

Investigating STDP and LTP in a spiking neural network

Daniel Bush*, Andrew Philippides, Phil Husbands and Michael O'Shea

Centre for Computational Neuroscience and Robotics
University of Sussex,
Brighton, BN1 9QG

{daniel.bush, andrewop, philh, m.o-shea}@sussex.ac.uk
www.cogs.susx.ac.uk/ccnr

Abstract. The idea that synaptic plasticity holds the key to the neural basis of learning and memory is now widely accepted in neuroscience. The precise mechanism of changes in synaptic strength has, however, remained elusive. Neurobiological research has led to the postulation of many models of plasticity, and among the most contemporary are spike-timing dependent plasticity (STDP) and long-term potentiation (LTP). The STDP model is based on the observation of single, distinct pairs of pre- and post- synaptic spikes, but it is less clear how it evolves dynamically under the input of long trains of spikes, which characterise normal brain activity. This research explores the emergent properties of a spiking artificial neural network which incorporates both STDP and LTP. Previous findings are replicated in most instances, and some interesting additional observations are made. These highlight the profound influence which initial conditions and synaptic input have on the evolution of synaptic weights.

1. Introduction

The ability of the brain to translate ephemeral experience into enduring memories has long been attributed by neuroscientists to activity-dependent changes in synaptic efficacy. One of the first to suggest a mechanism that could govern this plasticity was Donald Hebb, who hypothesised that ‘when an axon of cell A is near enough to excite a cell B, and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place ... such that A’s efficiency as one of the cells firing B, is increased’ (Hebb, 1949). This concept of ‘Hebbian’ learning has become a mainstay of neural theories of memory, but more precise rules of synaptic change have been difficult to elucidate.

It has become clear, however, that there are certain features which are crucial to a successful model of plasticity (Roberts and Bell, 2002 ; Song, Miller and Abbott, 2000 ; van Rossum, Bi and Turrigiano, 2000). It must generate a stable distribution of synaptic weights, and stimulate competition between inputs to a neuron, in order to account for the processes of activity-dependent development and forgetting, and to maximize the capacity for information storage (Miller, 1996). Pure Hebbian learning cannot achieve this, not least because it fails to make any mention of synaptic weakening processes, but also because those inputs which correlate with post-synaptic firing are repeatedly strengthened, thus growing to infinitely high values. This creates an inherently unstable, bimodal distribution of synaptic weights. Earlier plasticity models have had to resort to a variety of means in order to solve this problem. Often these promoted competition through the use of global signalling mechanisms, such as limiting the sum of strengths of pre-synaptic inputs to a cell, but the biophysical realism of such protocols can be questioned. The exact nature of the additional constraints used can also strongly influence the behaviour of the model (Miller and McKay, 1994).

In considering the neural basis of memory, it is long-lasting alterations in synaptic strength that are of most interest. Experimental evidence for such changes was first found in the hippocampus – a region of the brain long identified with learning – when it was shown that repeated activation of excitatory synapses by high frequency spike trains caused an increase in synaptic strength which lasted for hours, or even days (Lomo and Bliss, 1973). This phenomena - known as long-term potentiation (LTP) - has since been the subject of a great deal of investigation, because it exhibits several features which make it an attractive candidate as a neural learning mechanism (see Malenka and Nicol, 1999, for a review). It is synapse specific, vastly increasing the potential storage capacity of individual neurons. It is also associative, in that the repeated stimulation of one set of synapses can simultaneously facilitate LTP at adjacent sets of synapses. This has often been viewed as analogous to the process of classical conditioning.

The wealth of research into LTP has helped to inform and inspire new plasticity models which are more easily reconcilable with the tenets outlined earlier. The 'BCM' model, named after its creators (Bienenstock, Cooper and Munro, 1982) and based on their consideration of input selectivity in the visual cortex, is a good example. It is Hebbian, but achieves stability through the existence of a 'threshold' firing rate, a crossover point between depression and potentiation which is itself slowly modulated by post-synaptic activity. This makes the strengthening of a synapse more likely when average activity is low, and vice versa, thus generating competition between inputs.

