

The Trouble with Weight-Dependent STDP

Dominic Standage and Thomas Trappenberg

Abstract— We fit a weight-dependent STDP rule to the classic data of Bi and Poo (1998), showing that this rule leads to slow learning in a simulation with an integrate-and-fire neuron. The slowness of learning is explained by an inequality between the range of initial weights in the data and the largest relative potentiation. We show that slow learning can be overcome with an increased learning rate, but that this approach leads to rapid forgetting in the presence of realistic levels of background spiking. Our study demonstrates that weight-dependent STDP rules, commonly used in neural simulations, have biologically unrealistic consequences. We discuss the implications of this finding for several interpretations of weight-dependent plasticity and STDP more generally, and recommend directions for further research.

I. INTRODUCTION

Activity-dependent change in synaptic strength, or synaptic plasticity, is widely believed to provide the basis for learning and memory, as originally proposed by Hebb [1]. Countless physiological data reveal Hebbian plasticity, and similarly vast numbers of neural simulations show associative learning from activity-dependent rules. While early experiments showed potentiation [2] and depression [3] of synaptic strength as a function of pre-synaptic firing rates, more recent experimental methods reveal similar changes based on the timing of pre- and post-synaptic activity [4], [5], [6]. Experiments showing spike-time dependent plasticity (STDP) reveal that the direction (potentiation or depression) and magnitude of synaptic change depend on the order and latency of pre- and post-synaptic activity, within a time window on the order of tens of milliseconds [6], [7].

Weight-dependent STDP rules provide the additional constraint that changes in synaptic strength or *weight* are a function of the initial strength of a synapse. Such rules have been widely studied [8], [9], [10], [11], largely because they provide a cap on synaptic strength while leading to asymptotic weight distributions that compare favourably with experimental data [9] and that are believed to preserve the statistics of input activity [10]. As such, weight-dependent STDP rules are now commonly used in modelling studies.

While several experimental studies have shown weight-dependent plasticity [6], [12], [13], [14], only Bi and Poo's (1998) study [6] has done so under the STDP pairing protocol. In recent work [15] we fit a weight-dependent STDP rule to the data of Bi and Poo, showing how asymptotic weights depend on the correlation and rate of pre- and post-synaptic activity. Here, we use this rule in a neural

simulation, demonstrating that self-organisation of weights and associative learning are achieved, but that learning is extremely slow. While an increased learning rate accelerates the learning process, it also leads to rapid forgetting in the presence of biologically realistic background spiking.

Our study shows that the common formulation of weight-dependent STDP rules [8], [16], [17], [10] has biologically unrealistic consequences in rules fit to data. These consequences result from the assumption that weights may traverse the full range of initial values shown in weight-dependent studies (Bi and Poo's weight-dependent data are shown in Figure 1B). This assumption is inconsistent with Bi and Poo's data, and also with weight-dependent data resulting from other plasticity protocols [12], [14]. Given the common inclusion of weight-dependent terms in learning rules, this issue is an important one, and poses questions about experimental methods, data and their interpretation.

II. A WEIGHT-DEPENDENT STDP RULE FIT TO DATA

For simplicity, we assume that weight- and spike-time-dependencies are independent. Under this assumption, the general form of a weight-dependent STDP rule may be expressed as

$$\Delta w_{\{p,d\}} = k f_{\{p,d\}}(w) e^{-c_{\{p,d\}} \Delta t}, \quad (1)$$

where $\Delta w_{\{p,d\}}$ is the change in weight (p for potentiation, d for depression), $\Delta t = t_{post} - t_{pre}$ is the time between pre- and post-synaptic firing, c captures the timescale of the STDP window, and k is a learning rate.

A log-linear fit to Bi and Poo's weight-dependent STDP data, hypothesized by Bi and Poo and reproduced in Figure 1B, yields the following form of $f(w)$:

$$f_{\{p,d\}}(w) = (a_{\{p,d\}} - b_{\{p,d\}} \log w) w \theta(w),$$

where $\theta(w)$ is the Heaviside function, keeping weights positive, and the parameters a and b are different for potentiation and depression.

