

# Fast homeostatic neural oscillators induce radical robustness in robot performance

Ezequiel A. Di Paolo

School of Cognitive and Computing Sciences, University of Sussex  
Brighton, BN1 9QH, UK, ezequiel@cogs.susx.ac.uk

## Abstract

A network of relaxation oscillators is evolved to produce phototaxis in a simulated robot. Oscillations are faster than the timescale of performance, and are designed to maintain the same average activation value independently of sensory or synaptic input. Neural activation cannot correlate with any action-relevant sensory information, but must be continuously modulated by sensorimotor coupling. Radical sensor robustness is shown by inverting the position of the sensors and also by removing either of them in turn – operations that do not alter the success of the strategy. Slowing down the timescale of oscillations results in less robustness.

This paper explores some of the issues that arise from evolving oscillatory neural controllers which operate timescales faster than that of performance, and whose elements compensate for long terms input patterns by keeping their average activation as close as possible to a middle range, thus making it difficult for action relevant information to be stored in individual elements.

Continuous-time recurrent neural networks (CTRNNs) (Beer, 1990) are extended to transform each neuron into a centre-crossing relaxation oscillator that maintains a constant average activation of 0.5. The equations describing traditional CTRNNs are:

$$\tau_i \dot{y}_i = -y_i + \sum_j w_{ji} z_j + I_i; \quad z_j = \frac{1}{1 + \exp[-(y_j + b_j)]},$$

where  $y_i$  represents the membrane potential,  $\tau_i$  the decay constant (range [0.4,4]),  $b_i$  the bias,  $z_i$  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 on the sensor node. These equations are extended by turning the neuron bias into a responsive variable that “keeps track” of the opposite value of the membrane potential with a  $\tau_i^b$  greater than  $\tau_i$ :

$$\tau_i^b \dot{b}_i = -(b_i + y_i); \quad \tau_i^b = \tau_i G_i;$$

where  $G_i$  ranges from 1 to  $T$  and is genetically set. The consequence of adding this equation is that the bias term

is no longer constant but adapts so as to maintain a long term average firing rate of 0.5, regardless of all the other parameters and input pattern. The parameter  $T$  indicating the range of allowed values for  $\tau_i^b$  was set for different series of runs at 2, 5, and 10.

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. Motors can drive the robot backwards and forwards in a 2-D unlimited arena. Tests were run with fully connected networks of 4, 6, 8, 10 and 20 nodes with similar results. More details in (Di Paolo, 2002).

Robots are run for 4 independent evaluations, each consisting on the sequential presentation of 6 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 [300,500]. 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 [500,1500].

A population of 30 robots is evolved using a generational GA with truncation selection. All parameters are encoded in a real-valued vector. Only vector mutation is used with a standard deviation of vector displacement of 0.25. Fitness is calculated according to:

$$F = \frac{1}{T_S} \int f dt; \quad f = 1 - \frac{d}{D_i}$$

if the current distance to the source  $d$  is less than the initial distance  $D_i$ , otherwise  $f = 0$ .

In all cases the GA was run for 2000 generations. Control runs were performed using standard CTRNNs with genetically set but fixed bias terms from the range [-3,3]. The following results correspond to 6-node neural controllers. Each data point is taken by averaging the fitness over 10 trials of 5 independently evolved robots over a series of 50 light sources. The final distance to the source was in all cases between 3 and 10 units showing that all controllers perform phototaxis equally well.

A series of radical sensor perturbations has been performed. Figure 1 shows the proportional decay in performance for sensor inversion, consisting in swapping

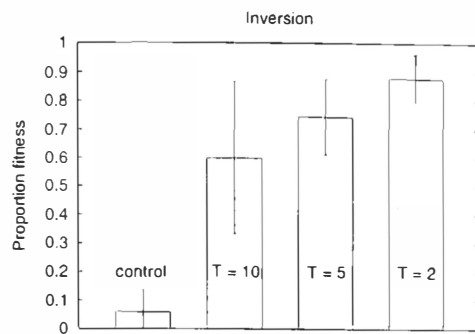


Figure 1: Average relative robustness (measured as proportion of unperturbed performance) for sensor inversion.