Another contemporary plasticity model, based on the more straightforward empirical observation of distinct pairs of pre- and post- synaptic action potentials (Roberts and Bell, 2002 ; Bi and Poo, 1998), has also generated a great deal of interest. It is known as spike timing dependent plasticity (STDP), because it dictates that the direction and degree of changes in synaptic efficacy are determined by the relative timing of pre- and post- synaptic spiking. Only pre-synaptic spikes which provoke post-synaptic firing within a short temporal window potentiate a synapse, while those which arrive after post-synaptic firing cause depression. Those inputs with shorter latencies or strong mutual correlations are thus favoured, at the expense of others.

The most pertinent feature of STDP is that it implicitly generates competition between synapses, and experiments with artificial neural networks (ANNs) have shown that this precipitates inherently stable weight distributions. The shape of the resulting distribution is dependent on the exact nature of the STDP implementation, and the values of parameters used. Some researchers, for example, include the experimental observation that stronger synapses seem to undergo relatively less potentiation than weaker synapses, or an activity dependent scaling mechanism such as that outlined by the BCM model (van Rossum, Bi and Turrigiano, 2000). These features help to generate a weight distribution that more closely resembles the stable, unimodal, and positively skewed distribution found *in vivo* (see *fig 1*). Their omission tends to produce a bimodal distribution (Song, Miller and Abbott, 2000; Iglesias et al. 2005) more similar to that produced by pure Hebbian learning, but stabilised by innate competition and the inclusion of hard limits on the maximum achievable strength of a synapse.

The analysis of STDP is based on isolated pairs of pre- and post- synaptic action potentials, while observations of LTP are mediated by the application of prolonged spike trains more characteristic of normal brain activity. It is not clear how the STDP model causes synaptic weights to develop with such input, which involves many possible spike pairings. We can presume that both forms of plasticity arise from the same underlying biophysical mechanisms, and some recent work has attempted to reconcile both models within a single theoretical framework (Izhikevich and Desai, 2003). By making a few biologically plausible assumptions, this research has demonstrated that the parameters of STDP can be linked directly with the sliding threshold of the BCM model.

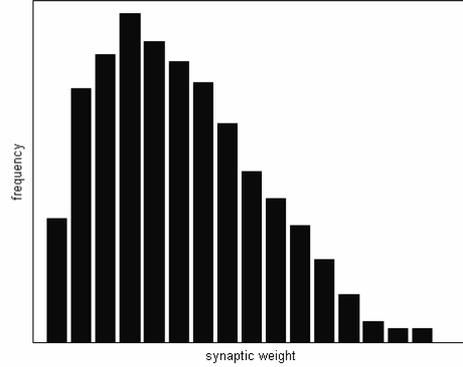


Fig 1. – Synaptic weight distribution found in vivo, taken from Bekkers et al., 1990

This paper explores the emergent properties of an artificial neural network which implements spike timing dependent plasticity. The form of STDP used is compatible with the BCM model of long-term potentiation, and thus the value of the threshold firing rate can be directly manipulated. The effects this has on synaptic weight distributions and dynamics are examined. Size-dependent potentiation is also introduced into the model, and results obtained from the input of random uncorrelated or partially correlated spike trains are compared with those generated by the performance of two simple, embodied, sensorimotor tasks. The latter will have temporal patterns that are perhaps more representative of firing regimes found *in vivo*, and which STDP has previously been shown to make use of (Izhikevich, Gally and Edelman, 2004).

2 Methods

2.1 Neural Controller

In the majority of tests, the neural network consists of 20 neurons, which are divided into 9 sensory, 9 intermediate and 2 motor neurons. During the phototaxis task, however, the network has only 2 sensory neurons, and thus a total of 13 neurons. The network is realistic of the mammalian cortex in that these are 80% excitatory and 20% inhibitory, and that each has a randomly chosen axonal delay in the range [1ms, 20ms]. Each neuron has 5 randomly assigned post-synaptic connections. Motor neurons have no post-synaptic connections, and sensory neurons have no pre-synaptic connections.