Bi and Poo controlled spike timing in their weight-dependent experiment by limiting their spike latencies Δt . We used the midpoints of these latencies in our weight-dependent fit ($\Delta t_p = 10, \Delta t_d = 17.5$) and assumed that for a given Δt , the largest absolute values for $\Delta w/w$ represent synapses with the smallest initial weights. We therefore assign an initial weight $w = 30\text{pA}$ to the largest STDP data (open symbols in Figure 1A) and only include these data in our spike-time fit for c . The resulting parameter values are $a_{\{p,d\}} = \{208, -54\}$, $b_{\{p,d\}} = \{26.4, 3.5\}$, $c_{\{p,d\}} = \{0.054, 0.042\}$. Because these parameters were fit to data expressed in percentage terms, and because we make the

Manuscript received January 31, 2007. This work was supported in part by the NSERC grant RGPIN 249885-03.

Dominic Standage and Thomas Trappenberg (corresponding author) are with the Faculty of Computer Science, Dalhousie University, 6050 University Avenue, Halifax, NS, B3H 1W5, Canada (phone: (902) 494-3087; fax: (902) 492-1517; email: {standage,t}@cs.dal.ca).

common assumption that each of the 60 pairings in Bi and Poo’s experiment contributed equally to synaptic change [9], we use a learning rate of $k = 1/6000$.

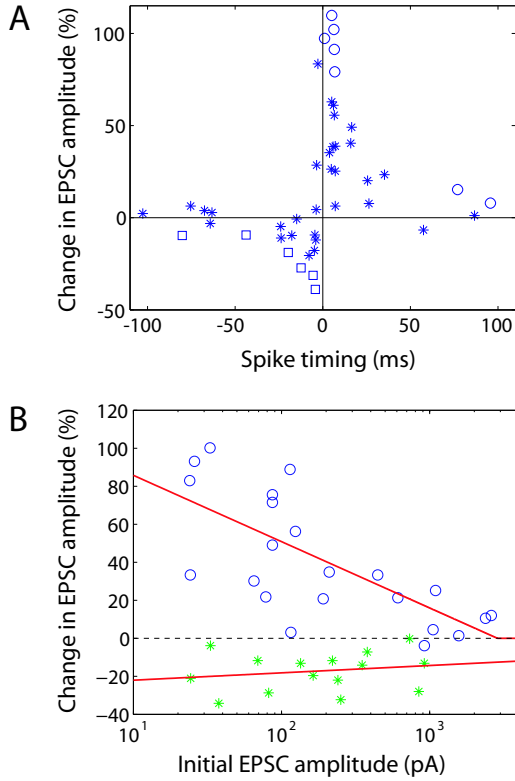


Fig. 1. (A) Reproduction of Bi and Poo’s (1998) spike-time dependent data. Circles and squares represent estimates of $w = 30$ pA for potentiation and depression respectively. Stars represent estimates of $w > 30$ pA. (B) Log-linear fit to Bi and Poo’s weight-dependent STDP data, imposing a maximum weight where the upper solid line meets the x-axis.

III. STDP AND SPIKE INTERACTIONS

STDP rules are fit to data from experiments in which pairings of pre- and post-synaptic spikes occur at low frequency, typically 0.5-5Hz. The long inter-spike intervals (ISI’s) of these spike trains effectively isolate spike pairings from the effects of previous and subsequent spikes. How to apply STDP rules to spike trains more realistic than temporally isolated pairings of pre- and post-synaptic spikes is the subject of great research interest [18], [19], [20], [17], [14], [21]. Here, we use a model we call *closest pair* interactions, where only the pair of pre-synaptic spikes closest to the surrounding post-synaptic spikes contributes to plasticity, as shown in Figure 2.

When exactly one pre-synaptic spike falls between two post-synaptic spikes, this model is equivalent to the nearest neighbour model suggested in [17]. These two interaction models are depicted in Figure 2. Because the analysis in [17] does not consider spike trains beyond this simplified case, we compare asymptotic weights calculated with their analytical method to asymptotic weights from simulations under closest pair interactions.




Fig. 2. Cartoons depict spike interaction models described in the text. It is unclear how the *nearest neighbour* model extends to the case where multiple pre-synaptic spikes fall between two post-synaptic spikes. We define the *closest pair* model for this case.

