Evolving robust robots using homeostatic oscillators

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

CSRP No 548

Abstract

A network of homeostatic relaxation oscillators is evolved to produce non-rhythmic phototactic behaviour in a simulated robot. Neural oscillations take place at a faster timescale than that of performance, and are designed to maintain an average activation value which is independent of sensory or synaptic input. In this way, neural activation cannot correlate directly with any action-relevant sensory information, but must be continuously modulated by sensorimotor coupling. Evolution finds robust controllers which work by altering their central oscillation patterns. Robot are evolved with a fixed set of body parameters, including sensor positions. 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. Fast dynamics and long-term homeostasis seem to be required for robustness; slowing down the timescale of oscillations results in less robustness. The need for long-term homeostasis is shown both by modifying the oscillators and by running control experiments using a network of FitzHugh-Nagumo neurons. In none of these cases robustness is obtained. A general hypothesis is proposed according to which removing functional specificity from the lower-level mechanisms is likely to result in robust performance at the global level.

1. Introduction

There is an increasing interest within the context of autonomous robot design in biologically inspired controllers capable of acting over a large range of timescales (Husbands et al., 1998) as well as in controllers that combine medium and slow timescales (Beer, 1990, Yamauchi and Beer, 1994, Floreano and Urzelai, 2000, Fujii et al., 2001, Di Paolo, 2000) and medium and fast timescales (Floreano and Mattiussi, 2001). Harvey (1997) has emphasized the importance of robot controllers implementing a wide range of timescales in the context of dynamical accounts of adaptive behaviour. Despite this interest the majority of work in this area has concentrated on control architectures operating at a single preferred timescale: the typical timescale of behavioural performance. The average time it takes for an element of performance to be achieved does not differ significantly from the average time the robot controller takes to change from one state to another. This is more apparent in neural network controllers undergoing no plastic change, where the neural elements do not have any intrinsic time constants differentially affecting their dynamics. Controllers in such cases are 'driven' Braitenberg-vehicle style by the sensorimotor coupling obtained while the robot is behaving. The notable exception to this trend is the growing body of research into oscillatory controllers of limbs for legged locomotion (Beer et al., 1992, Fujii et al., 2001), rhythmic arm movement (Williamson, 1998) and multi-segmented robot control (Ijspeert et al., 1998). There is, however, little research on the use of controllers producing autonomous and central dynamic patterns for tasks that are not necessarily rhythmic.

Data from neuroscience indicate that not only do many neuronal processes operate at a range of timescales significantly faster than that of behavioural performance, but also real neurons tend to adapt their firing properties homeostatically as a consequence of their recent history of activation (Turrigiano, 1999). Put together, these properties signify that it is increasingly difficult to think of single neuronal activation as playing the role of a 'record' of a relevant action-oriented event (the legacy of a computational view of cognition), since for many classes of neurons firing patterns will tend to rapidly lose any transient information at a timescale which is faster than that of behavioural performance. In contrast, global variables such as relative phase between neuronal assemblies and mode of oscillation become more relevant (Rodriguez et al., 2001, Varela et al., 2001). Even though real nervous systems can get around these problems by resorting to different means (e.g. slower processes, bistable neurons, or synergistic effects between interacting neuronal assemblies), the question of what sort of performance would obtain if only fast oscillatory neurons with homeostatic average activation were used becomes one of conceptual and practical interest for evolutionary robotics.

This paper explores some of the issues that arise from evolving neural controllers which operate under these conditions, i.e. at timescales faster than that of performance, and whose elements tend to compensate for long term patterns of activation 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 such individual elements. In order to facilitate understanding of the results and comparative analysis (and for other reasons explained below) a simple task of phototaxis is chosen. The next section further discusses the conceptual and practical motivations of this work. Section 3 describes the experimental setup and the neuron model which is a simple extension of a continuous-time network architecture. The results are presented in section 4 which shows that evolved controllers are highly robust to radical sensor perturbations such as exchange of sensor position and removal of one sensor. For every single case studied robots were able to perform the desired task as long as they had at least one sensor in the frontal half of the body. Robustness decreases as the allowed timescale of oscillation is made closer to that of performance. An analysis of the evolved strategy is also presented in this section. It is suggested that fast oscillations are not sufficient for robustness but that long-term homeostatic behaviour of neural activation is also necessary. This claim is supported by evolving a network of fast *non*-homeostatic FitzHugh-Nagumo oscillators which turn out to be much less robust. The final section discusses the implications of these results.

