Incremental Self-Organization of Spatio-Temporal Spike Pattern Detection

Mohammad Dehghani-Habibabadi[1,*], Lenny Müller[1], and Klaus Pawelzik[1]

[1]Institute for Theoretical Physics, University of Bremen, 28359 Bremen, Germany
[*]mohammad@neuro.uni-bremen.de

ABSTRACT

Brains learn new information while retaining already acquired contents. It is not known by what mechanisms synapses preserve previously stored memories while they are plastic and contribute to further learning. To understand how this notorious stability-plasticity dilemma might be resolved, we investigate a one layer self-supervised neural network that incrementally learns to detect new spatio-temporal spike patterns without loosing existing selectivities. A plausible combination of Hebbian mechanisms, hetero-synaptic plasticity, and synaptic scaling enables unsupervised learning of spatio-temporal input patterns by single neurons. In networks acquisition of different patterns one after the other is achieved by including pre-synaptic hetero-synaptic plasticity which enforces differentiation of the neurons' selectivities. We found, however, that only when additionally the training spikes are both, jittered and stochastic past memories persist despite ongoing learning. This input variability selects a subset of weights and drives them into a regime where synaptic scaling induces self-stabilization. Thereby our model provides a novel explanation for the stability of synapses related to preexisting contents despite ongoing plasticity, and suggests how nervous systems could incrementally learn and exploit temporally precise Poisson rate codes.

Significance Statement

Activity-dependent changes in synaptic efficacy are thought to underlie learning. While ongoing synaptic plasticity is necessary for learning new content, it is detrimental to the traces of previously acquired memories. Here, we show how memories for spatio-temporal patterns can be protected from overwriting. A combination of biologically plausible synaptic plasticity mechanisms turns single neurons into robust detectors of statistically dominant input patterns. For networks, we find that memory stability is achieved when the patterns to be learned are temporally sloppy and noisy, as opposed to being frozen. This variability drives the relevant synaptic weights to large efficacies, where they become self-reinforcing and continue to support the initially learned patterns. As a result, such a network can incrementally learn one pattern after another.

Introduction

Memory persistence depends on the stability of synaptic weight patterns in the neuronal networks that encode memories[1]. In order to learn from and adapt to new experiences, however, a brain is required to be plastic as well. In fact, plasticity is a permanent mechanism, and neurons constantly modify their synapses also to fire at the desired biological firing rate[2,3]. For understanding learning, memory consolidation and retrieval it is essential to investigate how neural networks maintain stability in the face of ongoing plasticity. Deeper insights into the balance between stability and plasticity in biological neural networks also have the potential to lead to the development of more robust and efficient artificial neural networks. In particular, stable networks may fail to learn new information, while hyperplastic networks may suffer from forgetting, which occurs if previously learned information is "forgotten" when new information is learned[4–6]. Therefore, it is a matter of question what mechanisms are required for a neural network to be both stable and plastic, the notorious so-called stability-plasticity dilemma (SPD)[7].
In the past biologically more or less plausible algorithms were proposed for learning spatio-temporal spike patterns. A basic Tempotron can distinguish between patterns based on their duration and synaptic time constants[8], however, utilizes a supervisory signal. While the Tempotron provides no spike timing information, supervised learning algorithms for Chronotrons can train neurons to fire at specific times during a pattern[9–11]. Toward a biological learning algorithm, as a network always receives input, it is required to be quiescent when receiving random input and to fire when a specific pattern is present, which is embedded in background activity. This can be achieved through correlation based learning based on the N-methyl-D-Aspartate (NMDA) receptors[12–14] by selecting synapses with a sufficiently significant correlation between input and membrane potential[15]. The selection of synapses can be done also completely unsupervised based on pre- and post-synaptic hetero-synaptic plasticity[16]. In fact, a combination of fundamental but realistic mechanisms, including synaptic scaling, Hebbian mechanisms, and hetero-synaptic plasticity leads to a balance of excitatory and inhibitory inputs, and allows individual neurons to become detectors of

repetitive patterns in the input without any supervision. It was already found that the memory for the learned pattern is retained if the pattern is not shown again and only noise is presented afterwards which enables learning also of patterns repeating at low rates. If the learned embedded patterns are not presented to the neuron and the scaling term is large enough, all weights become scaled up until the neuron fires at the desired rate. This explains why memory is maintained when learning continues with random input patterns[16]. In stark contrast, however, if a different embedded pattern is presented after learning a first pattern, this single neuron model encounters catastrophic forgetting: the memory for the original pattern becomes lost, and the new pattern is learned.

A system with several output neurons has the potential to detect and identify several patterns[16]. But it is unknown by which mechanisms the stability-plasticity dilemma (SPD) might be alleviated for neuronal networks that serve spike pattern detection. In this study, we first present realistic synaptic mechanisms enabling neurons to learn spatio-temporal spike patterns. Then, we investigate the conditions for which a learned weight distribution remains particularly stable during ongoing plasticity when the learned patterns are absent in the input. Finally, we identify conditions for which new patterns are indeed learnable in assemblies of output neurons such that the weights related to previously learned patterns are maintained, which allows for incremental learning.

Results

Neurons receive input from a large number of excitatory and inhibitory neurons. For simplicity, in the following we demonstrate our results with the simple example of 400 excitatory and 100 inhibitory neurons which all fire with Poisson statistics and fixed rates of 5Hz and 20Hz, respectively. Embedded patterns are short epochs with the same statistics, however contain consistent patterns which repeat either with frozen spike times or are noisy versions of a fixed pattern.

The learning algorithm (see Materials and Methods) combines plausible synaptic plasticity mechanisms for which the signs of the weights remain unchanged throughout the learning process, i.e., Dale's law is enforced.

We incorporate homeostatic plasticity that is applied to excitatory synapses and works to regulate the firing rate to match the biological firing rate[2,3]. This means that when the firing rate of a neuron exceeds the desired rate on long time scales, all excitatory input synapses undergo a process of down-scaling, and they are scaled up when the firing rate is lower than desired. This synaptic scaling is crucial for the maintenance of system stability. However, because this mechanism is universally applied to all excitatory synapses, regardless of their efficacy, it is not sufficient to identify embedded patterns, that is, the neuron's spike timing does not become pattern specific under synaptic scaling.

To determine whether a neuron has learned or memorized a particular pattern, we look at the timing of the spikes in the post-synaptic neurons. If all output spikes occur during the epoch when the embedded pattern is present in the input, we can conclude that the neuron has learned that pattern. To evaluate how well the learning mechanism and the memory recall work in the recognition of random patterns, we determine the average percentage of spikes that correctly identify the patterns, which is denoted by R,

$$R = \left\langle \frac{n_s^p(L)}{n_s + \zeta} \right\rangle_\mu , \tag{1}$$

where n_s is the total number of spikes observed during a test period and n_s^p is the number of spikes related to the presence of the pattern to be detected. To account also for spikes that may occur shortly after the pattern due to the finite decay time of the excitatory synaptic kernel, the test window is extended by L ms after the pattern ends. To ensure a definitive result ($R = 0$) when no spikes are observed at all, a small arbitrary positive value $\zeta << 1$ is added to the denominator. The ratio is then averaged over an ensemble of inputs denoted by μ, where each input contains the embedded patterns and independent random background spikes.

Furthermore we consider Hebbian mechanisms of synaptic plasticity and their potential role in the self-organization of spike pattern selectivity. Hebbian mechanisms affect both excitatory and inhibitory synapses and depend on correlations between input kernels and deviations of the membrane potential from a given threshold. For changes in inhibitory synapses, weight enhancement is driven by positive deflections, while negative deflections contribute to their attenuation. For excitatory synapses, we allow only positive deflections to contribute, mimicking NMDA-dependent processes. It is well understood that without any further constraints, this approach would lead to an instability of excitatory synaptic efficacy. This is avoided by a combination of two simple but biologically plausible mechanisms.

First: By capping excitatory synaptic weights at their upper limits, unbounded growth is effectively prevented.

Second: Decreases in excitatory synaptic weights are induced by hetero-synaptic plasticity, where pre- and post-synaptic components of each synapse are modified differently. Hetero-synaptic plasticity couples the weight changes of different synapses in such a way that the resources required for significant increases in efficacy of a component are drawn from the components of other synapses which otherwise would increase only weakly. Thereby, competition for synaptic increases is

