# Evolving spike-timing dependent synaptic plasticity for robot control

**Ezequiel A. Di Paolo**
School of Cognitive and Computing Sciences, University of Sussex
Brighton, BN1 9QH, UK, ezequiel@sussex.ac.uk

## Abstract

Plastic spiking neural networks are evolved as controllers for simulated phototatic robots. Synaptic plasticity depends on the precise relative time between presynaptic and postsynaptic spikes and on longer-term activity dependent regulatory scaling. Comparative studies have been carried out for non-plastic and different kinds of plastic neural networks. The evolved controllers are highly robust against sensorimotor perturbations such as sensor swapping, as well as against internal synaptic decay. The importance of the precise timing of spikes is demonstrated by randomizing the spike trains which results in degraded performance. Regulatory scaling plastically hypersensitizes the neural networks during periods of low activity, allowing robots to make use of single random spikes to trigger an active scanning strategy. Weight values undergo rhythmic changes at the mesoscale, but during periods of high activity they are finely regulated at the microscale by synchronous or entrained firing.

## 1. Introduction

Evolutionary robotics is a direct inheritor of the cybernetic research tradition on the origins of autonomous behaviour epitomised by the robotic work of W. Grey Walter. It is almost axiomatic in this kind of research that models of adaptive behaviour should be studied in integrated agents that are situated within, and in dynamic interaction with their environments. The daunting task of modelling anything resembling the simplest animal brain in its whole complexity in this way was anticipated by Grey Walter (1953), yet his hypothesis was that adaptive behaviour stemmed from the organisation and internal interactions within an organism and that these were partly replicable. If so, simple biologically-inspired controllers giving rise to adaptive, lifelike robot behaviour can turn out to be very valuable despite, or perhaps because of, their simplicity. In contrast, contemporary work in computational neuroscience, it all its sophistication, often lacks such a whole-agent dimension. Typically, functional properties of single neurons in specific subsystems are modelled and studied under idealised conditions (e.g., random uncorrelated inputs, etc.) and it is hard to see how the sensorimotor loops might ever be closed, valuable though this work is. Even for small networks the number of dimensions to optimize make a hand-designed approach very difficult. Recent studies in evolutionary robotics have aimed at harnessing the power of automatic evolutionary design to try to cross the gap between these two modes of research. For the time being these studies have tended to be mainly exploratory, drawing inspiration from neuroscience to enrich the building blocks used for evolutionary design – but the potential is there for feeding useful information back to neuroscience. A good example of this is the work by Husbands and colleagues using gaseous diffusion of neuromodulators as part of their evolved robot controllers (Husbands et al., 1998).

This paper follows a similar path. Recent work in evolutionary robotics has began to explore the use of spike-based neural controllers (Floreano and Mattiussi, 2001). Here the exploration continues by addressing the evolvability and properties of plastic spiking neural networks where synaptic plasticity depends on the precise timing of spikes. Parameters regulating plasticity in light-seeking robots are evolved in simulation. It is found that not only such controllers are evolvable, but also that they produce a rich variety of behaviours and desirable properties such as sensorimotor and synaptic robustness. The evolved controllers rely heavily on precise timing information and the relative firing activity between neurons to achieve adaptive behaviour and as this investigation will show they also open a series of important questions for future research.

## 2. Spike-timing dependent plasticity

Experimental neuroscientific evidence suggests that the degree and direction of change in the strength of a synapse subject to repeated pairings of pre- and postsynaptic action potentials depends on the relative timing between them (Markram et al., 1997, Bi and Poo, 1998). See (Bi and Poo, 2001) for a review. Synaptic modification depends on whether the pre- and postsynaptic spikes are separated in time in

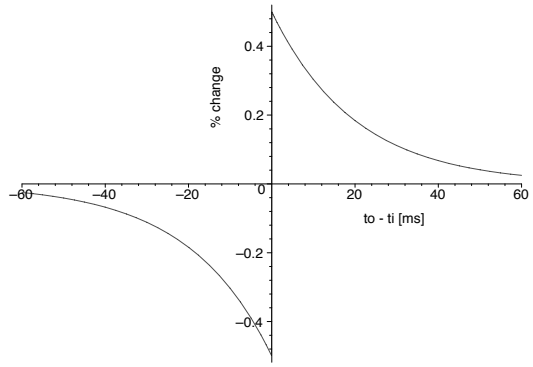


Figure 1: Time window for spike-timing dependent plasticity. The percentage and direction of synaptic change is given by time difference between pre-synaptic ($ti$) and postsynaptic ($to$) spikes.

less than a critical window of the order of a few tens of milliseconds. In most cases studied, if a presynaptic spike precedes the postsynaptic spike the synapse is potentiated, whereas the opposite relation leads to depression of the synapse. This results in a temporally asymmetric plasticity rule (Fig. 1) that deserves the name Hebbian because of its tendency to strengthen causal correlations between spikes. There is empirical evidence, however, for non-Hebbian STDP (Abbott and Nelson, 2000, Bi and Poo, 2001). Many theoretical studies have concerned themselves with this rule and its desirable properties such as a trend towards inherent stability in weight distribution and neural activity, unlike purely rate-based Hebbian rules which often require additional constrains (Kempter et al., 1999, Song et al., 2000, Rubin et al., 2001). The rule can be expressed as:

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

where $s = to - ti$ is the time difference between a postsynaptic and presynaptic spike and $A^+$ and $A^-$ are positive constants. Other filters may be used instead of the exponential decay, but this form is particularly suitable for implementation in and evolutionary robotics context.