2. Motivations

This is an exploratory piece of work aiming at generating hypotheses. The motivations are conceptual as well as practical.

An animal nervous system is a complex network of relational patterns of electrochemical activity which is coupled with the rest of the organism and its medium through its sensorimotor surfaces. Neural dynamics present spontaneous extended patterns, characterized by their own measures of coherence. Such patterns are constantly happening in real brains. One of the legacies of the information processing perspective on cognition, typical of cognitivism but which can be traced to the behaviourist movement it purported to replace, is that such inherent brain dynamics are driven by sensory input and drive motor output in an almost linear causal manner. This view has had its strong opponents in different periods: Dewey (1896) and Merleau-Ponty (1942/1963) in their respective critiques of the concept of the reflex arc and S-R explanations in psychology, proponents of central views on neural population activity (Skarda and Freeman, 1987) and, more recently, alternative, dynamical views of cognition (Beer, 1995, Varela et al., 1991). According to such views, dynamic patterns in the brain are autonomous and sensory input is best understood as parametric perturbation or modulation of the intrinsic dynamics (Maturana and Varela, 1980). Technological improvements that permit non-invasive studies of brain patterns in behaving animals, as well as complex modelling techniques have increased the appeal of this point of view, not only in neuroscience (for instance in research into cell assemblies and large scale synchronization, Varela et al., 2001), but also amongst roboticists sympathetic with the dynamical approach.

Still, the latter have not been able to shake off entirely the computational legacy. This is a point that is difficult to defend, since often a dynamical perspective is outspokenly argued for by researchers in autonomous robotics. But even in such cases (the present author included) there is a tendency to think of a robot controller mostly as a task-driven device rather than as an autonomous system (as the very word "controller" implies)[1]. One reason for this is that under certain circumstances descriptions in terms of computations (e.g., nodes in a neural network representing or keeping track of action-relevant information) and dynamics seem to fit equally well to the case at hand. Such is the case of Brainteberg-like or behaviour-based architectures (Braitenberg, 1984, Brooks, 1986) which demonstrate just how interesting performance can be achieved by exploiting relevant environmental couplings but are limited to mostly reactive behaviour.

There is some justification, therefore, for the project of trying to purposefully devise robot controllers in which the contrast between a computational and a dynamical approach is maximal. Such a project would help illuminate the degree or kind of difference between these points of view in very concrete terms. Of course, notable examples already exists, typically when controllers (often neural networks) exhibit an explicit variety of spatial and temporal scales (Husbands et al., 1998). In those cases a dynamic interpretation of how they work is typically easier to envision. One of the aims of this paper is to explore such a contrast in a simple setup.

Real nervous systems operate at a large range of timescales. The faster end of this range can roughly be described as the 1/10, and 1 timescales. The fast 1/10 timescale include intrinsic cellular rhythms (e.g., in bursting interneurons) ranging from 10 to 100 milliseconds. The 1 timescale corresponds to durations of the order of a few seconds, roughly the timescale of the formation of large-scale coherent assemblies, perceptual integration and behavioural acts. Slower timescales are also active for instance in activity-dependent regulation of intrinsic neural properties and modifications to synaptic efficacy.

---

[1] The issue of organismic autonomy, clearly, extends beyond endogenous dynamic patterns in the nervous system, not only to other physiological processes, but also to behavioural and bodily intentional dimensions.

