

STDP Pattern Onset Learning Depends on Background Activity

James Humble, Sue Denham and Thomas Wennekers

Abstract Spike-timing dependent plasticity is a learning mechanism used extensively within neural modelling. The learning rule has been shown to allow a neuron to find the beginning of a repeated spatio-temporal pattern among its afferents. In this study we adduce that such learning is dependent on background activity, and is un-stable when in a noisy framework. We also present insights into the neuron's encoding.

1 Introduction

Spike-timing dependent plasticity (STDP) is a well established mechanism that permits spike time differences between pre- and postsynaptic activities to affect changes in synaptic efficacy [1, 2, 3, 4, 11, 16, 22]. If presynaptic activity precedes postsynaptic activity, the conjoining synapse is strengthened—usually according to some STDP function—and if the reverse is observed, the synapse is depressed. Moreover, STDP can increase the mutual information between inputs and outputs of simple networks ([14] used information theory to quantify learning performance), provides a function for Hebbian learning and development and captures the causality of determining the direction of synaptic change that is implied by Hebb's original

James Humble
School of Computing and Mathematics, University of Plymouth, UK
e-mail: james.humble@plymouth.ac.uk

Sue Denham
School of Psychology, University of Plymouth, UK
e-mail: s.denham@plymouth.ac.uk

Thomas Wennekers
School of Computing and Mathematics, University of Plymouth, UK
e-mail: thomas.wennekers@plymouth.ac.uk

postulate. STDP has accordingly been studied extensively and is commonly used as a substrate of forms of learning [9, 10, 20].

STDP has been applied successfully to simple pattern learning and complicated competitive pattern learning [7, 12, 13, 15, 19, 6]. Masquelier et al. [12] demonstrated that a neuron equipped with STDP can learn repeated spatio-temporal spike patterns even when embedded in a statistically identical distractor signal [12]. Such learning takes advantage of the view that a pattern is a succession of spike coincidences, and these coincidences combined with STDP's pre- and postsynaptic temporal considerations form the basis of many STDP-based learning rules.

Envisage a neuron that has learnt to fire after a subset of its afferents present a pattern as in [12]. In the present work we adduce this neuron's high sensitivity and dependence on background activity unrelated to the pattern; we also show that the temporal accuracy of the learning can be un-stable. Subsequently, insights into the neuron's encoding are also possible.

2 Methods

The simulations were performed using custom made C software. Source code is available from the authors upon request.

2.1 *Network Structure and Afferent Input*

The network structure was identical to that of Masquelier et al. [12]: 2000 afferents converge on one neuron. These 2000 afferents carry Poisson spike trains (64 Hz) which were produced 'on the fly', with 1000 of these afferents occasionally conveying a repeated spatio-temporal pattern of 50 ms duration.

To produce this input a simulation is segmented into 50 ms windows. During each subsequent bin, afferents communicating the pattern have a 0.25 probability of displaying the pattern; the remaining afferents consistently projected random Poisson activity. The pattern was not presented consecutively: this ensured each pattern presentation was preceded and followed by random activity. This protocol is similar to Masquelier et al. but allows on-line production of afferents' inputs. Analysis was carried out to ensure that the pattern was statistically identical to the random distractor signal. As in [12], Poisson spike trains were 54 Hz initially, with 10 Hz further noise added to all afferents including those of the pattern; however no jitter was added to the pattern.

A full simulation ran for up to 3000 s. It was found that any singular weight could converge (change from one weight limit to another) in 200 s, therefore weight values were recorded every 2 s allowing sufficient visualisation of weight trajectories and distributions.