The neurons operate using the Izhikevich (2004) spiking model, which dynamically calculates the membrane potential (v) and a membrane recovery variable (u), based on the values of four dimensionless constants (a, b, c and d) and a dimensionless applied current (I), according to the equations below.

$$\begin{aligned}
 v' &= 0.04v^2 + 5v + 140 - u + I \\
 u' &= a(bv - u) \\
 \text{if } v \geq +30 \text{ mV then } &\begin{cases} v \leftarrow c \\ u \leftarrow u + d \end{cases}
 \end{aligned} \tag{1}$$

This model was chosen for two main reasons. Firstly, it uses very few floating point operations, and so is computationally advantageous. Secondly, it can exhibit firing patterns of all known types of cortical neurons, by variation of the parameters $a - d$.

The values used for a standard excitatory neuron are $[0.02, 0.2, -65, 6]$ respectively, and those for an inhibitory neuron are $[0.02, 0.25, -65, 2]$.

In order to introduce neural noise into the system, one neuron is selected at random each time step, and a small current applied to it. A value of $I=10$ was used in most tests, although this was varied to assess the effects of neural noise. When distributed randomly over 20 neurons, an applied current of this size produces a spiking rate of approximately 3Hz per neuron.

2.2 STDP

Mathematically, with $s = t_{post} - t_{pre}$ being the time difference between pre- and post-synaptic spiking, the change in the weight of a synapse (Δw) due to spike timing dependent plasticity can be expressed as:-

$$\Delta w = \begin{cases} A^+ \exp(-s / \tau^+) & \text{if } s > 0 \\ A^- \exp(s / \tau^-) & \text{if } s < 0 \end{cases} \quad (2)$$

The method of implementing this plasticity is outlined by Song et al (2000) and Di Paolo (2003). Two recording functions (P_+ and P_-) are kept for each synapse. These values decay exponentially according to the time constants of potentiation and depression, except when pre-synaptic spikes arrive or post-synaptic spikes are fired, in which case the values are reset to A_+ or A_- respectively. This means that only those spikes which are temporally adjacent affect the degree of synaptic weight change, and hence this is known as the ‘nearest neighbour’ model of STDP. Research has shown that this implementation allows the reconciliation of the BCM model with STDP (Izhikevich and Desai, 2003). It also outlines a formula for the calculation of the threshold firing rate, which is given by eqn. 3 below. The expressions $A_+ > |A_-|$ and $|A_- \tau_-| > |A_+ \tau_+|$ must be satisfied during experiments, to ensure that the threshold has a positive value at all times.

$$v = - \frac{A_+ / \tau_- + A_- / \tau_+}{A_+ + A_-} \quad (3)$$

Previous research (Bi and Poo, 1998) has shown that an inverse exponential relationship between the level of potentiation and initial synaptic weight may also exist *in vivo*. The modified formula which governs increases in synaptic weight when such ‘size dependent potentiation’ is examined is given in eqn. 4. It should be noted that there is no evidence for any such size-dependent effects in synaptic weakening.

$$w_{ij}(t) = w_{ij}(t) + P_+ e^{-kw_{ij}} \quad (4)$$

2.3 Tasks

The network was first examined using uncorrelated Poissonian spike trains of varying frequencies as input. In later experiments, correlated spike trains and two simple robotics tasks were used to assess how temporal patterns and more widely varying spike frequencies may affect the behaviour of the network. The tasks chosen were a simple phototaxis exercise similar to that used by Di Paolo (2003), and a falling block task which has been employed previously by Goldenberg et al (2004).

The correlated input was generated by creating a set number of Poissonian spike trains of a certain frequency, and distributing these amongst the 9 sensory neurons. Each time step, the spike trains were re-distributed amongst the inputs. The number of trains that exist thus determine the ‘strength’ of the correlation between inputs.

In the ‘falling block’ task, an agent of radius 15 moves horizontally in an arena which is 400 units wide. The agent has 9 sensory neurons with a range of 205 units, which are distributed evenly over a visual angle of $\pi/6$. These sensory neurons each have a randomly determined bias in the range [0.6:1.0] which is used to scale an applied current, relative to the distance of any object in their direct line of vision.