A large part of current work in understanding central pattern-generating circuits (CPGs) is focused on their role in the generation of rhythmic behaviour such as locomotion and respiration (Marder and Bucher, 2001). This is also true in robotics (Beer et al., 1992, Fujii et al., 2001, Ijspeert et al., 1998, Williamson, 1998). Rhythmic neural activity (not necessarily associated with CPGs) may also be involved in the generation of patterns of behaviour or perception that are non-rhythmic and happen at significantly longer timescales that those of oscillations (Rodriguez et al., 2001). This aspect has been less explored but it should be of considerable practical interest in robotics. If a system is synthesized to produce a large scale pattern with a typical timescale which is much longer than the timescale of its micro-components, then certain degree of robustness of performance should be expected, as, by design, no single micro-component can take a large share in the control of the overall system – the faster micro-timescale would not allow this – and so the system must make use of long range synergies that tend to be highly robust. Similar phenomena have been demonstrated in different contexts, (Di Paolo, 2001, Thompson, 1996) but apparently has not been applied in robotics so far.

Whether such robustness could also happen in robots is one of the main angles of investigation of this work. For this purpose, a task that is not intrinsically rhythmic has been chosen deliberately. Phototactic behaviour in a wheeled robot can easily be achieved by reactive, non-oscillatory controllers exhibiting the same internal and behavioural timescale. Thus phototaxis provides an ideal initial testbed for the exploring homeostatic oscillatory neurons as building blocks not only because of ease of analysis but also because of the maximum contrast in comparison with reactive controllers.

Tied to the issue of robustness is the issue of stability. Previous investigations have addressed the importance and benefits of balancing stability and change when trying to design plastic robot controllers. Based on Ashby's idea of ultrastability (Ashby, 1960), adaptation to radical sensorimotor perturbation has been obtained when single neural elements in the controller are evolved with the further requirement of long term homeostatic behaviour (Di Paolo, 2000). These biologically-inspired neural controllers (Turrigiano, 1999) have been evolved in order to maintain their elements within allowed ranges of activation while the robot performs its task. Local plasticity is induced by unstable elements until a new stable state is reached. This process can lead to sensorimotor adaptation to perturbations in sensor and motor gains and inversion of the visual field. The whole process operates at a slow timescale (of the order of 10 to 100 times the timescale of a behavioural event). Of the many questions that derived from this piece of work, one is the issue of what sort of behaviour would be obtained from evolving neural elements that were intrinsically homeostatic. Although in the present context controllers will not be plastic, the relation between long-term homeostasis of components relates similarly to the potential robustness of the whole system as the use of fast timescales. Neurons that always compensate for input currents will not be able to keep long-term records of any kind, and so this forces the evolution of strategies based on global patterns of activation, and so, hopefully, robustness and stability of performance.

3. Methods

3.1 *Homeostatic relaxation oscillators*

A simple modification is made to the equations governing the activation of continuous-time recurrent neural networks (CTRNN) (Beer, 1990) to transform each neuron into a relaxation oscillator which maintains a constant average activation over long periods. The equations describing traditional CTRNNs are:

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

where, using terms derived from an analogy with real neurons, y_i represents the membrane potential, τ_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 sensory node (modelled here as an incoming current, this term is always 0 for the other nodes). These equations are extended by turning the bias in each neuron into a responsive variable that "keeps track" of the opposite value of the membrane potential with a τ_i^b which greater than τ_i:

$$\tau_i^b \frac{db_i}{dt} = -(b_i + y_i); \quad \tau_i^b = \tau_i G_i;$$

where G_i ranges from 1 to T and is genetically set. The response variable always has a longer time constant than the membrane potential. 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. This modification transform each node in a special case of a centre-crossing dynamical neuron (Mathayomchan and Beer, submitted).

3.2 *Experimental setup*

Simulated robots are evolved to perform phototactic behaviour on a series of light sources. Robots are modelled as solid circular bodies of radius $R_0 = 4$ with two diametrically opposed motors and two light sensors. The angle between sensors is always of 120 degrees (60 degrees each to the body central midline).

Motors can drive the robot backwards and forwards in a 2-D unlimited arena. Robots have a very small mass, so that the motor output is the tangential velocity at the point of the body where the motor is located. The translational movement of the whole robot is calculated using the velocity of its center of mass (the vectorial average of the motor velocities), and the rotational movement by calculating the angular speed (the difference of the tangential velocities divided by the body diameter). There is no inertial resistance to either form of movement.

Light from point sources impinges on sensors with a local intensity proportional to the source intensity and the inverse square of the distance from sensor to source. The model includes shadows on sensors produced when light is occluded by the body (i.e., a sensor angle of acceptance of 180 degrees). Input current from each sensor are fed to a corresponding neuron.

