

Inhibitory feedback enables predictive learning of multiple sequences in neural networks

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

Anticipating future events is a key computational task for neuronal networks. Experimental evidence suggests that reliable temporal sequences in neural activity play a functional role in the association and anticipation of events in time. However, how neurons can differentiate and anticipate multiple spike sequences remains largely unknown. We implement a learning rule based on predictive processing, where neurons exclusively fire for the initial, unpredictable inputs in a spiking sequence, leading to an efficient representation with reduced post-synaptic firing. Combining this mechanism with inhibitory feedback leads to sparse firing in the network, enabling neurons to selectively anticipate different sequences in the input. We demonstrate that intermediate levels of inhibition are optimal to decorrelate neuronal activity and to enable the prediction of future inputs. Notably, each sequence is independently encoded in the sparse, anticipatory firing of the network. Overall, our results demonstrate that the interplay of self-supervised predictive learning rules and inhibitory feedback enables fast and efficient classification of different input sequences.

Introduction

The prediction of future events is a key mechanism for the organization of behaviour [1], and is arguably supported by anticipatory neuronal firing observed in various brain regions [2, 3, 4, 5, 6]. In order to fire ahead of predictable events, neurons must learn to credit synapses that carry information about future inputs, and downregulated the ones that are predictable. However, this prediction process is difficult when different spiking patterns influence the activity of the same neuronal populations. For instance, synaptic inputs carrying relevant information for predicting one stimulus may also be predictable by other synaptic inputs during another stimulus. To address this issue, neurons must accomplish two tasks: (a) learn predictive relationships in the input spikes patterns, and adjust synaptic weights accordingly; (b) discriminate different components of input statistics and participate in the prediction of specific input patterns. Currently, it remains unclear how neurons can show anticipatory firing that is specific to the particular temporal relationships in the input.

Theoretical works indicate that spike-timing-dependent plasticity (STDP) leads, in some cases, to a reduction in the post-synaptic spike latency when spike patterns are systematically repeated [7, 8]. Thereby, neurons can fire ahead of future sensory stimuli and can give rise to predictions at the network level [9], suggesting a mechanism for anticipatory firing in the brain [2]. Another recent work addressed the problem directly, hypothesizing that neurons adjust their synaptic weights proportionally to the predictability of pre-synaptic inputs [10]. This plasticity mechanism enables neurons to anticipate and signal future events and unifies several STDP phenomena from the perspective of predictive processes. However, this mechanism alone cannot generate anticipatory firing when the temporal relationships between spikes are ambiguous, e.g. when the same input predicts other inputs in one spike sequence but is fully predictable during another sequence. Introducing competition among neurons offers a solution to this issue, as it allows neurons to become selective for specific sequences. Winner-take-all-like models capture the fundamental principle of competition, where single neurons encode specific inputs while other neurons remain silent [11]. A possible competitive mechanism is a recurrent inhibition, which exerts inhibitory control over neurons receiving the same input [12, 13]. Previous studies suggest that local inhibitory interactions, observed in the neocortex [14] and hippocampus [15], are essential for decorrelated neuronal firing and independent representation of input components [16]. However, how the interplay

26 between inhibitory mechanisms and local learning rules can support the prediction of future inputs is still not well understood.

28 In this work, we propose that the combination of a predictive learning rule (PLR) as in [10] with a competitive mechanism at the network level (inhibitory feedback; IF) can lead to selective anticipation of input patterns. We demonstrate that neurons in the network suppress the firing of the entire network while adjusting their synaptic weights to anticipate predictable inputs. This inhibitory mechanism allows neurons to become selective for specific sequences in the input patterns and to show anticipatory firing for those sequences. Moreover, the identity of different input sequences can be read out from a few early spikes in the network activity, enabling fast and energy-efficient anticipation of future events.

Results

Inhibitory feedback enables selective anticipation of spike sequences in neural networks

The present work builds further on a predictive learning rule from a previous work [10] (see Methods for a detailed description of the learning rule). The learning rule potentiates synaptic weights associated with predictive spikes, that is, inputs that anticipate other pre-synaptic inputs, while it suppressed inputs that are anticipated by other pre-synaptic spikes. For example, we explored the scenario where a neuron endowed with the predictive learning rule received a temporal sequence composed of N pre-synaptic neurons firing sequentially with fixed delays (Figure 1a). In this case, the pre-synaptic inputs occurring earlier in the sequence should be potentiated as they anticipate the subsequent inputs. We trained the neuron by repeating the input sequence for 750 epochs. As anticipated, during training, the neuron progressively potentiated the inputs occurring earlier in the sequence. Eventually, it assigned the highest credit to the first inputs in the sequence and fired upon their occurrence (Figure 1b). Notably, our results remained consistent regardless of the size of the pre-synaptic population (Figure S1).

46 According to the predictive learning rule described above [10], synaptic weights undergo potentiation or depotentiation based on the predictability of the pre-synaptic inputs. However, in cortical networks, neurons may receive many different input sequences. As a result, one would expect that neurons do not selectively fire at the beginning of one particular sequence, but rather diffusely throughout multiple sequences. This raises the question, of how the predictive learning rule can be integrated with competition mechanisms to enable neuronal selectivity for specific spike sequences.

52 We simulated a case where neurons received input patterns consisting of n_{class} distinct sequence classes (Figure 1c). In each sequence, the N_{presyn} pre-synaptic neurons fired in sequential order with fixed delays. The total sequence duration was 200 ms. The pre-synaptic neurons participated in each of the n_{class} input sequences, and each sequence class $c \in \{1, \dots, n_{\text{class}}\}$ was characterized by a unique order of pre-synaptic input spikes. During each training epoch, the input also included randomly drawn sequences whose firing times varied at each iteration, representing irrelevant input patterns (these random sequences comprised 10% of all sequences). We compared two network configurations: In the first configuration, there were N_{nn} neurons, but without recurrent connections (“PLRnoIF”). In the second configuration, the neurons were recurrently connected through inhibitory feedback (“PLR+IF”) (see Methods). Specifically, each neuron provided inhibition onto all the neurons in the network, which represents a simplified implementation of the lateral global inhibition in cortical networks [12] (see Methods). In both network configurations, the weight updates followed the predictive learning rule [10]. Furthermore, the synaptic connections from the pre-synaptic inputs were randomly assigned at the start of training, while the weights of the recurrent connections remained fixed.

We exposed the neural network to the input patterns, and we examined the network activity across epochs. In the absence of inhibitory feedback (PLRnoIF network), the neurons showed prolonged firing during the whole duration of each sequence (Figure 1d). The dynamics of the synaptic weights did not show any structure, as all the pre-synaptic inputs became potentiated, which is due to the fact that many sequences were presented. By contrast, we observed that including inhibitory feedback (PLR+IF network) between neurons led to the anticipation of input sequences (Figure 1e). Specifically, neurons in the PLR+IF network became selective to one sequence class and learned to potentiate the early inputs of that sequence (Figure 1f). Accordingly, the neurons in the PLR+IF network fired shortly after the onset of the selected sequence (Figure 1c), and typically did not fire when exposed to the other input sequences (Figure S2). The percentage of neurons that became selective to specific sequences increased across epochs (Figure 2a). By contrast, in the PLRnoIF, no sequence selectivity was developed. As a result of the sequence selectivity, the average firing activity of the

