

A triplet spike-timing–dependent plasticity model generalizes the Bienenstock–Cooper–Munro rule to higher-order spatiotemporal correlations

Julijana Gjorgjieva^{a,1,2}, Claudia Clopath^b, Juliette Audet^c, and Jean-Pascal Pfister^{d,e}

^aDepartment of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge CB3 0WA, United Kingdom; ^bLaboratory of Neurophysics and Physiology, Université Paris Descartes, 75270 Paris, France; ^cLaboratory of Computational Neuroscience, Ecole Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland; ^dComputational Neuroscience Laboratory, Department of Physiology, University of Bern, CH-3012 Bern, Switzerland; and ^eDepartment of Engineering, University of Cambridge, Cambridge CB2 1PZ, United Kingdom

Edited* by Leon N. Cooper, Brown University, Providence, RI, and approved September 30, 2011 (received for review April 14, 2011)

Synaptic strength depresses for low and potentiates for high activation of the postsynaptic neuron. This feature is a key property of the Bienenstock–Cooper–Munro (BCM) synaptic learning rule, which has been shown to maximize the selectivity of the postsynaptic neuron, and thereby offers a possible explanation for experience-dependent cortical plasticity such as orientation selectivity. However, the BCM framework is rate-based and a significant amount of recent work has shown that synaptic plasticity also depends on the precise timing of presynaptic and postsynaptic spikes. Here we consider a triplet model of spike-timing–dependent plasticity (STDP) that depends on the interactions of three precisely timed spikes. Triplet STDP has been shown to describe plasticity experiments that the classical STDP rule, based on pairs of spikes, has failed to capture. In the case of rate-based patterns, we show a tight correspondence between the triplet STDP rule and the BCM rule. We analytically demonstrate the selectivity property of the triplet STDP rule for orthogonal inputs and perform numerical simulations for nonorthogonal inputs. Moreover, in contrast to BCM, we show that triplet STDP can also induce selectivity for input patterns consisting of higher-order spatiotemporal correlations, which exist in natural stimuli and have been measured in the brain. We show that this sensitivity to higher-order correlations can be used to develop direction and speed selectivity.

Synaptic plasticity depends on the activity of presynaptic and postsynaptic neurons and is believed to provide the basis for learning and memory (1, 2). It has been shown that low-frequency stimulation (1–3 Hz) (3) or stimulation paired with low postsynaptic depolarization (4) induces synaptic long-term depression (LTD), whereas synapses undergo long-term potentiation (LTP) after high-frequency stimulation (100 Hz) (5). Such findings are consistent with the well-known Bienenstock–Cooper–Munro (BCM) learning rule (6). This BCM model has been shown to elicit orientation selectivity and other aspects of experience-dependent cortical plasticity (6, 7). Furthermore, in this model the modification threshold between LTP and LTD varies as a function of the history of postsynaptic activity, a prediction that has been confirmed experimentally (8).

Despite its consistency with experimental data and its functional relevance, the BCM framework is still limited experimentally and functionally. Experimentally, because the learning rule is expressed in terms of firing rates, it cannot predict synaptic modification on the basis of the timing of pre- and postsynaptic spikes (9, 10). This form of plasticity, called spike-timing–dependent plasticity (STDP), uses the timing of spike pairs to induce synaptic modification (11, 12). The presynaptic spike is required to shortly precede the postsynaptic spike to elicit LTP, whereas the reverse timing of pre- and postsynaptic spikes leads to LTD (9, 10). Functionally, the BCM model cannot segregate input patterns that are characterized by their temporal spiking structure. STDP provides a possible solution, but how STDP relates to BCM remains debated (13–15).

Here, we consider a spike-based learning rule, “the triplet STDP model” (15, 16), and show that it overcomes those two important limitations of the BCM rule and thus generalizes the BCM framework. This triplet model uses sets of three spikes (triplets)—instead of pairs of spikes as in the case of classical STDP—to

induce potentiation. More precisely, LTP depends on the interval between the pre- and postsynaptic spikes and on the timing of the previous postsynaptic spike (Fig. 1A). Furthermore, this triplet learning rule has been shown to explain a variety of synaptic plasticity data (17, 18) significantly better than pair-based STDP (15) (Fig. 1B). Plasticity induced by multiples of spikes has also been the focus of other studies (19, 20); despite using the same spike combinations some differences have been observed, most likely due to the different (extracellular or intracellular) stimulation protocols used in these studies (21).

Computationally, it has been shown that under some rather crude assumptions—when the input and output neurons have independent Poisson statistics—the triplet STDP model can be mapped to the BCM learning rule (16). In this paper, we take a more biologically plausible approach by incorporating contributions from input–output spiking correlations in inducing synaptic plasticity. Consistent with results from the BCM theory, we demonstrate that in the presence of orthogonal rate-based patterns, the maximally selective fixed points of the weight dynamics induced by the triplet rule are stable. Furthermore, we show that the triplet rule acts as a generalized BCM rule in the sense that postsynaptic neurons become selective not only to rate-based patterns of the inputs, but also to patterns differentiated only by their spiking correlation structure. The mathematical simplicity of the triplet model allowed us to characterize the explicit dependence of the weight dynamics on higher-order input correlations. We believe this study is of great relevance given the ubiquity of higher-order correlations in the brain (22, 23) and their relevance for neural coding (24).

Model and Methods

Neuronal Dynamics. We considered a feedforward network with N input neurons $x_i(t)$ as Dirac delta spike trains connected to a single output neuron through the weights $w_j(t)$ and giving rise to the postsynaptic spike train $y(t)$ (SI Text). The input spike trains had average firing rates $\rho_i(t)$.

We assumed that the membrane potential of the postsynaptic neuron $u(t)$ increased with the spike times of each input by the excitatory postsynaptic potential (EPSP) scaled by the corresponding weight

$$u(t) = \sum_{k=1}^N w_k(t) \int_0^{\infty} \varepsilon(r) x_k(t-r) dr. \quad [1]$$

The function $\varepsilon(r)$ denoted the EPSP kernel, taken to be a decaying exponential with a membrane time constant of 11 ms. For spatio-temporal

Author contributions: J.G., C.C., and J.-P.P. designed research; J.G. performed research; J.G., J.A., and J.-P.P. contributed new reagents/analytic tools; J.G. and C.C. analyzed data; and J.G., C.C., and J.-P.P. wrote the paper.