One of the key concerns when studying rules for synaptic plasticity is their regulatory properties. Hebbian learning on its own leads to runaway processes of potentiation and cannot account for the stability of neural function. Additional elements, such as directional damping of synaptic change (Rubin et al., 2001) or longer-term stabilizing regulation based on postsynaptic activity (Turrigiano, 1999) may come to the rescue. These can lead to unsaturated distributions of synaptic strengths in the first case and to regulated neuronal firing in the second and will also be investigated in this work.

## 3. Methods

**Integrate-and-fire model.** The time evolution of the membrane potential $V$ of an integrate-and-fire neuron with reversal is given by:

$$\tau_m \frac{dV}{dt} = V_{rest} - V + g_{ex}(t)(E_{ex} - V) + g_{in}(t)(E_{in} - V)$$

where $\tau_m$ is the membrane time constant (range [10 ms, 40 ms]), $v_{rest}$ = -70 mV is the rest potential, the excitatory and inhibitory reversal potentials are respectively $E_{ex} = 0$ mV and $E_{in}$ = -70 mV.

A noisy threshold value, $V_{thres}$, is given by a normal distribution with a genetically set mean value (range [-65 mV, -50 mV]) and a deviation of 1 mV. When the membrane potential reaches this threshold, an action potential is fired and $V$ is reset to $V_{rest}$. An absolute refractory time of 4 ms prevents the neuron from firing another spike within this period.

Every time a spike arrives to neuron $j$ from an excitatory presynaptic neuron $i$ the excitatory conductance of $j$ is increased by the current value of the synaptic strength ($w_{ij}(t)$): $g_{ex}(t) \rightarrow g_{ex}(t) + w_{ij}(t)$. The inhibitory conductance $g_{in}$ is similarly affected by spikes coming from inhibitory neurons. Conductances otherwise decay exponentially:

$$\tau_{ex} \frac{dg_{ex}}{dt} = -g_{ex}; \quad \tau_{in} \frac{dg_{in}}{dt} = -g_{in}$$

with $\tau_{ex}$ and $\tau_{in}$ genetically set for each neuron from the range [4 ms, 8 ms]. The firing rate $z$ of a neuron is estimated by a similar leaky integration:

$$\tau_z \frac{dz}{dt} = -z + \sum_i \delta(t - t^{(f)})$$

where $t^{(f)}$ are the times when the neuron emits a spike and $\tau_z = 100$ ms.

**Spike-timing dependent plasticity (STDP).** The properties of plastic windows (Fig. 1) are evolved for each synapse in the neural network controller. Following (Song et al., 2000), synaptic change is implemented using two functions per synapse $P^-(t)$ and $P^+(t)$. Every time a spike arrives at the synapse $P^+(t)$ is incremented by $A^+$, and every time the postsynaptic neuron fires $P^-(t)$ is decremented by $A^-$. Otherwise, these functions decay exponentially with time constants $\tau^-$ and $\tau^+$ respectively. $P^-(t)$ is used to decrease the synaptic strength every time the presynaptic neuron fires: $w_{ij}(t) \rightarrow w_{ij}(t) + w_{max} \times P^-(t)$. Analogously, $P^+(t)$ is used to increase the synaptic strength every time the postsynaptic neuron fires: $w_{ij}(t) \rightarrow w_{ij}(t) + w_{max} P^+(t)$. In all the cases studied here the maximum synaptic strength is $w_{max} = 1$. This method facilitates the computational implementation of STDP by eliminating the

need of keeping track of spike trains or calculating other response functions which could be more costly.

The values for $A^+$ and $A^-$, and $\tau^+$ and $\tau^-$ are genetically set per synapse from the ranges [0.0001, 0.05] and [10 ms, 40 ms] respectively. In all the experiments reported here the plastic windows are Hebbian, that is, spikes arriving before a postsynaptic action potential always potentiate a synapse and those arriving after always depress it. Experiments relaxing this constraint, i.e., allowing anti-Hebbian or purely potentiating or depressing windows, have also been carried out successfully, but are not reported here.