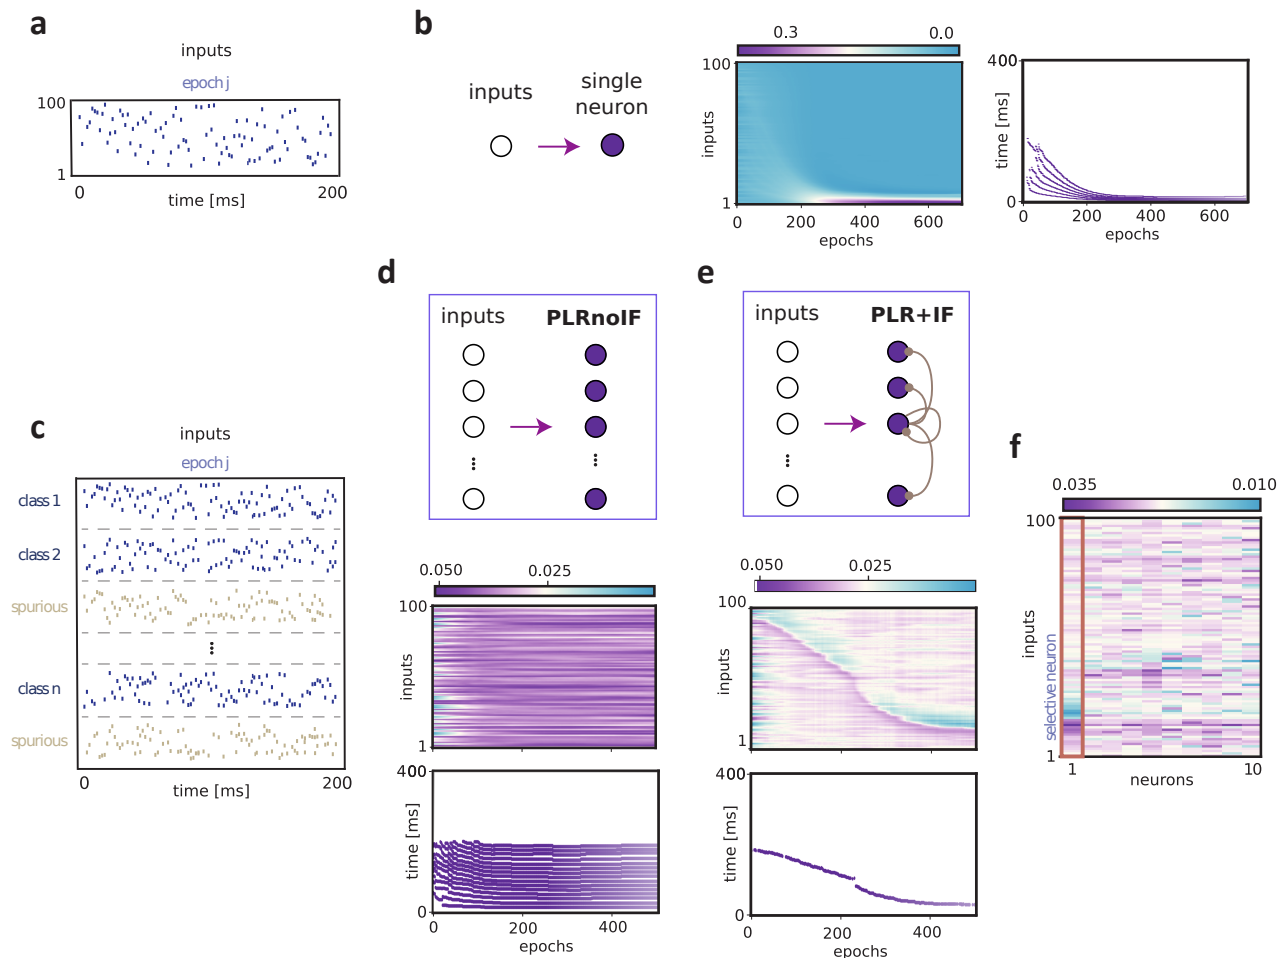


Figure 1: Anticipation of multiple sequences in a neural network with inhibitory feedback. **a)** The input was composed of one spike sequence, given by the correlated firing of $N=100$ pre-synaptic neurons. The N pre-synaptic neurons fired sequentially with relative delays of 2 ms, resulting in a total sequence length of 200 ms. **b)** Left: Illustration of the connections between pre-synaptic inputs and the post-synaptic neuron. Middle: Dynamics of the synaptic weights associated with each pre-synaptic input as a function of training epochs. The synaptic weights are ordered along the y-axis from 1 to 100 following the temporal order of the sequence. Right: Dynamics of the neuron spiking activity as a function of training epochs, as in panel b but for 100 pre-synaptic inputs. **c)** The input was composed of sequences belonging to 30 different classes ($n_{\text{class}} = 30$) and of n_{random} random spiking sequences. **d)** Top-Left: illustration of the PLRnoIF network. Here, we illustrated 5 neurons and all the connections from the neuron in the center as an example. The network was composed of $N_m = 10$ neurons. Each neuron receives feedforward inputs and is not connected to the other neurons. Bottom-Left: Dynamics of the synaptic weights \vec{w} for a neuron in the PLnoIF network (neuron 1) in function of the training epochs, and dynamics of the spiking activity of the same neuron. The synaptic weights are ordered along the y-axis from 1 to 100 following the temporal order of a sequence class (class $c=1$). PLR+IF: Each neuron inhibits all the neurons in the network via recurrent inhibition. **e)** Same as in **d)** for the PLR+IF network. Each neuron receives feedforward inputs and is connected to the other neurons through recurrent inhibition. **f)** The synaptic weight matrix of the PLR+IF network at the end of training (epoch 500). The synaptic weights are ordered along the y-axis from 1 to 100 following the temporal order of the input spikes in sequence class $c = 1$. We highlighted the neuron that developed selectivity to class $c = 1$ (neuron 1), thereby assigning credit to the first inputs of the specific class.

neurons in the PLR+IF network decreased substantially (Figure 2b). Importantly, the average duration of the network activity decreased, as many neurons learned to fire at the beginning of distinct sequences (Figure 2c). Additional analyses showed that the network was able to develop sequence selectivity for a vast range of weight initialization (Figure S3a), for different numbers of random sequences in the input pattern (Figure S3b), and for multiple sets of model parameters (Figure S3c).

Together, our results demonstrate that combining the predictive learning rule with inhibitory feedback leads neurons to fire at the start of particular sequences in the input spike train. The increased selectivity is accompanied by the efficient encoding of input sequences.

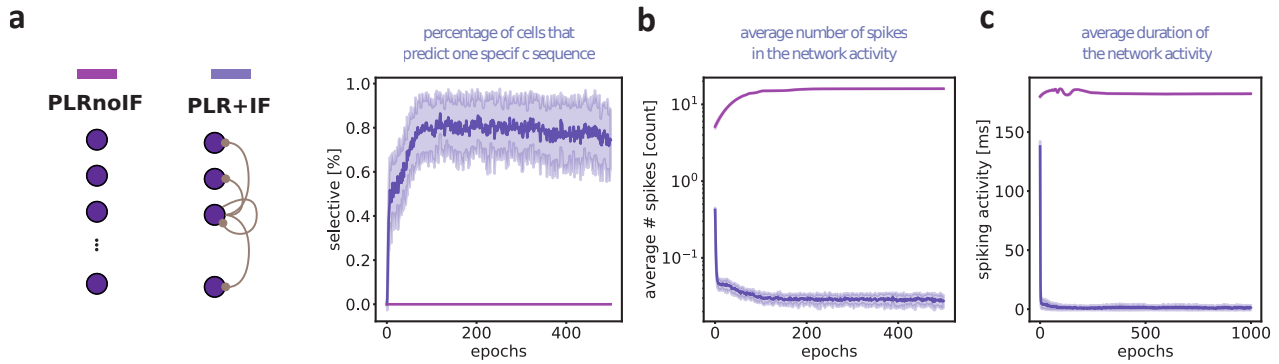


Figure 2: Inhibitory feedback support network selectivity and efficient firing. **a**) Left: Illustration of the two networks and color code. Right: Percentage of selective neurons in the network as a function of the training epochs. A neuron in the network is labeled as *selective* if it fires only for one specific class. The panel shows the percentage of selective neurons for the PLRnoIF and PLR+IF networks, see the color code on the left. **b**) The average number of spikes in the network activity as a function of the training epochs. The total number of spikes is averaged across neurons in the network and across sequence classes ($n_{\text{class}} = 30$). **c**) The average duration of network activity as a function of the training epochs. We computed the temporal difference between the last spike and the first spike of the network activity to estimate the total duration of the network’s activity. The duration of network activity is averaged across neurons and across sequence classes, as in the middle plot. Each panel shows the mean and standard deviation computed over 100 different simulations, and have the same color code as in panel **a**.

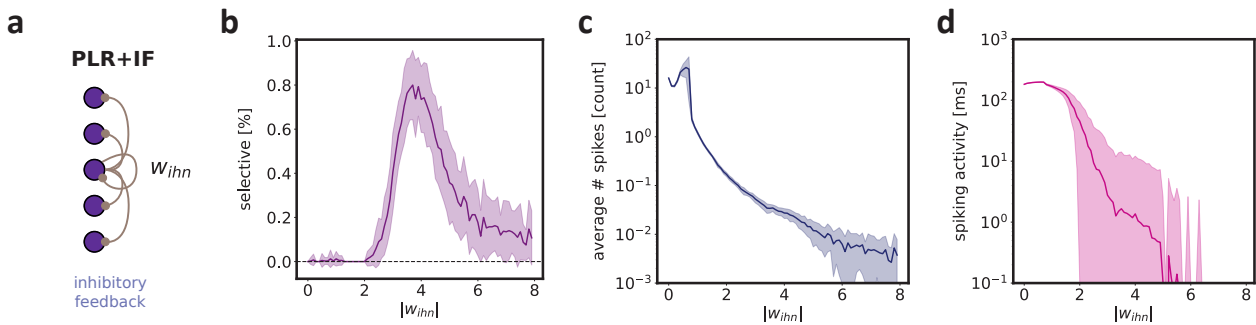


Figure 3: Selectivity and anticipation depend on the strength of inhibition. **a**) Illustration of the global inhibition. The parameter w_{ihn} defines the strength of the recurrent connections (see Methods). **b**) Percentage of selective neurons as a function of the inhibition strength. A neuron is labeled as *selective* if it fires only for sequences belonging to one specific class. **c**) The average number of spikes in the network activity as a function of the inhibition strength. The total number of spikes is averaged across neurons in the network and across sequence classes. **d**) The average duration of network activity as a function of the inhibition strength. The duration of network activity is averaged across neurons and across sequence classes. Each panel shows the mean and standard error of the mean computed over 1000 different simulations.

Intermediate levels of inhibition maximize network selectivity and sequence anticipation

84 In the previous section, we studied the effect of inhibitory feedback in a network of neurons that receives a
 85 feedforward input consisting of many different spike sequences. In the PL+GI network, the neurons developed
 86 selectivity for specific spike sequences, leading to the anticipation of the selected sequences [10]. These results
 87 likely depend on the specific type of interaction between the neurons, in particular on the overall strength of
 88 inhibitory feedback. Thus, our next question was how the level of inhibition influences the development of
 sequence selectivity.

90 We conducted simulations similar to those in the previous section and systematically varied the strength
 of inhibitory feedback (Fig 3a). We then examined the percentage of neurons selective to a single sequence
 92 at the end of training. The percentage of selective neurons showed an inverted U-curve dependence on the
 strength of inhibition. Specifically, we observed that the highest percentage of selective neurons was obtained
 94 for intermediate values of the inhibition strength, whereas weak or strong inhibition resulted in minimal network
 selectivity (Fig 3b). The average number of spikes and the duration of network activity decreased progressively
 96 as the inhibition strength increased (Fig 3c). The network activity was primarily concentrated at the beginning
 of each sequence within the same range of inhibition strength (Fig 3d).

Thus, we found that without inhibitory feedback, neurons fired for every sequence, while excessive inhibitory feedback caused the network activity to collapse. However, we observed that an intermediate level of inhibition was optimal for achieving maximum decorrelation among neurons and enabling the anticipation of predictable inputs. Our results suggest that precisely tuned inhibition is a powerful mechanism to maximize the capacity of neural networks.