2.2 Neuron Model

A leaky integrate-and-fire (LIF) neuron model was used such that membrane potential was modeled by (1) where $\tau_m=10$ ms is the membrane time constant and I is total input to the cell from all afferents. Input activity consists of pulses of current 1 for duration dt . A threshold of $\theta = 1$ was used.

$$\tau_m \frac{dV}{dt} = -V + I \quad \text{if } V \geq \theta \text{ then reset } V = 0 \quad (1)$$

2.3 Spike-Timing Dependent Plasticity

An exponential STDP rule was used such that $W_{t+\Delta t} = W_t + f(\tau)$. Equation (2) describes the STDP function used where $\tau_p = \tau_m=20$ ms. These values of τ_p and τ_m were chosen similar to those observed experimentally [16, 5, 8, 21, 23], where the strongest synaptic modifications occurred within a window of ± 20 ms. These constants vary from other studies such as Masquelier et al. [12, 13] who used $\tau_p=16.8$ ms and $\tau_m=33.7$ ms.

$$\begin{aligned} f(\tau) &= A_p \times \exp\left(\frac{-\tau}{\tau_p}\right) & \text{if } \tau \geq 0 \\ f(\tau) &= A_m \times \exp\left(\frac{\tau}{\tau_m}\right) & \text{if } \tau < 0 \end{aligned} \quad (2)$$

Learning rates were assigned as follows: $A_p = 0.002 \times W_{\max}$ and $A_m = -A_p \times \left(\frac{\tau_p}{\tau_m}\right) \times 1.05$. The assignment of A_m is in accordance with the finding by Song et al. [19] that $\frac{A_m}{A_p}$ needs to be slightly larger than one to maintain reasonable postsynaptic activity.

2.4 Maximum Synaptic Weight

Many studies into STDP pattern learning use weight bounding to control synaptic change; such saturation limits stop unstable forms of STDP based learning changing synaptic efficacies ad infinitum. Furthermore, the maximum synaptic weight can have a great impact on the learning process: if too low, postsynaptic activity may not reach threshold, and if too high, postsynaptic activity may be erratic and uninformative. It is thus implemented with (3), where Δv is the difference in membrane potential required to go from resting to threshold, $\langle r(t) \rangle$ is the average firing rate of afferents, N_{input} is the number of afferents in the pattern (here 1000) and A is an additional constant with arbitrary units.

$$W_{\max} = \frac{\left[\frac{\Delta v}{\tau_m \times \langle r(t) \rangle \times \Delta t} \right] + A}{N_{\text{input}}} \quad (3)$$

Therefore, after learning one envisages that all synapses participating in a pattern are at W_{\max} and the rest are at 0, and thus the afferents in the pattern will evoke an action potential from the output neuron. However, as all afferents in the pattern may not be learnt through STDP, and the afferents are not guaranteed to fire at the average firing rate or sufficiently within the membrane time constant, A can be adjusted to allow fluctuations in the total membrane potential to reach the neuron's threshold. To that end, an appropriately selected value can allow STDP to be affected by positive fluctuations in input and membrane potential—found to be crucial by Song et al. [19] and designed into a learning rule by Senn et al. [18].

Crucially, (3) is a ‘best guess’ approach to assigning maximum synaptic weights because a ‘best’ value will depend on how many afferents are fully potentiated at the end of learning, and this changes between simulations depending on the pattern and chance.

2.5 Initial Synaptic Weights

The initial strength of synaptic weights is as important as their maximum efficacy. Several different approaches are used to set initial weights; one approach is to set weights \mathbf{w} with random values drawn from a Uniform distribution between 0 and W_{\max} as in (4). This initially results in the output neuron firing frequently.

$$0 < \mathbf{w} \leq W_{\max} \quad (4)$$

Initial weights should be set such that fluctuations from all afferents cause the output neuron to fire. This is important because it allows all synapses contributing to postsynaptic activity to be affected by STDP [19].




Fig. 1 Postsynaptic latency—relative to the pattern start—as a function of discharges. When the neuron discharges outside of the pattern a latency of 0 is shown. The STDP clearly learns the pattern and has similar periods to those observed by Masquelier et al.: 1) when the neuron is non-selective to the pattern and most weights are being depressed due to $\int_{\tau_m}^{\tau_p} f(\tau) < 0$; 2) when the neuron is training to the beginning of the pattern; and 3) when the neuron consistently fires to a part of the pattern.