**Activity-dependent synaptic scaling (ADS).** Some of the mechanisms used by neurons to regulate their firing rate homeostatically are thought to affect all incoming synapses scaling them up or down independently of the presynaptic activity, (Turrigiano, 1999). If the postsynaptic activity is above a certain target, excitatory synapses are scaled down, otherwise, they are scaled up, thus preventing sustained levels of activity that are too high or too low. Following (van Rossum et al., 2000) excitatory synapses are modified according to:

$$\tau_{ads}\frac{dw_{ij}}{dt} = w_{ij}(z_{goal} - z_j)$$

where $z_{goal} = 50$ Hz and $\tau_{ads}$ is genetically set from the range [1 s, 10 s]. In real neurons, this is a mechanism that acts over long timescales (over hundreds to thousands of seconds) (Turrigiano et al., 1998), but due to computational limitations (the very long evaluation runs that would be required) the chosen timescale ($\sim \tau_{ads}$) is faster than this but still significantly slower than the rest of the timescales in the system. Even though the above mechanism acts on excitatory synapses, in the current context it has also been applied when the presynaptic neuron is inhibitory by multiplying the right hand side above by $-1$.

**Directional damping.** Synaptic weights are constrained within the range [0, 1]. This can be achieved simply by a stop condition at the boundaries or by means of damping factors that vanish as the weight value approaches a boundary. The choice can have important consequences. No damping leads to a bimodal distribution of weights under random stimulation (Song et al., 2000), where most weights adopt the minimum or maximum values in the range, but few values in-between. The same happens with purely positional damping, i.e., factors that slow down weight change near the boundaries, but depending only on the current weight value. A biologically plausible alternative is directional damping whereby if a weight value is near a boundary, changes that push this value towards the boundary are slowed down, but changes that push it away from the boundary are not. The equilibrium weight distribution in this case tends to be unimodal and centred around the point where potentiation and depression equilibrate, (Rubin et al., 2001). Directional damping is supported empirically by the observation that spike-driven potentiation is more pronounced than the expected linear variation at synapses of relatively low initial strength in cultured hippocampal cells (Bi and Poo, 1998).

Linear directional damping is simply implemented by transforming a weight change (as resulting from STDP or ADS or both): $\Delta w_{ij} \rightarrow (1 - w_{ij})\Delta w_{ij}$ if $\Delta w_{ij} > 0$ and $\Delta w_{ij} \rightarrow w_{ij}\Delta w_{ij}$ if $\Delta w_{ij} < 0$ for $w_{ij} \in [0, 1]$.

**Synaptic decay.** In order to test the robustness of the evolved controllers to perturbations in their internal configuration, synaptic weights are allowed to decay to 0 with a time constant that can be as fast as 100 ms. Synaptic decay is not affected by directional damping and is not applied during evolution.

**Poisson filters and randomized delays.** In order to test the reliability of the evolved controllers on the precise timing of spikes, the simple expedient of filtering the output of a neuron with a Poisson process emitting random spikes at the same instantaneous rate has been used. Information about firing rate is conserved, but precise spike-timing is disrupted. Because, the rate $z$ is estimated using only previous spikes, it is only possible to approximate the instantaneous firing rate of the neuron in this manner. It is expected, however, that if controllers rely heavily on firing rates, the disruption in performance should not be too strong. A more sophisticated method consist in introducing artificial random delays in the firing time of single or multiple neurons. This is done by keeping a short sub-train corresponding to the last $T$ ms of activity and swapping the current fire state of a neuron with a randomly selected state in the sub-train thus conserving the number of spikes. These tests are not applied during evolution.

**CTRNN.** Control runs using rate-based continuous-time, recurrent neural networks (Beer, 1990) have been performed. These are defined by:

$$\tau_i\frac{dV_i}{dt} = -V_i + \sum_j w_{ji}z_j + I_i; \quad z_j = \frac{1}{1+\exp[-(V_j+b_j)]},$$

where, $V_i$ represents the membrane potential of neuron $i$, $\tau_i$ the decay constant (range [0.4 s, 4 s]), $b_i$ the bias (range [-3, 3]), $z_i$ the firing rate, $w_{ij}$ the strength of synaptic connection from node $i$ to node $j$ (range [-8, 8]), and $I_i$ the degree of sensory perturbation for sensory nodes.

**Robot sensors and effectors.** The simulated robots use photoreceptors that are activated by the light intensity corresponding to their current position if the light source is directly visible (i.e., an angle of acceptance of 180 degrees). This intensity is multiplied by the sensor gain (genetically set from range [0.1, 50]) and clipped for values beyond of maximum of 20. A spike train is gen-

erated using a Poisson process with variable rate (maximum 200 Hz) by linearly transforming the sensor value into the instantaneous rate. Additionally, uniform noise is present in the sensors (and motors) with range 0.2 – this results in spikes that fire randomly with very low probability when the sensor is not stimulated.