Two blocks of radius 13 fall from a height of 198 at randomly assigned angles and from randomly assigned horizontal start positions, constrained only by the criteria that it must be possible for the agent to catch them both. The first object has a random velocity in the range [0.03:0.04] and the second object in the range [0.01:0.02]. The agent’s horizontal velocity is determined by the sum of the two opposing motors outputs, its maximum velocity being set at 0.05 units/ms. The two motor neurons are leaky integrators, operating according to eqn 5 below, where t^o is the time at which a spike was last received. Each has a randomly assigned gain in the range [0.01:0.05] and a decay constant (τ) in the range [20ms : 40ms].

$$v = ve^{-(t-t^o)/\tau} \quad (5)$$

In the phototaxis task, an agent of radius 2 is placed at a random angle of orientation and randomly determined distance in the range [60 : 80] from a light source, in an arena of unlimited size. The agent has two sensory neurons, which are connected to light sensors separated by an angle of $2\pi/3$ on the agents body, plus or minus a random displacement of $\pi/36$. These light sensors have an angle of acceptance of π , and a randomly assigned bias in the range [10:50], which is used to scale the intensity of any light into an applied current. The intensity of the light source is assigned randomly in the range [3000:5000]. Two motors are placed diametrically opposite on the agents body, and driven in a forwards-only direction by the two motor neurons. The speed of each motor is limited at 0.5 units/ms, and thus this is also the maximum achievable forwards velocity of the agent. As before, the motor neurons are leaky integrators with gains in the range [0.01:0.5] and decay constants in the range [40ms:100ms]. It is important to note that the capacity of the network to learn how to perform this task is not being tested in this paper. The embodiment is needed only to provide realistic sensorimotor input, which has correlated temporal properties that are considered important in assessing the properties of the plasticity model and network.

2.4 Stability

After each 100ms of experimental time, a histogram of synaptic weights is generated. If the values in each bin (which are of size 1) do not vary by more than ± 1 for 10 of the 100ms steps (i.e. 1 second), then the network is considered to have achieved a stable synaptic weight distribution. In order to test that this criteria was adequate, 30 tests were performed, with random initial conditions and parameter values, and network operation was continued for 100 seconds of simulated time after stability was flagged. In all cases, no further discernible change in the synaptic weight distribution occurred. In each experiment, 30 random incarnations of the neural network are created, and each one is run twice, in each case until stability is achieved. Thus, the results presented in this paper are a conglomeration of 60 individual tests, or a total of 5400 synaptic weights (3300 in the case of the phototaxis task).

3 Results and Discussion

3.1 Manipulation of threshold firing rate

Figure 2 represents a typical synaptic weight distribution generated when the network

was operated with purely uncorrelated input at a rate of 30Hz , and the results replicate previous research findings (Song et al., 2000). The values of STDP parameters used in this case correspond to a threshold firing rate of approximately $\nu=17\text{Hz}$. The effects of moving the threshold firing rate (by varying any of the four main STDP parameters) are intuitive, and demonstrated by figures 3 ($\nu=350\text{Hz}$) and 4 ($\nu=6.25\text{Hz}$). A higher threshold for long-term potentiation allows fewer synapses to reach the maximum possible strength, and a lower threshold has the reverse effect.

However, results suggest that the relationship between weight distribution and STDP parameters is dictated by more complex factors than simply the position of the BCM threshold. Figure 5 shows a weight distribution for an identical threshold firing rate as 2 ($\nu=17\text{Hz}$), but with different STDP values. The number of synapses which have been potentiated to saturation are fewer, and those which have been persistently depressed are larger, in frequency. The value of $|A_+ \tau_+|$ is identical in both cases, but the longer temporal window for potentiation that existed in fig. 5 clearly had a lower overall strengthening effect on weight values, compared with the higher degree of synaptic strengthening per spike which was present in the results for 2. Further investigation demonstrated that the ratio of $A_+ : A_-$ is particularly important in determining the shape of the stable weight distribution. Results generated with identical values of this ratio are consistently very similar, more so than results with equal values of the modification threshold.

3.2 Firing rates

It is clear that the key to a good plasticity model, and one of the reasons why STDP is so highly regarded, is that it regulates network output in the face of wide fluctuations in input. In previous research (Song, Miller and Abbott, 2000) an increase in input firing rate has been observed to cause a decrease in the number of synapses saturating at the uppermost weight values, a finding that was replicated in these experiments. One may expect that fewer strong synapses would correlate with lower post-synaptic activity, but previous work has shown that the STDP model actually exhibits a ‘damping’ effect – increasing the input firing rate precipitating a much smaller increase in post-synaptic firing rate. An analysis of firing rates in the intermediate, excitatory neurons during this investigation, however, led to a finding which, at least to some extent, contradicts this previous research (Song, Miller and Abbott, 2000). Figure 6 illustrates the correlation between input and intermediate firing rates for four different sets of STDP parameters.