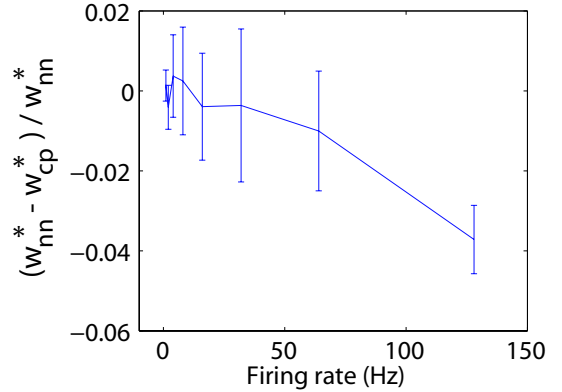


Fig. 3. The relative difference between mean weights w^* under *nearest neighbour* (analytic) and *closest pair* (numeric) spike interactions (see text). There is a very small (but significant) relative difference at high rate. Error bars show standard deviations.

To this end, we numerically approximate the means of equilibrium weight distributions (equilibrium weights w^*) for closest pair interactions, comparing them to analytic values for the nearest neighbour case. In the simulations, post-synaptic spikes were time-locked to Poisson-distributed pre-synaptic spikes at $\Delta t = 4$ ms. Weights were averaged over 5000 pairings following 5000 equilibrating pairings at spike rates $r = \{1, 2, 4, 8, 16, 32, 64, 128\}$ Hz. Analytic calculations for the nearest neighbour case with the same spiking statistics were calculated in [15]. The relative difference between numeric and analytic asymptotic weights for these spike interaction models is shown in Figure 3. There is a systematic difference that increases with increasing spike rate, but the relative difference is very small.

IV. SIMULATIONS WITH A NEURON

To test the effectiveness of Equation 1 for associative learning, we drive an integrate-and-fire (IF) node with two groups of synapses. Before learning, synchronous input from one group is strong enough to sporadically drive the node in combination with background activity from all synapses. The other group cannot drive the node. We test whether repeated synchronous activation of the high-strength group allows self-organisation of these weights, and whether synchronous activation of both groups (following self-organisation) is sufficient for the low-strength group to ‘piggyback’ the

high-strength group. If so, the low-strength group will have learned to drive the node by associativity.

A. Model and parameters

The membrane potential of the IF node is described by

$$\tau_m \frac{dv(t)}{dt} = V_{rest} - v(t) + G_e(V_e - v(t)) + G_i(V_i - v(t))$$

with membrane time constant $\tau_m = 20$ ms and reversal potentials $V_{rest} = -70$ mV, $V_e = 0$ mV and $V_i = -70$ mV where subscripts e and i refer to excitatory and inhibitory potentials respectively. When the membrane potential crosses threshold Θ , it is reset to -60 mV with an absolute refractory period of 2mS. A relative refractory period is implemented by increasing $\Theta = -54$ mV by $\gamma = 20$ mV when the neuron fires, after which γ decays exponentially with half width 10mS. Excitatory and inhibitory conductances G_e and G_i are described (as in [22]) by

$$\tau_e \frac{dG_e(t)}{dt} = -G_e(t) + \tau_e \sum_j g_j \delta(t - t_j^f)$$

and

$$\tau_i \frac{dG_i(t)}{dt} = -G_i(t) + \tau_i \sum_j g_j \delta(t - t_j^f),$$

where $\tau_e = \tau_i = 5$ ms, δ refers to the Dirac delta function and t_j^f is the time of firing of inputs mediated by conductance synapses g_j (excitatory for G_e and inhibitory for G_i).

EPSC's as conductances: There is no obvious way to translate Bi and Poo's EPSC measurements to synaptic conductances. To begin with, we do not know the amplitudes of corresponding EPSP's and so cannot determine conductances with Ohm's law. More problematic is the variation in these data ($\sim 25 - 2500$ pA) and the size of the larger EPSC's. Because our objective is to study the implications these data for STDP rules, we cannot simply ignore the larger currents. Rather than use current synapses (independent of v), we use a linear transformation, allowing a range of 30-3000pA to correspond to 10-150pS. A maximum conductance of 150pS is used in [22]. Inhibitory synapses are not subject to plasticity and are uniformly allocated conductances of 500pS (also from [22]).