3 Results

During any typical simulation several phases are visible—first observed by Masquelier et al. [12]. The first two of these phases are similar to [12]; the last phase slightly differs.

3.1 Typical Results

It was found that STDP trained best with $A = 20$ (cf. (3)). Figure 1 depicts a typical simulation of 1000 s. Our results reinforce those of Masquelier et al., viz. that a neuron equipped with STDP is able to effectively find and train to a repeated spatio-temporal pattern embedded within a statistically identical signal. This is not surprising as the mechanism presented by Masquelier et al. is very general.

Masquelier et al. found that once a neuron had trained to a pattern it had “converged towards a fast and reliable pattern detector.” We did not find this; instead, we found that a learnt neuron’s firing latency appeared to drift backward through the pattern, see Fig. 1.

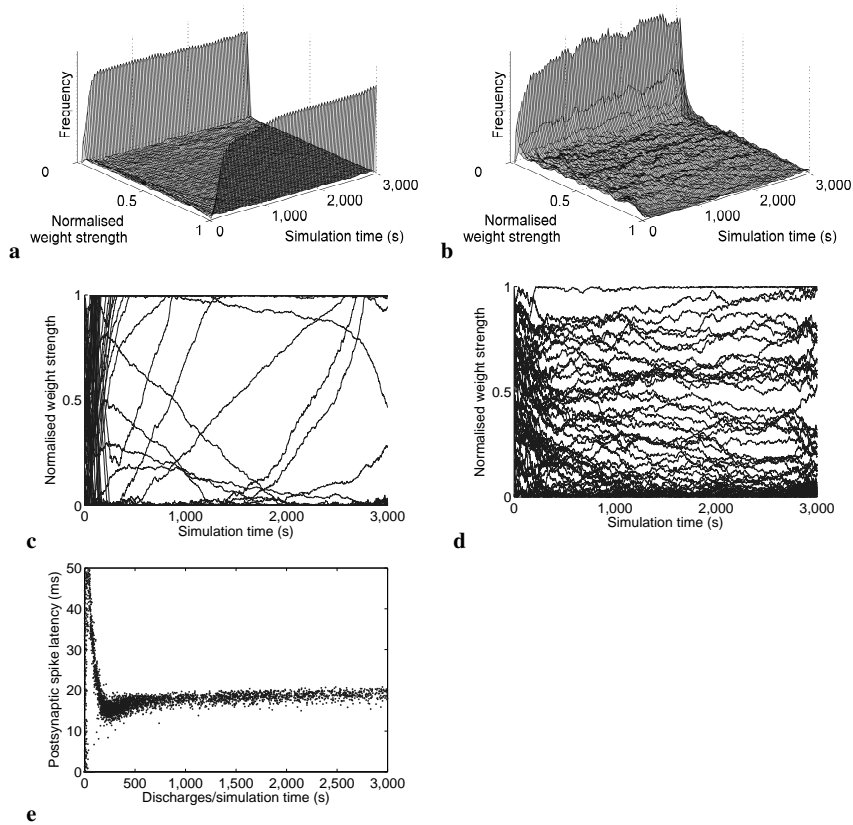


Fig. 2 Results of a longer simulation (3000 s) including weight distributions and trajectories. Figure 2a depicts the evolution of the weight distribution of afferents in the pattern and Fig. 2b of those not in the pattern. Figure 2c shows the weight trajectories of 50 afferents in the pattern and Fig. 2d 50 of those not in the pattern. Figure 2e depicts discharge latency as a function of simulation time (cf. Fig. 1).

3.2 Long Simulation

To study this drift in more detail longer simulations (3000 s) were ran (Fig. 2). Figure 2e depicts a neuron's discharges as a function of simulation time. The typical phases described earlier (Sect. 3.1 and Fig. 1) are clearly visible again, together with their time scales. The neuron's latency stabilizes at ≈ 18 ms however does fire as early as 13 ms. This drift backward through the pattern has a longer time scale than the the second period where STDP reduces the latency.