The data demonstrates that, if any relationship exists between these two variables, then it is very complex, and could depend on many factors. As the figure shows, in some cases there seems to be an inverse relationship between the two firing rates, while in others previous research has been replicated and a simple damping effect can be seen. Once again, it seems that the ratio of $A_+ : A_-$ has a pronounced effect on post-synaptic firing rates. In the data presented here, similar values of this ratio do seem to produce similar relationships between input and intermediate firing rates. Further investigation will be required to elucidate the nature of this relationship.

3.3 Varying network input

It is useful to make a comparison between the weight distributions arising from uncorrelated input, correlated input, and those generated by input from closed-loop sensorimotor tasks. Figures 7 to 9 illustrate these results – in each case, identical parameter values to figure 2 were used, but in each case the stable synaptic weight distributions are markedly different.

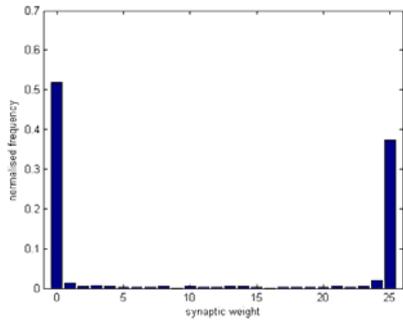


Fig. 2 –
 $A_+=0.16$; $A_-=-0.1$; $\tau_+ = 20\text{ms}$; $\tau_- = 40\text{ms}$

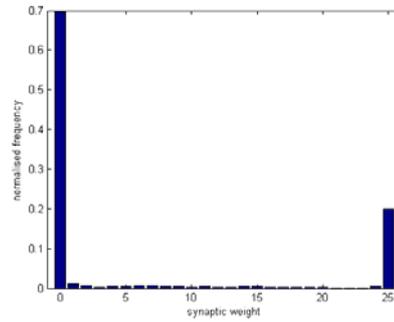


Fig. 3 -
 $A_+=0.12$; $A_-=-0.$; $\tau_+ = 10\text{ms}$; $\tau_- = 40\text{ms}$

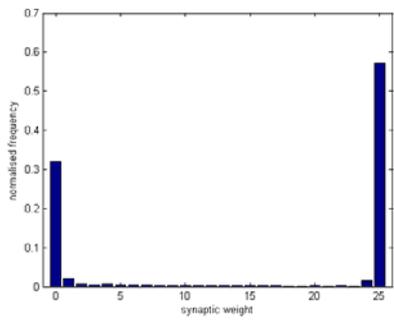


Fig. 4 –
 $A_+=0.18$; $A_-=-0.$; $\tau_+ = 20\text{ms}$; $\tau_- = 40\text{ms}$

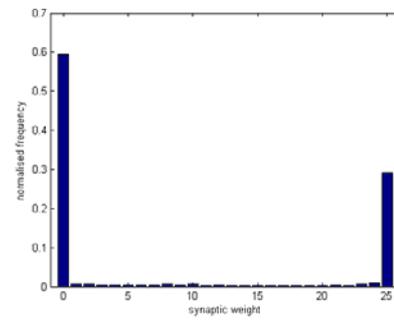


Fig. 5 -
 $A_+=0.12$; $A_-=-0.1$; $\tau_+ = 30\text{ms}$; $\tau_- = 40\text{ms}$

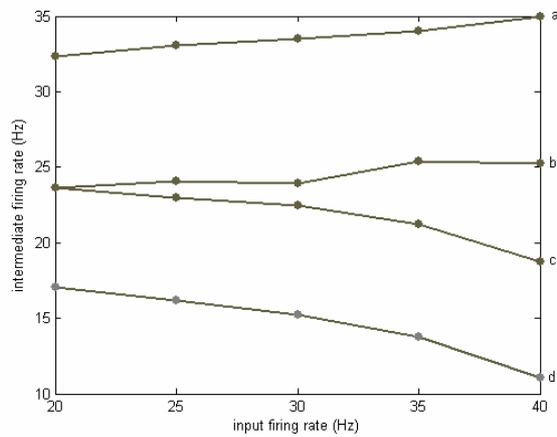


Fig. 6 – The relationship between input and intermediate firing rates

- a -** $A_+=0.2$; $A_-=-0.1$; $\tau_+ = 10\text{ms}$; $\tau_- = 40\text{ms}$
- b -** $A_+=0.2$; $A_-=-0.1$; $\tau_+ = 20\text{ms}$; $\tau_- = 40\text{ms}$
- c -** $A_+=0.2$; $A_-=-0.15$; $\tau_+ = 20\text{ms}$; $\tau_- = 40\text{ms}$
- d -** $A_+=0.2$; $A_-=-0.15$; $\tau_+ = 10\text{ms}$; $\tau_- = 40\text{ms}$

It seems that input in the sensorimotor tasks caused more synapses to adopt intermediate weight values, rather than be pushed to the bounds. In the phototaxis task there are also a much greater frequency of synapses at maximum strength and fewer at zero weight, in contrast to the falling block task. Other results showed that distributions generated by input from the robotics tasks are generally much more consistent in shape. The effects of manipulating the threshold rate can still be seen, but rather than simply altering the size of the bimodal peaks (as seen in figures 2 – 4), it is the frequency and distribution of the intermediate strength synapses that are most affected. Correlated input also produced a markedly different weight distribution, with much more similarly sized modal peaks and a more uniform intermediate distribution. As with uncorrelated input, the intermediate weight values are more sparsely populated.