The authors declare no conflict of interest.

*This Direct Submission article had a prearranged editor.

See Commentary on page 19103.

¹Present address: Center for Brain Science, Harvard University, Cambridge, MA 02138.

²To whom correspondence should be addressed. E-mail: gjorgjieva@fas.harvard.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1105933108/-DCSupplemental.

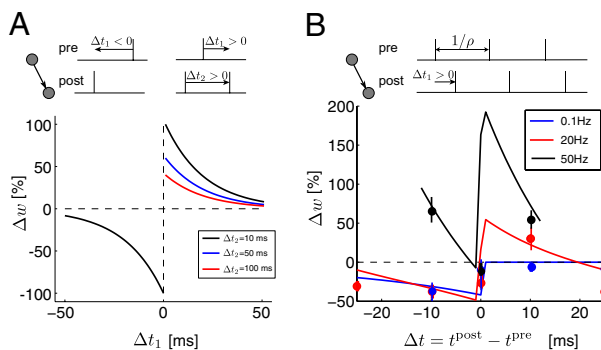


Fig. 1. The triplet STDP rule. (A) Synaptic depression is induced as in classical pair-based STDP using spike pairs separated by $\Delta t_1 = t^{\text{post}} - t^{\text{pre}} < 0$. Synaptic potentiation is induced using triplets of spikes consisting of two postsynaptic spikes and one presynaptic spike on the basis of the timing interval between them $\Delta t_1 = t^{\text{post}} - t^{\text{pre}} > 0$ and $\Delta t_2 = t^{\text{post}} - t^{\text{post}} > 0$. (B) Synaptic change as a function of the time between pre- and postsynaptic spikes in a protocol where 60 pairs were presented at different frequencies $\rho = 0.1, 20,$ and 50 Hz. Depression predominated at low frequency, whereas potentiation was more prevalent at high frequencies. The data points are experiments from ref. 17 and the lines were generated with the triplet STDP rule with the parameters taken from ref. 15.

receptive field development, an inhibitory postsynaptic membrane potential kernel with a membrane time constant of 20 ms was also used. Postsynaptic spikes were generated stochastically from the membrane potential, with a probability density of firing a spike at time t given by the transfer function $\langle y(t) \rangle = g(u(t))$. For simplicity, we used linear neurons where the transfer function was approximated by

$$g(u(t)) \approx g(u_0) + g'(u_0)(u(t) - u_0), \quad [2]$$

where the averaged membrane potential was u_0 . We also used $\nu = g(u_0)$ to denote the mean postsynaptic firing rate.

Synaptic Dynamics and Input Selectivity. Following the approach of ref. 25, we expressed the weight change as a Volterra expansion of both pre- and postsynaptic spike trains and the two learning rules: pair-based STDP

$$W_2(\Delta t) = \begin{cases} A_2^+ e^{-\Delta t/\tau_+}, & \Delta t \geq 0 \\ -A_2^- e^{\Delta t/\tau_-}, & \Delta t < 0, \end{cases} \quad [3]$$

where $\Delta t = t^{\text{post}} - t^{\text{pre}}$ denotes the timing difference between a post- and a presynaptic spike, τ_+ is the potentiation time constant, and τ_- is the depression time constant and triplet STDP

$$W_3(\Delta t_1, \Delta t_2) = A_3^+ e^{-\Delta t_1/\tau_+} e^{-\Delta t_2/\tau_+}, \quad \Delta t_1 \geq 0, \Delta t_2 \geq 0 \quad [4]$$

and 0 otherwise, where spike triplets $(t^{\text{pre}}, t^{\text{post}}, t^{\text{post}})$ affect synaptic potentiation depending on their timing difference $\Delta t_1 = t^{\text{post}} - t^{\text{pre}}$ and $\Delta t_2 = t^{\text{post}} - t^{\text{post}}$. The parameters used throughout this paper were those of the minimal triplet rule (15), i.e., $A_2^+ = 0$, $A_2^- = -6.5 \times 10^{-3}$, $A_3^+ = 7.1 \times 10^{-3}$, $\tau_+ = 16.8$ ms, $\tau_- = 33.7$ ms, $\tau_y = 114$ ms. Assuming slow learning dynamics (25), we derived Eq. 6 (Results) to describe the weight dynamics (SI Text).

We considered M input patterns, where pattern i had mean firing rate $\rho^{(i)}$, and pairwise and triplet correlation terms $A^{(i)}$ and $B^{(i)}$, respectively. Each input pattern i was associated with a probability p_i of occurrence and gave rise to an average postsynaptic firing rate $\nu^{(i)} = w^T \rho^{(i)}$. The selectivity of the postsynaptic neuron was $\text{Sel}(w) = 1 - (\sum_i p_i w^T \rho^{(i)}) / (\max_i w^T \rho^{(i)})$ (6).

To match the triplet rule to the BCM model, we set $A_2^- \rightarrow A_2^- \bar{\nu} / \rho_0^2$, where the expectation of the p^{th} power of the postsynaptic firing rate can be expressed as $\bar{\nu} = \sum_{i=1}^M p_i \nu^{(i)p}$. This quantity was approximated by low-pass filtering the p^{th} power of the instantaneous postsynaptic firing rate $\nu(t) = g(u(t))$ with a time constant of $\tau_p = 5$ s. For all of the calculations in this paper we took $p = 2$.

In the case of orthogonal rate-based patterns modeled as independent Poisson inputs, we proved that the maximally selective fixed points of the weight dynamics are stable (SI Text). For the development of selectivity in the case of correlation-based patterns, we calculated the fixed points of maximal selectivity only in the case of a reduced 2D system

(SI Text). In this case, two patterns were presented to the feedforward network, each consisting of two groups (or pools) of input neurons. Extensions to more than two correlated patterns are currently possible only with numerical simulations.

Numerical Simulations with Multiple Patterns. For all numerical simulations we simulated the triplet learning rule given by Eq. 4 that can also be expressed in differential form (SI Text). A lower bound of 0 and an upper bound of 3 were imposed on the weights. The methods for generating correlated spike trains and the correlation strength used in each figure are described in the SI Text. The phase plane diagrams for the 2D systems in Figs. 2D and 3E were plotted using the MATLAB software pplane written by John Polking (Rice University, Houston, TX).