To elucidate the mechanism behind this backward drift, synaptic weight values were recorded every 2 s. Weight distributions and trajectories are fairly stable after 1000–2000 s and are bimodal (as found by [12, 17]). The period when STDP

learns the pattern can be seen in the weights distributions and trajectories (Figs. 2a and 2c—these plots include afferents that are within the pattern). The convergence of weights to the pattern is mostly complete before 250 s with a small fraction of weights still varying; in contrast however weights not in the pattern take relatively longer to converge: $\approx 1000\text{--}2000$ s.

These non-pattern weights are depressed over this longer period. This is the desired effect of STDP as they are not conveying the repeated pattern. However, the longer time scale has a possible side-effect: the latency drift described above. During the first phase the neuron fires non-specifically, thus STDP depresses most weights. This continues until pattern-driven fluctuations in the neuron’s membrane potential facilitates STDP learning [19]. These fluctuations are based on a background activity level consisting of the synaptic weights of *all* afferents at that time. Therefore, when STDP depresses non-pattern weights *after* the initial pattern learning has completed, this background level is consequently modified: it is reduced. This reduction in background activity increases the time the neuron takes to temporally sum incoming spikes to threshold. It is hypothesized that this increase in latency visible in Fig. 2e is due to the relatively longer time scale of weight depression.

3.3 Analysis of Drift

To test the hypothesized reliance of the output neuron on background activity, an additional (constant) electrical current was ‘injected’ into the output neuron after it had learnt a pattern. If a trained neuron has no significant reliance on the activity and responds *only* to a short succession of coincidences from a subset of afferents, this additional electrical current—if small relative to threshold—should not affect when (and if) the neuron fires. Figure 3 shows the results of such a simulation.

When the additional current is applied to the neuron at 2000 s there is a significant change in weights (Fig. 3a–d). Moreover, Fig. 3e shows a significant change in the neuron’s firing latency during this additional current: the latency is reduced. Interestingly the neuron’s latency then drifts again. The synapses are plastic for the entirety of the simulation; if the synapses are fixed at their current strengths from 2000 s onwards this drift does not occur and the neuron consistently fires at the initial reduced latency (results not shown herein).

An analysis of the weight distributions and trajectories from 2000 s helps to postulate as to this second drift: many afferents which convey the pattern change their efficacy significantly—either going from 0 to W_{\max} or vice-versa. Accompanying these changes is a reduction in weights for afferents not conveying the pattern—similar to the initial training phases. The neuron is adjusting to its new background activity.

To further clarify this reliance on background activity Fig. 4 depicts a neuron’s membrane potential. Figure 4a is just after the neuron has trained to the pattern, and Fig. 4b is later in the simulation after some drift has occurred. The background