The variations in the size and shape of the distributions seen can be loosely explained by the slight differences in the nature of the input presented to the neural network. However, the main issue is that these discrepancies support the intuitive hypothesis that the nature of input to an ANN has a pronounced effect on the evolution of synaptic weights in that network. Much of the previous research in this area has made exclusive use of uncorrelated input, but results found here show that care must be taken in generalising from these findings. The development of a network directed by any plasticity model is at least partially defined by the nature of the input it receives – and there are gross differences between uncorrelated and more realistic sensorimotor input.

3.4 Size-dependent potentiation

The introduction of size-dependent potentiation into the plasticity model also has a pronounced effect on synaptic weight distributions. Figure 10 (which was generated using a value of $k=50$ and identical parameter values to fig. 1) illustrates this, and more closely resembles results found *in vivo*. The peak at $w=0$ has been omitted, as these ‘silent’ synapses are not considered (and cannot be detected) when biological appraisals of weight distributions are made. It is interesting to note, however, that the frequency of synapses found at the lower bound was generally consistent between experiments with and without size-dependent potentiation. This implies that the larger number of synapses adopting intermediate weights was simply a product of the fact that fewer synapses were able to saturate to the upper bounds. By tuning the value of k appropriately, the hard limit on synaptic weights can be completely removed, giving a more biophysically realistic plasticity model, which in turn will generate a more biophysically realistic weight distribution.

3.5 Effect of initial weight values

The results obtained also demonstrated that the initial synaptic weight values have some considerable influence on the appearance of the stable weight distribution. The spread of initial weights has little effect on the shape of the distribution, although it does make the network slower to converge to stability. Experiments in which initial weights were uniformly distributed from 0 to w_{\max} took the most time to reach stability, while those runs in which all weights were initially set at or near the maximum value were quick to converge. Uniform and Gaussian distributions produce very similar results, as do tests where all synapses begin with the exact same weight. The value of this weight, however, whether it be the point around which initial strengths are (relatively narrowly) distributed, or that which all synapses originate with, does have a considerable effect on the shape of the final distribution. This is illustrated by figures 11 – 13. In these instances, synaptic weights were given a Gaussian distribution around values of 20, 25 and 30 respectively, with a standard deviation of 5. The differences in the shape of the stable weight distributions are clear.