Two motors control the robot wheels. Each motor is controlled by two neurons, one that drives it forwards and the other one backwards, using a spike-based leaky integrator. The current motor value is stored in a variable $M_{L,R}$ which is directly translated into the left and right velocities respectively.

$$\tau_{mot}\frac{dM_{L,R}}{dt} = -M_{L,R}+M_G(\sum \delta(t-t^{(f)}_{forw})-\delta(t-t^{(f)}_{back}))$$

with $\tau_{mot}$ genetically set from the range [40 ms, 100 ms] and $M_G$ from [0.1, 50]. Both motors have a same value for their gains and decay constants. This approach marks a difference with previous work on the evolution of spiking controllers which have used a neural rate estimation method for driving the motor (Floreano and Mattiussi, 2001).

**Evaluation.** Simulated robots are evolved to perform phototaxis on a series of light sources. Robots have circular bodies of radius $R_0 = 4$ with two motors and two light sensors. The angle between sensors is of 120 degrees but a small random displacement between -5 to 5 degrees is added at the start of each evaluation. Motors can drive the robot backwards and forwards in a 2-D unlimited arena.

The neural network consists of 6 nodes, fully-connected except for self-connections. The left motor is controlled by neurons 0 and 4, and the right by neurons 1 and 5; the Poisson spike trains coming from the left and right sensors are fed into neuron 2 and 3 respectively. Neurons can be either excitatory or inhibitory and this is set genetically. Trials with larger number of neurons have been carried out successfully, but not systematically studied.

The whole system is simulated using an Euler integration method with a time step of 1 ms [25% of the minimum timescale]. Robots are run for 2 independent evaluations, each consisting on the sequential presentation of 2 distant light sources. Only one source is presented at a time for a relatively long period $T_S$ chosen randomly for each source from the interval [7.5 s, 12.5 s], (each evaluation consists therefore of an average of $2 \times 10^4$ update cycles). The initial distance between robot and new source is randomly chosen from [60, 80], the angle from [0, $2\pi$) and the source intensity from [3000, 5000]. The intensity decays in inverse proportion to the square of the distance to the source.

**Evolutionary Algorithm.** A population of 30 robots is evolved using a generational GA with truncation selection. In the plastic scenarios described below initial weights are randomly chosen at the start of each evaluation from the interval [0, 1] while the parameters for the plasticity windows and scaling constants are evolved. In the non-plastic scenario, synaptic strengths are encoded genetically. Other genetically set parameters include sensor and motor gains, motor decay constant and whether neurons are inhibitory or excitatory. All parameters are encoded in a real-valued genotype, each gene assuming a value within [0, 1]; each parameter is linearly scaled to the corresponding range of values, except for sensor and motor gains which are scaled exponentially. Only vector mutation (Beer, 1996) is used with a standard deviation of vector displacement of 0.5 (maximum genotype length is 220), genetic boundaries are reflective. Fitness is calculated according to:

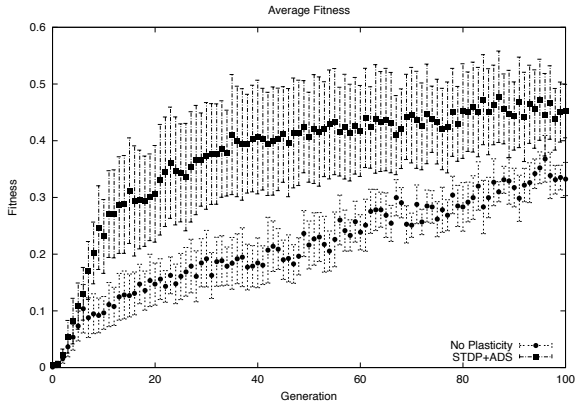


Figure 2: Mean population fitness averaged over five independent runs for the 2 evolutionary scenarios.

$$F = \frac{(1-M^2)}{T_S}\int f dt; \quad f = 1 - \frac{d}{D_i}$$

if the current distance to the source $d$ is less than the initial distance $D_i$, otherwise $f = 0$. $M$ measures the average difference in activity between the motors divided by the motor gain.

$$M = \frac{0.125}{T_S}\int \frac{(M_L - M_R)}{M_G} dt.$$

Near optimal fitness will be obtained by robots approaching a source of light rapidly and with minimal integrated angular movement.

## 4. Results

Robots using spiking controllers were successfully evolved for four different scenarios: 1. No plasticity (evolution of fixed weights), 2. STDP No Damping (evolution of STDP windows without directional damping, random initial weights), 3. STDP (evolution of STDP windows with directional damping, random initial weights), and 4. STDP+ADS (evolution of STDP windows and ADS with directional damping, random initial weights).