Results

Triplet STDP Induces Selectivity with Rate-Based Patterns. Orientation-selective neurons in the primary visual cortex respond with higher firing rates when a bar is presented in a particular orientation and with lower rates when the bar is presented in a different orientation (26). This orientation selectivity is learned during receptive field development, and normal patterns of sensory experience are important for receptive field maturation (27). Bienenstock et al. (6) proposed a model for how orientation selectivity, or more generally pattern selectivity, is learned by a neural network: the BCM learning rule. In the BCM framework, a randomly chosen input pattern i (of M possible patterns) with rates $\rho^{(i)}$ is presented with probability p_i to a feedforward network with N inputs. The postsynaptic neuron responds with a firing rate $\nu^{(i)} = w^T \rho^{(i)}$, where w is the weight vector. The weight change induced by the BCM rule is proportional to the input firing rate

$$\dot{w} = \phi(\nu, \bar{\nu}) \rho \quad [5]$$

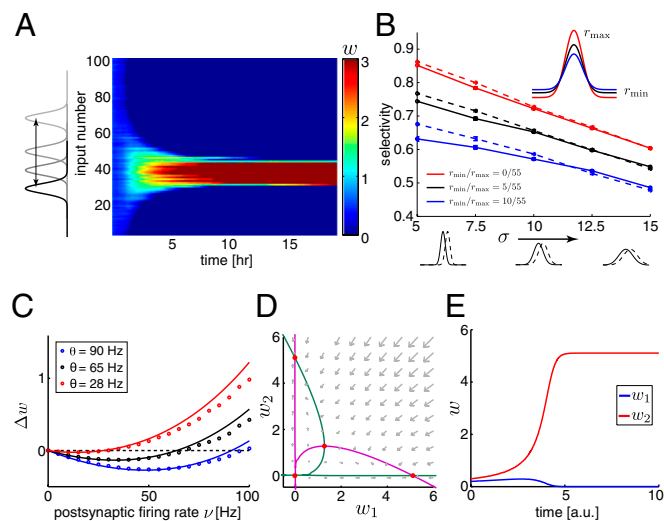
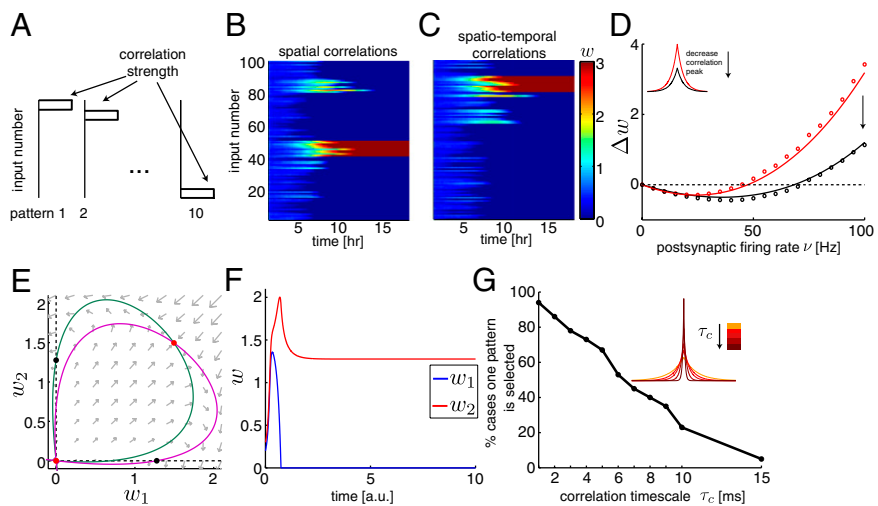


Fig. 2. Triplet STDP induces selectivity with rate-based patterns. (A) Evolution of the weights (Right) for 10 rate-based patterns (uniformly spaced Gaussian profiles across the 100 inputs) determining the inputs' firing rates (Left) presented to a feedforward network. The selected pattern corresponds to a Gaussian profile with $r_{\text{min}}/r_{\text{max}} = 5/55$ and $\sigma = 15$ Hz. (B) Mean (\pm SEM) selectivity (for 10 trials) as a function of the Gaussians' SD σ and for Gaussian profiles with different ratios of background to peak firing rates $r_{\text{min}}/r_{\text{max}} = \{0/55, 5/55, 10/55\}$ (solid lines, the triplet STDP rule; dashed lines, the BCM rule). The Gaussian profiles below illustrate the amount of overlap for two neighboring Gaussians. Numerical simulations implementing the differential form of the triplet STDP were performed in A and B. (C) Weight change Δw as a function of postsynaptic activity for three different input firing rates, which determine the threshold θ for weight modification. Symbols denote numerics and lines analytics. (D) 2D phase plane analysis for the analytically derived weight equation with orthogonal rate-based patterns. Nullclines in green and purple intersect at the equilibria shown in red. (E) An example trajectory for the two weights attracted to one of the stable nodes in D. (a.u., arbitrary units).

Fig. 3. Triplet STDP induces selectivity with correlation-based patterns. (A) Ten correlation-based patterns that have the same firing rates, but different correlation strength. (B) Evolution of the weights illustrates selectivity in the case of 10 correlation-based patterns. The firing rate of each of the 100 inputs was set to 10 Hz: 90 inputs had no correlations and 10 neighboring inputs (one of 1–10, 11–20, ..., 91–100 for each pattern) had strong spatial correlations (90% identical spikes). (C) Same as B except for the 10 correlated inputs in each pattern, for which exponentially decaying correlations with a time constant of 5 ms were used. Numerical simulations implementing the differential form of the triplet STDP were performed in B and C. (D) The average weight change Δw (for 100 weights) was computed for different initial conditions w_0 after 100 s. The symbols denote numerical results obtained by simulating the differential form of the triplet rule, and the lines indicate a semi-analytic solution by numerically solving Eq. 6 given an initial condition w_0 for 100 s. The average weight change was plotted as a function of the postsynaptic firing rate given by $\nu = w_0\rho$, where ρ was the input firing rate. Here we simulated two networks where the inputs had the same firing rate (10 Hz) and exponentially decaying correlations with a timescale of 10 ms. The correlation peak for the curve in black (SI Text) was half of the correlation peak for the curve in red ($\gamma = 9.09$, $\lambda = 9.09$; see Inset). (E) 2D (two groups of inputs) phase plane analysis for correlation-based patterns. Nullclines in green and purple intersect at the unstable fixed points shown in red. Imposing a lower bound at 0 resulted in stable maximally selective fixed points on the axes shown in black. (F) An example trajectory for the two weights attracted to one of the black equilibria in E. (G) Percentage of cases (over 100 trials) where all of the synapses from 1 of 10 input patterns (each consisting of 10 inputs) potentiate (Eq. 6). Same scenario as C, but for different correlation time constants τ_c . Correlations were symmetric and exponentially distributed (Inset).