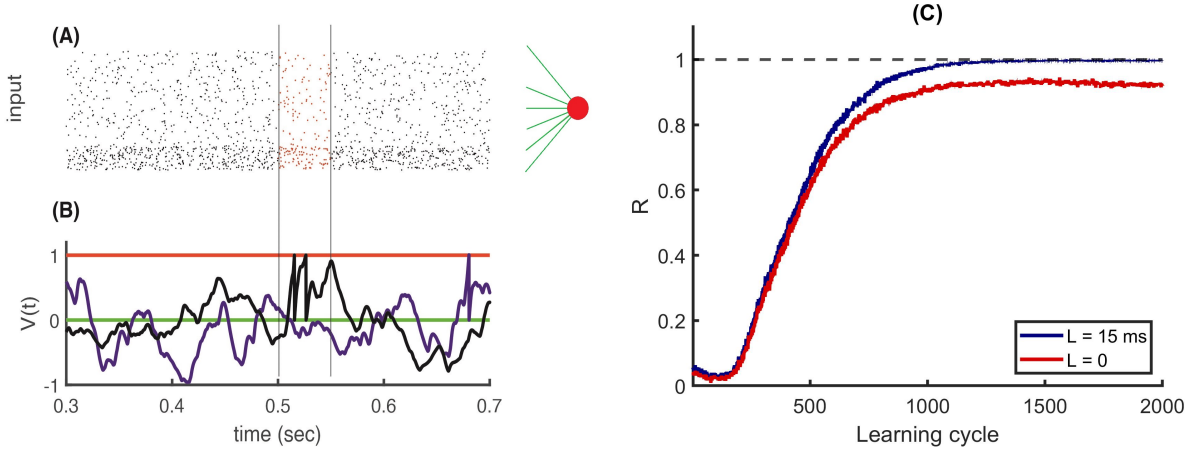

Figure 1. Learning of one embedded pattern. A: the input activity and B: membrane potential of a neuron in response to input from 500 afferents, with 80% excitatory and 20% inhibitory neurons. The excitatory neurons fired at a rate of 5 Hz, while the inhibitory neurons fired at 20 Hz. A 50 ms random pattern is embedded in the red area between two black vertical lines. The membrane potential of the neuron is in black trace showing that learning reduced fluctuations in the membrane potential outside the embedded pattern, compared to the case where the weights were learned from the random background alone. As a result, two spikes occurred during the embedded pattern time, between the vertical black lines. The green line represents the resting potential, the red line represents the threshold, the blue line represents the membrane potential after weight initialization by learning with random patterns only, and the black line represents the membrane potential after learning. C: Learning performance R versus learning cycle, for $L = 0$ and $L = 15$ ms (based on an average of 500 simulations). Note that in this example, the original embedded pattern is shown for each step, and no noise is applied during the process.

induced with the consequence that synapses receiving relatively small signals for increases are weakened. In pre-synaptic hetero-synaptic plasticity a large increase of a pre-synaptic component constrains or even inverts a potential increase of all pre-synaptic components originating from the pre-synaptic neuron to which these components belong. In post-synaptic hetero-synaptic plasticity it is the post-synaptic neuron that provides the limited resources for efficacy increases such that a substantial increase in one post-synaptic component (e.g. receptor density) comes at expense of the post-synaptic components of other synapses that target the same neuron. Note that these implementations of hetero-synaptic plasticity do not imply normalization of excitatory weights (see Materials and Methods).

1 Single Post-Synaptic Neuron

Here, we first reproduce the result from[16] demonstrating that the mechanisms mentioned above enable a single post-synaptic neuron to learn embedded patterns. Since here only one postsynaptic neuron is taken into account pre-synaptic hetero-synaptic plasticity doesn't participate (for details see Material and Methods). We first show that the resulting balance of excitatory and inhibitory inputs into the target neuron yields robustness of pattern detection against a range of stochastic perturbations, which is particularly significant when the training patterns are generated from Poisson processes with temporally modulated rates. Subsequently, we demonstrate that the trained neuron retains its selectivity when after learning a specific pattern only random inputs are presented. This changes dramatically when instead of random spikes a different pattern is repeatedly presented: here the neuron looses its selectivity for the first pattern and becomes selective for the second.

Figure 1A shows the membrane potential of the neuron before and after learning a random fixed (frozen) pattern of length 50 ms that has the same statistics as the random background. Note that spikes are elicited already before learning the pattern because we always initialize the weights by applying the plasticity mechanisms with random spike patterns with the same statistics (Poisson process with fixed firing rates: 5Hz for the 400 excitatory input neurons, and 20 Hz for the 100 inhibitory input neurons). Thus, the post-synaptic neurons already fire with average rate r_0, mainly due to synaptic scaling before the training patterns are presented. The embedded patterns used for the subsequent learning are taken from the same processes for inhibitory and excitatory neurons, but are fixed and repeated in every stimulation. Figure 1C shows that the R-value, used as a measure of convergence, is approaching one. This shows that the neurons learn to fire only when presented with the embedded pattern and remain inactive otherwise. When the test window is extended by L =15 ms, the performance becomes perfect.

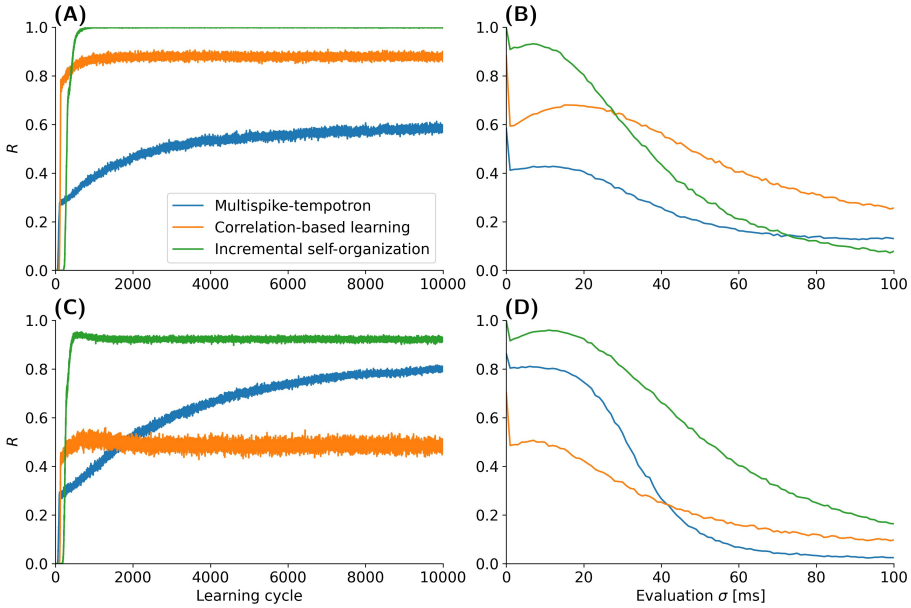

Figure 2. Comparison of performance. A: Performance during training for 10000 learning cycles with fixed (frozen) input spike patterns. B: Evaluation of the networks trained in (A) with modulated spike pattern of different σ used for the convolution. C: Performance during training for 10000 learning cycles with modulated spike patterns with $\sigma = 2\,\text{ms}$. D: Evaluation of the networks trained in (C) with modulated spike pattern of different σ used for the convolution. (Results based on an average of 500 simulations.)

This model for synaptic plasticity is particularly robust with respect to a range of perturbations of the patterns, including jitter of spike times, failures and additional stochastic spikes (not shown). Most noteworthy is its performance for temporally modulated Poisson spike rates since they appear to represent cortical codes. Here it shows superior performance when compared with previous approaches as e.g. the learning algorithm described in a 2016 study on spiking neurons[15] (Figure 2).

To generate a temporally modulated Poisson spike rate pattern from a basic frozen spike pattern, the given fixed time-coded spike input of each afferent was converted into rate-coded input by first convolving each spike with a Gaussian distribution characterized by a mean of zero and a standard deviation of $\sigma = 2\,\text{ms}$. The resulting function then served as a modulated firing rate for a Poisson point process. This procedure generates for each basic pattern an ensemble of spatio-temporal patterns with the statistics of a Poisson point process including failures and a Fano factor equal to 1. Patterns from this ensemble were then used for learning, plasticity and determination of the robustness of pattern detection.

In Figure 2 the left column shows the performance during the learning process and the right column presents an evaluation of the performance after learning for a range of jitter width of the Poisson modulation of the spike patterns (right column).

We compared our model to two algorithms from[15], pure correlation-based learning and multispike-tempotron learning. In contrast to our approach those models employ also a momentum heuristic and furthermore some attenuation[15] in addition to the learning rule. Even with these features these alternative models fail to reach the same performance level as the incremental self-organization for the case of frozen input spikes (i.e. no blurred Poisson inputs). When evaluating the network with the modulated input spike trains (Figure 2B), the performance drops visibly for $\sigma > 0$ for all models. The incremental self-organization, however, allows the network to perform consistently throughout the evaluated parameter range. In contrast, the biologically most realistic alternative correlation-based learning exhibits a significant decline in the performance and the multispike-tempotron learning rule does not reach a satisfactory performance in the first place. This demonstrates that our approach provides a superior noise robustness.

When exposed to modulated spike input trains, the neuron takes longer to learn the embedded pattern, and the performance becomes less distinct, while still remaining close to $R = 1$. Compared to the other models, however, learning is faster. When exposed to higher levels of temporal blur of the patterns than present during learning, our model allows for a much greater noise robustness of the neuron (see Figure 2B & D).

The reason for this stable performance lies in the balance between excitatory and inhibitory inputs which leads to more sparsely distributed weights (maintaining membrane potential fluctuations around the resting state, V_0, (Figure 1B)). This then enforces a stronger robustness against additional spikes fired or transmission failures which are mimicked by the missing of input spikes in the Poisson process.

To evaluate the robustness and stability of the memory trace over time we investigate the consequences of ongoing plasticity while the embedded pattern used for training is not presented to the neuron. Figure 3A shows that after 20,000 learning cycles with only random input input spikes and without the learned embedded pattern, still over 75% of the spikes would occur during the embedded pattern, exceeding the chance level of 0.025 by far. With realistic parameters the memory in this simulation would persists for over 10 hours without any decay or diffusion of the weight vector. This striking stability is explained by synaptic scaling that scales the excitatory synaptic efficacies up until a predetermined long-term firing rate is reached also for the random patterns (as shown by the red line in figure 3B). Weights are only scaled because the random inputs lack structure and the neurons therefore do not lose their selectivity for the learned pattern. That is, the weight vector mainly becomes larger and when the learned pattern is presented (only for testing) the neurons fire far more spikes, resulting in a much higher firing rate (as shown by the green line in figure 3B).