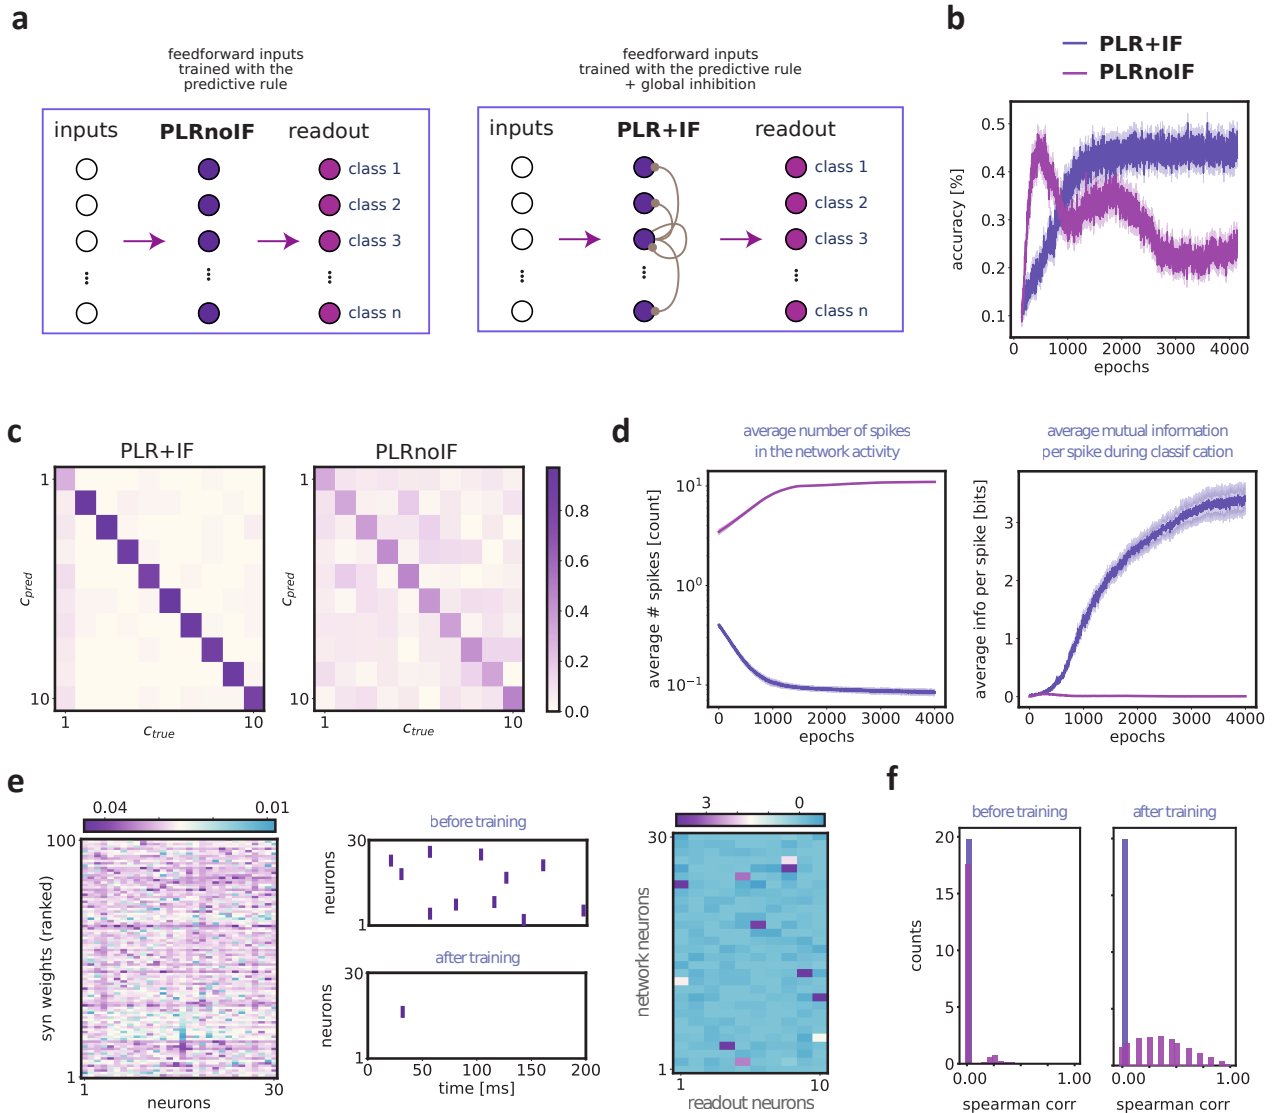


Figure 4: Classification of spike sequences with the predictive learning rule. **a**) The system was composed of an intermediate network (either a PLRnoIF or a PLR+IF network) with $N_m = 30$ neurons and a readout layer with $N_{out} = 10$ neurons. We updated the synaptic connections from the input to the network following the predictive learning rule, while we updated the connections from the network to the readout layer with backpropagation-through-time (BPTT), see Methods. **b**) Classification accuracy of the readout layer. The panel shows the mean and standard error of the mean computed over 100 different simulations. **c**) Confusion matrix at the end of training. **d**) Left: The average number of spikes in the intermediate network activity as a function of the training epochs. The total number of spikes is averaged across neurons in the network and across sequence classes ($n_{class} = 10$). Right: The average mutual information per spike in the classification task. The mutual information is computed with the conditional probability obtained from the confusion matrix (see Methods). All plots show the mean and standard error of the mean computed over 100 different simulations and have the same color code as in **b**. **e**) Left) Synaptic weight matrix of the PLR+IF intermediate network after training (epoch 10000) in one example trial. The y-axis represents the ordering of synaptic weights from 1 to 100, following the temporal order of sequence class $c = 9$. The neuron $n = 16$ exhibits selectivity for sequence class $c = 9$, assigning credit to its first inputs. Center) Network activity associated with sequence class $c = 9$ before training (first epoch) and after training (epoch 10000). Right) The synaptic weight matrix of the readout layer at the end of training (epoch 10000). **f**) Distribution of pairwise Spearman correlation coefficients between spike trains of the intermediate network, respectively at the beginning of training (left, epoch 0) and at the end of training (right, epoch 4000). The correlation coefficients are computed from spike trains obtained with 100 different simulations. The color code is as in **b**.

Inhibition leads to sequence classification with decorrelated and sparse network activity

We found that when combined with a predictive learning rule, inhibitory feedback promotes sparse activity in the network, and allows individual neurons to fire only for specific sequences. This finding would suggest that a decoder should be able to read out the activity of the recurrent network of neurons and correctly classify the spiking sequences.

To address this, we considered a neural network consisting of two modules: (1) An intermediate PLRnoIF or a PLR+IF layer, and (2) an additional readout layer, where each neuron was assigned to a specific sequence class in the input (Figure 4a). The weights from the intermediate layer to the readout layer were trained using the Adam optimizer with a cross-entropy objective function on the membrane potentials of the readout neurons (see Methods). We exposed the network to sequences belonging to one of ten different classes and evaluated its performance on a dataset where half of the examples did not belong to any specific sequence class (Figure 4b). The input was composed of sequences belonging to $n_{\text{class}} = 10$ different classes, and of n_{random} random spiking sequences. We used a training dataset with $n_{\text{random}}/n_{\text{class}} = 1$ and a testing dataset with $n_{\text{random}}/n_{\text{class}} = 10$.

For the PLRnoIF network, we first observed an increase in the classification accuracy and later a decay in performance during training (Figure 4b). This decrease in performance indicates a loss of information in the recurrent network about the sequence inputs. Conversely, the PLR+IF network showed a stable increase in the classification accuracy which was greater than for the PLRnoIF network (Figure 4b-c). This increase in classification accuracy was accompanied by a rise in the percentage of selective cells in the network (Figure S4a). The classification accuracy was driven by a simultaneous decrease in both the prediction error in the intermediate network and the readout layer's classification error, during both the training and the testing phase (Figure S4b). The average firing activity was two orders of magnitude lower in the PLR+IF network, such the average amount of information per spike was considerably higher than in the PLRnoIF case (Fig 4d).

We further illustrated how synaptic plasticity in both the intermediate network and the readout layer leads to fast sequence classification (Figure 4f). We found that in the PLR+IF network neurons fired in response to the initial inputs of specific sequences, allowing for early classification of the sequence identity (Figure 4f). As a result, only a few early spikes from a small group of neurons were sufficient to achieve maximal classification accuracy, resulting in a sparse connectivity matrix at the readout (4f).

The observation that the PLR+IF network had substantially higher classification performance despite much sparser activity suggests that the input was encoded with a much higher degree of statistical independence between neurons. Previous work has indeed suggested that inhibitory feedback leads to a strong reduction in firing rate correlations among excitatory neurons [16]. We, therefore, quantified the pairwise (Spearman) correlations between all the neurons in the intermediate network module. Firing-rate correlations were much stronger in the PLRnoIF network, and were close to zero in the PLR+IF network (Fig 4e). Hence, inhibitory feedback led to the maximization of channel capacity and efficient encoding of the sequential inputs.

In sum, these analyses show that the predictive learning rule combined with inhibitory feedback leads to improved classification of sequence inputs as compared to a network without inhibitory feedback. Furthermore, predictive learning with inhibition leads to sparse coding with a high degree of information per spike and decorrelated firing between neurons.

Fast and efficient classification during predictive plasticity

Finally, we compared the (unsupervised) predictive learning rule with two different supervised learning algorithms: BPTT and BPTT with fixed inhibitory connections (BPTT+IF). In the BPTT network, all synaptic connections in both the intermediate network module and the readout layer were trained based on a global loss (cross-entropy) via backpropagation-through-time [17] (Figure 5a). In the BPTT+IF network, the recurrent network module had a fixed level of global inhibition (as in the PLR+IF network, Figure 5a).

We found that the BPTT and the BPTT+IF networks reached a higher classification performance than the PLR+IF network (about 10% higher) (Figure 5b), which is expected considering that both networks were trained using a supervised learning rule. Nevertheless, the PLR+IF network exhibited several features that were absent in the other networks. First, the average firing activity was orders of magnitude lower in the PLR+IF network, and the amount of information per spike was considerably higher (Figure 5c). Second, the readout layer consistently attained maximum accuracy at progressively earlier times in the trial (Figure 5d). In other words, the PLR+IF network signaled the sequence information earlier in time as compared to the network trained with backpropagation-through-time. Third, the number of spikes required to achieve maximal accuracy