and scales with a nonlinear function ϕ , which depends not only on the postsynaptic firing rate ν , but also on the average (over all patterns) of a nonlinear function of the postsynaptic rate $\bar{\nu} = \sum_i p_i (\nu^{(i)})^2$. The nonlinear function ϕ must be negative when the postsynaptic firing rate is below a given threshold θ —which itself depends on $\bar{\nu}$ —and positive when it is above it (Fig. 2C).

Interestingly, assuming a linear transfer function, the average weight change under the triplet rule can be written precisely as the BCM term plus some perturbation terms due to the input correlations (SI Text)

$$\dot{w}_j = \phi(\nu)\rho_j + \sum_{k=1}^N \Delta A_{jk} w_k + w^T \Delta B_j w \quad [6]$$

where $\phi(\nu) = \bar{W}_2 \nu + \bar{W}_3 \nu^2$ is the BCM term (Fig. 2C) with \bar{W}_2 and \bar{W}_3 being the overall area under the pair-based and triplet STDP rules, respectively. ΔA and ΔB describe the contributions from the input statistics (SI Text). To get depression at low postsynaptic firing rate and potentiation for higher firing rate, pairs of spikes must have an overall depressive effect ($\bar{W}_2 < 0$) and triplets of spikes must induce potentiation ($\bar{W}_3 > 0$), as is the case for the minimal triplet STDP model considered here (15).

There are two differences between the original BCM rule in Eq. 5 and the triplet model in Eq. 6. First, in the triplet model, the function ϕ depends only on the temporally averaged postsynaptic activity ν , whereas in the BCM model, ϕ also depends on the postsynaptic activity averaged over all patterns, $\bar{\nu}$. However, if we redefine the amplitude parameter for pair-based depression A_2^- as $A_2^- \bar{\nu} / \rho_0^2$, where ρ_0 is a constant denoting the target rate of the postsynaptic neuron (15), then both ϕ and ΔA in Eq. 6 will depend on $\bar{\nu}$.

The second difference is the presence of the two additional terms (ΔA and ΔB) in Eq. 6. If the inputs are Poisson neurons, we can rewrite Eq. 6 as

$$\dot{w} = (\phi(\nu, \bar{\nu}) \mathbf{1} + \Lambda) \rho, \quad [7]$$

where $\mathbf{1}$ denotes the identity matrix and Λ is a diagonal matrix (SI Text). If we now assume that the patterns are orthogonal, we can show that the condition $\dot{w} = 0$ in Eq. 7 gives rise to 2^N fixed points. Moreover, the N maximally selective fixed points, $w^{*(n)} = (0, \dots, 0, w_n^*, 0, \dots, 0)$, are stable fixed points (Fig. 2D and E shows an illustration in two dimensions), which is consistent with results of the BCM theory.

The general problem of deriving selectivity analytically has not yet been solved; however, numerical simulations suggest that the triplet rule successfully drives selectivity even when the rate-based inputs are nonorthogonal. We designed an experiment to examine the level of selectivity as 10 Gaussian input patterns were presented to the network with varying amounts of overlap (Fig. 2A). The Gaussian profiles were uniformly spaced across the input neurons, and we varied the ratio of the background firing rate (r_{\min}) to the peak firing rate (r_{\max}) of each Gaussian (Fig. 2B, Inset) and the SD (σ). The Gaussian profiles were closest to orthogonal for small σ and $r_{\min}/r_{\max} = 0$, whereas their amount of overlap increased as either σ or r_{\min}/r_{\max} increased. We computed the amount of selectivity of the postsynaptic neuron at the end of a simulation when the weights reached a steady state (Fig. 2A) as a function of the Gaussians' SD (σ). We observed that for the case of nearly orthogonal Gaussian profiles ($r_{\min}/r_{\max} = 0/55$ and small σ , Fig. 2B, red lines), the achieved selectivity was close to the maximally attainable selectivity of 0.9 for 10 orthogonal patterns (6). The selectivity dropped as σ or r_{\min}/r_{\max} increased (Fig. 2B). We compared the performance of the triplet to the BCM rule using the same Gaussian input profiles, while keeping the weights nonnegative during the entire simulation (Fig. 2B, dashed lines). As expected, we obtained similar results to those of the triplet rule.

Triplet STDP Induces Selectivity with Correlation-Based Patterns. In addition to mapping the triplet to the BCM rule for rate-based patterns, the triplet rule further generalizes the BCM model: In Eq. 6, ΔA and ΔB depend on the second- and third-order input correlations, respectively (SI Text); therefore, we expected triplet STDP to be sensitive to spatiotemporal correlations in the inputs.

To examine our hypothesis, we presented 10 “correlation-based” patterns to the feedforward network with 100 inputs. The correlation-based patterns were determined by different pairwise and third-order correlations (Fig. 3A), but had the same input firing rates. Therefore, the response of the postsynaptic neuron to each pattern was the same, which prevented us from using the same measure of selectivity as for rate-based patterns. Instead, selectivity was defined in terms of the selective potentiation of a group of correlated inputs. Fig. 3B shows a simulation with purely spatial correlations that had no temporal structure (the input correlations were due to identical spikes in the neurons). The weights from one pattern potentiated (inputs 41–50), whereas the other weights depressed. When we presented spatiotemporal correlations with an exponentially decaying

correlation function, selectivity was also achieved: A set of 10 weights (81–90) characteristic of one pattern potentiated, whereas the other weights depressed (Fig. 3C).