To further evaluate this result, we consider the cosine of the angle between the weight vector from the last learning cycle with patterns embedded in the input and each further learning cycle with only random spikes (i.e. without the learned pattern in the input, figure 3C). The results show that the cosine remains constantly high, implying that the weight distribution does not change significantly, which demonstrates the stability of the memory. The same procedure is used to calculate the norm of the weight vector, showing that the unstructured data leads to a gain in synaptic strength (Figure 3D). However, the memory is stored in the same synapses because when we subsequently use the embedded pattern again for relearning, the weights return to their fixed point (Figure 3E). This system, however, is not able to preserve the memory for the first pattern if another different embedded pattern is repeatedly presented (Figure 3F).

2 Neural Network

In principle, a single neuron is capable of learning more than one embedded pattern along the lines of the Tempotron[8]. However, a system with only one post-synaptic neuron cannot identify when and which pattern is present in the input. This would require an ensemble of output neurons which acquire different selectivities. It was already shown that for this purpose pre-synaptic hetero-synaptic plasticity is essential[16] since it induces a competition among out-going synapses. For investigation of the conditions under which plasticity and stability of memories meet, we consider a very simple network with several post-synaptic neurons and include also pre-synaptic hetero-synaptic plasticity.

We find that including this establishes approximately orthogonal weight vectors. As a result, when a new pattern is introduced, it adds an additional dimension to the system. As a result, the synaptic competition induced by pre-synaptic hetero-synaptic plasticity is sufficient to faithfully represent pattern identity and order in an ensemble of target neurons with shared input from a set of input neurons if, during learning, all patterns are presented in each epoch.

To quantify this, we use a matrix-based approach to evaluate the performance of the population of neurons. The matrix has rows representing neurons and columns representing patterns, with matrix elements set to 1 if a neuron is active for a particular pattern and 0 otherwise. The rank of this matrix reflects the number of linearly independent row vectors. We propose a performance measure, denoted Ω, which is the ratio of the matrix rank to the number of patterns When the rank of the matrix is equal to the number of patterns, it indicates that the population of neurons accurately represents the presence and order of the patterns in a stimulus. In this study, Ω is utilized as a criterion for evaluating learning, plasticity, and stability.

As a simple example, we trained a network with seven post-synaptic neurons using four embedded patterns in each training epoch. We found that when the pre-synaptic competition was present, all four patterns were selectively represented by the neurons, regardless of whether the patterns were shown in every learning cycle or only stochastically (Figure 4). In the absence of pre-synaptic competition, the selectivity for the patterns was not effectively distributed among the neurons[16]. Therefore, hetero-synaptic plasticity promotes the functional specialization of output neurons to different subsets of patterns, facilitating the self-organization of a precise representation of the "which" and "when" aspects of input patterns by the neuronal ensemble. The fact that specialization ($\Omega \to 1$) takes place also when each pattern is shown intermittently (i.e., not in each epoch, the green line in Figure 4) indicates the robustness of memories not only against the absence of patterns but also against the presence of different patterns. This raises the question under which conditions true incremental learning is possible without overwriting previous memories. We find that this is indeed the case, however only when the patterns are given as Poisson rate modulations (as explained above), as opposed to precise pattern repetitions.

In figure 5A, the initial learning phase involves $M = 7$ target neurons learning two embedded patterns. Subsequently, in figure 5B, learning continues with a third embedded pattern, without presenting the previously learned patterns again. The system's ability to recognize all patterns is tested every 50 cycles (with plasticity switched off) to assess memory retention. In figure 5C, we start with the initial conditions from the last learning cycle in figure 5A, but with only one of the two embedded pattern for ongoing learning. Every 50 learning cycles, the learning process is paused and the system's ability to recognize both initial patterns is tested.

To understand the superior stability of memories for patterns when learned from temporally modulated Poisson rate patterns we investigated the weight matrix $w = ab$ with its pre- and postsynaptic components a and b respectively (see Materials and

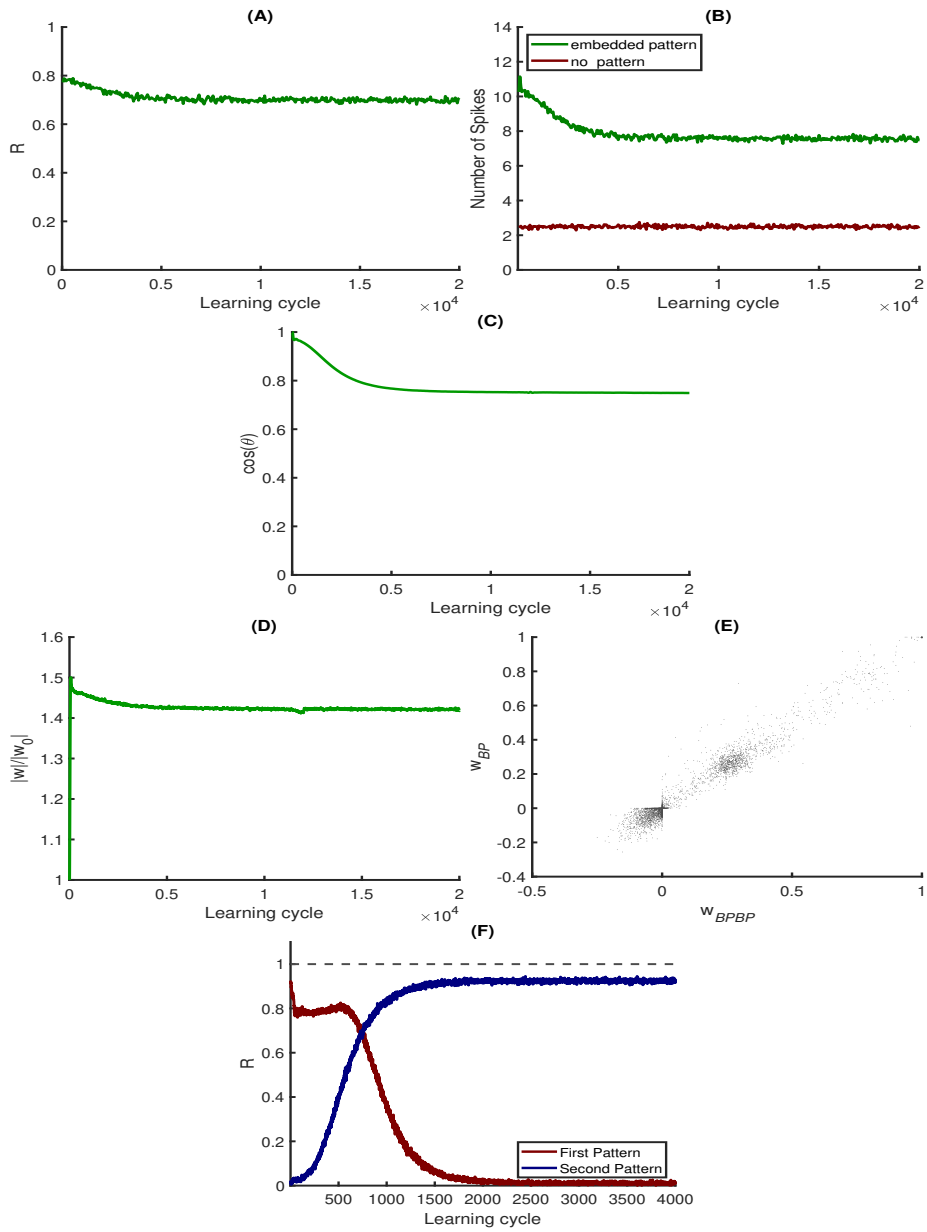

Figure 3. Memory of networks trained in with modulated spike pattern ($\sigma = 2$ ms). A: R as a memory criterion versus the learning cycle without learned patterns in afferents. Plasticity continues, and every 50 cycles, it is paused, and the learned pattern is placed in the background. Then R is computed for testing memory retention. B: Number of spikes versus learning cycle; the post-synaptic neuron elicits more spikes when the embedded pattern is present in the afferents (green) than in the absence of the pattern (red). (C and D) Without learned patterns in afferents, learning continues. C: Average of the cosine between the current weight vector and w_0. D: Average of the ratio of the norm of the current weight vector divided by the norm of the learned pattern weight vector (w_0). E: First, neurons learned to fire at 2 Hz, while there was no embedded pattern in afferents until learning cycle 5000 (W_B). Then, in each simulation, an embedded pattern is present in the afferents for 10,000 learning cycles yielding the weights W_{BP}. Next, there is again no embedded pattern in afferents for 20000 learning cycles for which we obtain W_{BPB}. Finally, there is again the original embedded pattern in the afferents for another 20000 learning cycles leading to W_{BPBP}. The scatter plot W_{BP} versus W_{BPBP} from all simulations demonstrates that most of the weights remain roughly the same, i.e., learning a pattern leads to a fixed point in weight space. F: The neurons first learn the first embedded pattern for the 10000 learning cycles, and then there is a different embedded pattern (the second pattern) in the afferents for another 10000 learning cycles. This figure shows the R-value for the second 10000 learning cycles. The results are based on an average of 500 independent simulations.

