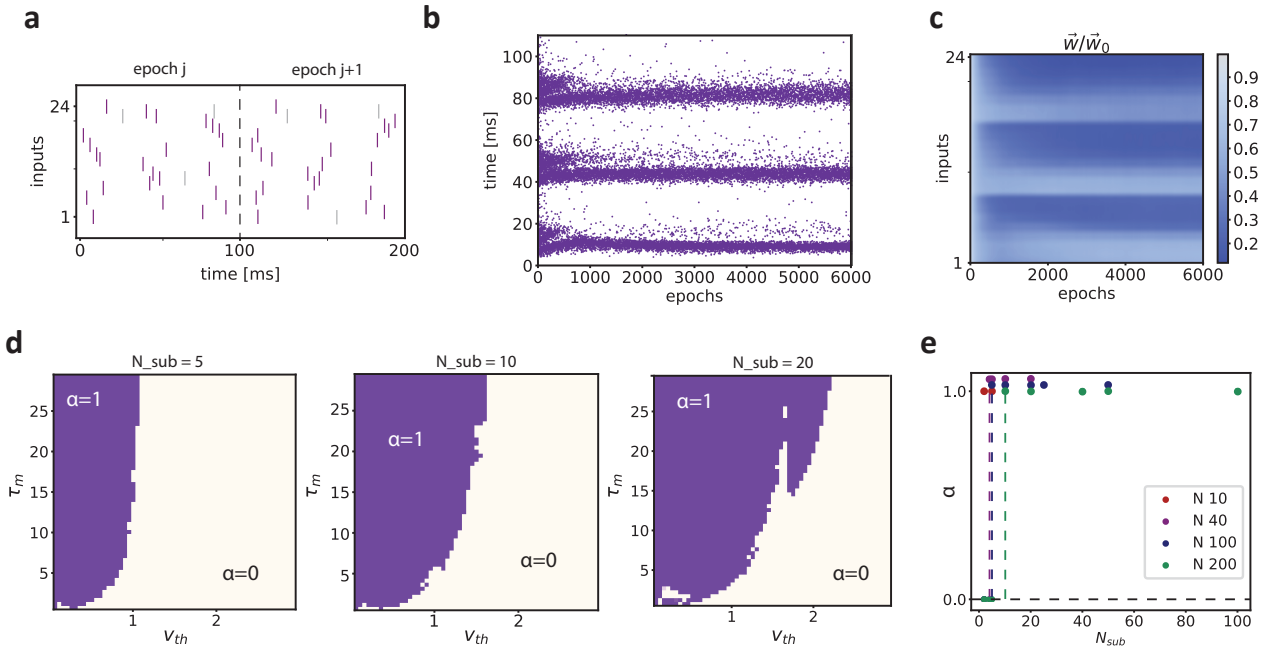

Figure S5: Learning multiple independent sequences and the capacity of the model. Relates to Figure 2 of the main text. **a**) Example spike sequence during different training epochs. The input spike trains contain a sequence given by the correlated activity of 24 pre-synaptic neurons that firex sequentially with relative delays of 2 ms (purple spike pattern). Each epoch contains two different sources of noise: (1) jitter of the spike times (random jitter between -2 and 2 ms, which is applied randomly in each epoch); and (2) random firing following an homogeneous Poisson process with rate $\lambda = 5$ Hz. The total sequence is divided into 3 sub-sequences of $N = 8$ separated by 20 ms. **b**) Dynamics of the post-synaptic spiking activity during learning for the four different sub-sequences. The neuron learns to represent several components of the input structure by potentiating the unpredictable part, i.e. the start, of each sub-sequence. **c**) Dynamics of the synaptic weights during training. The synaptic weights are ordered from 1 to 24 following the temporal order of the sequence. The synaptic weights corresponding to the first inputs of each sub-sequence are maximally potentiated. **d**) Capacity of the model α for different values of model parameters τ_m (membrane time constant) and v_{th} (spiking threshold). The capacity α is defined as the percentage of sub-sequences that the neuron anticipates at the end of training (same performance criterion as in Figure S3). The sequence was composed by the subsequent firing of $N = 100$ pre-synaptic inputs, divided in N_{sub} sub-sequences with $N_{sub} = 5$, $N_{sub} = 10$ and $N_{sub} = 20$ from left to right. In each condition, the neuron anticipates every sub-sequence in the input for a broad range of model parameters. **e**) Capacity of the model as a function of the number of sub-sequences N_{sub} for a different total number of pre-synaptic neurons N. Here we fixed $v_{th} = 1$ and $\tau_m = 20$ ms. Each color corresponds to a different total number of pre-synaptic neurons. We then quantified the capacity for different lengths of the sub-sequences (i.e. N_{seq}). For example, there are 2 sub-sequences of length N_{seq} for a total number of 200 pre-synaptic neurons. The figure shows that full capacity (i.e. $\alpha = 1$) can be reached even for small sub-sequences, and for both small and large numbers of pre-synaptic neurons. Values higher than $\alpha = 1$ correspond to values of $\alpha = 1$ and are shown as such for visualization purposes.