In the case of correlation-based patterns, the sliding threshold depends both on the input firing rates and correlations. Although the additional terms ΔA and ΔB in Eq. 6 prevent us from deriving an explicit expression for the modification threshold, we illustrated the dependence of the threshold on the correlation strength in Fig. 3D. Here we computed the average weight change for 100 weights as a function of the postsynaptic firing rate for two different input correlations: In both cases the firing rate of the 100 inputs was the same (10 Hz), and the correlation function was a decaying exponential with a timescale of 10 ms (Fig. 3D, *Inset*); however, the two functions differed in the correlation peak (black peak was one-half of the red peak). The network with the higher correlation peak had a lower modification threshold, resulting in a larger potentiation region.

Due to the increased complexity of the system when the inputs are correlated, we derived the fixed points of maximal selectivity (w_1^* , 0) and (0, w_2^*) in a small network of two groups of input neurons and analyzed their stability (*SI Text*). A lower bound had to be introduced to prevent the weights from becoming negative (in agreement with Dale's law). For the 2D network, we found that the maximally selective fixed points were always stable. Fig. 3E shows the 2D phase plane, where the two unstable fixed points (red symbols) drive the weight trajectories toward the axes where the stable maximally selective fixed points are located (black symbols). Example weight trajectories are shown in Fig. 3F for one choice of initial condition.

We extended the simulation in Fig. 3C to examine how the temporal correlation structure of the inputs influences the selective potentiation of synaptic weights corresponding to different patterns. We studied a particular example of a symmetric spatiotemporal correlation: an exponential function decaying in time, which was the same for all pairs and the same for all triplets of inputs (*SI Text*). Therefore, while preserving the correlation strength, we examined the role of the correlation timescale on the selective potentiation of synaptic weights (Fig. 3G, *Inset*). Increasing the correlation timescale had a similar effect as “diluting” the correlation strength. The triplet STDP rule failed to consistently potentiate the weights of one input pattern, and often two or three patterns were simultaneously selected (Fig. 3G). Therefore, correlations over broad timescales fail to evoke selective potentiation of correlation-based input patterns and could be used to understand the implications of different correlation structures in different brain regions.

Triplet STDP, but Not Pair-Based STDP, Can Induce Selectivity Driven by Third-Order Correlations. Despite the advantage of triplet STDP over classical pair-based STDP to capture a large variety of experimental plasticity data (for instance, frequency dependence) (17), we asked whether triplet STDP can do computations that pair-based STDP cannot. Previous studies have shown that for correlation-based patterns pair-based STDP selects the correlated groups of inputs in the case of static patterns (where the correlations are always presented to the same group of inputs), but have not addressed the case of dynamic patterns (28). We hypothesized that triplet STDP will be able to select patterns determined by the inputs' third-order correlations, whereas pair-based STDP will not be able to distinguish any higher-than-pairwise correlations.

For this task, we designed a selectivity scenario consisting of two correlation-based patterns presented to a feedforward network of six input neurons. The inputs in the two patterns consisted of the same firing rates and the same pairwise correlations, but differed in the presence or absence of third-order correlations in half of the inputs (Fig. 4A). Pattern 1 consisted of third-order correlations in inputs 1–3 (denoted as group 1) and no third-order correlations in inputs 4–6 (denoted as group 2). Pattern 2 consisted of third-order correlations in inputs 4–6 of group 2 and no third-order correlations in inputs 1–3 of group 1. Next, we presented each pattern to the network with a fixed probability; for instance, in Fig. 4A we illustrate a scenario in which, of 10 pattern presentations, pattern 1 was presented on average eight times (with probability 0.8). As the probability of presenting pattern 1

varied between 0.5 and 1 (Fig. 4C), we estimated the probability that pattern 1 wins in 200 simulation runs. An example of pattern 1 winning is illustrated in Fig. 4B, where inputs 1–3 potentiate, and inputs 4–6 depress to 0. We estimated the probability that pattern 1 wins for both pair STDP (Fig. 4C, red symbols) and triplet STDP (Fig. 4C, black symbols). When only pattern 1 was presented to the network (third-order correlations only in inputs 1–3), almost all simulations resulted in the potentiation of these inputs under the triplet STDP rule. As the probability of presenting pattern 1 decreased to 0.5, the probability that pattern 1 wins also decreased. However, pair-based STDP was not sensitive to the third-order input correlations and it treated both patterns equally, selecting each pattern randomly with equal probability of 1/2 regardless of how frequently pattern 1 was presented.

This result demonstrates that the triplet STDP rule can distinguish between inputs solely on the basis of the higher-order correlation structure, which pair-based STDP ignores. As a result, triplet STDP will be computationally more powerful in systems where such higher-order correlations have been characterized (22–24) and where firing rates and pairwise correlations are of similar magnitude.

These studies demonstrate that because pair-based STDP uses only pairs of spikes to induce synaptic plasticity, it is sensitive only up to pairwise correlations. Thus, we suspected that the triplet STDP rule, which evokes plasticity using triplets of spikes, will be sensitive only up to third-order correlations. To confirm this, we repeated the simulation scenario above for a network of two groups of five neurons each. In each case, the two input groups had the same lower-order correlations, but differed in the presence or absence of higher-order correlations in each group. We studied correlations with highest order of five. The triplet