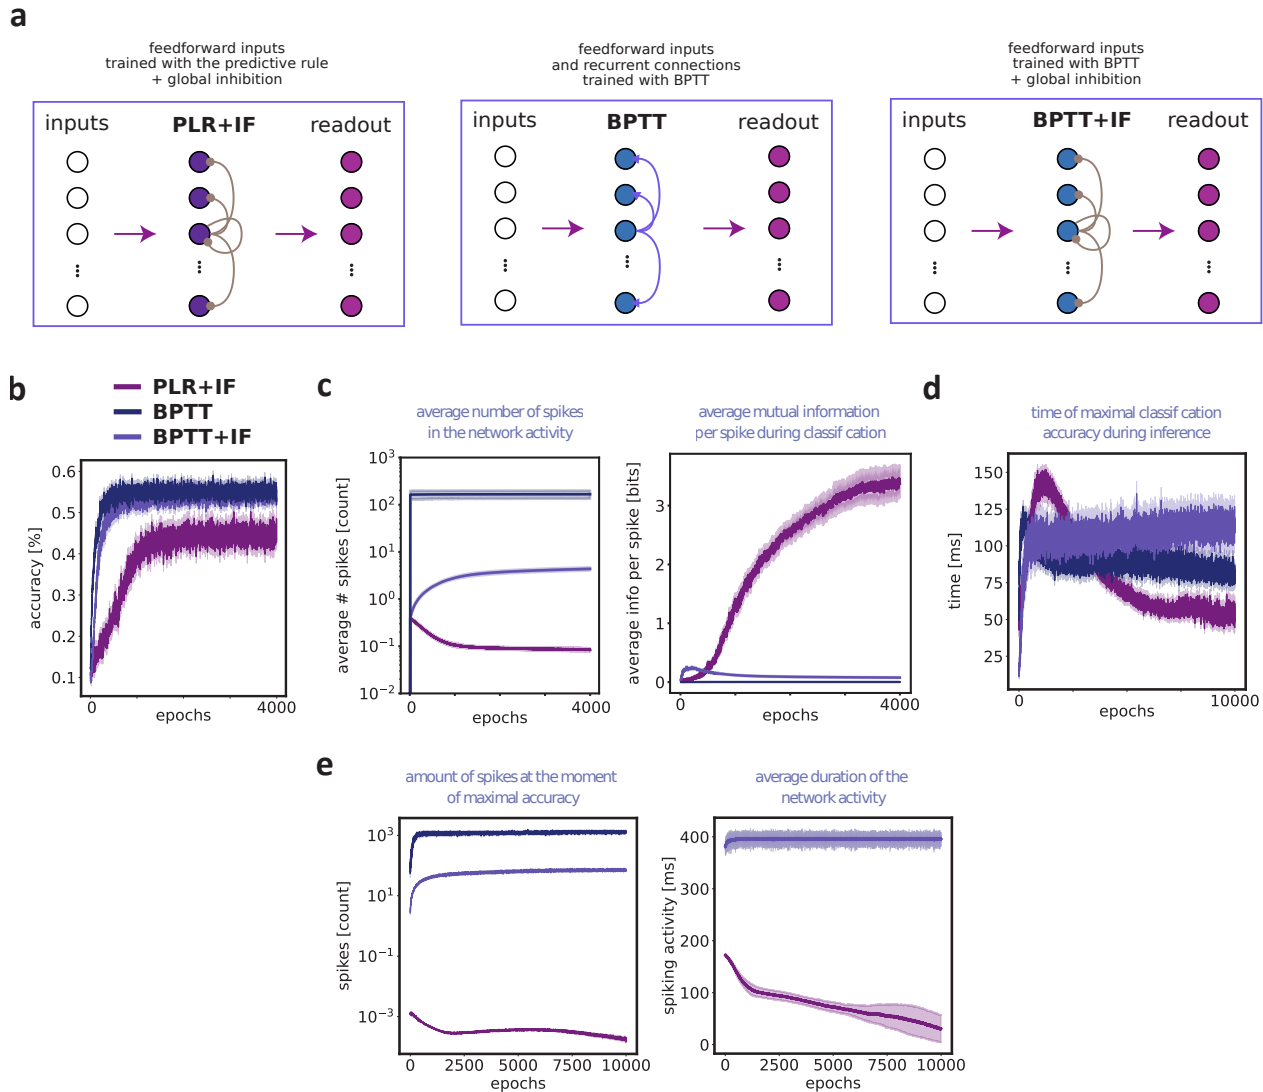


Figure 5: Comparison of the predictive learning rule with back-propagation-through-time algorithms. **a**) Illustration of the three systems. PLR+IF: the system was composed of an intermediate PLR+IF network and a readout layer. We updated the synaptic connections from the input to the network following the predictive learning rule, while we updated the connections from the network to the readout layer with backpropagation-through-time. BPTT: the system was composed of an intermediate network with all-to-all recurrent connections and a readout layer. We updated the synaptic connections from the input to the network, the recurrent connections in the intermediate layer, and the connections from the network to the readout layer with backpropagation-through-time. BPTT+GI: the system was composed of an intermediate network with recurrent connections as in the PLR+IF network. We updated the synaptic connections from the input to the network and the connections from the network to the readout layer with backpropagation-through-time. As in the PLR+IF network, the recurrent connections in the intermediate layer were kept fixed. **b**) Classification accuracy of the readout layer across training. The legend indicates the different training procedures employed for the network. The panel shows the mean and standard deviation computed over 100 different simulations. **c** Left: The average number of spikes in the network activity as a function of the training epochs. The total number of spikes is averaged across neurons in the network and across sequence classes ($n_{class} = 10$). Right: The average mutual information per spike in the classification task. The mutual information is computed with the conditional probability obtained from the confusion matrix (see Methods). **d**) The average time step at which the classification accuracy reached its maximum value. **e**) Left: The average number of spikes observed by the readout layer prior to reaching the maximum accuracy. Right: The average duration of network activity. In panel **c**, **d**, and **e**, all plots show the mean and standard error of the mean computed over 100 different simulations, and have the same color code as in **b**.

varied significantly among the different networks (Figure 5d, left). In particular, the number of spikes in the PLR+IF network rule was about 4 orders of magnitude lower than in the other two networks. This result even held true when we analyzed a network trained with BPTT that included regularization (Figure S5). Finally, the duration of network activity was substantially lower in the PLR+IF network (Figure 5d, right).

Discussion

In this study, we investigated how ensembles of neurons can selectively and efficiently encode multiple temporal input sequences. We implemented a predictive learning rule that adjusts synaptic strengths based on the predictability of pre-synaptic inputs [10]. The main contribution of the present work is to combine the predictive learning rule with recurrent inhibition in a network of neurons. We show that including recurrent inhibition leads to sparse firing in the network, where neurons selectively fire at the beginning of specific sequences. The selectivity of the network was maximized with an intermediate level of inhibition. Moreover, recurrent inhibition led to decorrelated firing between neurons with a high degree of information per spike. Finally, the network activity could correctly reveal the sequence identity, even with just a few early spikes when compared to end-to-end training. Overall, our results demonstrate that fast and energy-efficient classification of different input components can be achieved through the interplay of self-supervised predictive learning rules and inhibitory feedback. The combination of these two mechanisms leads to distinct network activities, which can be used to classify temporal patterns and anticipate their occurrence.

The two main components of our computational model are the predictive learning rule and global inhibition. Neurons endowed with this learning rule learn to fire for the initial, unpredicted inputs in a spiking sequence, resulting in an efficient representation with reduced post-synaptic firing. Importantly, this rule also describes various synaptic plasticity mechanisms that are observed experimentally [10]. Indeed, the synaptic updates crucially depend on changes in the post-synaptic membrane potential that occur at different moments in time, which are believed to trigger LTP and LTD in biological neurons [18, 19]. The predictive learning rule can thereby reproduce STDP mechanisms observed *in-vitro* and describe network phenomena that are usually modeled with STDP rules [2, 4]. For example, the standard asymmetrical STDP window and symmetrical (Hebbian) learning windows come about as a result of predictive plasticity. The impact on synaptic plasticity of pairing frequency [20], initial synaptic strength [21] and higher-order spike patterns (beyond pre-post pairing) [22, 23] are also described by the predictive plasticity rule. Importantly, this rule was derived from an optimization problem based on predictive processes, thus suggesting that STDP is a consequence of a general learning rule given the particular state of the system and the stimulation protocol [24]. We refer to [10] for a detailed analysis of the relation between the predictive plasticity rule and STDP models and experiments.

The inclusion of global inhibition in our model is well-supported by experimental evidence, indicating patterns of massive convergence and massive divergence from excitatory neurons to interneurons, e.g. fast-spiking parvalbumin-positive (PV) cells [25], without connection specificity. These interneurons span several cortical columns and provide fast and unspecific inhibition to pyramidal neurons [26]. Our model captures the essence of an excitatory-inhibitory network model, where single neurons adjust their synaptic weights based on local learning rules. At present, our model did not include recurrent excitation, which is a hallmark of cortical networks [27, 28, 29] but is typically absent in subcortical areas. Thus, in the present form, the model presented here may best approximate unsupervised learning in subcortical networks. Future works should include recurrent excitation in the model and investigate its impact on neuronal selectivity and anticipation of spike sequences.

The recurrent inhibition implemented in this study resembles the mechanism known as *k*-Winner-Take-All (*k*-WTA) with $k = 1$, where only one neuron emerges as the winner, and represents a specific pattern [11]. Experimental evidence demonstrates that lateral inhibition in cortical networks enables this type of operation [30] while also promoting decorrelation of neuronal activity [16]. In agreement with experimental and theoretical findings [31, 32, 16], we found that global inhibition decreases the correlation between neurons, and provides a mechanism for independent encoding of multiple spike sequences. Our results also indicated that an intermediate level of inhibition is optimal for decorrelated firing in the network with high information content per spike. This suggests that there might be a vulnerability of functional cortical activity to GABAergic disturbances or excitation-inhibition balance [33]. Indeed, experimental findings showed that extreme levels of inhibition are associated with psychiatric diseases and that intermediate inhibitory strength correlates to functional brain activity [34, 33].

However, in cortical networks, different cell types contribute to lateral inhibition, exhibiting diverse neuronal properties and connectivity patterns with excitatory neurons [35, 36, 37]. This heterogeneity gives rise to various excitation-inhibition motifs [38, 39], transient and oscillatory activity [40, 41], and significantly affects synaptic plasticity [42]. Therefore, it is important to note that our current work employs a simplified implementation of recurrent inhibition in the network. Future works will incorporate inhibitory neurons into the network

and explore the role of different connectivity schemes in anticipating multiple sequences. This will also provide insights into how the prediction of multiple sequences relates to the balance of excitation and inhibition in neural networks. Furthermore, the emerging evidence of inhibitory plasticity [43, 44] and its diverse consequences on cortical computations [45] introduces another interesting direction for exploration. It is e.g. possible that inhibitory plasticity contributes to reaching an optimum, i.e. intermediate level of global inhibition.

Previous studies have investigated the combination of local learning rules and competition mechanisms. For instance, certain forms of spike-timing-dependent plasticity (STDP), when combined with global winner-take-all mechanisms, can result in neurons representing different segments of a repeating pattern [13, 46]. Additionally, specific STDP windows can produce a reduction in output latency for frequently occurring inputs [8, 47]. Other spike-based learning rules, when combined with WTA mechanisms, have also shown high performance in classification tasks [48, 49]. It is worth noting that the predictive learning rule encompasses multiple STDP phenomena, including multiple STDP windows and higher-order effects (beyond pre-post pairing) [10]. Thereby, our findings provide a general description of sequence anticipation as a result of predictive processes based on synaptic plasticity, whereby STDP mechanisms come about as a consequence of predictive plasticity. Differently from previous works, we also performed a direct comparison of the predictive learning rule with standard training algorithms, and we showed that the combination of self-supervised learning rules and global inhibition led to fast, and energy-efficient encoding of stimulus identity compared to backpropagation-through-time (BPTT).

Our model relies on the temporal relationships between input spikes, and it has been tested on sequential pre-synaptic spike trains. Experimental evidence from various brain regions supports the existence of reliable spike sequences [3, 50]. For instance, the serial firing of neurons can represent the sequential nature of external stimuli. Moving objects can trigger the sequential firing of neurons in the visual system [4], resulting in waves of activity that can contain predictive information about the object's trajectory [51, 52]. In the hippocampus, the sequential activity of place cells encodes the motion of the animal entering different locations [53]. Moreover, spike sequences can also emerge as internally generated activities, e.g. during goal-oriented behaviour [54]. Neurons can engage in structured sequential firing during internal simulations and memory consolidation [55], motor planning, and online control of actions [54]. Interestingly, these sequential activities exhibit experience-dependent facilitation [2], and can be replayed at compressed time scales [56, 5]. Moreover, reliable spike sequences can be induced only by the activation of a few spikes neurons [57, 58]. Our findings support the idea that even a few spikes can initiate a cascade of neuronal activity and that learning can flatten internal activity by assigning credit to inputs with predictive power. Thereby, our model provides insight into the neuronal mechanisms underlying reliable spike sequences and their experience-dependent compression at the network level. Our results suggest a functional role of sequential spiking activity for predictive processes and for the engagement of networks in anticipating future events and organizing behaviour [2].

Finally, the results presented in this study have implications for training spiking neural networks (SNNs) with biologically-inspired algorithms, particularly for assigning credit to local synapses based on non-local information [59]. Indeed, the credit-assignment problem is particularly hard in SNNs, where the errors are propagated both in time and in space and through non-differentiable activation functions [60, 61]. Previous research has approached this problem by drawing insights from neuroscience, including the interplay between different neuronal types [62, 63], the multiplexing of signals through distinct firing patterns [63] and the involvement of different neuronal compartments [64, 65, 66]. Here, we showed that the combination of a local, self-supervised learning rule and an arguably simple inhibitory mechanism leads to distinguishable network activities and successful classification of neuronal inputs. Moreover, the predictive learning rule enables gradient propagation through differentiable functions. Indeed, the loss functions at the single neuron level depend solely on the membrane potential of each respective cell. It is important to note that our model was specifically tested on classifying spike sequences with direct temporal relationships between input spikes. An interesting direction for future research would be to apply our model to different classification tasks and systematically evaluate its performance in various scenarios.