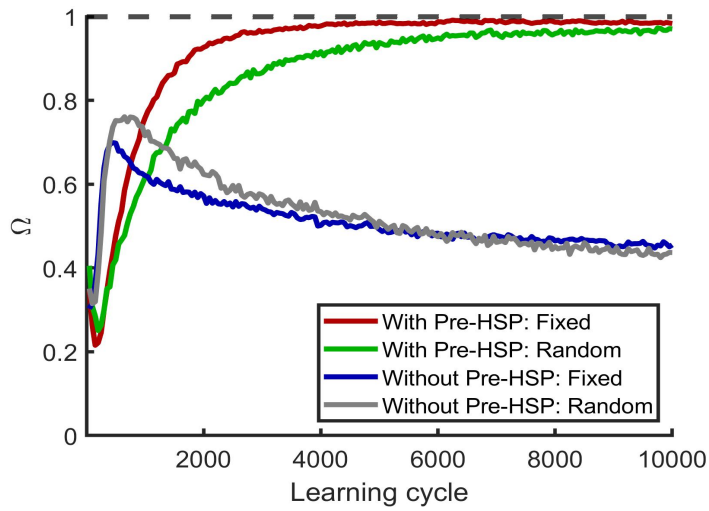

Figure 4. Representation of four embedded patterns by seven post-synaptic neurons. Ω versus learning cycle (There are 7 post-synaptic neurons). Red and Blue lines: Each pattern is shown in each learning cycle in a fixed position. Green and Gray lines: Each pattern is shown in the epoch with the probability of $1/4$ at a random position. Note that in this case some learning cycles have no embedded pattern in the epoch. Also, we randomly chose locations from $([85\ 85+L_{em}),[175\ 175+L_{em}),...[895\ 895+L_{em}]$ ms) to position the embedded patterns. (The data are from every 50'th learning cycles. This figure is based on the average of 500 simulations.)

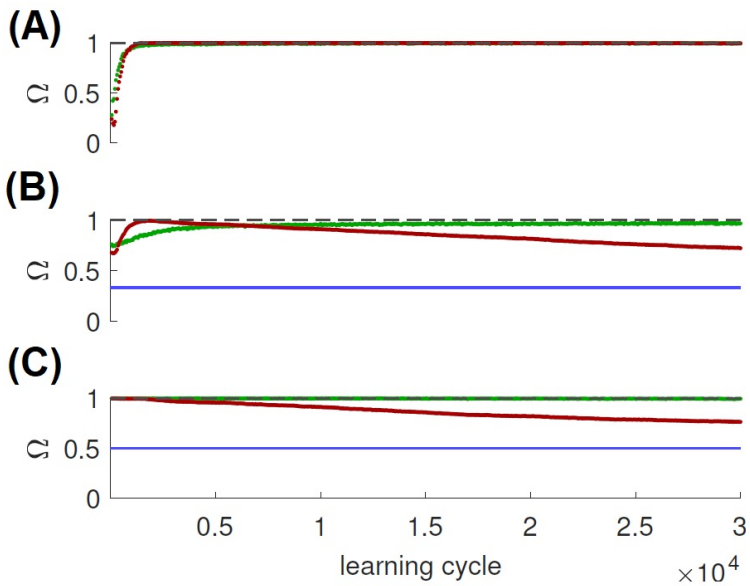

Figure 5. Learning, Plasticity and Stability. Green: Poisson spike rate. Red: Initial fixed pattern. Blue: Baseline. A, B, and C represent learning (from 0 to 2 patterns learned), plasticity (from 2 to 3 patterns learned), and stability (from 2 to 1 pattern shown), respectively. A: After learning with random input (noise), the network continues learning with two different embedded patterns for 30,000 learning cycles. B: The network continues learning with the third embedded pattern, while the first two are no longer shown. C: The network continues learning with only one of the embedded patterns, with initial conditions from the last learning cycles of A. Every 50 learning cycles, learning is stopped and Ω is computed for all two embedded patterns in A and C, and for all three embedded patterns in B. There are 7 post-synaptic neurons and the results are averages over 500 simulations.

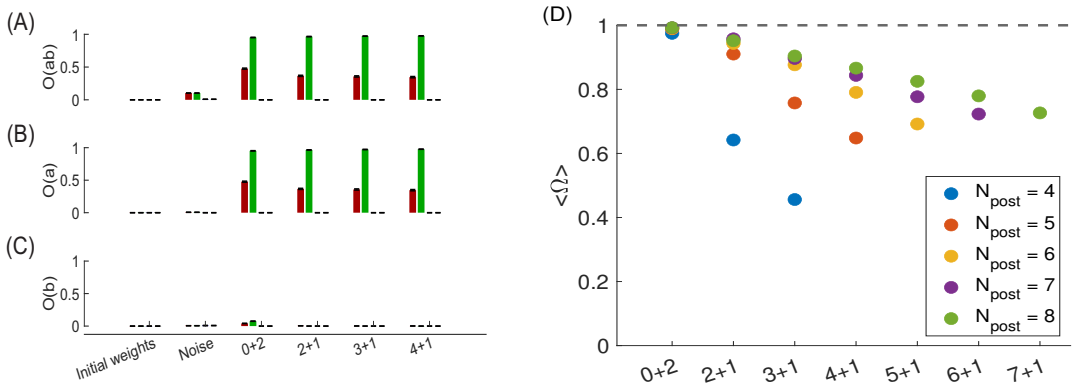

Figure 6. Orthogonality with and without pre-synaptic hetero-synaptic plasticity for original fixed patterns and Poisson spike rates. Red: with pre-synaptic hetero-synaptic plasticity and original patterns. Green: with pre-synaptic hetero-synaptic plasticity and Poisson spike rates. Blue: without pre-synaptic hetero-synaptic plasticity and original patterns. Gray: without pre-synaptic hetero-synaptic plasticity and Poisson spike rates. (A) The orthogonality measure is applied to the overall synaptic efficacy (i.e. $w = ab$), (B) to the pre-synaptic part (i.e. a), and (C) to the post-synaptic part (i.e. b). Initial weights: Random weight vectors in the first learning cycle. Noise: inputs are random spikes for 2000 learning cycles. 0+2: After learning noise, the network continues learning with two distinct embedded patterns for 30,000 learning cycles. 2+1: The network continues learning with the third embedded pattern, while the first two are no longer shown (30,000 learning cycles) and so on. D: Averaged Ω with pre-synaptic hetero-synaptic plasticity and Poisson spike rates with respect to the number of post-synaptic neurons. A, B. and C: There are 7 post-synaptic neurons. The results are averages over 100 independent simulations.

Methods). In particular, we tested the hypothesis that stability relates to the orthogonality of the weight vectors. As a graded measure O of 'orthogonality' we consider the volume spanned by the M weight vectors in the N^E-dimensional input space. We normalized the weight vectors leading to a matrix Y. The volume is then given by the product of the square roots of the M largest eigenvalues λ_i of $X = YY^T$:

$$O = \prod_{i=1}^{M} |\lambda_i(X)|^{\frac{1}{2}} \tag{2}$$

If the M normalized vectors are all mutually orthogonal this volume is $O = 1$, if one or more weight vectors are linearly dependent on the others this volume collapses to zero ($O = 0$). Figure 6 A, B and C show that without patterns in the input this measure is close to zero which is expected for random positive vectors. However, when then two rigid patterns are learned the weight vectors become more orthogonal while the orthogonality is close to maximum for jittered noisy (Poisson rate) patterns and increases even further when subsequently a third pattern is incrementally learned. Also, we see that without pre-synaptic hetero-synaptic plasticity the weights stay practically linearly dependent for both cases. Figure 6D shows that the averaged Ω at each incremental learning step increases with respect to the number of post-synaptic neurons.
To illustrate why the system maintains stability and plasticity simultaneously, we show the weight distribution of a target neuron at (0+2, 2+1, and 3+1 stages) using weight sorting when the system has learned two patterns (Figure 7: the red line is (0+2). The averaged sorted weight vectors, $w = ab$, for the (0+2) stage are represented by the red line. The green and blue lines, respectively, represent the averaged weight vector for the (2+1) and (3+1) stages (using the same sorting as the red line). As shown in figure 7, the decrease in the weight vector amount is much smaller in the learning with rate code than in the original patterns. However, some intermediate weight vectors are amplified in noise learning. This shows that pre-synaptic hetero-synaptic plasticity is essential but insufficient for memory retention. The advantage of introducing noise is that some important spikes in certain learning cycles are missed, resulting in an identical input-output correlation across all synapses on average. As a result, the effects of synaptic scaling overshadow hetero-synaptic plasticity. As a result, large weight vectors scale up more than when the network learns the original pattern; therefore, post-synaptic neurons search for the new embedded pattern within smaller and zero-weight vectors.