Each motor is controlled by the firing rate of a single motor neuron (mapped so as to produce both positive and negative values). Alternatively, each motor can be controlled antagonistically by two neurons whose firing rates are summed with opposite signs. Results were the same in both cases and will be presented only for the former. All neurons are fully connected. Tests were run with 4, 6, 8, 10 and 20 neurons in the network with similar results.

Motor and sensor signals are multiplied by a genetically set gain which is left-right symmetrical and taken from the range [0.1,10]. Uniform noise with range 0.1 is applied to both input and output signal priors to gain scaling.

Each time step the network is updated using an Euler integration method with a time step of 0.2.

A population of 30 robots is evolved using a generational genetic algorithm with truncation selection. The population is ranked according to fitness, the top third is copied to the next generation and used as parents for individuals that replace the bottom two thirds of the population. All parameters are encoded in a real-valued vector, each component encoding a single parameter in the neural network. These components belong to the range [0,1] and are linearly scaled to the range corresponding to the parameter encoded (with the exception of gain values which are exponentially scaled). A global mutation operator is used (Beer, 1996) which adds a small random vector to the whole genotype each time a mutation occurs. The mutation rate is of 2% per locus and the standard deviation of vector displacement is of 0.25. Recombination was not used.

Each individual robot is run for 4 independent evaluations. Fitness is calculated by averaging the fitness obtained in each evaluation minus one standard deviation to favour low variability. Each evaluation consists of the sequential presentation of 6 distant light sources which the robot must approach in turn and remain close to. 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]. After T_S the source is extinguished and another one appears at a random distance and angle. The initial distance between robot and source is uniformly chosen from the interval [60,80], the angle from [0,2π) and the intensity of each source from the interval [500,1500], so that, on average, the intensity of light at the initial stages is of the same order of magnitude as the noise on the sensors.

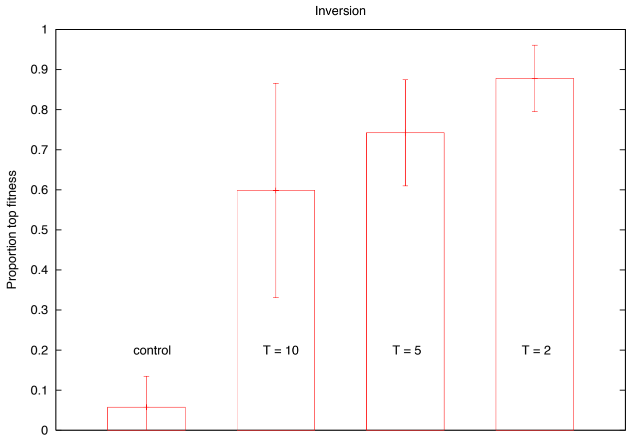

Figure 1: Average relative robustness (measured as proportion of unperturbed performance) for left-right inversion of sensors and different values of T.

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 $f = 0$.

The parameter T indicating the range of allowed values for τ_i^b was set for different series of runs at 2, 5, and 10 in order to study the effect of increasingly longer timescales of oscillation.

4. Results

4.1 Robustness

In all cases the genetic algorithm was run for 2000 generations; fitness values stabilizing at around 1000 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 performance over 10 trials of 5 independently evolved robots over a series of 50 light sources. Robots are evolved that approach light sources equally well –

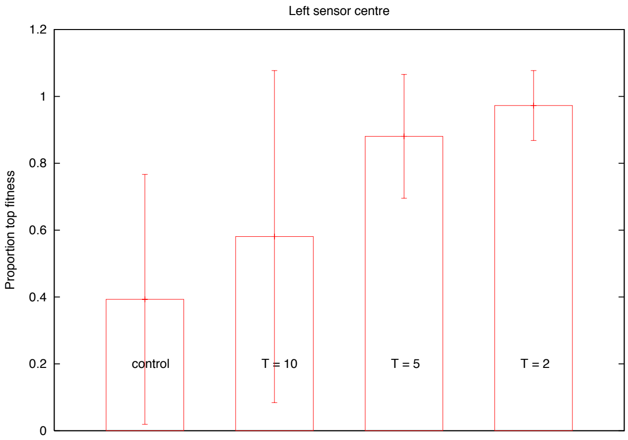

Figure 2: Average relative robustness (measured as proportion of unperturbed performance) for removal of right sensor.