Methods

Network model

We implemented a recurrent neural network composed by N_m neurons, where each i -th neuron in the network received external inputs from N_{presyn} pre-synaptic neurons $\vec{x}_i \in \mathbb{R}^{N_{\text{presyn}}}$. The synaptic weights asso-

ciated with the recurrent connection and with the external inputs were determined by the connectivity matrix
 260 $W_{\text{rec}} \in \mathbb{R}^{N_{\text{nn}}, N_{\text{nn}}}$ and $W \in \mathbb{R}^{N_{\text{presyn}}, N_{\text{nn}}}$, respectively. Each i -th neuron in the network obeyed a discrete-time model of the form

$$\begin{cases} v_{i,t} = \alpha v_{i,t-1} - v_{(\text{th})} s_{i,t-1} + \vec{w}_i^\top \vec{x}_{i,t} + \vec{w}_{\text{rec},i}^\top \vec{s}_t \\ s_{i,t} = H(v_{i,t} - v_{(\text{th})}). \end{cases} \quad (1)$$

262 Here, $v_{i,t} \in \mathbb{R}$ is the membrane potential of the i -th neuron at timestep t , $\alpha \equiv 1 - h/\tau_m$ where τ_m is the membrane time constant and h is the timestep size, $\vec{x}_{i,t} \in \mathbb{R}^{N_{\text{presyn}}}$ is the pre-synaptic input to the i -th neuron at timestep
 264 t , $v_{(\text{th})}$ is the spiking threshold (the subscript (th) : “threshold”), and $H(\cdot)$ is the Heaviside function. The weight vector of the external inputs $\vec{w}_i \in \mathbb{R}^{N_{\text{presyn}}}$ and of the recurrent connections $\vec{w}_{\text{rec},i} \in \mathbb{R}^{N_{\text{nn}}}$ are the i -th column
 266 of the connectivity matrix W and W_{rec} , respectively. The two connectivity matrices W and W_{rec} were defined depending on the specific type of simulations, see the section *Optimization and training schemes*. The variable
 268 $s_{i,t} \in \{0, 1\}$ 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
 270 a constant value $v_{(\text{th})}$ at the next time step. This implementation of the membrane potential reset relates our model to the spike response model [67]. We set $h = 1$ ms in all numerical simulations.

272 We also implemented a readout layer composed of n_{class} neurons, one for each class c . Each i -th neuron in the layer received the activity of the recurrent neural network \vec{s}_t , and the associated synaptic weights were
 274 determined by the connectivity matrix $W_{\text{out}} \in \mathbb{R}^{N_{\text{nn}}, n_{\text{class}}}$. The readout neurons obeyed a discrete-time equation of the form

$$u_{i,t} = \beta u_{i,t-1} + \vec{w}_{\text{out},i}^\top \vec{s}_t. \quad (2)$$

276 Here, $\beta \equiv 1 - h/\tau_{m,\text{out}}$ where $\tau_{m,\text{out}}$ is the membrane time constant of the readout neurons and h is the timestep size, and $\vec{w}_{\text{out},i}$ is the i -th column of the connectivity matrix W_{out} .

Predictive learning rule

278 The predictive learning rule used in this study was adapted from a previously published paper [10]. This rule is derived from an optimization problem in time at the single neuron level. Specifically, each i -th neuron
 280 in the network is assigned with an objective function \mathcal{L}_i as follows,

$$\mathcal{L}_i \equiv \sum_{t=0}^T \mathcal{L}_{i,t} \equiv \sum_{t=0}^T \frac{1}{2} \|\vec{x}_{i,t} - v_{i,t-1} \vec{w}_i\|_2^2, \quad (3)$$

where $\|\cdot\|_2$ is the l_2 -norm, $\vec{x}_{i,t}$ are the external inputs to the i -th neuron in the network, $v_{i,t-1}$ is the membrane
 282 potential of the i -th neuron and \vec{w}_i is the synaptic weight vector of the external inputs, see Equation (1). Here, the objective is to obtain the minimal difference between the input $\vec{x}_{i,t}$ received by the neuron, and its prediction
 284 via $v_{i,t-1}$ and \vec{w}_i . To derive the predictive learning rule, one computes the gradient of \mathcal{L}_i w.r.t. \vec{w}_i ,

$$\vec{\nabla}_{\vec{w}} \mathcal{L}_i = \sum_{t=0}^T \frac{1}{2} \left(\vec{\nabla}_{\vec{w}_i} \mathcal{L}_{i,t} + \frac{\partial \mathcal{L}_{i,t}}{\partial v_{i,t-1}} \vec{\nabla}_{\vec{w}_i} v_{i,t-1} \right). \quad (4)$$

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

$$\vec{\nabla}_{\vec{w}_i} \mathcal{L}_{i,t} = -2(\vec{x}_{i,t} - v_{i,t-1} \vec{w}_i) v_{i,t-1}, \quad (5)$$

while the second term is given by

$$\frac{\partial \mathcal{L}_{i,t}}{\partial v_{i,t-1}} = -2(\vec{x}_{i,t} - v_{i,t-1} \vec{w}_i)^\top \vec{w}_i, \quad (6)$$

288 and by

$$\vec{\nabla}_{\vec{w}_i} v_{i,t} = \left(\alpha - v_{(\text{th})} \frac{\partial s_{i,t-1}}{\partial v_{i,t-1}} \right) \vec{\nabla}_{\vec{w}_i} v_{i,t-1} + \vec{x}_t. \quad (7)$$

We define an influence vector $\vec{p}_{i,t} \equiv \vec{\nabla}_{\vec{w}_i} v_{i,t}$ such that it obeys the recursive equation (see Equation (7))

$$\vec{p}_{i,t} = J_{i,t} \vec{p}_{i,t-1} + \vec{x}_t, \quad (8)$$

where $J_{i,t}$ is the Jacobian from the recurrent equation of the membrane voltage of the i -th neuron in Equation (1) 290

$$J_{i,t} = \frac{\partial v_{i,t}}{\partial v_{i,t-1}} = \alpha - v_{i,t-1}^{(\text{th})} \frac{\partial s_{i,t-1}}{\partial v_{i,t-1}}. \quad (9)$$

The gradient of \mathcal{L}_i w.r.t. \vec{w}_i is then given by 292

$$\vec{\nabla}_{\vec{w}_i} \mathcal{L}_i = - \sum_{t=0}^T \left[(\vec{x}_{i,t} - v_{i,t-1} \vec{w}_i) v_{i,t-1} - 2(\vec{x}_{i,t} + v_{i,t-1} \vec{w}_i)^\top \vec{w}_i \vec{p}_{i,t-1} \right] \quad (10)$$

We define the prediction error $\vec{\epsilon}_{i,t}$ at timestep t as

$$\vec{\epsilon}_{i,t} \equiv \vec{x}_{i,t} - v_{i,t-1} \vec{w}_i, \quad (11)$$

that determines the sign and amplitude of synaptic plasticity, and the global signal $\mathcal{E}_{i,t}$ at timestep t as 294

$$\mathcal{E}_{i,t} \equiv \vec{\epsilon}_{i,t}^\top \vec{w}_i, \quad (12)$$

thereby the gradient of \mathcal{L}_i w.r.t. \vec{w}_i is given

$$\vec{\nabla}_{\vec{w}_i} \mathcal{L}_i = - \sum_{t=0}^T \left[\vec{\epsilon}_{i,t} v_{i,t-1} + \mathcal{E}_{i,t} \vec{p}_{i,t-1} \right], \quad (13)$$

After the end of the period $[0, T]$, the exact gradient can be used to update the weight vector via gradient descent 296

$$\begin{aligned} \vec{w}_{i,k} &= \vec{w}_{i,k-1} - \eta \vec{\nabla}_{\vec{w}_i} \mathcal{L}_i = \vec{w}_{i,k-1} - \eta \sum_{t=0}^T \Delta \vec{w}_{i,t} \\ &= \vec{w}_{i,k-1} + \eta \sum_{t=0}^T \left[\vec{\epsilon}_{i,t} v_{i,t-1} + \mathcal{E}_{i,t} \vec{p}_{i,t-1} \right]. \end{aligned} \quad (14)$$

where η is the learning rate, and $\Delta \vec{w}_i$ is the weight update obtained with the predictive learning rule [10] 298

$$\Delta \vec{w}_{i,t} = \vec{\epsilon}_{i,t} v_{i,t-1} + \mathcal{E}_{i,t} \vec{p}_{i,t-1}. \quad (15)$$

Each neuron in the network receives inputs from other neurons, and the associated recurrent connections are defined by the specific connectivity scheme. The contributions to the gradient of the recurrent connections from neuron j to neuron i in the network are proportional to 300

$$\propto \frac{\partial s_{j,t}}{\partial s_{i,t}} \frac{\partial v_{i,t}}{\partial v_{i,t}}. \quad (16)$$

These contributions have a discontinuous effect at the time of output spikes, and thereby we neglected these contributions to the gradient. Conforming to previous works [68, 61, 69], we also neglected the contribution of the reset mechanisms in Equation (9), that is 304

$$J_{i,t} \simeq \alpha. \quad (17)$$

In the predictive learning rule, the membrane potential $v_{i,t}$ serves as the input variable for the loss function, while $s_{i,t}$ acts as a hidden variable. Thus, the predictive learning rule circumvents the issue of back-propagating through discrete output variables. 306

In this work, we also considered weight updates that take place in real time with the prediction of the pre-synaptic inputs. To do so, we approximated the Equation (14) with the current estimate of the gradient at each timestep t , 308

$$\vec{w}_{i,t} = \vec{w}_{i,t-1} - \eta \vec{\nabla}_{\vec{w}_i} \mathcal{L}_i \Big|_{\vec{w}_i = \vec{w}_{i,t-1}} = \vec{w}_{i,t-1} + \eta \left[\vec{\epsilon}_{i,t} v_{i,t-1} + \mathcal{E}_{i,t} \vec{p}_{i,t-1} \right]. \quad (18)$$

Spike sequences datasets

Each dataset is composed of spike sequences belonging to n_{class} different sequence classes, and comprise the correlated activity of N_{presyn} pre-synaptic neurons. Each sequence class $c \in n_{\text{class}}$ is defined by a unique order of the firing times of the pre-synaptic neurons. The N_{presyn} pre-synaptic neurons fire one spike in each sequence with relative delays of 2 ms, resulting in a total sequence length of $2N_{\text{presyn}}$ ms. We also add to the dataset n_{random} random spiking sequences, that is, sequences whose order of firing was drawn randomly. We created the dataset as follows. First, we set the number of pre-synaptic neurons N_{presyn} and the number of sequence classes n_{class} , then we randomly drew the order of firing of the N_{presyn} pre-synaptic neuron in each class c . Second, we created b_{class} examples of the spike sequences belonging to each class c . Finally, the inputs were convoluted with an exponential kernel with $\tau_x = 2$ ms to replicate the fast dynamics of post-synaptic currents. Thereby, the dataset was composed by $N_{(\text{dataset})} = (b_{\text{class}} * n_{\text{class}} + n_{\text{random}})$ spike sequences. An example set of spike sequences from $N = 10$ pre-synaptic neurons and belonging to $n_{\text{class}} = 4$ classes is as follows:

	1	2	3	4	5	6	7	8	9	10
\mathbf{c}_1	10 ms	18 ms	2 ms	4 ms	12 ms	16 ms	8 ms	0 ms	20 ms	14 ms
\mathbf{c}_2	0 ms	8 ms	12 ms	18 ms	10 ms	6 ms	2 ms	20 ms	4 ms	14 ms
\mathbf{c}_3	14 ms	10 ms	20 ms	4 ms	2 ms	8 ms	16 ms	6 ms	0 ms	12 ms
\mathbf{c}_4	4 ms	2 ms	0 ms	12 ms	20 ms	18 ms	6 ms	16 ms	14 ms	10 ms

Table 1: Example set of spike sequences from $N = 10$ pre-synaptic neurons and belonging to $n_{\text{class}} = 4$ classes

In Figure 1a-b and Figure S1 the dataset was composed by one example ($b_{\text{class}} = 1$) of a single sequence class ($n_{\text{class}} = 1$), and without examples of random spiking sequences ($n_{\text{random}} = 0$).