Neural Decoding of Spiked Audio

The aim of this section is to assess the algorithm's ability to identify embedded patterns within signals that replicate natural conditions. We used auditory data from the Audio Signal dataset, which consists of spoken digits from 0 to 9 in English,

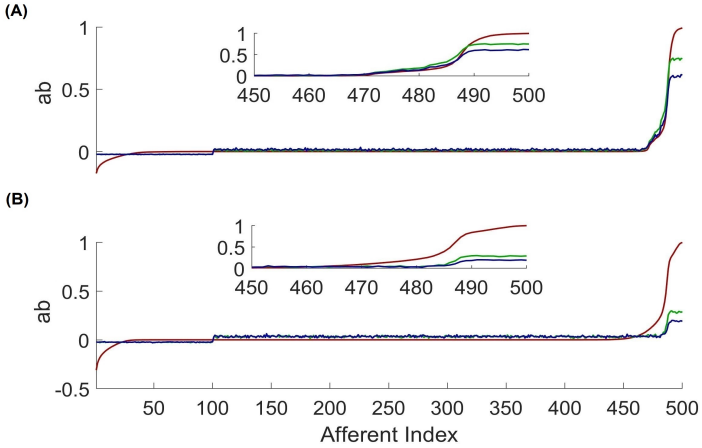

Figure 7. Weight vector distributions. Weight vector distributions for the Poisson rate (A) and for the original pattern (B). First, there is no embedded pattern in the afferents for 2000 learning cycles. Then, there are two embedded patterns in the afferent for 30,000 learning cycles (0+2). The red line shows the average sorted weights for this case ($w = ab$). Then, in each simulation, there is another distinct embedded pattern in the afferents for 30,000 learning cycles, and the first two learned patterns are not shown, and the average of the sorted weights is obtained in this learning cycle (2+1, the green line, and using the ranks from the red line). Finally, there is again another distinct embedded pattern in the afferents for 30,000 learning cycles, and the blue line shows the resulting weight vector at this learning cycle (using the ranks from the red line). There are seven post-synaptic neurons, and this figure is averaged over 500 simulations.

converted into spike patterns by the Lauscher artificial cochlea model[17]. This method allows us to analyze spikes across 700 input channels, taking advantage of the high-quality, aligned studio recordings in the Spoken HD (SHD) dataset. Moreover, this dataset includes recordings from 12 different speakers, providing a diverse basis for testing the model's pattern recognition capabilities under naturalistic conditions.[17]
Figure 8 shows the amplitude and spatiotemporal spike patterns of speaker number 11 when saying "*two*" (panels A and C) and "*three*" (panels B and D), as represented in the files *lang-english_speaker-11_trial-23_digit-2* and *lang-english_speaker-11_trial-23_digit-3*, respectively.[17]

These patterns, however, exhibit a high degree of synchronization, which poses the challenge of generating the same level of correlation between input and output across different afferents, potentially leading to identical synaptic efficacies. Therefore, the competition among excitatory neurons may need to be more sensitive to noise and small differences between afferents. To address this issue, we have improved the effectiveness of presynaptic heterosynaptic plasticity mechanisms (Materials and Methods). This adjustment enables a more subtle and effective modulation of synaptic strength, which is essential for discriminating between highly synchronized inputs.
Furthermore, it is challenging because a single word can contain multiple phonemes, increasing the possibility that the network interprets each segment as a separate pattern and does not learn another word at all. Consequently, the network may allocate all spikes to a single stimulus, thereby overshadowing the subtle differences between phonemes. To address this issue, we propose two solutions: increasing the number of postsynaptic neurons or using a larger value of the desired firing rate. These adjustments aim to improve the network's ability to distinguish between closely related phonemic patterns, resulting in a more precise and detailed representation of auditory inputs. With the exception of Speakers 5 and 6, training employs data from all other speakers. In contrast, the testing phase exclusively utilizes data from these two speakers. During the training phase, a randomly selected sequence of the paired indices ([0 1], [2 3], [4 5], [6 7], [8 9]) from speakers other than 5 and 6 is input into the neural network. Subsequently, the efficiency of the network, Ω, is assessed using data from speakers 5 and 6. During the training phase, the network initially processes pure noise input for 500 cycles to establish the target biological firing rate. Subsequently, two sequential numbers from the set ([0 1], [2 3], [4 5], [6 7], [8 9]) are randomly selected from one of the speakers. These two patterns are then employed as embedded patterns for the 10,000 learning cycles.
To assign the afferent type, we initially aggregated all spikes from each trial involving all speakers. A probability distribution was subsequently generated from the aggregated data, which was then used for randomly selecting inhibitory afferents. Specifically, channels exhibiting higher levels of activity were more likely to be classified as inhibitory afferents. Furthermore, this distribution was employed to randomly generate spikes (Figure 9 A). In each epoch, lasting 2.5 seconds, background noise was generated from the aforementioned distribution.

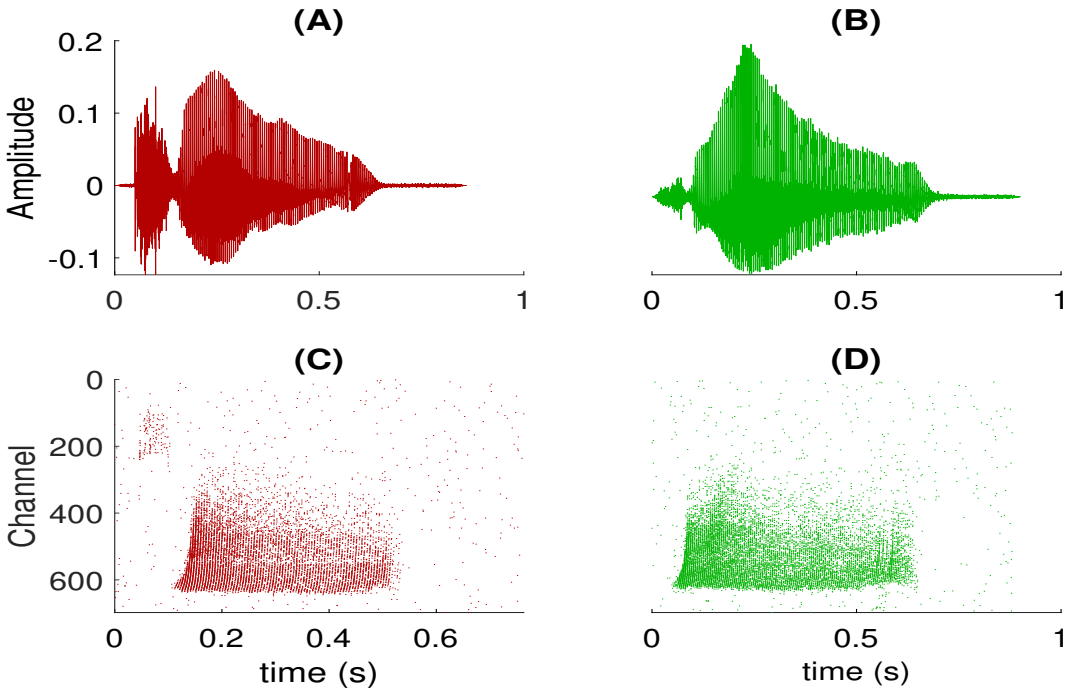

Figure 8. Spatio-temporal data in spike pattern of speech data In Panel A, Speaker 11 says"*two*" and "*three*" in Panel B. Panels C and D show the spatiotemporal spike patterns corresponding to A and B, respectively. lang-english_speaker-11_trial-23_digit-2 and lang-english_speaker-11_trial-23_digit-3.[17]

In a demonstration experiment, the speaker number 11 produced the numbers "*two*" and "*three*" 9 B. As it is shown in 9 C, the corresponding weight vectors demonstrate that distinct afferent inputs contribute to the generation of spikes.

Discussion

Selectivity for spatio-temporal patterns can self-organize in simple one-layer networks of integrate and fire neurons through a combination of Hebbian plasticity mechanisms, synaptic scaling for excitatory synapses and hetero-synaptic plasticity[16]. We found that the competition for weight increases of the pre-synaptic components of excitatory synapses which are controlled by the pre-synaptic neuron (termed 'pre-synaptic hetero-synaptic plasticity') have a tendency to orthogonalize the weight vectors of the target neurons in input space (figure 6, 0+2). When the patterns are supposed to be learned one after the other, our simulations demonstrate that the weights of the output neurons of a network stay fully orthogonal only if the patterns are stochastic variants of the original pattern. We showed this by generating the training patterns from a Poisson process with temporally modulated firing rates obtained from convolving the originally fixed pattern with a Gaussian function. In this case, orthogonality of the weights is preserved also when the patterns are learned incrementally. The orthogonality then supports minimal interference with previously learned weight vectors[18,19]. This novel mechanism realizes the balance of stability of memories and plasticity, in particular, it protects already present memories against deterioration during ongoing learning of further patterns (5).

Taken together, we proposed a biologically realistic solution to the stability-plasticity dilemma for networks that learn to detect spatio-temporal spike patterns without any supervision. This approach relies on the balance between excitatory and inhibitory inputs that naturally emerges during learning[16] which was suggested to be necessary for maintaining memory already before[20]. However, the dependence of the balance between stability of memories and plasticity on the balance between excitation and inhibition remains to be investigated.

To our knowledge this is the first approach towards a biologically realistic model for incremental self-supervised learning of spatio-temporal spike patterns. Previous approaches were proposed for recurrent networks[21,22] and rate-based models[18] which relied also on systematic orthogonalization of weight vectors, however by different mechanisms. Furthermore, our approach relies heavily on mechanisms for synaptic scaling which has also been proposed before as a mechanisms for stabilizing long term memories, however in rate based models and with a w^2 dependence of scaling[22] which in our case was found to be detrimental for performance of learning (not shown).