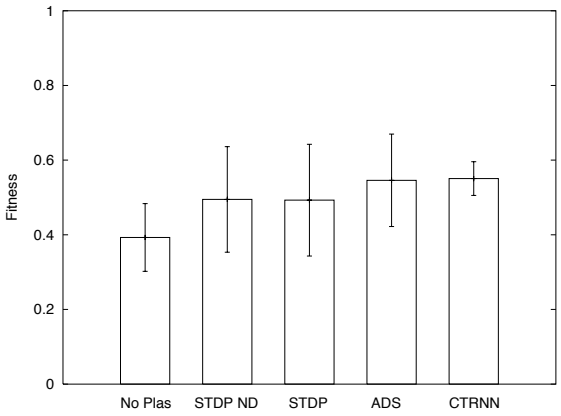


Figure 3: Average fitness of the best individual in the last generation.

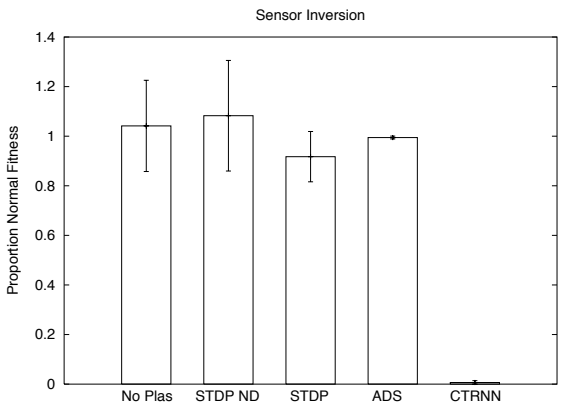


Figure 4: Robustness against right-left inversion of visual field. Proportion of original fitness for each scenario.

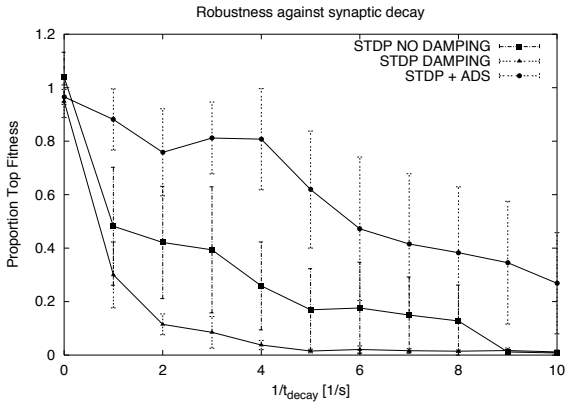


Figure 5: Robustness against synaptic decay; $t_{decay}$ indicates the speed with which weights decay exponentially to 0.

**Evolvability.** Five 100-generation independent runs were made for each of the four scenarios above. It was qualitatively found that, contrary to expectations, the more complex case (in term of the dimensions of the search space and the additional features of the mechanism), that is STDP+ADS, was easier to evolve, particularly during the initial generations, than the simpler case of no plasticity. This is observed in figure 2 where the mean population fitness, averaged over the five runs, is plotted for these cases (error bars indicate standard deviation). There is little observable distinction between using STDP with or without damping (not shown) but there is a significant difference between STDP+ADS and no plasticity.

However, the long term fitness of the best individual of the last generation is not significantly different between the cases (Fig. 3). These were obtained by running the best individual of each of the 20 runs for 10 independent evaluations. For comparison purposes, a CTRNN controller has been evolved and tested under the same conditions and is also shown in the figure.

**Sensor robustness.** A number of tests have been performed to explore the sensorimotor robustness of the evolved controllers. Most of them show significant robustness to perturbations in motor and sensor gains and and the introduction of motor and sensor offsets. Here only robustness to sensor inversion is reported. This perturbation consists in swapping the left and right sensors. All four scenarios have evolved highly robust controllers, approaching 100 % of their original fitness under the perturbed condition (Fig. 4). The very small standard deviation of the STDP+ADS should be noticed. In contrast, CTRNN evolved under the same condition shows no robustness under sensor inversion.

**Synaptic decay.** Plastic controllers maintain their functionality dynamically as a consequence of their own activity. In order to assess their reliability a disruptive decay of synapses was introduced as described above. Figure 5 shows results for the three plastic set-ups (again using 5 independent runs for each and testing the best individual 10 times). It is apparent that STPD+ADS controllers are able to perform quite reliably even for decay times of up to 250 ms, while controllers using STDP with or without damping are unable to maintain their performance. It is possible to explain this as a consequence of the compensatory nature of the ADS mechanism, which is able to alter synapses as a consequence of longer term changes in neural activity in ways that tend to maintain this activity and therefore the functionality of the controller.