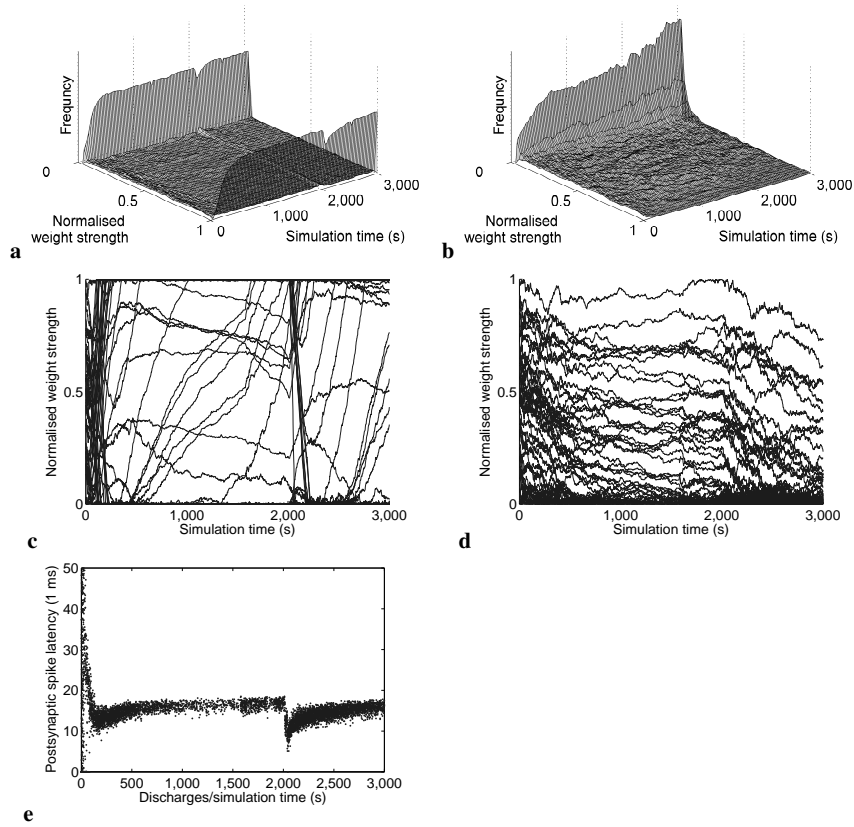


Fig. 3 Results of a long simulation (3000 s) with an additional injection of constant electrical current at 2000 s. As with Fig. 3, weight distributions (a and b) and weight trajectories (c and d) were recorded every 2 s. e clearly shows a decrease in latency at the current onset; the latency then increases and levels off.

activity is higher in Fig. 4b than Fig. 4a; this leads to a longer integration time to threshold, thus increasing the firing latency by ≈ 5 ms.




Fig. 4 A neuron’s membrane potential depicting a lower background activity and consequently a longer integration time. a reaches threshold with a latency of ≈ 10 ms, whereas b with ≈ 15 ms. Patterns started at 251.25 s and 952.35 s.

3.4 Impact of Noise

To this point simulations trained to the beginning of the pattern; however if noise was added to the neuron’s membrane potential, this is no longer the case. Figure 5 depicts two typical results with noise: in Fig. 5a the shortest latency is ≈ 15 ms, and in Fig. 5b it is ≈ 30 ms. It appears the noise interferes with the learning phase (Fig. 1); in fact, with this noise we found that a neuron’s learnt latency varied greatly.




Fig. 5 Two typical results when Gaussian noise ($\mu = 0.05$ and $\sigma = 0.013$) is added to the membrane potential. The neuron starts training to the beginning of the pattern but stops before reaching it.

4 Discussion

The results we present strengthen the findings by Masquelier et al. [12]: even though we used—among other things—a different STDP window and different learning rates, STDP was able to distinguish a repeated pattern from random activity, learn

it and fire consistently within it. However, we found that with long simulations the latency of a neuron's firing increased with respect to the beginning of the pattern. It is not clear if Masquelier et al. [12] ran their simulations for longer than 450 s.

To eliminate this drift as an artefact of the model, parameters were adjusted as follows:

- Different learning rates had no affect. As rates are assigned with $A_m = -A_p \times \left(\frac{\tau_p}{\tau_m}\right) \times 1.05$, $A_m \propto A_p$ so $\int_{\tau_m}^{\tau_p} f(\tau)$ was always < 0 thus different learning rates changed the convergence/learning time only.
- Different values of W_{\max} had no affect on the drift.

It would be advantageous for resources if synapses which are depressed—and determined not to carry information—were eliminated. However if this was to occur, a neuron may either not fire at all or fire with a much greater latency; thus the STDP-based learning schema used in this study may have less biological plausibility than previously postulated.

An offshoot of the findings presented in this study is a glimpse at the pattern encoding used by its neurons. A neuron encodes the coincidence of spikes from a subsection of the pattern's afferents, and when these spikes occur the neuron fires. However the temporal precision of a response depends on some background activity, which must be similar to that present during learning. Consequently, in a noisier framework—perhaps with greatly fluctuating afferent activity—a neuron may no longer respond to its learnt pattern. In addition we found that a relatively small amount of noise added to the neuron's membrane potential produced less stable learning; in fact the earliest trained latency appears to be highly un-stable.