In Figure 1, Figure 2, Figure 3 and Figure S2 the dataset was composed of several examples ($b_{\text{class}} = 500$) for each sequence class ($n_{\text{class}} = 30$), and by random spiking sequences ($n_{\text{random}} = 500$). In Fig S3 we used the same ratio between class examples and random sequences, but with $n_{\text{class}} = 10$.

In Figure 4, Figure 5, Figure S4, Figure S5 and Figure S2 the dataset was composed of several examples ($b_{\text{class}} = 50$) for each sequence class ($n_{\text{class}} = 10$), and by random spiking sequences ($n_{\text{random}} = 500$).

Objective functions and regularizations

We defined the predictive objective function of a network as

$$\mathcal{L}_{\text{pred}} = \frac{1}{N_{nn}} \sum_{i=1}^{N_{nn}} \mathcal{L}_i, \quad (19)$$

where \mathcal{L}_i is the objective function associated to each i -th neuron in the network, see Equation (3).

We trained the readout layer using supervised learning. During each iteration in the training and testing phase, the actual label y_s of every random sequence was drawn randomly from the n_{class} possible classes, resulting in the following dataset

$$\{(X_s, y_s) | s = 1, \dots, N_{(\text{dataset})}; y_s \in \{1, \dots, n_c\}\}. \quad (20)$$

Here, X_s is an example sequence in the dataset and y_s is the associated label. We applied a cross-entropy loss to the activity of the readout layer \vec{u}_t ,

$$\mathcal{L}_{\text{class}} = -\frac{1}{N_{(\text{dataset})}} \sum_{s=1}^{N_{(\text{dataset})}} \mathbb{1}(i = y_s) \log \left[\frac{\exp(u_{i,t_{\text{max}}})}{\sum_{c=1}^{n_c} \exp(u_{i,t_{\text{max}}})} \right]. \quad (21)$$

Here $\mathbb{1}$ is the indicator function and $u_{i,t_{\text{max}}}$ is the maximal membrane potential of the i -th readout neuron across timesteps, where $t_{\text{max}} = \text{argmax}_t u_{i,t}$.

In Figure S5 we also applied an L1-regularization to the total number of spikes emitted by the network, as to penalize excessive firing

$$\mathcal{L}_{\text{reg}} = \lambda_{L1} \sum_{i=1}^{N_{nn}} \sum_{t=0}^T s_{i,t} \quad (22)$$

where λ_{L1} is a hyperparameter controlling the influence of the regularization term.

Optimization and training schemes

In the following, we provide a description of the optimization techniques and the training schemes used in this study. 342

In Figure 1, Figure 3, Figure S1, Figure S2 and Figure S3, we minimized $\mathcal{L}_{\text{pred}}$ using the online approximation of the predictive learning rule. That is, each i -th column of W was updated following Equation (18). 344

In Figure 4, Figure 5, Figure S4, and Figure S5 we used different networks with different optimization schemes. For the case of the predictive neural network, we employed backpropagation-through-time (BPTT) to minimize both the network's objective $\mathcal{L}_{\text{pred}}$ and the readout layer's objective $\mathcal{L}_{\text{class}}$. For the LIF network with recurrent inhibition, we used BPTT to minimize the objective $\mathcal{L}_{\text{class}}$. We updated all the synaptic weights in the networks, except for the connectivity matrix W_{rec} that was kept fixed. Finally, in the LIF network without recurrent inhibition, we also employed BPTT to minimize the objective $\mathcal{L}_{\text{class}}$ of the readout layer. Unlike the previous case, we updated all the synaptic weights in the networks. 348
350

For all the spiking neural networks that were not optimized with the predictive learning rule, we applied a surrogate gradient to replace the spiking non-linearity with a differentiable function [61]. Specifically, we chose a fast sigmoid function of the form 352
354

$$-v^{(\text{th})} \frac{\partial s_{i,t}}{\partial v_{i,t}} \simeq \frac{1}{\gamma |v_{i,t}| + 1} \quad (23)$$

where γ is a parameter controlling the steepness of the sigmoid function. We optimized the model parameters with the Adam optimizer [70] in all the simulations shown in this study. 356

Weights initialization

In Figure 1, Figure 2, Figure 3, and Figure S1 the synaptic weights from the input dataset to the network were determined by the connectivity matrix $W \in \mathbb{R}^{N_{\text{presyn}} \times N_{nn}}$. Before training, the entries of W were independently drawn from a uniform distribution $\mathcal{U}(0, w_{\text{init}})$, with $w_{\text{init}} > 0$. The connections between neurons in the network were defined by the connectivity matrix $W_{\text{rec}} \in \mathbb{R}^{N_{nn} \times N_{nn}}$. Before training, the connectivity matrix W_{rec} was defined as $W_{\text{rec}} = w_{\text{ihn}} (\mathbb{I}_{nn} + N_{nn}^{-1/2} \mathbb{U}_{nn})$, where \mathbb{I}_{nn} is the identity matrix, N_{nn} is the number of neurons in the network, \mathbb{U}_{nn} is a random matrix drawn from a uniform distribution between $[0, 1]$, and $w_{\text{ihn}} < 0$ defines the overall strength of the recurrent inhibition. The connectivity matrix W_{rec} was fixed during the simulations, that is, the associated synaptic weights were not subject to plasticity. 358
360
362
364

In Figure 4, Figure 5, Figure S4, and Figure S5 we used different types of networks with different initialization schemes. The connections between the network and the readout layer were determined by the connectivity matrix $W_{\text{out}} \in \mathbb{R}^{N_{nn} \times n_{\text{class}}}$. Before training, the entries of W_{out} were independently drawn from a normal distribution $\mathcal{N}(0, 1/\sqrt{N_{nn}})$. For the case of the predictive neural network and the LIF network with inhibition, we initialized the connectivity matrices W and W_{rec} as in Figure 1 and Figure 3. For the case of the LIF network without inhibition and the random neural network, the entries of W and of W_{rec} were independently drawn from the normal distribution $\mathcal{N}(0, 1/\sqrt{N_{\text{presyn}}})$ and $\mathcal{N}(0, 1/\sqrt{N_{nn}})$, respectively. 366
368
370

Data analysis

We calculated each entry C_{ij} of the confusion matrix in Figure 4 as the average number of observations known to belong to sequence class i and predicted by the readout to belong to sequence class j . We calculated the average entry C_{ij} at the end of training (epoch 4000) and over 100 different simulations. 372
374

To calculate the Mutual Information (MI) in Figure 4, we first assume a uniform prior distribution $p(j)$ of the sequence class, and we calculated the joint probability $p(i, j)$ as the product of the prior distribution $p(j)$ and the conditional distribution $p(i|j)$ obtained from the confusion matrix, 376

$$p(i, j) = p(i|j)p(j). \quad (24)$$

Then, we calculated $p(j)$ as the sum of the joint probability $p(i|j)p(j)$ over the expected sequence class $p(j)$, and we finally calculate the MI as 378

$$\text{MI} = \sum_i \sum_j p(i, j) \log \left(\frac{p(i, j)}{p(i)p(j)} \right). \quad (25)$$

380 We averaged the MI in each epoch over 100 different simulations.

382 We calculated the average Spearman correlation $\text{corr}(k, l)$ in Figure 4 between each pair of spike train k and l as

$$\text{corr}(k, l) = \frac{\langle s_k - \mu_k, s_l - \mu_l \rangle}{\sqrt{\langle s_k - \mu_k, s_k - \mu_k \rangle \langle s_l - \mu_l, s_l - \mu_l \rangle}}, \quad (26)$$

384 where s is the spike train, μ is the mean of the respective spike train, and $\langle \cdot \rangle$ is the dot product. We calculated $\text{corr}(k, l)$ of the spike trains of any pair of neurons (k, l) , for each sequence class c and for 100 different simulations.

Acknowledgments

386 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 Richard Gao for helpful comments, edits, and insightful discussions.

Authorship contributions

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

Code availability

392 The code to reproduce all the Figures in the main text and in the Supplementary Materials is freely available at *we will insert GitHub repo link here*.

Declaration of Interests

394 The authors declare no competing interests.

References

- [1] Jakob Hohwy. *The predictive mind*. OUP Oxford, 2013.
- [2] 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. 396
- [3] Xiaofeng Lu and James Ashe. “Anticipatory activity in primary motor cortex codes memorized movement sequences”. In: *Neuron* 45.6 (2005), pp. 967–973. 398
- [4] Shengjin Xu et al. “Activity recall in a visual cortical ensemble”. In: *Nature neuroscience* 15.3 (2012), pp. 449–455. 400
- [5] 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. 402
- [6] Georg B Keller and Thomas D Mrsic-Flogel. “Predictive processing: a canonical cortical computation”. In: *Neuron* 100.2 (2018), pp. 424–435. 404
- [7] 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. 406
- [8] 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. 408
- [9] Pau Vilimelis Aceituno et al. “Learning Cortical Hierarchies with Temporal Hebbian Updates.” In: *bioRxiv* (2023), pp. 2023–01. 410
- [10] Matteo Saponati and Martin Vinck. “Sequence anticipation and spike-timing-dependent plasticity emerge from a predictive learning rule”. In: *Nature Communications* 14.1 (2023), p. 4985. 412
- [11] Tomoki Fukai and Shigeru Tanaka. “A simple neural network exhibiting selective activation of neuronal ensembles: from winner-take-all to winners-share-all”. In: *Neural computation* 9.1 (1997), pp. 77–97. 414
- [12] Robert Coultrip, Richard Granger, and Gary Lynch. “A cortical model of winner-take-all competition via lateral inhibition”. In: *Neural networks* 5.1 (1992), pp. 47–54. 416
- [13] Timothée Masquelier, Rudy Guyonneau, and Simon J Thorpe. “Competitive STDP-based spike pattern learning”. In: *Neural computation* 21.5 (2009), pp. 1259–1276. 418
- [14] Jeffrey S Isaacson and Massimo Scanziani. “How inhibition shapes cortical activity”. In: *Neuron* 72.2 (2011), pp. 231–243. 420
- [15] M Megias et al. “Total number and distribution of inhibitory and excitatory synapses on hippocampal CA1 pyramidal cells”. In: *Neuroscience* 102.3 (2001), pp. 527–540. 422
- [16] Alfonso Renart et al. “The asynchronous state in cortical circuits”. In: *science* 327.5965 (2010), pp. 587–590. 424
- [17] Timothy P Lillicrap et al. “Backpropagation and the brain”. In: *Nature Reviews Neuroscience* 21.6 (2020), pp. 335–346. 426
- [18] 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. 428
- [19] Claudia Clopath et al. “Connectivity reflects coding: a model of voltage-based STDP with homeostasis”. In: *Nature neuroscience* 13.3 (2010), p. 344. 430
- [20] Per Jesper Sjöström, Gina G Turrigiano, and Sacha B Nelson. “Rate, timing, and cooperativity jointly determine cortical synaptic plasticity”. In: *Neuron* 32.6 (2001), pp. 1149–1164. 432
- [21] 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. 434
- [22] Robert C Froemke and Yang Dan. “Spike-timing-dependent synaptic modification induced by natural spike trains”. In: *Nature* 416.6879 (2002), pp. 433–438. 438