**Behavioural strategies.** Evolved robots show a rich variety of behavioural strategies and not all can be discussed here. Practically all of the observed strategies are active, involving scanning behaviour, which is not surprising given that sensors saturate rather easily. Figure 6 shows a trajectory for an STDP+ADS robot. Unlike what is commonly observed using rate-based neural

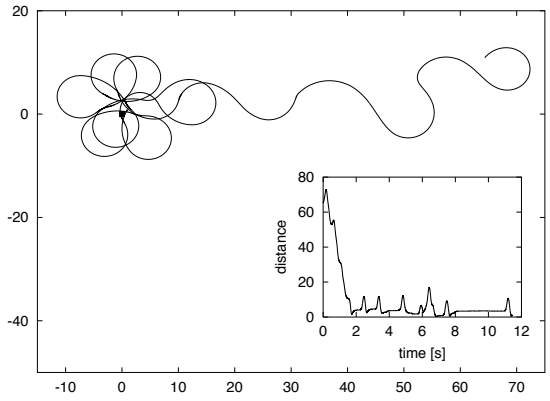


Figure 6: Trajectory of an evolved robot using STDP+ADS, inset: distance to light source.

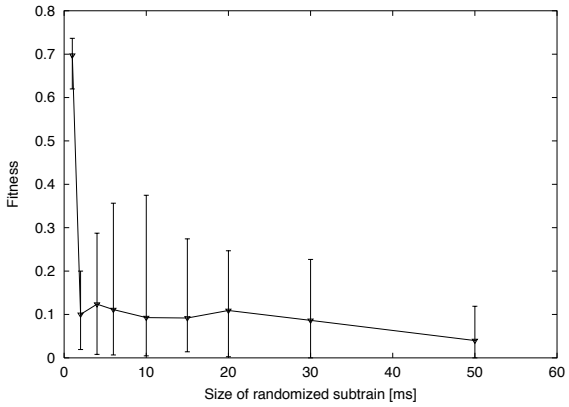


Figure 7: Fitness effect of randomizing spike sub-trains in all neurons for an STDP+ADS controller. Bars indicate maximum and minimum values.

controllers, robots are often able to come to a full stop near a source of light (but not facing it) – this can be observed in the inset showing the distance to the light source. During these periods there is very little neural activity, until a random spike from the sensors triggers a round of activity and the robot starts scanning again (see below). In the absence of light robots scan their surroundings using this mechanism.

**Timing dependence.** Applying a Poisson filter test to all the neurons in a controller, whereby rate information is conserved but not spike timing, results in total loss of fitness. This was tested in different controllers in the four scenarios with the same result. A more detailed study involves the application of random delays to single neurons. As explained above, this method conserves the number of spikes but randomizes the timing within a sub-train of a given size corresponding to the last $T$ ms in the simulation. Figures 7 and 8 show how fitness is affected for a best individual in one STDP+ADS run when random delays are applied to all the neurons and to single neurons respectively (similar results were obtain for the other classes). There is a sharp reduction in fitness in the first case for the randomization of subtrains as short as 2 ms, but applying the test to single neurons shows that the controller is crucially dependent on the precise timing of only 3 of the six neurons (corresponding in this case to one motor and the two sensor neurons, but to different neurons in other cases).

**Synaptic dynamics and internal regulation.** Weight values also exhibit rich dynamical behaviour. Figure 9 shows three of the synaptic weights affecting neuron 0 for the same run corresponding to figure 6. The other neurons show similar qualitative dynamics at this timescale, consisting of rhythmic periods of activity (during which the robot rotates left and right) punctuated by silent periods (during which the robot gently comes to a full stop facing away from the light). While

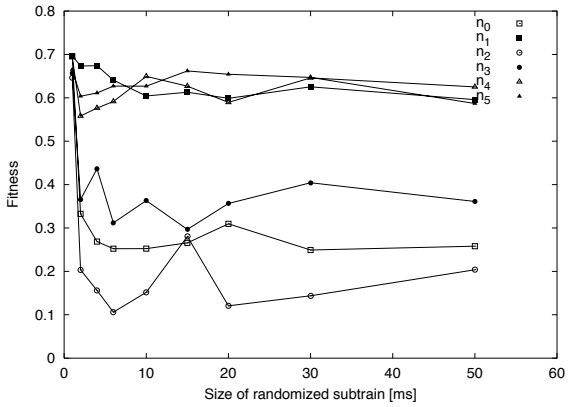


Figure 8: Fitness effect of randomizing spike sub-trains in individual neurons for an STDP+ADS controller.

neurons are firing, STDP drives the weights always to a same area within the range, the rest of the time ADS takes control and increases the synaptic efficacies so as to compensate for the lack of firing activity in the neuron. This has the effect of hyper-sensitizing the whole network, so that even a single spike arriving from the sensors is often capable of triggering a new round of activity and the robot starts moving again.