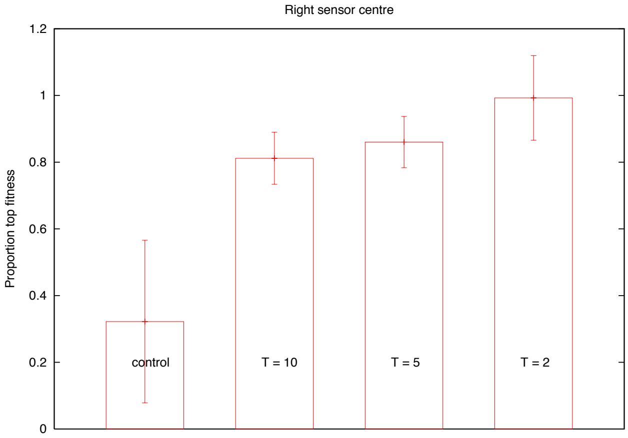

Figure 3: Average relative robustness (measured as proportion of unperturbed performance) for removal of left sensor.

the range of final distances to the source was the same in all cases (between 3 and 10 units). However, because the fitness benefits fast trajectories absolute performance was significantly better for the control runs (0.56 ± 0.06) than for the oscillatory controllers (0.35 ± 0.04 for $T = 2$, 0.42 ± 0.03 for T = 5 and 0.35 ± 0.08 for $T = 10$) which take a longer time to reach the source.

Figures 1, 2 and 3 show the proportional decay in performance for a series of radical sensor perturbations. Inversion consists in swapping the sensor positions left and right. Sensor removal consists in making the corresponding incoming current equal to 0 and placing the other sensor at an equal angle between the motors (central midline). Error bars indicate standard deviation. In all cases the leftmost bar corresponds to control runs. For all three sensor perturbations oscillating neural controllers show high robustness which increases as T is

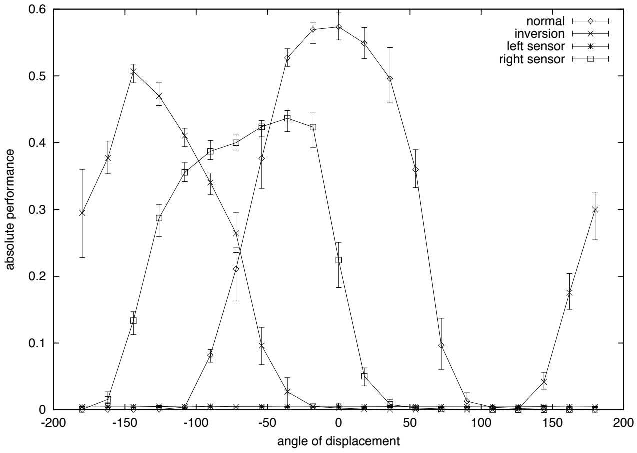

Figure 4: Absolute performance vs. angular sensor displacement for a control individual.

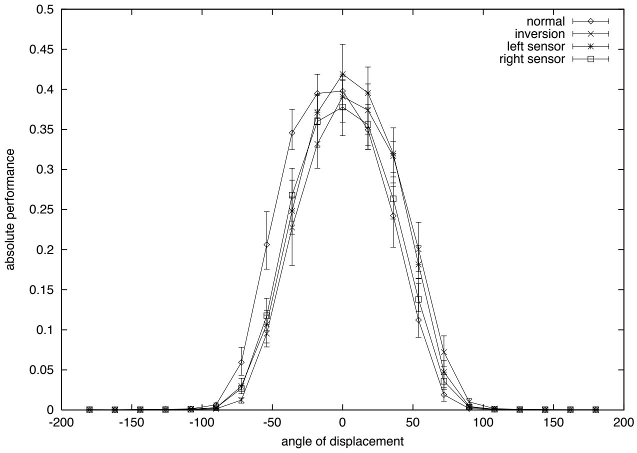

Figure 5: Absolute performance vs. angular sensor displacement for individual evolved using fast homeostatic oscillators.

made smaller, i.e., when the range of allowed oscillations is constrained to higher frequencies. Observed behaviour is unchanged by these radical perturbations for $T = 2$. On application of the sensor distortion, the robot continues moving along its trajectory as if the perturbation had not been applied. This justifies the description in terms of robustness rather than adaptation since no maladapted period has ever been observed in any of the trials.