- 440 [23] Thomas Nevian and Bert Sakmann. “Spine Ca²⁺ signaling in spike-timing-dependent plasticity”. In: *Journal of Neuroscience* 26.43 (2006), pp. 11001–11013.
- 442 [24] Harel Z Shouval, Samuel S-H Wang, and Gayle M Wittenberg. “Spike timing dependent plasticity: a
444 consequence of more fundamental learning rules”. In: *Frontiers in computational neuroscience* 4 (2010),
p. 19.
- [25] Adam M Packer and Rafael Yuste. “Dense, unspecific connectivity of neocortical parvalbumin-positive
446 interneurons: a canonical microcircuit for inhibition?” In: *Journal of Neuroscience* 31.37 (2011), pp. 13260–
13271.
- 448 [26] Frédéric Pouille and Massimo Scanziani. “Enforcement of temporal fidelity in pyramidal cells by somatic
feed-forward inhibition”. In: *Science* 293.5532 (2001), pp. 1159–1163.
- 450 [27] Rodney J Douglas et al. “Recurrent excitation in neocortical circuits”. In: *Science* 269.5226 (1995),
pp. 981–985.
- 452 [28] Ho Ko et al. “Functional specificity of local synaptic connections in neocortical networks”. In: *Nature*
473.7345 (2011), pp. 87–91.
- 454 [29] Wolf Singer. “Recurrent dynamics in the cerebral cortex: Integration of sensory evidence with stored
knowledge”. In: *Proceedings of the National Academy of Sciences* 118.33 (2021).
- 456 [30] Zhi-Hong Mao and Steve G Massaquoi. “Dynamics of winner-take-all competition in recurrent neural
networks with lateral inhibition”. In: *IEEE transactions on neural networks* 18.1 (2007), pp. 55–69.
- 458 [31] Jessica A Cardin. “Inhibitory interneurons regulate temporal precision and correlations in cortical cir-
cuits”. In: *Trends in neurosciences* 41.10 (2018), pp. 689–700.
- 460 [32] Molly M Huntsman et al. “Reciprocal inhibitory connections and network synchrony in the mammalian
thalamus”. In: *Science* 283.5401 (1999), pp. 541–543.
- 462 [33] Jacob B Ruden, Laura L Dugan, and Christine Konradi. “Parvalbumin interneuron vulnerability and
brain disorders”. In: *Neuropsychopharmacology* 46.2 (2021), pp. 279–287.
- 464 [34] Brielle R Ferguson and Wen-Jun Gao. “PV interneurons: critical regulators of E/I balance for prefrontal
cortex-dependent behavior and psychiatric disorders”. In: *Frontiers in neural circuits* 12 (2018), p. 37.
- 466 [35] Kenneth A Pelkey et al. “Hippocampal GABAergic inhibitory interneurons”. In: *Physiological reviews*
97.4 (2017), pp. 1619–1747.
- 468 [36] Henry Markram et al. “Interneurons of the neocortical inhibitory system”. In: *Nature reviews neuro-
science* 5.10 (2004), pp. 793–807.
- 470 [37] Xiangmin Xu, Keith D Roby, and Edward M Callaway. “Immunochemical characterization of inhibitory
mouse cortical neurons: three chemically distinct classes of inhibitory cells”. In: *Journal of Comparative
472 Neurology* 518.3 (2010), pp. 389–404.
- [38] Thilo Womelsdorf et al. “Dynamic circuit motifs underlying rhythmic gain control, gating and integra-
474 tion”. In: *Nature neuroscience* 17.8 (2014), pp. 1031–1039.
- [39] Martin Vinck et al. “Principles of large-scale neural interactions”. In: *Neuron* 111.7 (2023), pp. 987–
476 1002.
- [40] Agostina Palmigiano et al. “Flexible information routing by transient synchrony”. In: *Nature neuro-
478 science* 20.7 (2017), pp. 1014–1022.
- [41] Irene Onorato et al. “Distinct roles of PV and Sst interneurons in visually-induced gamma oscillations”.
480 In: *bioRxiv* (2023), pp. 2023–04.
- [42] Juan Burrone, Michael O’Byrne, and Venkatesh N Murthy. “Multiple forms of synaptic plasticity trig-
482 gered by selective suppression of activity in individual neurons”. In: *Nature* 420.6914 (2002), pp. 414–
418.
- 484 [43] Henning Sprekeler. “Functional consequences of inhibitory plasticity: homeostasis, the excitation-inhibition
balance and beyond”. In: *Current opinion in neurobiology* 43 (2017), pp. 198–203.

- [44] Christoph Miehl and Julijana Gjorgjieva. “Stability and learning in excitatory synapses by nonlinear inhibitory plasticity”. In: *PLOS Computational Biology* 18.12 (2022), e1010682. 486
- [45] Auguste Schulz et al. “The generation of cortical novelty responses through inhibitory plasticity”. In: *Elife* 10 (2021), e65309. 488
- [46] James Humble, Susan Denham, and Thomas Wennekers. “Spatio-temporal pattern recognizers using spiking neurons and spike-timing-dependent plasticity”. In: *Frontiers in computational neuroscience* 6 (2012), p. 84. 490
492
- [47] Pau Vilimelis Aceituno, Masud Ehsani, and Jürgen Jost. “Spiking time-dependent plasticity leads to efficient coding of predictions”. In: *Biological cybernetics* 114.1 (2020), pp. 43–61. 494
- [48] Xiaohui Xie, Richard Hahnloser, and H Sebastian Seung. “Learning winner-take-all competition between groups of neurons in lateral inhibitory networks”. In: *Advances in neural information processing systems* 13 (2000). 496
- [49] Raphaela Kreiser et al. “On-chip unsupervised learning in winner-take-all networks of spiking neurons.” In: *BioCAS*. 2017, pp. 1–4. 498
- [50] Andrew J Zimnik and Mark M Churchland. “Independent generation of sequence elements by motor cortex”. In: *Nature neuroscience* 24.3 (2021), pp. 412–424. 500
- [51] Michael J Berry et al. “Anticipation of moving stimuli by the retina”. In: *Nature* 398.6725 (1999), pp. 334–338. 502
- [52] Giacomo Benvenuti et al. “Anticipatory responses along motion trajectories in awake monkey area V1”. In: *BioRxiv* (2020), pp. 2020–03. 504
- [53] John O’Keefe and Jonathan Dostrovsky. “The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat.” In: *Brain research* (1971). 506
- [54] Giovanni Pezzulo et al. “Internally generated sequences in learning and executing goal-directed behavior”. In: *Trends in cognitive sciences* 18.12 (2014), pp. 647–657. 508
- [55] Yingxue Wang et al. “Theta sequences are essential for internally generated hippocampal firing fields”. In: *Nature neuroscience* 18.2 (2015), pp. 282–288. 510
- [56] Kamran Diba and György Buzsáki. “Forward and reverse hippocampal place-cell sequences during ripples”. In: *Nature neuroscience* 10.10 (2007), pp. 1241–1242. 512
- [57] Mike Hemberger et al. “Reliable sequential activation of neural assemblies by single pyramidal cells in a three-layered cortex”. In: *Neuron* 104.2 (2019), pp. 353–369. 514
- [58] Juan Luis Riquelme et al. “Single spikes drive sequential propagation and routing of activity in a cortical network”. In: *Elife* 12 (2023), e79928. 516
- [59] Timothy P Lillicrap and Adam Santoro. “Backpropagation through time and the brain”. In: *Current opinion in neurobiology* 55 (2019), pp. 82–89. 518
- [60] Friedemann Zenke and Surya Ganguli. “Superspike: Supervised learning in multilayer spiking neural networks”. In: *Neural computation* 30.6 (2018), pp. 1514–1541. 520
- [61] 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. 522
524
- [62] Nicolas Perez-Nieves et al. “Neural heterogeneity promotes robust learning”. In: *Nature communications* 12.1 (2021), p. 5791. 526
- [63] Alexandre Payeur et al. “Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits”. In: *Nature neuroscience* 24.7 (2021), pp. 1010–1019. 528
- [64] Jordan Guerguiev, Timothy P Lillicrap, and Blake A Richards. “Towards deep learning with segregated dendrites”. In: *Elife* 6 (2017), e22901. 530
- [65] João Sacramento et al. “Dendritic cortical microcircuits approximate the backpropagation algorithm”. In: *Advances in neural information processing systems* 31 (2018). 532

- 534 [66] Alexander Meulemans et al. “Credit assignment in neural networks through deep feedback control”. In:
Advances in Neural Information Processing Systems 34 (2021), pp. 4674–4687.
- 536 [67] Wulfram Gerstner et al. *Neuronal dynamics: From single neurons to networks and models of cognition*.
Cambridge University Press, 2014.
- 538 [68] Guillaume Bellec et al. “Long short-term memory and learning-to-learn in networks of spiking neurons”.
In: *arXiv preprint arXiv:1803.09574* (2018).
- 540 [69] 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.
- 542 [70] Diederik P Kingma and Jimmy Ba. “Adam: A method for stochastic optimization”. In: *arXiv preprint*
arXiv:1412.6980 (2014).

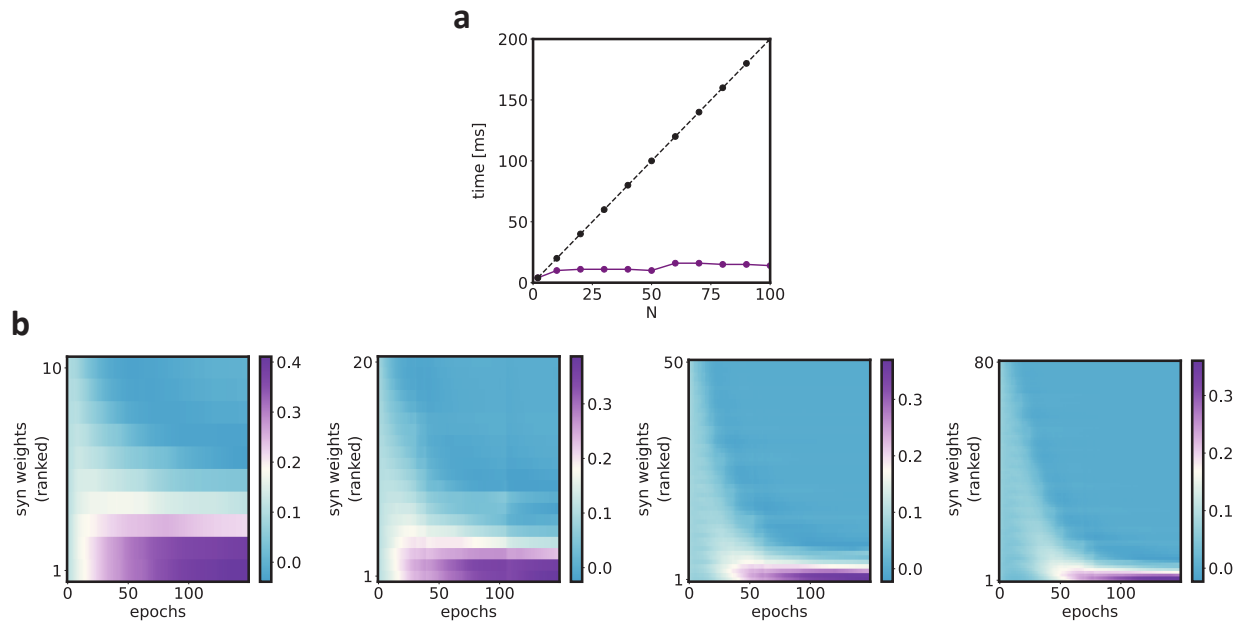


Figure S1: Dependence of sequence anticipation to model parameters - relates to Figure 1 in the main text. a) The purple line represents the duration of the neuron activity at the end of training (epoch 150) for different numbers of pre-synaptic neurons N . We performed the same simulations with the same model parameters for each value of N . The black line represents the duration of the associated pre-synaptic spike sequence. As in Figure 1, the N pre-synaptic neurons fired sequentially with relative delays of 2 ms, resulting in a total sequence length of $2 * N$ ms in each case. The neuron learns to fire at the beginning of each sequence, regardless of the number of pre-synaptic neurons N . b) Dynamics of the synaptic weights \vec{w} as a function of the training epochs, and for different numbers of pre-synaptic inputs N . From left to right: $N = 10$, $N = 20$, $N = 50$, $N = 80$. In each plot, the synaptic weights are ordered along the y-axis from 1 to N following the temporal order of the sequence.