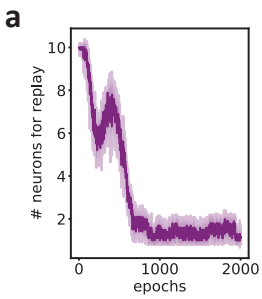

Figure S6: Number of neurons in the network that need to be activated for sequence recall to occur. Relates to Figure 3 of the main text. **a**) The y-axis corresponds to the number of neurons in the network that need to be activated by the input sequence in order for recall to occur. For each epoch, we tested how many neurons needed to be activated by the corresponding pre-synaptic sequence to obtain full recall. First, we presented the pre-synaptic sequence corresponding to the first neuron in the network. We observed how many neurons in the network were active after the sequence presentation. Note that if a neuron is not activated by the pre-synaptic sequence, the pre-synaptic neurons still exhibit background firing. Then, we systematically increased the number of neurons in the network which received the corresponding pre-synaptic sequence. We examined the minimum number of neurons in the network that needed to be active such that every other neuron in the network also fired sequentially. We repeated this analysis for each training epoch and for 100 simulations with different noise realizations. At the beginning of training, each sequence input to each neuron in the network is needed to obtain a full recall of the sequence. At the end of training, only the inputs to the first neuron are required to trigger a full recall of the input sequence in the network.