The simulation was run with 1000 excitatory synapses and 200 inhibitory synapses. We initialised the high-strength synapses (trainers) to uniformly distributed random values between 800 and 1200pA. The low-strength synapses (piggybackers) were initialised to 50pA, providing sub-threshold input. All other excitatory synapses were given uniformly distributed random values between 500 and 900pA. Poisson background firing with mean 10Hz was mediated by all synapses throughout the simulation. Synchronous inputs mediated by the trainers and piggybackers were given a uniformly distributed random jitter of 1-5ms, superimposed on the Poisson activity. Under our transformation to conductances, these weights are too weak for background activity to drive the node, but are able to elicit occasional firing when the trainers fire synchronously on top of background activity. Axonal delays AD were assigned uniformly distributed

random values (to the nearest integer) from $AD = 4 - 17$ ms for excitatory synapses and $AD = 3 - 6$ ms for inhibitory synapses. Equation 1 was applied under the closest pair spike interaction model.

B. Associative learning under Equation 1

The IF node was first driven by Poisson activity alone, establishing a baseline membrane potential. This activity was followed by a period of self-organisation by the trainer synapses, driven by periodic synchronous activity at 10Hz. Periodic 10Hz activity provides a sufficiently long ISI to isolate neural activity in each cycle from the effects of the previous cycle. Following this activity, the piggybackers were simultaneously activated with the trainers for the associative task, again at 10Hz. The node was subsequently driven by Poisson activity for comparison with the pre-training response to noise.

Figure 4A shows the trajectories of three trainer synapses and three piggyback synapses over the course of the simulation. The lower solid curve depicts a fast trainer synapse with a small initial weight [$AD = 4$ ms, $w(t = 0) = 800$ pA], the upper solid curve shows a slow trainer with a large initial weight [$AD = 17$ ms, $w(t = 0) = 1200$ pA] and the middle curve shows an intermediate case [$AD = 9$ ms, $w(t = 0) = 1000$ pA]. Early during self-organisation, the trainer group is not strong enough to regularly fire the node, but in combination with background activity, input synchrony generates enough post-synaptic firing for weights to increase until the group regularly drives the node.

Post-synaptic firing during early and late self-organisation is depicted by the upper and lower insets of the figure respectively. All three weights initially increase because they all contribute to post-synaptic activity. As faster weights gain in strength, EPSC onset of slower weights occurs after post-synaptic firing, subjecting them to LTD. As expected, this effect is first exhibited by the slow synapse and later the by mid-latency synapse. The former is effectively inactivated, with final value $w(t = 4 \cdot 10^5 \text{ms}) = 2.4$ pA. The latter assumes an intermediate value [$w(t = 4 \cdot 10^5 \text{ms}) = 1009$ pA] determined by an ongoing cycle of 'promotion' and 'relegation' to and from the effective input group. This cycle is governed by the interaction of background firing and post-synaptic activity, and the overlap between the jitter in synchronous input and axonal delay. The fast synapse reaches an asymptotic value close to the log rule cut-off.

The dashed curves in Figure 4A show the trajectories of three piggyback synapses. The arrow at the top of the figure shows onset of associative learning. At this point, these small synapses are climbing toward equilibrium weights determined by their background firing and the now-periodic post-synaptic firing generated by the trainer group. With the onset of associative activity, the fast piggyback synapse follows the fast trainer synapse toward an asymptotic weight, the mid-latency synapse assumes a value similar to that of the mid-latency trainer synapse, and the slow piggyback synapse is effectively inactivated.