Last but not least, our results align with previous research[23] emphasizing the critical role of noise in the input for incremental

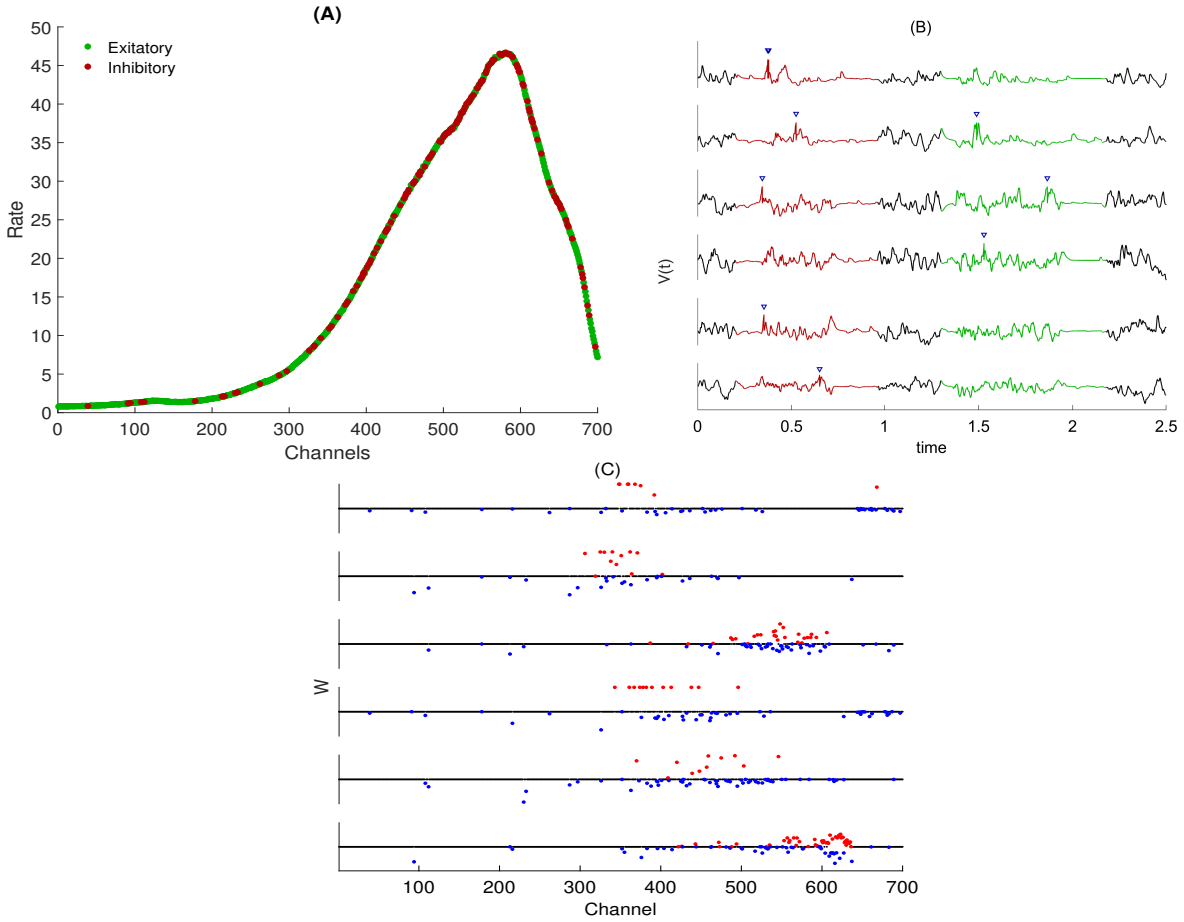

Figure 9. Application to speech data. Each embedded pattern was placed between background noise generated from the distribution in panel A. The network was trained to recognize only the spike patterns for "*two*" and "*three*" from speaker 11, trial number 23. B shows the membrane potential of each neuron after 10000 learning cycles and the corresponding weight vector in panel C. The same training process was repeated for the other speakers and the other embedded patterns. During the evaluation, the network demonstrated an ability to remember and discriminate these patterns with input from the other speakers (Speakers 4 and 5), achieving an accuracy rate of 94% ($\Omega = 0.94 \pm 0.02$). The mean network performance across all speakers, after 10,000 learning cycles, is 98% ($\Omega = 0.98 \pm 0.002$).

learning by driving weights to become orthogonal and serving to minimize their overlap for different output neurons.

Materials and Methods

Signal transmission is more effective when pre-synaptic neurons release more neurotransmitters and fire more frequently, and also when post-synaptic neurons have more receptors and sensitivity. Therefore, the efficacy of a synapse,w_{ji} (pre-synaptic neuron i targeting a post-synaptic neuron j), can be expressed in terms of the combination of pre-synaptic and post-synaptic components.

$$w_{ji} = a_{ij} b_{ji}. \tag{3}$$

where the resources for a_{ij} are provided by the pre-synaptic neuron i and b_{ji} is supported by the post-synaptic neuron j. In the following we assume N pre-synaptic neurons ($i = 1, ..., N$) and M post-synaptic neurons ($j = 1, ..., M$).
The efficacy of a synapse can be modified separately in the pre- and the post-synaptic components through synaptic plasticity, resulting in changes in the total efficacy of the synapse in transmitting signals between neurons.

$$\frac{dw_{ji}}{dt} = a_{ij}\frac{db_{ji}}{dt} + \frac{da_{ij}}{dt} b_{ji}. \tag{4}$$

Here, we first consider the homeostatic plasticity of synaptic strength to maintain the overall stability of neuronal activity, referred to as synaptic scaling, in order to modify the synapse. It has been shown that neurons regulate their excitatory synaptic weights to achieve the biologically desired firing rate, r_0[3,24]. It allows the neuron to decrease or increase synaptic effectiveness when the long-term firing rate, r_j, is greater or less than r_0.

$$\tau_E \frac{da^E_{ij}}{dt} = -a^E_{ij} + \alpha a^E_{ij} tanh(r_0 - r_j). \tag{5}$$

$$\tau_E \frac{db^E_{ji}}{dt} = -b^E_{ji} + \alpha b^E_{ji} tanh(r_0 - r_j). \tag{6}$$

where α and τ_E are the scaling factor and time constant, respectively. Note that the sub- and superscript E indicate excitatory synapses. The long-term firing rate r_j is determined by

$$\tau_R \frac{dr_j}{dt} = -r_j + \sum_{t^l_i} \delta(t - t^l_i). \tag{7}$$

where τ_R is the time constant for the long-term firing rate. At the time t^l_i there is a spike in the afferent i, and at that time the afferent i begins to send input to the post-synaptic neuron j equal to $w_{ji}K$.
The kernel shape, K, is an alpha function with the equation:

$$K(t - t^l_i) = I_{norm}\left[\exp\left(-\frac{t - t^l_i}{\tau_d}\right) - \exp\left(-\frac{t - t^l_i}{\tau_r}\right)\right]\theta(t - t^l_i). \tag{8}$$

In this equation, θ represents the Heaviside step function. τ_r and τ_d denote the time constants for rise and decay respectively associated with the synaptic current, with different values for excitatory (τ^E_d, τ^E_r) and inhibitory synapses (τ^I_d, τ^I_r). The term $I_{norm} = \frac{\eta^{\eta/\eta-1}}{\eta-1}$ normalizes the amplitude of K to unity, where $\eta = \frac{\tau_d}{\tau_r}$. Note that the parameters are chosen to be as close as possible to the biologically feasible range.
While synaptic scaling is effective in stabilizing the firing activity of a network, it is insufficient for learning the information encoded in the input. Therefore, a correlation-based learning rule inspired by the N-Methyl-D-Aspartate (NMDA) receptor[12–14] has been proposed in previous works[15,16]. In the weakly supervised learning rule[15], neurons are able to learn embedded patterns, as well as in[16], a version following Dale's rules, which operates in an unsupervised manner.
In more detail, the correlation between the input from pre-synaptic neuron i and output membrane potential of post-synaptic neuron j, denoted as $V_j(t)$, provides the signal for modifying the synaptic strength, denoted as $\varepsilon_{ij}(t)$ which is referred to as eligibility[15]:

$$\tau_\varepsilon \frac{d\varepsilon_{ij}(t)}{dt} = -\varepsilon_{ij}(t) + \vartheta_{ij}. \tag{9}$$

In this context, $\vartheta_{ij}(t)$ represents the correlation between the input from the afferent i and the output neuron j, reflected by the membrane potential V_j.

$$\vartheta_{ij}(t) = \sum_{l=1}^{n_i} K(t - t_i^l)\{[V_j(t) - V_0]_+ + q[V_j(t) - V_0]_-\}. \tag{10}$$

In the simulations discussed in this paper, V_0 represents the modification threshold[25] and is consistently set to zero. The parameter q takes a value of zero for excitatory afferents and a value of 1 for inhibitory afferents[16]. During each time step, all spikes n_i in the i-th afferent (elicited before time t) are employed to update the synapses. Note, $V_j(t)$ is described by the simple leaky integrate and fire neurons' differential equation[26,27]:

$$\tau_m \frac{d}{dt} V_j(t) = -V_j(t) + R_m I_j^{ext}(t), \tag{11}$$

where τ_m and R_m are the membrane time constant and resistance, respectively. The external current, I_j^{ext}, is determined by the synaptic strength of the excitatory and inhibitory afferents, denoted by w_{ji}^E and w_{ji}^I, respectively:

$$I_j^{ext}(t) = \sum_{i=1}^{N^E} w_{ji}^E \sum_{t_i^l < t} K^E(t - t_i^l) - \sum_{i=1}^{N^I} w_{ji}^I \sum_{t_i^l < t} K^I(t - t_i^l). \tag{12}$$

where N^E and N^I represent the number of excitatory and inhibitory afferents, respectively, and $N = N^E + N^I$. The external current is determined by the convolutions of the input spike trains with their respective synaptic kernels (Eq. 8), $K^E(t)$ and $K^I(t)$.