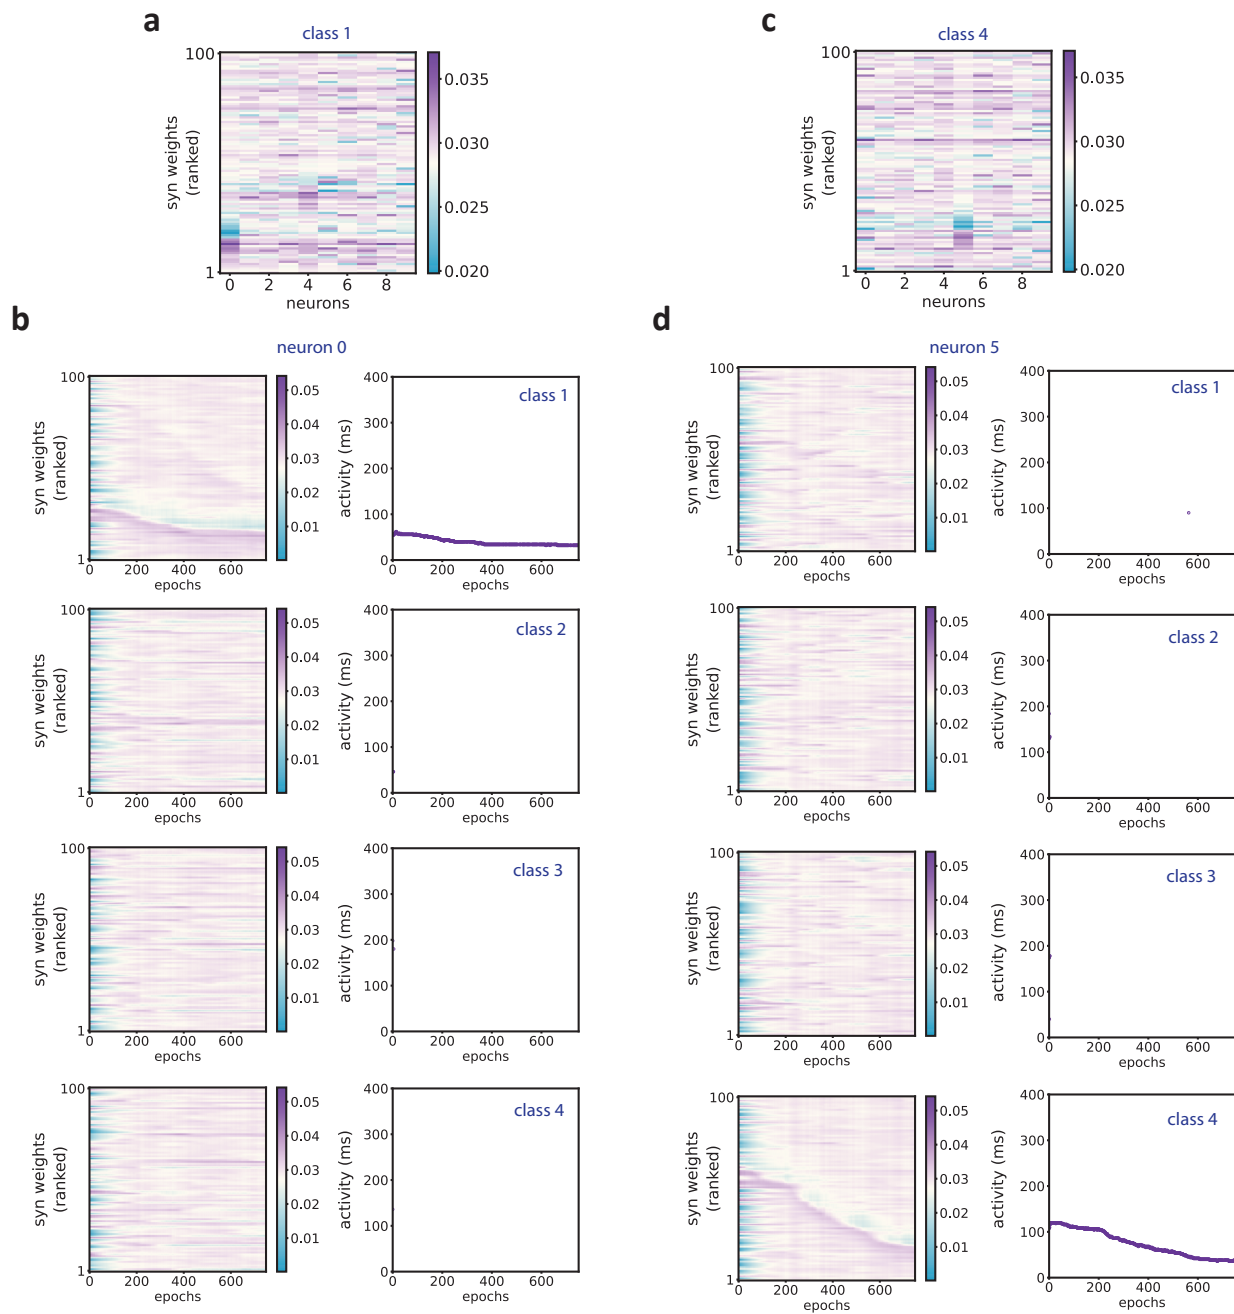


Figure S2: Evolution of the synaptic weights and the spiking activity for selective neurons - relates to Figure 1 in the main text. Simulations were performed as in Figure 1, but with different network initialization. **a)** The synaptic weight matrix of the network at the end of training (epoch 750). The input was composed of sequences belonging to 30 different classes ($n_{\text{class}} = 30$). The synaptic weights are ordered along the y-axis from 1 to 100 following the temporal order of the input spikes in sequence class $c = 1$. Neuron $n = 0$ developed selectivity to class $c = 1$, and assigned credit to the first inputs in the order of class $c = 1$. **b)** Left) Dynamics of the synaptic weights \vec{w} of the neuron $n = 0$ as a function of the training epochs. In each plot, the synaptic weights are ordered along the y-axis from 1 to 100 following the temporal order of the specific sequence class c . The synaptic weights dynamics resemble the one in Figure 1d only for class $c = 1$. The neuron developed selectivity due to the recurrent inhibition in the network, thereby assigning credit only to the first inputs of class $c = 1$. Right) Dynamics of the spiking activity of the neuron $n = 0$ as a function of the training epochs. Each plot shows the dynamics of the spiking activity across epochs when the neuron was presented to the specific sequence class c . The neuron developed selectivity by firing only when presented to class $c = 1$. The neuron grouped its activity earlier in time, such that it eventually learned to fire for the first inputs in the sequence class $c = 1$. **c)** Same as in **a)** for the sequence class $c = 4$. **d)** Same as in **b)** for the neuron $n = 5$.




Figure S3: Dependence of network selectivity to model parameters - relates to Figure 1 in the main text. Simulations were performed as in Figure 1, but with $n_{class} = 10$. **a)** Network selectivity is shown as a function of the mean of the initial weight matrix for the external input W . The network selectivity 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 1000 epochs. **b)** As in panel a, but as a function of the percentage of random sequences in the dataset. **c)** As in panel a, but as a function of the membrane time constant τ_m and of the spiking threshold $v_{(th)}$. Here we performed the simulations with $b_{class}/b_{random} = 1$. The network gets selective for a vast range of model parameters. The highest degree of selectivity is obtained at the border of the region where the neurons in the network start spiking.

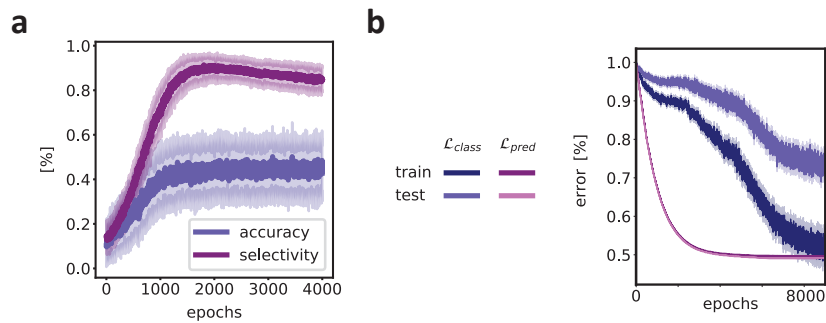


Figure S4: Model performances and network selectivity - relates to Figure 4 in the main text. **a)** Classification accuracy of the readout layer and percentage of selective neurons in the PLR+GI network, see the legend. The panel shows the mean and standard deviation computed over 100 different simulations. **b)** Evolution of the objectives \mathcal{L}_{class} and \mathcal{L}_{pred} during the training and testing phase. The random sequences comprised 10% of the dataset during the training phase, and 50% of the dataset during the testing phase. All the objective functions are normalized to the initial value at epoch 0.




Figure S5: Effect of regularization on the classification of spike sequences - relates to Figure 5 in the main text. Simulations were performed as in Figure 5, but here we compared the performance of the PLR+GI network with the BPTT network augmented with a $L1$ regularization on the network activity (reg_BPTT). Specifically, we trained the network end-to-end on the total objective $\mathcal{L}_{tot} = \mathcal{L}_{class} + \lambda_{L1}\mathcal{L}_{reg}$, where $\mathcal{L}_{reg} = \sum_{t=0}^T \sum_{i=1}^{N_n} |s_{i,t}|$ and λ_{L1} is a hyperparameter that controls the influence of the regularization term, see Methods. **a)** Evolution of the objectives \mathcal{L}_{class} for the train and test sets, and of the objective \mathcal{L}_{reg} for the train set. All the objective functions are normalized to the initial value at epoch 0. **b)** The classification accuracy of the readout layer during training is shown on the left. The reg_BPTT network achieved similar performances to the BPTT network in Figure 5d. The average time step at which the classification accuracy reached its maximum value is shown in the center. The PLR+GI network reaches its maximal performance in half the time compared to the reg_BPTT network. The average number of spikes observed by the readout layer prior to reaching the maximum accuracy is presented on the right. As expected, the regularization term limits the total amount of spikes in the network activity.