Robustness against sensor perturbation can also be appreciated by studying the dependence of absolute performance on angular sensor displacement along the circular body of the robot. In these tests, the position of both sensors is shifted by a same angle between -180 and 180 degrees. Figures 4 and 5 show this relation for the unperturbed robot and for the three sensor perturbations described above for the control case and an

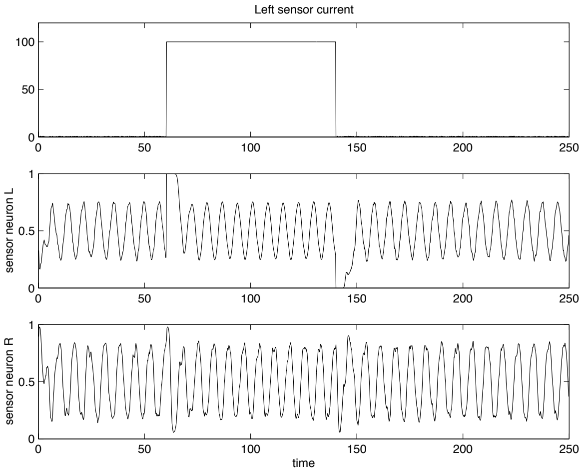

Figure 6: Compensatory behaviour of sensor neurons to step changes in left sensor current. Similar behaviour for steps in right sensor current.

oscillator-controlled robot evolved with $T = 2$ respectively. The control robot shows 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 practically the same angular distribution for all 4 cases, indicating that the strategy used by this robot will work as long as there is at least one sensor (any sensor) facing in the forward direction of the robot.

4.2 Analysis

Analysis of the evolved strategy has been performed for a 4-neuron network. The strategy itself and the behaviour of single neurons do not differ significantly from 6-neuron controllers.

All neurons behave as oscillators entrained to a common frequency of 0.11 cycles per unit of time. If all but self-connections are removed, neurons oscillate at their proper frequencies (0.094 for left motor neuron, 0.119 for right motor neuron, 0.211 for left sensor neuron and 0.237 for right sensor neuron). Homeostatic compensation to step changes in input current is shown in figure 6. At the start of an step, sensor neurons change their oscillation pattern but for no longer than one cycle. The perturbation is compensated for and an opposite effect obtains when an opposite step is applied. The effects on the other neurons is similar in duration and smaller in scale. Robot trajectories are affected only temporarily by such step functions and the way they are modified depend on precise timing of the onset of the step.

In all the cases tested the strategy used by the robot is a variation of the strategy shown in figures 7 and 8. Figure 7 corresponds to the robot behaving in the absence of light sources. The top right plot shows the two motor neurons. These oscillate with a same frequency and maintaining a roughly constant phase difference which means that for certain periods they will be both at their top range or both at their bottom range, while the rest of the time they will be at alternate high and low values of activations. These top and bottom values correspond to the motors moving forwards and backwards respectively. The trajectory shown at the top and bottom left in figure 7 can then be described as alternated segments where the robot moves forwards, rotates on the spot about three quarters of a full turn, moves backwards (but keeping in roughly the same general direction), rotates on the spot, moves forwards, and so on. Rotation points are marked in the figure. Because of the imposed limitations on the neural controller, the robot must use an active scanning strategy for phototaxis, as sensor neurons will lose any instantaneous information they might acquire from the input currents. Sensor neurons also oscillate in anti-phase (bottom left).