While the learning signal in excitatory synapses is positive, hetero-synaptic plasticity is used to provide negative changes, inducing selectivity and specificity. Here, we consider that the resources required to enhance the pre-and post-synaptic components a_{ij} and b_{ji} of the weights $w_{ji} = a_{ij} b_{ji}$ are inherently limited. As a result, we propose that these resources are distributed competitively, reflecting the limited availability and competitive nature of their allocation. The heterosynaptic plasticity of both pre- and post-synaptic regions is a function of the induced eligibility signal. The post-synaptic heterosynaptic plasticity, is implemented by subtracting the mean of the eligibilities. It affects changes in the post-synaptic components b_{ji}:

$$\tilde{\varepsilon}_{ij} = \varepsilon_{ij} - \frac{1}{N^E} \sum_{i=1}^{N^E} \varepsilon_{ij}. \tag{13}$$

Nevertheless, it is hypothesized that the signals governing the changes in the pre-synaptic components a_{ij} depend on the magnitude of the potentiation observed on the post-synaptic side[28]. The signal for pre-synaptic competition is given by the signal for postsynaptic potentiation $[\tilde{\varepsilon}_{ij}]_+$, diminished by the mean of the potentiation signals of all other postsynaptic synapses, however, only in case more than one postsynaptic component receives a signal for strengthening. This is implemented by the following:

$$B_{ji} = [\tilde{\varepsilon}_{ij}]_+ - \frac{\theta\left(\sum_{l \neq j} \theta([\tilde{\varepsilon}_{ij}]_+) - \zeta\right)}{\sum_l \theta([\tilde{\varepsilon}_{il}]_+) + \zeta} \sum_{l=1}^{M} [\tilde{\varepsilon}_{il}]_+. \tag{14}$$

where θ is the Heaviside function, ζ is a small positive number ($\zeta << 1$), and M is the number of post-synaptic neurons. When applying our model to highly synchronized data, such as speech recordings, we attempted to increase synaptic competition so that a reduced number of synapses undergo significant growth. This was done by making the following modification:

$$B_{ji} = [\tilde{\varepsilon}_{ij}]_+ - \left(\frac{\theta\left(\sum_{l \neq j} \theta([\tilde{\varepsilon}_{il}]_+) - \zeta\right)}{\sum_l \theta([\tilde{\varepsilon}_{il}]_+) + \zeta} \sum_{l=1}^{M} [\tilde{\varepsilon}_{il}]_+^{\psi} \right)^{\frac{1}{\psi}}. \tag{15}$$

where ψ is the competition enhancement parameter. Increasing the parameter ψ beyond 1 enforces competition between synapses. In other words, a ψ value greater than one results in a greater degree of synaptic competition than when ψ is equal to one (Jensen Inequality). With this we model the dynamics of the synaptic components a and b:

$$\tau_E \frac{da^E_{ij}}{dt} = -a^E_{ij} + \alpha a^E_{ij} \tanh(r_0 - r_j) + c^E B_{ji}. \tag{16}$$

$$\tau_E \frac{db^E_{ji}}{dt} = -b^E_{ji} + \alpha b^E_{ji} \tanh(r_0 - r_j) + c^E \tilde{\varepsilon}_{ij}. \tag{17}$$

where τ_E is the time constant for changes of excitatory synapses and c^E the learning rate. For inhibitory neurons we use:

$$\tau_I \frac{da^I_{ij}}{dt} = -a^I_{ij} + c^I \varepsilon_{ij} \tag{18}$$

$$\tau_I \frac{db^I_{ji}}{dt} = -b^I_{ji} + c^I \varepsilon_{ij}, \tag{19}$$

where τ_I is the time constant for inhibitory synapses and c^I is the learning rate. Note that the sub- and superscript I indicate inhibitory synapses. We consider c^E to be smaller than c^I; therefore, inhibitory neurons can adapt more quickly to maintain balance at both global and detailed levels, and excitation cannot conflict with the relatively slow limiting mechanism of synaptic scaling[16].

Here, the phenomenon of simultaneous Hebbian plasticity of inhibition is essential for maintaining a global balance within neural systems. The neuron is placed in a fluctuating regime by this balance, in which the excitatory and inhibitory weights exhibit significant strength, resulting in significant membrane potential fluctuations. Once this balance is achieved, further weight adjustments driven by the stochastic background restrict the weights to a fixed point balance via a random walk process. Furthermore, the integration of a frequent pattern induces an additional drift that systematically shifts the weights away from the fixed point until the target number of spikes is induced by the pattern alone, independent of the response of the background activity [16,29].

Simulation Details

In all simulations, 80% of the afferents are excitatory and 20% inhibitory (400 excitatory and 100 inhibitory neurons). Afferent and embedded pattern spikes are generated randomly with Poisson point processes in which the excitatory rate (r_E) is 5 Hz, and the inhibitory rate (r_I) is 20 Hz. In each learning cycle, the learning epoch length is 2000 ms when using a single output neuron. When investigating the network with multiple postsynaptic neurons, an epoch length of 1000 ms was used in order to reduce the computational effort. To optimize the simulations, changes over epochs e with a duration of T ms are considered instead of implementing weight changes at each integration step. Consequently, equation (10) is replaced by

$$\vartheta^l_{ij} = \int_0^T dt K(t - t^l_i)[V_j(t) - V_0]_+ + q[V j(t) - V_0]_-. \tag{20}$$

Differential equations of this form:

$$\tau \frac{dy}{dt} = -y + x \tag{21}$$

are replaced by moving averages:

$$y(e+1) = \gamma y(e) + (1-\gamma)\hat{x}(e), \tag{22}$$

where $\hat{x}(e)$ is the sum of contributions of x in each epoch e and $\gamma \simeq 1 - T/\tau$.

Synaptic scaling relies on firing rates averaged over long periods of time. In online learning, a low-pass filter is used to ensure that the rate estimate $r_j(c+1)$ of the post-synaptic neuron j for learning in epoch $e+1$ is based on the running average of the actual rates $\hat{r}$ from previous epochs:

$$r_j(e+1) = \gamma^* r_j(e) + (1-\gamma^*)\hat{r}_j(e). \tag{23}$$

The parameter γ^* is set to 0.9 for all figures, which corresponds to a time constant of 10s. Although smaller values do not change the asymptotic results, this value allows for faster convergence.

Plasticity is not immediate; it depends on the accumulation of signals for weight changes over time, called 'eligibility'[15]. A low-pass filter is used to implement this:

$$\varepsilon_{ij}(e+1) = \gamma\varepsilon_{ij}(e) + (1-\gamma)\hat{\vartheta}_{ij}(e). \tag{24}$$

The initial conditions are $r_j(0) = 0$ and $\varepsilon_{ij}(0) = 0$. The parameter γ is 0.99 for all figures, which corresponds to a time constant of about 100s. This parameter limits the learnability of infrequently presented patterns[16].
Due to structural limitations and other factors, synaptic strength cannot increase indefinitely. As a result, we restrict weight changes to prevent excitatory synapses from exceeding the upper limit of one.
Dale's rule states that excitatory and inhibitory synapses cannot convert into each other. Therefore, we consider synaptic weights to be zero if weight changes during learning would change their type. This ensures that subtracting the mean in Eq (13) and Eq (14) does not affect the type of a synapse. We use the Euler method with a time step of Δt for the numerical integration of Eq (12), and the associated parameters are available in table 1.
Note, for the network model, the initial synaptic efficacies a_{ij} and b_{ji} are drawn from a Gaussian distribution with a mean of 0.1 and 10^{-2} standard deviations, while negative values are set to zero. In the single post-synaptic neuron model, they are drawn from a Gaussian distribution with a mean of 10^{-2} and standard deviations of 10^{-3}, with negative values set to zero. In cases where there is only a single post-synaptic neuron, all values of a_{ij} are initially set to 1. Since there is no pre-synaptic hetero-synaptic plasticity, these a_{ij} values will remain at their initial value of 1, regardless of the number of learning cycles.
In the epoch-based approach, we use the following equations, which can be linearized to obtain the continuous version:

$$\delta w_{ij} = a_{ij}\delta b_{ji} + \delta a_{ij} b_{ji} + \delta a_{ij}\delta b_{ji}. \tag{25}$$

Inhibitory synapses

$$\delta a^I_{ij} = c^I \varepsilon_{ij}. \tag{26}$$
$$\delta b^I_{ji} = c^I \varepsilon_{ij}. \tag{27}$$

Excitatory synapses

$$\delta a^E_{ij} = -\chi a^E_{ij} + \alpha a^E_{ij}\tanh(r_0 - r_j) + c^E B_{ji}. \tag{28}$$
$$\delta b^E_{ji} = -\chi b^E_{ji} + \alpha b^E_{ji}\tanh(r_0 - r_j) + c^E \tilde{\varepsilon}_{ij}. \tag{29}$$

The additional factor χ in front of the decay term is given by the excitatory time constant τ_E. The other parameters α and c^E have been updated by that same factor as well.