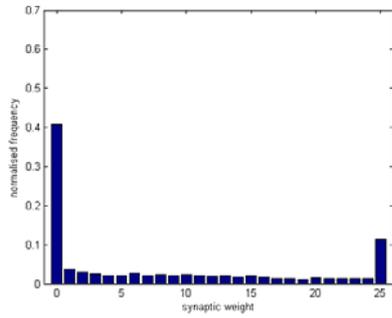


Fig. 7 – The falling block task

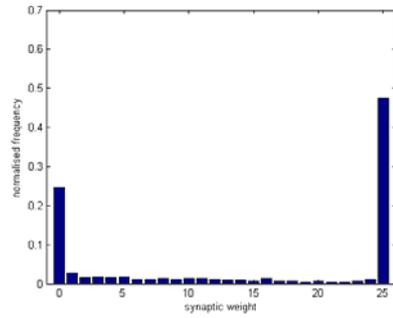


Fig. 8 – The phototaxis task

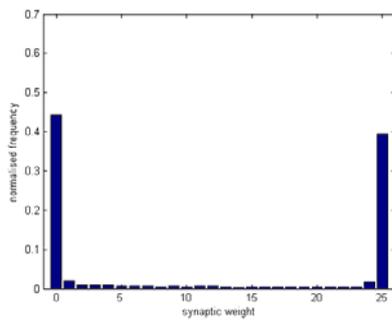


Fig. 9 – Correlated input

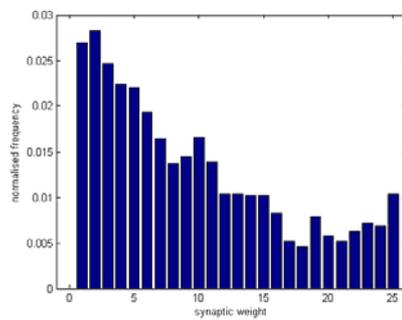


Fig. 10 – Size dependent potentiation

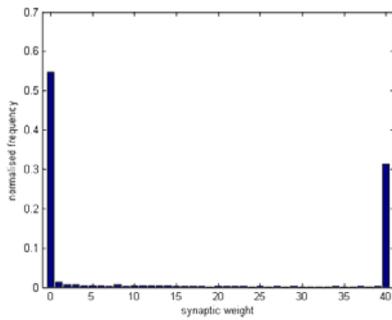


Fig. 11

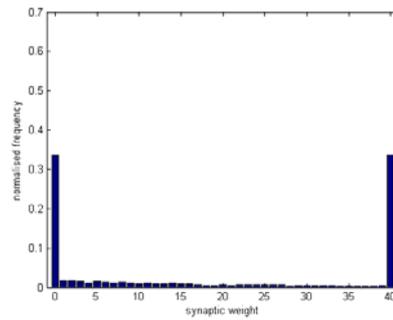


Fig. 12

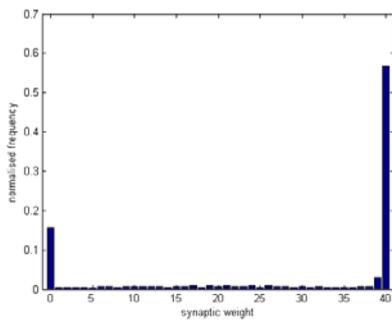


Fig. 13

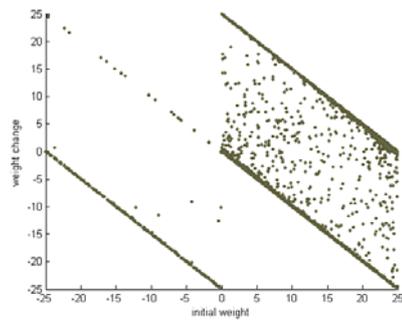


Fig. 14 – Changes in synaptic weight

It is interesting to note that the initial weight of a synapse bears no indication as to what its final weight value will be. In more simple Hebbian plasticity models, synapses above a certain strength would immediately be correlated with post-synaptic firing, and thus persistently potentiated, while those below that strength were persistently weakened. This implies that final weights could be predicted based on the initial configuration. With STDP, however, the competition generated between synapses allows no such predictions to be made. Figure 14 illustrates the relationship between the initial and final weights of a synapse. Although there is a slight residual tendency for weights which are originally strong to remain so, and likewise, for those which begin as very weak to remain at zero, the only clear correlation is caused by the hard limits on synaptic weight.