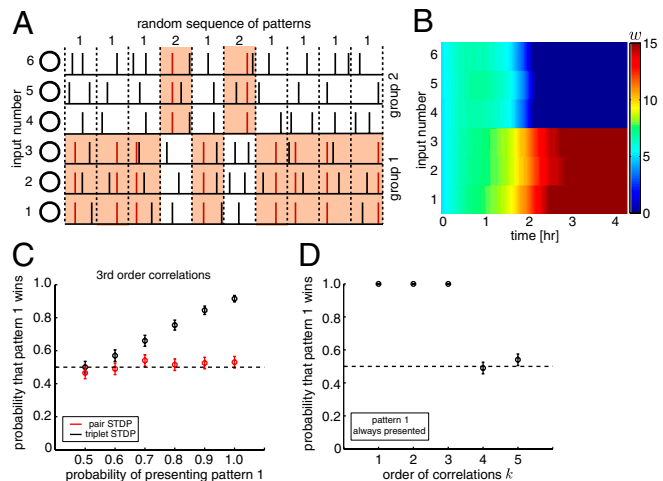


Fig. 4. Triplet STDP, and not pair-based STDP, can distinguish between patterns determined by third-order correlations. (A) Two patterns were randomly presented to a feedforward network: The inputs in the two patterns had the same firing rates and the same pairwise correlations. The patterns differed only by the presence or absence of third-order correlations in half of the inputs (illustrated with the red triplets of spikes and the colored background). The probability of presenting pattern 1 was varied, e.g. of 10 pattern presentations, pattern 1 was presented with probability 0.8. (B) The evolution of the weights under the triplet STDP rule demonstrating an example where pattern 1 (inputs 1–3) wins. (C) Pattern 1 was presented to the network at different probabilities and the mean \pm SEM of the probability that pattern 1 wins was computed for 200 simulation runs: triplet STDP (black symbols) and pair STDP (red symbols). (D) The triplet STDP rule is sensitive up to third-order correlations, but not to higher-order correlations. Pattern 1 was always presented and consisted of two groups of five neurons each. Both groups had the same correlations up to (but not including) order k (horizontal axis). Group 1 had nonzero $\geq k^{\text{th}}$ -order correlations, and group 2 had zero $\geq k^{\text{th}}$ -order correlations. The mean \pm SEM of the probability that pattern 1 wins (i.e., all of the weights of group 1 potentiate) was computed for 200 simulation runs. Dashed lines correspond to chance level.

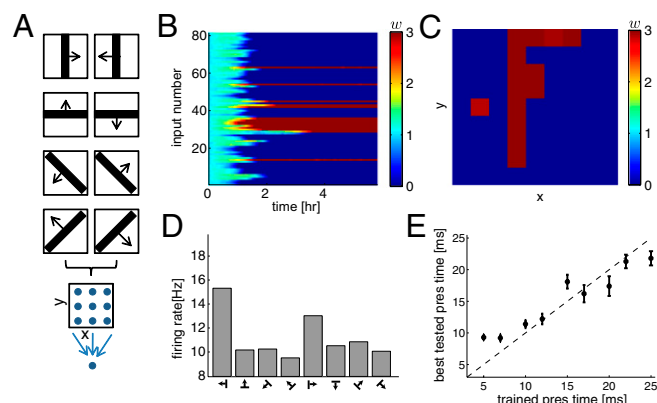


Fig. 5. Triplet STDP leads to spatiotemporal receptive field development. (A) Four different bars (horizontal, vertical, and the two diagonals) on a 9×9 -pixels image) were presented as inputs to a feedforward network with a single postsynaptic neuron; each bar can move in one of two directions, giving a total of eight patterns. (B) Time evolution of the 81 synaptic weights. (C) Final weights reordered in a grid corresponding to the input location. (D) Histogram of the postsynaptic firing rate plotted after convergence of the weights, at the end of the learning in B. The firing rates shown resulted from the presentation of the eight different patterns (four orientations and two directions) averaged over 200 s. (E) After learning with different training presentation times (5, 7, 10, 12, 17, 20, 22, and 25 ms), the weights were frozen during a testing phase. The pattern (of the eight patterns) that resulted in the highest firing rate at the training presentation time was presented again to the network at different tested presentation times (5, 7, 10, 12, 17, 20, 22, and 25 ms) while the firing rate was measured. The best tested presentation time for which the firing rate (averaged over 100 s) was the highest is plotted against the training presentation time (mean \pm SEM over 10 trials).

rule distinguished correlations up to third-order, but was insensitive to fourth- and fifth-order correlations (Fig. 4D). We expect that for a learning rule to be sensitive to higher than third-order correlations, the rule would need to incorporate more than three spikes, or the neural model would have to be nonlinear.

Spatiotemporal Receptive Field Development. Due to its sensitivity to higher-order correlations, we expected that the triplet rule would succeed in driving the development of spatiotemporal receptive field properties encountered in visual cortex, such as orientation and direction selectivity. Using the same feedforward framework we presented eight different patterns consisting of four bars at different orientations, each moving in one of two directions, as drawn in Fig. 5A. The input neurons in the network were organized in a 9×9 grid mapping to a single postsynaptic neuron. Each input spike produced both an EPSP and an IPSP with a longer time constant (*Model and Methods*) (29). The synapses in the network selectively refined under the triplet learning rule, some depressing to the lower bound of zero and others potentiating maximally to the upper bound (Fig. 5B). At the end of the simulation the postsynaptic neuron became selective to one of the bars, as shown in Fig. 5C by the strong weights corresponding to inputs arranged in the bar with a vertical orientation (although this selectivity was not as robust as for nonmoving bars, i.e., Fig. 2).

Recently, after training with moving stimuli, Li et al. (30) observed the emergence of direction selectivity in cortical neurons of visually naive ferrets for the trained directions of motion. Motivated by this experimental result, we examined the histogram of the postsynaptic firing rate for all patterns after learning (Fig. 5D). We found that the postsynaptic neuron responded with the highest firing rate for a bar in given (vertical) orientation and moving in one direction (left), thus becoming weakly selective to direction.

The receptive field in this scenario also developed a temporal structure as a result of the spatiotemporal correlations imposed by the moving bars. To visualize this effect, the weights were frozen after learning (at a given presentation time) and a test

phase was conducted. The pattern that produced the highest firing rate for the trained presentation time was tested at a range of testing presentation times. The best testing presentation time for which the firing rate was the highest was plotted against the training presentation time. There was a high correlation between training presentation time and best tested presentation time, demonstrating that the neuron became selective to the training speed (Fig. 5E).