Fig. 4. (A) Self-organisation and associativity under Equation 1 with an integrate-and-fire node. Curves show weights over time for three synapses with large initial weights (trainers, see text) and three synapses with small initial weights (piggybackers, see text). Curves to the left of the arrows show weights as the trainers self-organise during periodic synchronous activity (10Hz, uniform-distributed random jitter 1-5ms) on top of Poisson background activity (mean 10Hz). Arrows depict the onset of associative learning when the piggyback synapses fire synchronously with the trainers, also at 10Hz. A range of 30-3000pA is linearly compressed to 10-150pS for use by the node (see text). The upper solid curves show the weights of three trainers with axonal delays of $AD = 17$, $AD = 9$ and $AD = 4$ ms respectively, top to bottom on the left of the figure. During self-organisation, all synapses increase in weight until the faster trainers drive the node without the slower ones, rendering slower synapses subject to depression. The same effect is seen for the piggyback synapses during associative learning. Self organisation and associativity are achieved, but learning is very slow with learning rate $k = 1/6000$. The Fast ($AD = 4$ ms) piggyback synapse is still converging to the log rule cut-off after 200,000ms (2000 pairings). (B) An increased learning rate ($k = 1/200$) accelerates learning, but background noise soon causes forgetting.

Learning is extremely slow under Equation 1: Self-organisation and associativity have been achieved, as synchronous input by either group consistently drives the node. Furthermore, Equation 1 has selected the fastest weights within the groups and shown the ability to recruit and discard slower weights as and when needed, as shown in [22]. The mean membrane potential \bar{v} in response to background noise has increased from $\bar{v} = -62.7$ mV (standard deviation $sd = 1.1$) before learning to $\bar{v} = -60.7$ mV ($sd = 1.2$) after learning, rendering the node more responsive to the timing of its inputs.

Unfortunately, the figure also shows that learning is extremely slow. The fast synapses are still experiencing net potentiation after 400,000ms. While 400s may not be considered slow on a behavioural timescale, our simulations are based on *in vitro* data. This interval accounts for around 3000 pairings for the trainers and 2000 for the piggybackers, approximately 50 times the number of pairings in Bi and Poo's experiments. Relating plasticity data to behavioural learning is fraught with difficulty [23] but the frequency of periodic, synchronous input in our simulations is within the theta range, consistent with the hypothesis that the theta rhythm provides an optimal pairing frequency for LTP during learning tasks [24].

C. Fast learning implies fast forgetting under weight-dependent STDP

To combat the slowness of learning shown in Figure 4A we raise k from $k = 1/6000$ to $k = 1/200$. With this learning rate, weights begin to asymptote after around 60 pre- and post-synaptic pairings, the same number of pairings used in Bi and Poo's experiments [6]. It is unclear what a learning rate $k \neq 1/6000$ represents, but an accelerated learning rate is required due to the large variation in synaptic strength in Bi and Poo's weight-dependent data (Figure 1), allowing weights to make the transition from low to high values in a reasonable period of time.

Figure 4B shows the weight of a synapse mediating the same pre-synaptic spike statistics as the trainers did in the learning task, but over a much shorter timespan (10,000ms). The post-synaptic response is time-locked to $\Delta t = 2$ ms after pre-synaptic activity, again superimposed on Poisson spiking at 10Hz. Learning is stopped at time $t = 10,000$ ms (top arrow), after which pre- and post-synaptic spike trains are given independent Poisson distributions at 10Hz. The synapse forgets the learned association within 10s.

V. DISCUSSION AND CONCLUSIONS

It is possible to parameterize weight-dependent STDP rules to accomplish various learning tasks, but our study shows these rules have unwanted consequences when fit to data. The main consequence is due to the range of initial weights in the only weight dependent STDP data available [6], shown in Figure 1B. Weight-dependent STDP rules implicitly assume that a synapse can span this entire range of values, suggesting changes in synaptic efficacy around 10,000%. No synapse in these experiments, however, changed in strength by more than around 100% (see Figure 1A,B).

This issue suggests several possibilities. One possibility concerns the standard computational interpretation of weight-dependent plasticity, in which there is an implicit assumption that weight-dependent data do not reflect saturated synapses. As an example, consider a synapse in Bi and Poo's weight-dependent experiment with an initial strength of 25pA. According to their data, this synapse would not have been potentiated by more than around 100% of control following their potentiation protocol. It would therefore be no larger than around 50pA after 60 pairings, much less than their largest initial weights. If all synapses have the same maximum, then for all but the largest synapses, an additional 60 pairings in Bi and Poo's experiment would have led to weight changes on the same order as the 60 pairings they performed. The consequences of this assumption are far reaching. For instance, STDP data typically suggest a gradient of weight changes in either direction as a function of the time between pre- and post-synaptic spikes [6], [20]. In the case of potentiation, the standard interpretation dictates that given enough pairings, all weights will reach the same maximum, regardless of the latency between pre- and post-synaptic spikes.