The dynamics of the microstructure are also interesting. Figure 10 shows how a particular synaptic strength is regulated during a period of activity lasting about 0.6 s. Its value is maintained nearly constant. The insets show the pattern of firing of the pre- and postsynaptic neurons which undergoes a process of synchronization that persist while at least one sensor is active and finally de-synchronizes and stops. The plasticity window shown in the figure indicates that synchronous firing translates into net potentiation which is compensated by ADS resulting in an equilibrium (just after activity stops, ADS

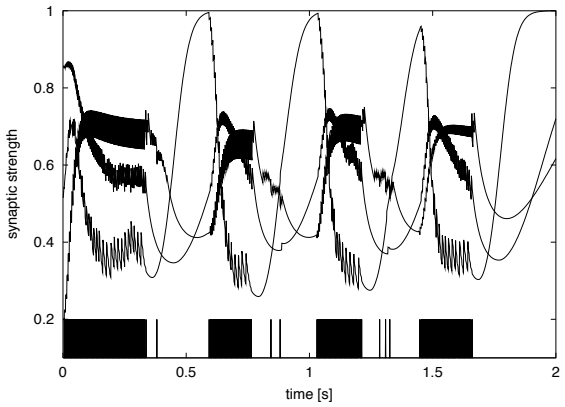


Figure 9: Example of weight dynamics shown for 3 synapses affecting node 0 together with its firing pattern.

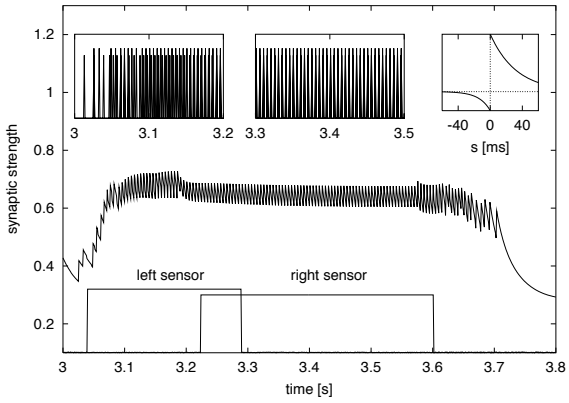


Figure 10: Weight regulation (synapse $w_{20}$) during a period of network activity and sensor activations (not to scale) – insets: firing patterns of corresponding neurons at the onset of activity round (left) and during highly ordered period (centre), evolved plasticity window for this synapse (right).

depresses the weight due to the inertia of the frequency estimation before re-sensitizing the network). This regulatory pattern has been observed in all of the cases studied for the STDP+ADS class.

Figures 11 and 12 show the spike trains for a same controller under normal and sensory inversion conditions. In both cases the robot's movement is practically the same, yet sensors become active in a different order as expected. It is hard to distinguish between the two cases by observing the spike trains. In both situations activity starts just at the end of a period of relaxation where the robot is coming to full stop. In the first case light comes into view before the robot stops and triggers another round of activity, but in figure 12 the robot has fully stopped and the network has reached its peak excitability. In this case activity is triggered by a random single spike coming from the left sensor (arrow).

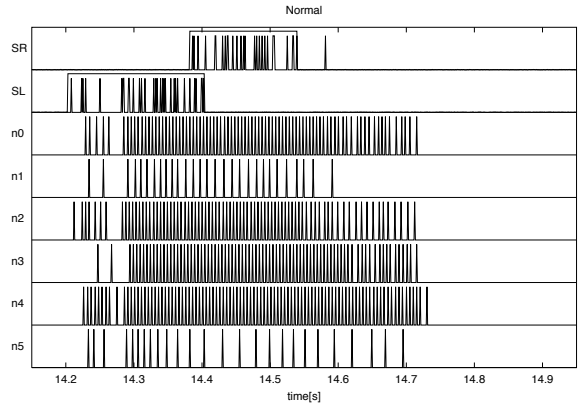


Figure 11: Network activity starting just after a period of relaxation. The Poisson trains generated by the sensors are shown at the top.

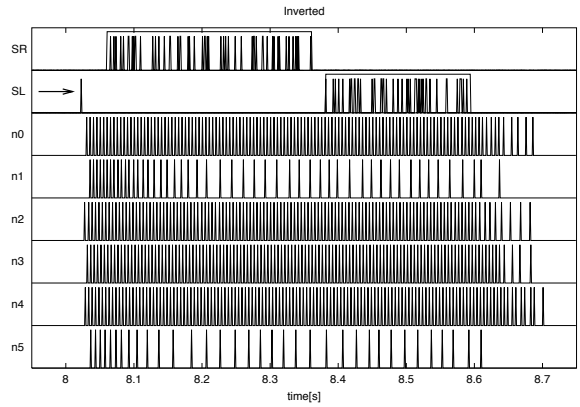


Figure 12: Same as in figure 11 for the same controller but with inverted sensors.

The neurons that drive the left motor (n0 and n4) fire at a same frequency, yet it is possible to observe that their phase relation changes non-randomly so as to produce a peak in motor output roughly towards the middle of the activity round. The right motor does something similar just after the left motor peaks, although its driving neurons (n1 and n5) are not clearly entrained, (data not shown).