In this scenario the development of direction selectivity relied on two elements. First, triplet STDP was sensitive to the spatiotemporal input correlations, and second, we used a modified postsynaptic potential kernel that included an inhibitory component with a longer time constant than for the excitatory current. This additional IPSP was necessary for obtaining direction selectivity because it made the postsynaptic neuron sensitive to a temporal derivative of the input currents. A similar assumption has been used in previous models of direction selectivity driven by pair-based STDP (29, 31), although in these models the postsynaptic neuron was usually trained in a single direction. The temporal sensitivity to obtain direction selectivity can be also implemented by the order of presentation of the stimuli, instead of modifying the dynamics of the postsynaptic neuron, as shown with the BCM rule (32). More robust direction selectivity as observed biologically (26, 33) can also be obtained by using a network with recurrent excitatory and inhibitory connections (29), where the source of inhibition could arise from neocortical interneurons, such as fast-spiking interneurons. Finally, we note that the receptive fields we observed depend on the choice of input statistics (here moving bars), the learning rule, and the neural model. Receptive fields with more complex structure (such as Gabor patches) could be developed with inputs of richer spatiotemporal contents, such as natural images.

Discussion

The BCM theory is attractive because it generates selectivity in a variety of scenarios and has been supported experimentally (6–8). Synaptic plasticity, however, has been shown to depend on the precise spike timing (9, 10) classically modeled by pair-based STDP (11, 12). In this paper, we show that a different spike-based rule, triplet STDP, known to accurately capture plasticity experiments (15), exhibits the computational properties of the BCM rule and is additionally sensitive to higher-order spatiotemporal input correlations.

We mapped the triplet STDP to the BCM learning rule for rate-based patterns, determined by the input firing rates. Consistent with the BCM theory, we showed that for nonoverlapping (orthogonal) patterns, the maximally selective fixed points of the weight dynamics under triplet STDP are always stable. For overlapping Gaussian patterns, numerical simulations demonstrated that the selectivity achieved with the triplet rule is similar to the selectivity achieved with the BCM rule.

We also showed that the triplet rule can generate selectivity in the case of correlation-based dynamic patterns, determined solely by the higher-order input correlations. However, because the rule uses triplets of spikes to induce plasticity, it is sensitive to higher-order correlations of maximum order three. This sensitivity led to the development of direction selectivity and speed selectivity. We observed that increasing the input correlation timescale dilutes the correlation strength, which prevents the cooperation of inputs necessary for the emergence of selectivity. Therefore, our results make experimental predictions about the types of correlation structure that lead to selectivity.

Higher-order correlations have not only been measured in the brain, but also shown to play an important role in visual coding and representing experimental data (22–24, 34). Higher-order correlations are ubiquitous in sensory stimuli, such as natural stimuli and speech signals (35, 36). These correlations have been previously used in learning rules to extract the independent components or features in natural images resulting in simple cell receptive fields as seen in V1 (35, 37). One such rule is the BCM rule, shown to perform projection pursuit that relies on higher-order correlations to find the most interesting projection, i.e., the one that minimizes the Gaussianity of the output distribution (38, 39) and is closely related to independent component analysis (ICA). Because of its mapping to the BCM rule, we can interpret

the triplet rule as a method for performing such ICA-like computations. In addition to spatial ICA (where the independent components are obtained from the input statistics at each fixed time point) (35), the triplet rule can also perform temporal ICA-like computations, which additionally rely on the temporal structure of the inputs (40).

Several other models have addressed the issue of spiking-based implementations of the BCM learning framework (13, 14). Izhikevich and Desai (14), for instance, proposed that by implementing classical pair-based STDP with nearest-neighbor spike interactions, the rule can be mapped to the BCM rule. However, their model failed to capture the frequency dependence of ref. 17 if pairs of spikes are presented at different frequencies (21) and considered the rather crude approximation that input and outputs are independent. Although the model of Senn et al. (13) captured the frequency dependence of the pairing protocol, it could not reproduce the triplet and quadruplet experiments of ref. 18 (see ref. 15) and the correspondence to the BCM rule is only approximate (the sliding threshold depends on the weights and not on the postsynaptic firing rate as in the BCM framework). Toyozumi et al. (41) derived an alternate spike-based learning rule designed to maximize the information transmission between an ensemble of inputs and the output of a postsynaptic neuron. Although such plasticity rules derived from the infomax principle can generalize the BCM theory to spiking neurons and can be reduced under some assumptions to the triplet STDP rule (42), the dynamics of these rules are rather complicated to be studied

analytically in contrast to the triplet STDP model. The same problem arises with biophysical models (43, 44) or with more elaborate phenomenological models (45), which have primarily been studied numerically. Therefore, the triplet STDP model is a good trade-off: It can reproduce a large set of electrophysiological data and yet has a relatively simple formulation so that we can study it analytically and generalize its functional properties to networks of different size and input statistics. Additionally, the triplet STDP model extends the BCM theory to correlation-based patterns with higher-order correlations, which were not considered by any of the above models.

The emergence of input selectivity in this work has been studied with a linear neuron. This analysis can be extended to consider nonlinear neurons and to relate it to a variant of the BCM theory that uses nonlinear neurons (38). We expect that with an appropriate nonlinearity, this variant will lead to a more robust orientation selectivity for natural image stimuli (46). Furthermore, because the triplet rule is also sensitive to the temporal structures of the inputs, we expect it to be able to capture relevant aspects of the spatiotemporal statistics of natural scene environments.

ACKNOWLEDGMENTS. We thank Stephen Eglen, Adrienne Fairhall, and Robert Froemke for helpful discussions and reading drafts of this manuscript. This work was funded by a Cambridge Overseas Research Studentship and Trinity College Internal Graduate Studentship (to J.G.), the Agence Nationale de la Recherche Grant ANR-08-SYSC-005 (to C.C.), the Wellcome Trust, and the Swiss National Science Foundation Grant 31-133094 (to J.-P.P.).