Alternatively, Bi and Poo's potentiation data may reflect saturated synapses. If so, to be consistent with the data, weight-dependent STDP rules must implement intrinsic synaptic maxima rather than a global maximum. For instance, a synapse with an initial weight of around 25pA should have an intrinsic limit of around 50pA, regardless of the size of the larger initial weights. We are unaware of any weight-dependent Hebbian rule that operates under this premise and the utility of such a rule is unclear. Certainly, large and small synapses would have different roles under this scheme.

These possibilities suggest that weight-dependent learning rules with a global maximum must be supported by weight-dependent data in which the largest relative change in weight is roughly equal to the relative difference between the largest and smallest initial weights. We may express this relationship as

$$\max\left(\frac{w_i^{fin} - w_i^{init}}{w_i^{init}}\right) \approx \frac{w_{max}^{init} - w_{min}^{init}}{w_{min}^{init}}, \quad (2)$$

where w_i^{init} and w_i^{fin} represent the weight of a synapse i before and after potentiation respectively, and w_{max}^{init} and w_{min}^{init} are the largest and smallest initial weights respectively. The data of Montgomery et al. (2001) [13], produced by a

pairing protocol with no spike-time-dependence, are roughly consistent with this equation for synapses with initial weights less than 50pA.

The weight-dependent data of Debanne et al. (1999) [12] are intriguing in this regard. These authors paired pre-synaptic spikes with post-synaptic bursts in organotypic hippocampal slice cultures. Their CA3-CA1 data are reminiscent of Bi and Poo's in that the relative difference in initial weights is much greater than the normalized weight change, suggesting (as above) that these synapses have widely varying intrinsic maxima or that they have not reached saturation. Their CA3-CA3 synapses, however, appear to show the opposite effect (the left side of Equation 2 is greater than the right). Their combined data fit roughly with Equation 2. The degree to which plasticity data from different protocols may be related under computational rules is an open question and is currently receiving considerable attention [25], [17].

We show in Section IV-B that a weight-dependent STDP rule fit to weight-dependent STDP data leads to slow learning. As shown in Section IV-C, speeding up learning with an increased learning rate leads to rapid forgetting in the presence of noise. Learning thresholds [26] and hysteretic weight dynamics [27] provide potential solutions to this problem, but our study focuses on the standard weight-dependent STDP formulation defined by Equation 1 and investigated in numerous studies [8], [9], [10], [11]. An unexplored solution to the problem may be grounded in pre- and post-synaptic mechanisms if we (quite happily) abandon the assumption that the effects of spike pairings sum linearly. Consider the insertion and removal of AMPA receptors (AMPA's) in the post-synaptic membrane, correlated with plasticity in hippocampal regions CA1 and CA3 [28]. It seems likely that potentiation should be harder to initiate as the number of AMPA's increases. Pre-synaptically, plasticity is expressed at least in part by changes in the probability of transmitter release [29]. Weight-dependence is implicit in such a mechanism, where synapses with a low probability of release have greater scope for potentiation, and synapses with a high probability of release have greater scope for depression. A convolution of pre- and post-synaptic plasticity mechanisms could protect learned associations if pre-synaptic plasticity were to precede post-synaptic plasticity in both directions. Under such a scheme, AMPA's may not be inserted in the post-synaptic density until a sufficiently high probability of release has been achieved, providing an initial resistance to post-synaptic learning. In reverse, AMPA's may not be removed from the post-synaptic density until the probability of release is sufficiently low, providing initial resistance to forgetting. Such a scheme requires further study.