sensor positions left and right. Oscillating neural controllers show high robustness which increases as  $T$  is made smaller, i.e., for higher frequencies. Observed behaviour is unchanged by inversion for  $T = 2$ . Similar results have been obtained for sensor removal. Robustness is also shown in angular displacement tests where the position of both sensors is shifted by a same angle (Di Paolo, 2002). Control networks show a heavy reliance on a single sensor. Displacement of this sensor results in very different angular ranges for high fitness, and removal of the same results in negligible performance. In contrast, the oscillator-controlled robot shows the same angular distribution for all perturbations, indicating that the strategy used will work as long as there is at least one sensor (any sensor) facing in the forward direction of the robot.

Analysis of the evolved strategy has been performed for a 4-neuron network. All neurons behave as oscillators entrained to a common frequency of 0.11 cycles per unit of time. Neural response to step-changes in input is confined to less than one cycle of oscillation. After that the oscillation pattern is recovered with a possible phase shift. In all the cases tested the strategy used by the robot is a variation of the strategy shown in figure 2. Motor neurons oscillate maintaining a roughly constant phase difference. The trajectory far from the light can be described as alternating segments where the robot moves forwards, rotates on the spot about three quarters of a full turn, moves backwards, rotates on the spot, moves forwards, and so on. Because of the imposed limitations on the controller, the robot must use an active scanning strategy as sensor neurons will lose any instantaneous information they might acquire from the input currents.

The controller uses input currents to modulate oscillations and make the network switch between different regimes. This is clear from observing the behaviour of the sensor neurons whose oscillations are unstable in the presence of the intermittent pattern of input currents caused by rotating near the light source. Analysis of the

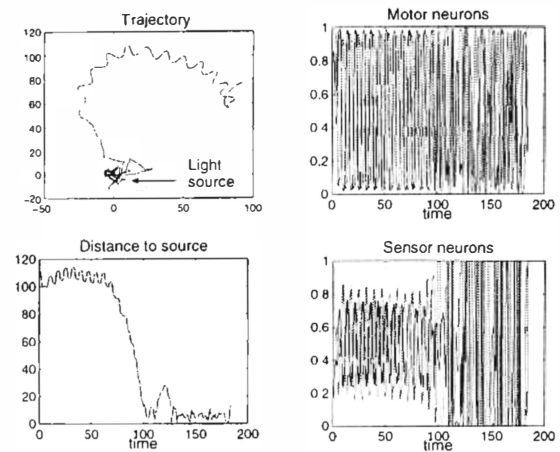


Figure 2: Robot behaviour, neural activation and distance to light source during phototaxis.

transition between the near and far regimes throws some light into the origins of robust behaviour. Even though low input currents do not alter the oscillation mode of the neurons, they do modulate the phase of oscillation (Di Paolo, 2002). The effect of this phase modulation correlates with whether, during rotations, the robot faces the general direction of the light source or the opposite direction. In the first case, the angle of rotation is made smaller as a consequence of the increase in input current. In the second case, the angle is not affected, resulting in a deviation of the trajectory towards the source of light. For such a strategy the precise location of the sensors is not a very sensitive parameter, as long as they are placed towards the front of the body.

The results show that when single neurons are prevented from storing long-term information in their activation values, good solutions evolve that make use of the relative coordination between neural oscillations to solve the desired task. These solutions are extremely robust to sensory perturbation. Two essential components seem to be needed for this: the timescale of oscillation must be faster than the timescale of performance and the long-term average activation of each neuron must conserve an undifferentiated average value independently of the history of inputs.

The author wishes to acknowledge the support of the Nuffield Foundation, (grant no. NAL/00274/G).

## References

- Beer, R. D. (1990). *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroscience*. San Diego: Academic Press.
- Di Paolo, E. A. (2002). Evolving robust robots using homeostatic oscillators. Cognitive Science Research Paper 548, COGS, University of Sussex.