- Bliss TV, Collingridge GL (1993) A synaptic model of memory: Long-term potentiation in the hippocampus. *Nature* 361(6407):31–39.
- Malenka RC, Nicoll RA (2009) Long-term potentiation—a decade of progress? *Science* 285:1870–1874.
- Dudek SM, Bear MF (1992) Homosynaptic long-term depression in area CA1 of hippocampus and effects of N-methyl-D-aspartate receptor blockade. *Proc Natl Acad Sci USA* 89:4363–4367.
- Artola A, Bröcher S, Singer W (1990) Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature* 347(6288):69–72.
- Bliss TV, Lomø T (1973) Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J Physiol* 232:331–356.
- Bienenstock EL, Cooper LN, Munro PW (1982) Theory of the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *J Neurosci* 2:32–48.
- Cooper LN, Intrator N, Blais BS, Shouval HZ (2004) *Theory of Cortical Plasticity* (World Scientific, Singapore).
- Kirkwood A, Rioult MC, Bear MF (1996) Experience-dependent modification of synaptic plasticity in visual cortex. *Nature* 381(6582):526–528.
- Markram H, Lübke J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic AP and EPSP. *Science* 275:213–215.
- Bi GQ, Poo MM (1998) Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci* 18:10464–10472.
- Gerstner W, Kempter R, van Hemmen JL, Wagner H (1996) A neuronal learning rule for sub-millisecond temporal coding. *Nature* 383(6595):76–78.
- Abbott LF, Nelson SB (2000) Synaptic plasticity - taming the beast. *Nat Neurosci* 3:1178–1183.
- Senn W, Markram H, Tsodyks M (2001) An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Comput* 13:35–67.
- Izhikevich EM, Desai NS (2003) Relating STDP to BCM. *Neural Comput* 15:1511–1523.
- Pfister JP, Gerstner W (2006) Triplets of spikes in a model of spike timing-dependent plasticity. *J Neurosci* 26:9673–9682.
- Pfister JP, Gerstner W (2006) Beyond pair-based STDP: A phenomenological rule for spike triplet and frequency effects. *Adv Neural Inf Process Syst* 18:1083–1090.
- Sjöström PJ, Turrigiano GG, Nelson SB (2001) Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32:1149–1164.
- Wang HX, Gerkin RC, Nauen DW, Wang GQ (2005) Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nat Neurosci* 8:187–193.
- Froemke RC, Dan Y (2002) Spike-timing dependent plasticity induced by natural spike trains. *Nature* 416(6879):433–438.
- Froemke RC, Tsay IA, Raad M, Long JD, Dan Y (2006) Contribution of individual spikes in burst-induced long-term synaptic modification. *J Neurophysiol* 95:1620–1629.
- Clopath C, Gerstner W (2010) Voltage and spike timing interact in STDP - A unified model. *Front Synaptic Neurosci* 2:25. doi:10.3389/fnsyn.2010.00025.
- Montani F, et al. (2009) The impact of high-order interactions on the rate of synchronous discharge and information transmission in somatosensory cortex. *Philos Trans R Soc A* 367:3297–3310.
- Ohiorhenuan IE, et al. (2010) Sparse coding and high-order correlations in fine-scale cortical networks. *Nature* 466(7306):617–621.
- Pillow JW, et al. (2008) Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454(7207):995–999.
- Kempter R, Gerstner W, van Hemmen JL (1999) Hebbian learning and spiking neurons. *Phys Rev E* 59:4498–4514.
- Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* 160:106–154.
- White LE, Fitzpatrick D (2007) Vision and cortical map development. *Neuron* 56:327–338.
- Song S, Abbott LF (2001) Cortical development and remapping through spike timing-dependent plasticity. *Neuron* 32:339–350.
- Shon AP, Rao RPN, Sejnowski TJ (2004) Motion detection and prediction through spike-timing dependent plasticity. *Network Comput Neural Syst* 15:179–198.
- Li Y, Hooser SDV, Mazurek M, White LE, Fitzpatrick D (2008) Experience with moving visual stimuli drives the early development of cortical direction selectivity. *Nature* 456(7224):952–956.
- Rao RP, Sejnowski TJ (2003) Self-organizing neural systems based on predictive learning. *Philos Trans A Math Phys Eng Sci* 361:1149–1175.
- Blais BS, Cooper LN, Shouval H (2000) Formation of direction selectivity in natural scene environments. *Neural Comput* 12:1057–1066.
- De Valois RL, Yund EW, Hepler N (1982) The orientation and direction selectivity of cells in macaque visual cortex. *Vision Res* 22:531–544.
- Luczak A, Barthó P, Marguet SL, Buzsáki G, Harris KD (2007) Sequential structure of neocortical spontaneous activity *in vivo*. *Proc Natl Acad Sci USA* 104:347–352.
- Olshausen BA, Field DJ (1996) Emergence of simple-cell receptive field properties by learning a space code for natural images. *Nature* 381(6583):607–609.
- Simoncelli EP, Olshausen BA (2001) Natural image statistics and neural representation. *Annu Rev Neurosci* 24:1193–1216.
- Bell AJ, Sejnowski TJ (1995) An information-maximization approach to blind separation and blind deconvolution. *Neural Comput* 7:1129–1159.
- Intrator N, Cooper LN (1992) Objective function formulation of the BCM theory of visual cortical plasticity: Statistical connections, stability conditions. *Neural Netw* 5:3–17.
- Blais BS, Intrator N, Shouval H, Cooper L (1998) Receptive field formation in natural scene environments. Comparison of single-cell learning rules. *Neural Comput* 10:1797–1813.
- Clopath C, Longtin A, Gerstner W (2008) An online Hebbian learning rule that performs Independent Component Analysis. *Adv Neural Inf Process Syst* 20:312–328.
- Toyozumi T, Pfister JP, Aihara K, Gerstner W (2005) Generalized Bienenstock-Cooper-Munro rule for spiking neurons that maximizes information transmission. *Proc Natl Acad Sci USA* 102:5239–5244.
- Hennequin G, Gerstner W, Pfister JP (2010) STDP in adaptive neurons gives close-to-optimal information transmission. *Front Comput Neurosci* 4:143. doi:10.3389/fncom.2010.00143.
- Castellani GC, Quinlan EM, Cooper LN, Shouval HZ (2001) A biophysical model of bidirectional synaptic plasticity: Dependence on AMPA and NMDA receptors. *Proc Natl Acad Sci USA* 98:12772–12777.
- Shouval HZ, Bear MF, Cooper LN (2002) A unified model of NMDA receptor-dependent bidirectional synaptic plasticity. *Proc Natl Acad Sci USA* 99:10831–10836.
- Clopath C, Büsing L, Vasilaki E, Gerstner W (2010) Connectivity reflects coding: A model of voltage-based STDP with homeostasis. *Nat Neurosci* 13:344–352.
- Shouval H, Intrator N, Cooper LN (1997) BCM network develops orientation selectivity and ocular dominance in natural scene environment. *Vision Res* 37:3339–3342.