Finally, it is possible that weight-dependent plasticity data reflect the statistics of *non-weight-dependent* plasticity in populations of synapses, as discussed by Debanne et al. (1999) [12]. This possibility suggests that tissue preparations and stimulation methods must be carefully considered when interpreting experimental data. Extracellular stimulation is commonly used to mimic pre-synaptic activity in STDP

paradigms [7], but this method may stimulate hundreds if not thousands of synapses [30]. While intracellular stimulation of pre-synaptic neurons (used in Bi and Poo's experiments) eliminates the potential involvement of multiple pre-synaptic cells, neurons in culture (also used in Bi and Poo's experiments) typically make multiple contacts with the post-synaptic cell [31]. It is safe to say that single synapses do not mediate currents on the order of Bi and Poo's larger data (~ 2500 pA). This issue overlaps with debates concerning pre- and post-synaptic contributions to plasticity [32], [33] and the degree to which plasticity is a graded phenomenon [34], [28], [30]. Varying the number of pairings in STDP experiments would directly address the latter. STDP protocols are clearly promising for the investigation of weight-dependent plasticity, but the potentially confounding effects of populations of synapses must be carefully addressed.

ACKNOWLEDGMENT

We thank Alan Fine and Stefan Krueger for helpful discussions.

REFERENCES

- [1] D. O. Hebb, *The Organisation of Behaviour*. John Wiley, New York, 1949.
- [2] T. V. P. Bliss and T. J. Lomo, "Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path," *Journal of Physiology*, vol. 232, pp. 331–356, 1973.
- [3] G. Lynch, T. Dunwiddie, and V. Gribkoff, "Heterosynaptic depression: a postsynaptic correlate of long-term potentiation," *Nature*, vol. 266, pp. 737–739, 1977.
- [4] W. B. Levy and O. Steward, "Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus," *Neuroscience*, vol. 8, pp. 791–797, 1983.
- [5] H. Markram, J. Lübke, M. Frotscher, and B. Sakmann, "Regulation of synaptic efficacy by coincidence of postsynaptic action potentials and epsps," *Science*, vol. 275, no. 5297, pp. 213–215, 1997.
- [6] G. Qiang Bi and M. Ming Poo, "Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type," *Journal of Neuroscience*, vol. 18, pp. 10464–10472, 1998.
- [7] G. M. Wittenberg and S. S.-H. Wang, "Malleability of spike-timing-dependent plasticity at the ca3ca1 synapse," *The Journal of Neuroscience*, vol. 26, no. 24, pp. 6610–6617, 2006.
- [8] W. M. Kistler and J. L. van Hemmen, "Modeling synaptic plasticity in conjunction with the timing of pre- and postsynaptic action potentials," *Neural Computation*, vol. 12, pp. 385–405, 2000.
- [9] M. C. W. van Rossum, G. Q. Bi, and G. G. Turrigiano, "Stable hebbian learning from spike timing-dependent plasticity," *Journal of Neuroscience*, vol. 20, no. 23, pp. 8812–8821, 2000.
- [10] R. Gutig, R. Aharonov, S. Rotter, and H. Sompolinsky, "Learning input correlations through nonlinear temporally asymmetric hebbian plasticity," *Journal of Neuroscience*, vol. 23, pp. 3697–3714, May 2003.
- [11] A. Burkitt, H. Meffin, and D. B. Grayden, "Spike-timing-dependent plasticity: The relationship to rate-based learning for models with weight dynamics determined by a stable fixed point," *Neural Computation*, vol. 16, pp. 885–940, 2004.
- [12] D. Debanne, B. H. Gähwiler, and S. M. Thompson, "Heterogeneity of synaptic plasticity at unitary ca3-ca1 and ca3-ca3 connections in rat hippocampal slice cultures," *Journal of Neuroscience*, vol. 19, no. 24, pp. 10664–10671, 1999.
- [13] J. M. Montgomery, P. Pavlidis, and D. V. Madison, "Pair recordings reveal all-silent synaptic connections and the postsynaptic expression of long-term potentiation," *Neuron*, vol. 29, pp. 691–701, 2001.