## 5. Discussion

Despite the exploratory nature of this work, it is possible to assert that the richness of behaviour in plastic spiking neural networks, not often found in other controllers, makes them extremely interesting candidates for further testing in adaptive behaviour research. It is perhaps inevitable that a synthetic design method should be used to approach their complexity. This paper has shown that it can be done successfully for a simple task. The suitability of these mechanisms in other scenarios, and ulti-

mately their potential for more complex cognitive performance, remains to be seen. One of the main disadvantages of the approach is obviously the extra computational cost involved in the longer evaluations. But this cost has passed from being prohibitive a few years ago to being acceptable nowadays if the benefits justify them.

The evolved controllers demonstrate a number of advantages as well as the use of uncommon mechanisms (in a robotics context) such as the precise timing information in spike trains. In all the cases studied in detail neural networks undergo rhythmic periods of activity during which pairs of neurons start un-correlated, then they reach a highly entrained state, and finally they lose their entrainment and become inactive. Such periods can be triggered by a sensory stimulus or if the inactivity has lasted long enough, even by single random spikes coming from the sensors, owing to the excitability that builds up thanks to the ADS mechanism. During these periods synaptic strengths are kept nearly constant and despite firing in strongly entrained mode motor neurons maintain enough variety to coordinate their relative timing and achieve functionally useful movement. Exactly how these coordinated microactivity is achieved remains an open issue. Single neuron randomization of spike trains has revealed that not all neurons are crucial in allowing the network to make use of timing information. And this is, so far, the strongest lead for further investigation. Future studies will also look at the effects of added noise (e.g., random refractory periods, baseline random firing).

As a contrast to timing disturbances, the evolved neural networks show an high measure of sensor robustness in both plastic and non-plastic cases, as compared with CTRNNs. Why would this be so is another open question. A similar effect has recently been observed in oscillatory rate-based controllers (Di Paolo, 2002) and may be related to the fact that neural controllers with intrinsic dynamics that are faster than the timescale of performance need to rely on integrative methods for making use of sensor stimuli (instantaneous information is rapidly lost within the endogenous pattern). In other words, such controllers need to function less as "sensory-driven" and more as "sensory-modulated" which means that they will rely on more global sensorimotor invariants (such as the fact that the periods of body occlusion of the light source become longer as the robot approaches it) as opposed to local invariants (e.g., the light is in the direction of the strongest sensor reading). This hypothesis could account for sensorimotor robustness. In order to further test it, more data involving multi-timescale controllers is necessary.

Thanks to Jianfeng Feng for first suggesting STDP might be worth looking into. The author wishes to acknowledge the support of the Nuffield Foundation, (grant no. NAL/00274/G).

## References

Abbott, L. F. and Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nature Neurosci*, **3**:1178–1183.

Beer, R. D. (1990). *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroscience*. San Diego: Academic Press.

Beer, R. D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In Maes, P., Mataric, M. J., Meyer, J.-A., Pollack, J. B., and Wilson, S. W., (Eds.), *From Animals to Animats 4: Proceedings of the 4th Int'l Conf on Simulation of Adaptive Behavior*, pages 421 – 429. Cambridge, MA: MIT Press.

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

Bi, G. Q. and Poo, M. M. (2001). Synaptic modifications by correlated activity: Hebb's postulated revisited. *Ann Rev Neurosci*, **24**:139–166.

Di Paolo, E. A. (2002). Evolving robust robots using homeostatic oscillators. Cognitive Science Research Paper 548, COGS, University of Sussex.

Floreano, D. and Mattiussi, C. (2001). Evolution of spiking neural controllers for autonomous vision-based robots. In Gomi, T., (Ed.), *Evolutionary Robotics IV*. Springer Verlag.

Husbands, P., Smith, T., Jakobi, N., and O'Shea, M. (1998). Better living through chemistry: Evolving GasNets for robot control. *Connection Science*, **10**:185–210.

Kempter, R., Gerstner, W., and van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. *Phys Rev E*, **59**:4498–4514.

Markram, H., Lubke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, **275**:213–215.

Rubin, J., Lee, D. D., and Sompolinsky, H. (2001). Equilibrium properties of temporally asymmetric Hebbian plasticity. *Phys Rev Lett*, **86**:364–367.

Song, S., Miller, K. D., and Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neurosci*, **3**:919–926.

Turrigiano, G. G. (1999). Homeostatic plasticity in neuronal networks: The more things change, the more they stay the same. *Trends Neurosci.*, **22**:221–227.

Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., and Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, **391**:892–896.

van Rossum, M. C. W., Bi, G. Q., and Turrigiano, G. G. (2000). Stable Hebbian learning from spike-timing dependent plasticity. *J Neurosci*, **20**:8812–8821.

Walter, W. G. (1953). *The living brain*. London: Duckworth.