3.6 Conclusions

The results obtained mostly support previous findings in this area. Manipulation of the BCM threshold firing rate directs synaptic weights in an intuitive manner. The STDP model has a strong regulatory effect on post-synaptic output, although in some cases it seems to reduce the intermediate firing rate in the face of an increased frequency of input. The initial conditions of the network underlying the plasticity model, and the nature of input used, seem to have a pronounced effect on the direction in which it develops, which is to be expected from any dynamical system. Although STDP implicitly generates competition between synapses, the weight distribution it creates is still not representative of that found *in vivo*, unless additional experimental observations such as size-dependent potentiation are included. It is left to future work to elaborate on the results presented here, and to assess how beneficial the phenomena identified by this paper are to developing simple learning behaviour in ANNs.

Acknowledgements

The authors would like to thank early reviewers for their input, and for the general help and support of the members of the CCNR at Sussex University.

References

1. Wickliffe Abraham et al. Heterosynaptic metaplasticity in the hippocampus *in vivo*: A BCM-like modifiable threshold for LTP. *PNAS*, 98 (19): 10924-10929, 2001.
2. J.M. Bekkers et al. Origin of variability in quantal size in cultured hippocampal neurons and hippocampal slices. *PNAS*, 87: 5359-5362, 1990.
3. Guo-qiang Bi and Mu-ming Poo. Synaptic modifications in cultured hippocampal neurons : dependence on spike timing, synaptic strength and post-synaptic cell type. *Journal of Neuroscience*, 18: 10464 – 10472, 1998.
4. Elie Bienenstock, Leon Cooper and Paul Munro. Theory for the development of neuron selectivity : Orientation specificity and binocular interaction in the visual cortex. *Journal of Neuroscience*, 2: 32 – 48, 1982.
5. Ezequiel Di Paolo. Evolving spike-timing-dependent plasticity for single-trial learning in robots. *Phil. Trans. R. Soc. London*, 361: 2299-2319, 2003.
6. Eldan Goldenberg, Jacob Garcowski and Randall Beer. May we have your attention : Analysis of a selective attention task. *Proc. Eighth Int. Conf. Sim. Adap. Behaviour*, 49-56, 2004.
7. Donald Hebb. *The Organisation of Behaviour: A Neuropsychological theory*. Wiley, New York, 1949.
8. Javier Iglesias et al.. Stimulus-driven unsupervised synaptic pruning in large neural networks. *Proceedings of BV & AI*, LNCS 3704: 59-68, 2005.
9. Eugene Izhikevich and Niraj Desai. Relating STDP to BCM. *Letters to Neural Computation*, 15: 1511 – 1523, 2003.
10. Eugene Izhikevich. Which model to use for Cortical spiking neurons? *IEEE Transactions on Neural Networks*, 15 (5): 1063 – 1070, 2004.
11. Eugene Izhikevich, Joseph Gally and Gerald Edelman. Spike timing dynamics of neuronal groups. *Cerebral Cortex*, 14: 933-944, 2004.
12. T. Lomo and T. Bliss. Long-lasting potentiation of synaptic transmission in the dentate area of the anesthetized rabbit following stimulation of the perforant path. *Journal Physiology*, 232: 331-341, 1973.
13. Robert Malenka and Roger Nicoll. Long-term potentiation – A decade of progress? *Science*, 285: 1870 – 1874, 1999.
14. K.D. Miller and D.J. McKay. The role of constraints in Hebbian learning. *Neural Computation*, 6: 100-126, 1994.
15. K.D. Miller. Synaptic economics: competition and co-operation in synaptic plasticity. *Neuron*, 17: 371-374, 1996.
16. Patrick Roberts and Curtis Bell. Spike timing dependent plasticity in biological systems. *Biological Cybernetics*, 87: 392 – 403, 2002.
17. Sen Song, Kenneth Miller and L.F. Abbott. Competitive Hebbian learning through spike timing dependent synaptic plasticity. *Nature Neuroscience*, 3: 919 – 926 , 2000.
18. M.C.W. van Rossum, G.Q. Bi and G.G. Turrigiano. Stable Hebbian learning from spike timing dependent plasticity. *Journal of Neuroscience*, 20 (23): 8812 – 8821, 2000.
19. M.C.W. van Rossum and G.G. Turrigiano. Correlation based learning from spike timing dependent plasticity. *Neurocomputing*, 38-40: 409-415, 2001.