- [14] H.-X. Wang, R. C. Gerkin, D. W. Nauen, and G.-Q. Bi, "Coactivation and timing-dependent integration of synaptic potentiation and depression," *Nature Neuroscience*, vol. 8, no. 2, pp. 187–193, 2005.
- [15] D. Standage, S. Jalil, and T. Trappenberg, "Computational consequences of experimentally derived spike-time and weight dependent plasticity rules," *Biological Cybernetics*, 2007, in press.
- [16] A. Kepecs, M. C. van Rossum, S. Song, and J. Tegner, "Spike-timing-dependent plasticity: common themes and divergent vistas," *Biological Cybernetics*, vol. 87, pp. 446–458, 2002.
- [17] E. M. Izhikevich and N. S. Desai, "Relating stdp to bcm," *Neural Computation*, vol. 15, no. 7, pp. 1511–1523, 2003.
- [18] P. J. Sjöström, G. G. Turrigiano, and S. B. Nelson, "Rate, timing, and cooperativity jointly determine cortical synaptic plasticity," *Neuron*, vol. 32, pp. 1149–1164, 2001.
- [19] G.-Q. Bi, "Spatiotemporal specificity of synaptic plasticity: cellular rules and mechanisms," *Biological Cybernetics*, vol. 87, pp. 319–332, 2002.
- [20] R. C. Froemke and Y. Dan, "Spike-timing-dependent synaptic modification induced by natural spike trains," *Nature*, vol. 416, no. 6879, pp. 433–438, 2002.
- [21] R. C. Froemke, I. A. Tsay, M. Raad, J. D. Long, and Y. Dan, "Contribution of individual spikes in burst-induced long-term synaptic modification," *Journal of Neurophysiology*, vol. 95, pp. 1620–1629, 2006.
- [22] S. Song, K. D. Miller, and L. F. Abbott, "Competitive hebbian learning through spike-timing-dependent synaptic plasticity," *Nature Neuroscience*, vol. 3, no. 9, pp. 919–926, 2000.
- [23] S. J. Martin, P. Grimwood, and R. G. M. Morris, "Synaptic plasticity and memory: an evaluation of the hypothesis," *Annual Review of Neuroscience*, vol. 23, pp. 649–711, 2000.
- [24] G. Buzsáki, "Theta oscillations in the hippocampus," *Neuron*, vol. 33, pp. 325–340, 2002.
- [25] H. Z. Shouval, M. F. Bear, and L. N. Cooper, "A unified model of nmda receptor-dependent bidirectional synaptic plasticity," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 99, no. 16, pp. 10831–10836, 2002.
- [26] W. Senn, "Beyond spike timing: the role of nonlinear plasticity and unreliable synapses," *Biological Cybernetics*, vol. 87, pp. 344–355, 2002.
- [27] M. V. Kiselev, "Statistical approach to unsupervised recognition of spatio-temporal patterns by spiking neurons," in *The International Joint Conference on Neural Networks 2003*, 2003, pp. 2843–2847.
- [28] J. M. Montgomery and D. V. Madison, "Discrete synaptic states define a major mechanism of synapse plasticity," *Trends in Neurosciences*, vol. 27, no. 12, pp. 744–750, 2004.
- [29] C. F. Stevens and Y. Wang, "Changes in reliability of synaptic function as a mechanism for plasticity," *Nature*, vol. 371, pp. 704–707, 1994.
- [30] D. H. O'Connor, G. M. Wittenberg, and S. S.-H. Wang, "Dissection of bidirectional synaptic plasticity into saturable unidirectional processes," *Journal of Neurophysiology*, vol. 94, pp. 1565–1573, 2005.
- [31] D. Debanne, B. H. Gähwiler, and S. M. Thompson, "Cooperative interactions in the induction of long-term potentiation and depression of synaptic excitation between hippocampal ca3-ca1 cell pairs in vitro," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 93, pp. 11225–11230, 1996.
- [32] N. J. Emptage, C. A. Reid, A. Fine, and T. V. Bliss, "Optical quantal analysis reveals a presynaptic component of ltp at hippocampal schaffer-associational synapses," *Neuron*, vol. 38, pp. 797–804, 2003.
- [33] R. A. Nicoll, S. Tomita, and D. S. Bredt, "Auxiliary subunits assist ampa-type glutamate receptors," *Science*, vol. 311, pp. 1253–1256, 2006.
- [34] C. C. Petersen, R. C. Malenka, R. A. Nicoll, and J. J. Hopfield, "All-or-none potentiation at ca3-ca1 synapses," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 95, pp. 4732–4737, 1998.