References

1. Bastrikova, N., Gardner, G., Reece, J., Jeromin, A., Dudek, S.: Synapse elimination accompanies functional plasticity in hippocampal neurons. *Proceedings of the National Academy of Sciences* **105**(8), 3123–3127 (2008)
2. Bell, C., Han, V., Sugawara, Y.: Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* **387**(6630), 278 (1997)
3. Bi, G., Poo, M.: Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* **18**(24), 10,464–10,472 (1998)
4. Bi, G., Poo, M.: Distributed synaptic modification in neural networks induced by patterned stimulation. *Nature* **401**(6755), 792 (1999)
5. Bi, G., Poo, M.: Synaptic modification by correlated activity: Hebb's postulate revisited. *Annual Review Neuroscience* **24**, 129–66 (2001)
6. Gerstner, W., Kistler, W.: *Spiking Neuron Models*. Cambridge University Press (2002)
7. Guyonneau, R., VanRullen, R., Thorpe, S.: Neurons tune to the earliest spike through STDP. *Neural Computation* **17**(859-879) (2005)
8. Levy, W., Steward, O.: Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience* **8**(4), 791–7 (1983)
9. Linsker, R.: From basic network principles to neural architecture: Emergence of spatial-opponent cells. *Proceedings of The National Academy of Science* **83**(19), 7508–7512 (1986)

10. Mackay, D., Miller, K.: Analysis of Linsker's application of Hebbian rules to linear networks. *Network: Computation in Neural Systems* **1**(3), 257–297 (1990)
11. Markram, L.: Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* **275**(5297), 213 (1997)
12. Masquelier, T., Guyonneau, R., Thorpe, S.: Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains. *PloS one* **3**(1), e1377 (2008)
13. Masquelier, T., Guyonneau, R., Thorpe, S.: Competitive STDP-based spike pattern learning. *Neural Computation* **21**(5), 1259 (2009)
14. Masquelier, T., Hughes, E., Deco, G., Thorpe, S.: Oscillations, phase-of-firing coding, and spike timing dependent plasticity: An efficient learning scheme. *The Journal of Neuroscience* **29**(43), 13,484–13,493 (2009)
15. Masquelier, T., Thorpe, S.: Unsupervised learning of visual features through spike timing dependent plasticity. *Plos Computational Biology* **3**(2), 31 (2007)
16. Pratt, K., Dong, W., Aizenman, C.: Development and spike timing-dependent plasticity of recurrent excitation in the xenopus optic tectum. *Nat Neurosci* **11**(4), 467–475 (2008)
17. Rubin, J., Lee, D., Sompolinsky, H.: Equilibrium properties of temporally asymmetric hebbian plasticity. *Physical Review Letters* **86**(2), 264–7 (2001)
18. Senn, W., Fusi, S.: Learning only when necessary: better memories of correlated patterns in networks with bounded synapses. *Neural Computation* **17**, 2106–2138 (2005)
19. Song, S., Miller, K., Abbott, L.: Competitive Hebbian learning through spike-timing dependent synaptic plasticity. *Nature Neuroscience* **3**, 919–926 (2000)
20. Swindale, N.: A model for the formation of orientation columns. *Proc R Soc Lond* **215**(1199), 211–230 (1982)
21. Wittenberg, G., Wang, S.: Malleability of spike-timing-dependent plasticity at the CA3-CA1 synapse. *Journal of Neuroscience* **26**(24), 6610–6617 (2006)
22. Zhang, L., Tao, H., C. H.: A critical window for cooperation and competition among developing retinotectal synapses. *Nature* **395**(6697), 37 (1998)
23. Zhang, X.: Long-term potentiation at hippocampal perforant path-dentate astrocyte synapses. *Biochemical and Biophysical Research Communications* **383**(3), 326 (2009)