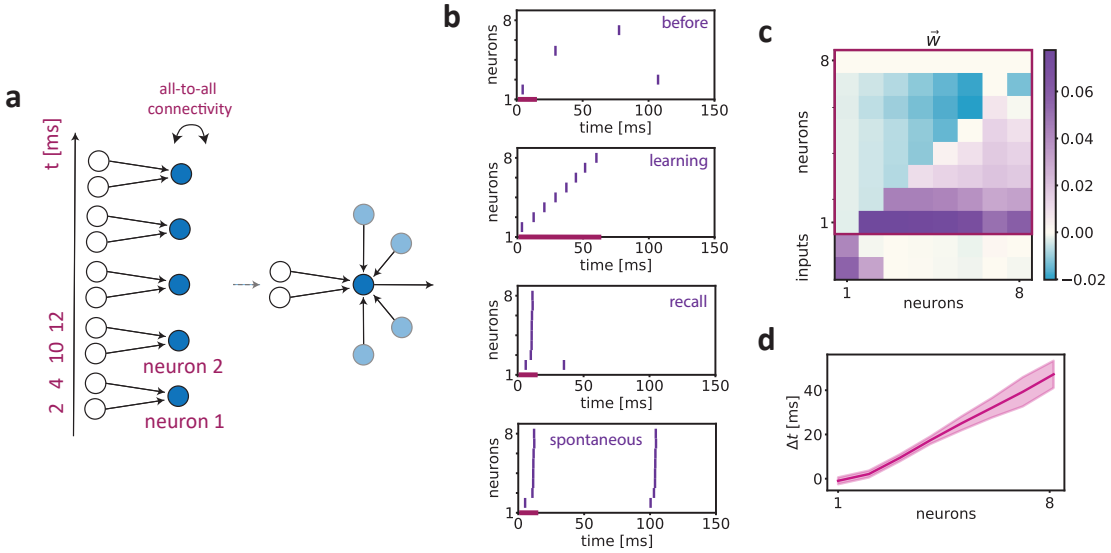

Figure S7: Anticipation and recall of sequences with all-to-all connected network. Relates to Figure 3 of the main text. We explored the dynamics of a network model where $N = 8$ neurons received an input sequence distributed across a total of 16 pre-synaptic neurons. Each neuron in the network received 2 pre-synaptic inputs The pre-synaptic neurons fired in a sequential manner with delays of 2 ms. Each epoch contains two different sources of noise: (1) jitter of the spike times (random jitter between -2 and 2 ms, which is applied randomly in each epoch); and (2) random firing following an homogeneous Poisson process with rate $\lambda = 10$ Hz. 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. The sequence onset of pre-synaptic inputs for the $n+1$-th neuron started 8 ms after the sequence onset for the n-th neuron in the network, etc. Accordingly, the pattern that each neuron tries to predict is composed by the external pre-synaptic input and by the internally generated activity of the network. **a**) Illustration of the network model, where only 5 neurons of the network are represented for simplicity. **b**) Top: Network spiking activity when presenting only part of the input sequence, namely the sequential input to the first 3 neurons in the network. Note that the other network neurons still received inputs from their corresponding pre-synaptic neurons due to stochastic background firing. Middle: Network spiking activity during the first training epoch, when the entire sequence was presented. Bottom: The first plot shows the network spiking activity at the end of training (epoch 2000) when we presented only part of the sequence (as in the top plot). This figure shows that after learning, the network learns to recall the entire sequence. The second plot shows the network spiking activity when we presented only part of the sequence, but also the end of the sequence (i.e. activating the last two network neurons). There is no effect of activating the neurons at the end of the sequence. Note that random background firing can also elicit spontaneous recall of the sequence, outside the stimulation period. **c**) The synaptic weight matrix obtained at the end of training (epoch 1000). The i-th column (and the top 10 rows) corresponds to the synaptic weights learned from the other neurons to the i-th neuron in the network. The bottom 2 entries correspond to the weights for the pre-synaptic inputs. **d**) Difference between the latency of the first spike of each single neuron in the first epoch of training and after training. Positive values indicate that an earlier latency due to training. The panel shows the mean and standard deviation computed over 100 different simulations.

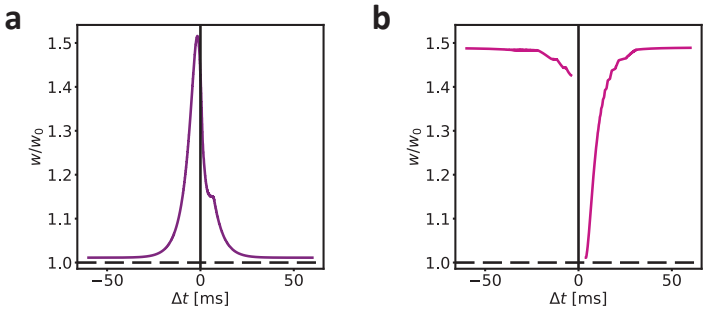

Figure S8: Different STDP kernels. Relates to Figure 4 of the main text. **a**) Weight change (in percentage) of the sub-threshold input as a function of the delay between the two input spikes (see Methods). We obtained the learning window by setting $\tau = 5$ ms, $v_{th} = 0.4$ and initial conditions $w_1 = 0.0007$ and $w_2 = 0.02$. **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 $w_1 = 0.02$ and $w_2 = 0.1$.

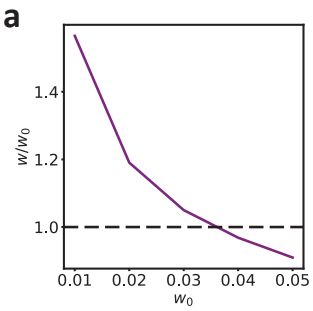

Figure S9: Dependence of STDP on the initial synaptic weight. Relates to Figure 4 of the main text. **a**) Simulations were performed as in Figure 4b, however we simulated different initial conditions for the synaptic weight corresponding to the sub-threshold input. Weight change (in percentage) as a function of different initial values for the synaptic weights.

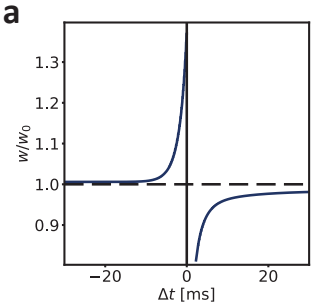

Figure S10: Spike-timing-dependent-plasticity with fixed synapses. Relates to Figure 4 of the main text. Simulations were performed as in Figure 3, however we fixed the synaptic weight of the supra-threshold input, that is the one eliciting a spike. Only the weight of the sub-threshold input was plastic. **a**) Weight change (in percentage) as a function of the delay between the two input spikes.