Large-angle on-the-spot rotation provides the basis for the active strategy. Figure 8 reveals more detail of behaviour in the presence of a light source. As action-relevant information cannot be maintained in levels of single neuron activation (neither in the short or long term) the neural controller resorts to using input currents to modulate intrinsic oscillations and make the network switch between oscillating regimes. This becomes clear by observing the behaviour of the sensor neurons (figure 8 bottom right) which show a change of regime from low amplitude oscillations to rapid “flip-flop” behaviour covering the whole range of activation when the robot is near the source of light (distance to the source is shown in the bottom left plot). The oscillations of sensor neurons turn out to be unstable in the presence of the intermittent pattern of input currents caused by rotating near the light source. The new regime in turn modulates the relative phase of the motor neurons. The same transitions can be observed by estimating the relative phase of the sensor neurons and plotting it in relation to periods when the robot is far away from a source of light or near it. Figure 9 shows precisely this; the line indicates when the robot is further (top) or closer (bottom) than 20 distance units. The phase relation changes from a very noisy anti-phase to a highly ordered state where only a few discrete values are taken when the robot is near the source.

As mentioned earlier, the effect of step-wise increments or decrements in input currents is concentrated in the transients after the onset of each step and their effect on the robot trajectory depends on their precise timing (which depends obviously on the robot’s current vector). The complex relation that maintains the robot close to the light source has yet to be fully elucidated. Closer analysis of the transition stage between the near

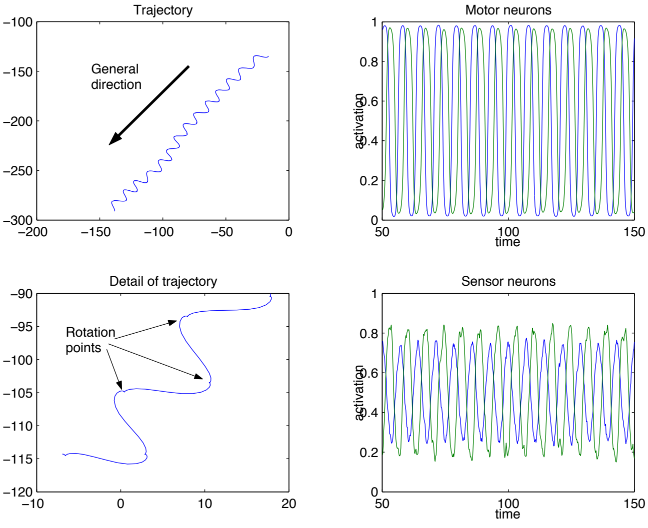

Figure 7: Robot behaviour and neural activation in the absence of light sources.

and far regimes throws some light into the approaching strategy and the origins of robust behaviour. Even though very low input currents do not alter the oscillation mode of the neurons, they do modulate the phase of oscillation. The effect of this phase modulation seems to correlate 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 the 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. Interestingly, 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.

4.3 Fast oscillations are not enough

The above results show that when single neurons are prevented from storing long-term information in their activation values, good solutions can nevertheless evolve that make use of the relative coordination between neural oscillators to solve the desired task. Additionally, these solutions tend to be 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 (otherwise transient information in the activation could be used) and the long-term average activation of each neuron must conserve an undifferentiated average value independently of the history of inputs.

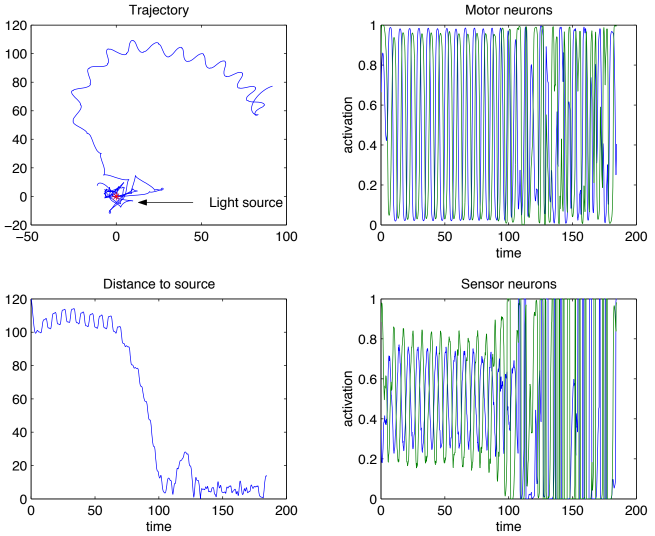

Figure 8: Robot behaviour, neural activation and distance to light source during phototaxis.