Noise robustness

Generation of noise

Noise was introduced in the form of Poisson spike rates. These were generated by computing the convolution of the initial (frozen) spike train with a Gaussian, where the standard deviation σ of the Gaussian represents the strength of the noise. This convolution is then used to generate the new spike train by drawing values from a uniform distribution and generating a spike wherever these values are less than the corresponding value from the convolution. This process is visualised in Figure 10.

Comparison to other models

When comparing our model to the approach of the multispike-tempotron and correlation-based learning the parameters listed in Tab. 2 were used. In addition to the learning rules of the respective algorithms, there was also a momentum heuristic and an additional attenuation implemented as in[15].

Code Availability

The source code supporting the findings of this study is available in the GitHub repository. The repository can be accessed at the following URL:
https://github.com/MohammadDehghaniH/Incremental_Self-Organization_of_Spatio-Temporal_Spike_Pattern_Detection

Table 1. List of parameters.

Symbol	Description	Value
τ_m	Membrane time constant	15 ms
R_m	Membrane resistance	1
τ_r^E	Rise time of excitatory currents	0.5 ms
τ_r^I	Rise time of inhibitory currents	1 ms
τ_d^E	Decay time of excitatory currents	3 ms
τ_d^I	Decay time of inhibitory currents	5 ms
r^E	Rate of excitatory neurons	5 Hz
r^I	Rate of inhibitory neurons	20 Hz
N	Number of pre-synaptic neurons	500
N^E	Number of pre-synaptic excitatory neurons	400
N^I	Number of pre-synaptic inhibitory neurons	100
c^I	Inhibitory learning rate	10^{-3}
c^E	Excitatory learning rate	0.9×10^{-3}
Δt	Time step	0.1 ms
α	Scaling factor	0.01
α	Scaling factor (real data)	0.04
T	Epoch duration (single neuron)	2000 ms
T	Epoch duration (Neural Network)	1000 ms
T	Epoch duration (real data)	2500 ms
χ	Decay term	10^{-4}
ψ	competition enhancement parameter	12

Table 2. Parameters used for MST & CBL.

Symbol	Description	Value
τ_m	Integration time constant	20 ms
τ_s	Decay time constant	5 ms
r	Neuron's firing rate	5 Hz
N	Number of pre-synaptic neurons	500
λ	Learning rate	10^{-5}
μ	Momentum heuristic	0.99
α	Additional attenuation	0.995
Δt	Time step	0.1 ms
T	Epoch duration (single neuron)	2000 ms

Contributions

Conceptualization: M.D-H; K.P.
Formal Analysis: M.D-H, K.P, L.M (specifically for Figure 2).
Funding Acquisition: K.P.
Investigation: M.D-H, L.M (specifically for Figure 2).
Methodology: M.D-H, K.P.
Project Administration: K.P.
Software: M.D-H, L.M (specifically for Figure 2).
Supervision: K.P.
Validation: M.D-H, K.P.
Visualization: M.D-H, K.P, L.M (specifically for Figure 2).
Writing – Original Draft: M.D-H, K.P, L.M.
Writing – Review & Editing: M.D-H, K.P, L.M.

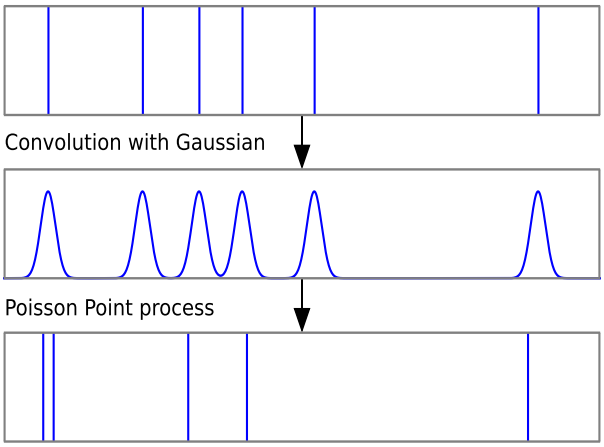

Figure 10. Schematic structure of generation of the Poisson spike rates. Top: Initial spike train. Center: Convolution with the Gaussian. Bottom: New spike train drawn from the distribution in the center graph.

Acknowledgements

We would like to thank David Rotermund for his fruitful discussions and for his assistance in providing the spiked audio data. We acknowledge support by the following grants: DFG - grant (PA 569/5-1 to KP)
During writing we took advantage from special tools including Bing Chat: https://www.bing.com, DeepL: https://www.deepl.com and ChatGPT.

References

1. Hebb, D. O. *The organization of behavior: A neuropsychological theory* (Psychology Press, 2005).
2. Turrigiano, G. G. The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell* **135**, 422–435 (2008).
3. Van Rossum, M. C., Bi, G. Q. & Turrigiano, G. G. Stable hebbian learning from spike timing-dependent plasticity. *J. neuroscience* **20**, 8812–8821 (2000).
4. Abraham, W. C. & Robins, A. Memory retention–the synaptic stability versus plasticity dilemma. *Trends neurosciences* **28**, 73–78 (2005).
5. French, R. M. Catastrophic forgetting in connectionist networks. *Trends cognitive sciences* **3**, 128–135 (1999).
6. Kirkpatrick, J. *et al.* Overcoming catastrophic forgetting in neural networks. *Proc. national academy sciences* **114**, 3521–3526 (2017).
7. Parisi, G. I., Kemker, R., Part, J. L., Kanan, C. & Wermter, S. Continual lifelong learning with neural networks: A review. *Neural Networks* **113**, 54–71 (2019).
8. Gütig, R. & Sompolinsky, H. The tempotron: a neuron that learns spike timing–based decisions. *Nat. neuroscience* **9**, 420–428 (2006).
9. Florian, R. V. The chronotron: A neuron that learns to fire temporally precise spike patterns. *PNAS* (2012).
10. Memmesheimer, R.-M., Rubin, R., Ölveczky, B. P. & Sompolinsky, H. Learning precisely timed spikes. *Neuron* **82**, 925–938 (2014).
11. Albers, C., Westkott, M. & Pawelzik, K. Learning of precise spike times with homeostatic membrane potential dependent synaptic plasticity. *PloS one* **11**, e0148948 (2016).
12. Artola, A., Bröcher, S. & Singer, W. Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature* **347**, 69–72 (1990).
13. Cummings, J. A., Mulkey, R. M., Nicoll, R. A. & Malenka, R. C. Ca2+ signaling requirements for long-term depression in the hippocampus. *Neuron* **16**, 825–833 (1996).
14. Malenka, R. C. & Nicoll, R. A. Long-term potentiation–a decade of progress? *Science* **285**, 1870–1874 (1999).
15. Gütig, R. Spiking neurons can discover predictive features by aggregate-label learning. *Science* **351**, aab4113 (2016).

16. Dehghani-Habibabadi, M. & Pawelzik, K. Synaptic self-organization of spatio-temporal pattern selectivity. *PLOS Comput. Biol.* **19**, e1010876 (2023).

17. Cramer, B., Stradmann, Y., Schemmel, J. & Zenke, F. The heidelberg spiking data sets for the systematic evaluation of spiking neural networks. *IEEE Transactions on Neural Networks Learn. Syst.* **33**, 2744–2757 (2020).

18. Leibold, C. Neural kernels for recursive support vector regression as a model for episodic memory. *Biol. Cybern.* **116**, 377–386 (2022).

19. Chenani, A. *et al.* Hippocampal ca1 replay becomes less prominent but more rigid without inputs from medial entorhinal cortex. *Nat. communications* **10**, 1341 (2019).

20. Koolschijn, R. S. *et al.* Memory recall involves a transient break in excitatory-inhibitory balance. *Elife* **10**, e70071 (2021).

21. White, O. L., Lee, D. D. & Sompolinsky, H. Short-term memory in orthogonal neural networks. *Phys. review letters* **92**, 148102 (2004).

22. Tetzlaff, C., Kolodziejski, C., Timme, M., Tsodyks, M. & Wörgötter, F. Synaptic scaling enables dynamically distinct short-and long-term memory formation. *PLoS Comput. Biol.* **9**, e1003307 (2013).

23. Schug, S., Benzing, F. & Steger, A. Presynaptic stochasticity improves energy efficiency and helps alleviate the stability-plasticity dilemma. *Elife* **10**, e69884 (2021).

24. Bi, G.-q. & Poo, M.-m. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. neuroscience* **18**, 10464–10472 (1998).

25. Sjöström, P. J., Turrigiano, G. G. & Nelson, S. B. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* **32**, 1149–1164 (2001).

26. Burkitt, A. N. A review of the integrate-and-fire neuron model: I. homogeneous synaptic input. *Biol. cybernetics* **95**, 1–19 (2006).

27. Burkitt, A. N. A review of the integrate-and-fire neuron model: Ii. inhomogeneous synaptic input and network properties. *Biol. cybernetics* **95**, 97–112 (2006).

28. Rasmussen, C. E. & Willshaw, D. J. Presynaptic and postsynaptic competition in models for the development of neuromuscular connections. *Biol. cybernetics* **68**, 409–419 (1993).

29. Yger, P., Stimberg, M. & Brette, R. Fast learning with weak synaptic plasticity. *J. Neurosci.* **35**, 13351–13362 (2015).