Figures 1, 2 and 3 show that when the timescale of oscillation is allowed to approach that of behaviour, robustness decreases, thus demonstrating that robustness relies on fast oscillations. The necessity of long-term homeostatic and uniform activation remains to be shown. In order to do this, two series of experiments using non-homeostatic fast relaxation oscillators was carried out.

In the first series, the modified CTRNN equations were slightly altered. The second equation now reads:

$$\tau_i^b \frac{db_i}{dt} = -(b_i + K_i y_i);$$

with K_i a genetically determined constant within [0,1] that specifies for each neuron the proportion of recovery for b_i. When K_i is near 1 the neuron will maintain an average activation in the middle range. At the other extreme ($K_i = 0$), b_i will approach zero and remain there in the long term. Neurons can differentiate in terms of the mean value of their oscillations in this case.

In the second series the experimental set-up was maintained with the exception of the neural activation which is now governed by a set of FitzHugh-Nagumo equations suitable coupled (Kanamaru and Okabe, 2000)

$$\tau_i \frac{du_i}{dt} = -v_i + u_i - u_i^3/3 + \sum_j w_{ji} u_j + I_i;$$

$$\frac{dv_i}{dt} = u_i - \beta_i v_i + \gamma_i;$$

where u_i is the fast variable representing the activation of the neuron and v_i is the response variable. β_i and γ_i are genetically set with values taken from the range [0.8,1.2], the range for weight values is [-1,1] and for τ_i [0.4,4.0].

In all of these cases (5 runs in each series) high fitness scores were rapidly evolved (e.g., 0.4 ± 0.04 for the F-N

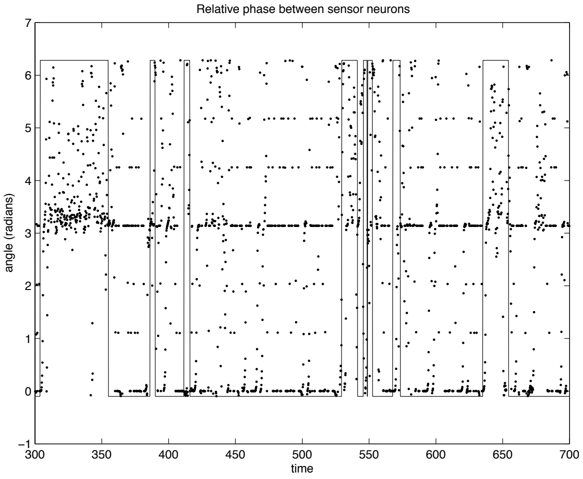

Figure 9: Relative phase of oscillation between the two sensor neurons during phototaxis. The solid line indicates when the distance to the source is more (up) or less (down) than 20 units. Highly ordered modes are observed when the robot is near the source.

network), but in none of them robustness against sensor perturbation was obtained (e.g., proportion of original fitness under inversion: 0.008 ± 0.004 for the F-N network). This result confirm that fast oscillations are not sufficient for robustness and gives support to the hypothesis that undifferentiated homeostatic behaviour is necessary. Obviously further analysis and tests will be required for understanding why this seems to be the case.

5. Conclusions: A hypothesis

By constraining the class of neural controllers that evolution can make use of, the solution space is, in principle, similarly constrained. Such a move would appear as contrary to the objective of enhancing evolvability that is embraced by a good part of research in evolutionary robotics. Yet the quality of solutions so constrained is significantly different and interesting to deserve more exploration. This work is part of a more general trend towards using synthetic techniques not just to arrive at more interesting robot behaviour, but to understand what are the global effects of biologically inspired micro-mechanisms when these are implemented in controllers for well understood tasks. Examples of this trend include visual shape discrimination using diffusible neuromodulators (Husbands et al., 1998), navigation and obstacle avoidance using spiking neurons (Floreano and Mattiussi, 2001), and phototaxis using homeostatic regulation of local plasticity (Di Paolo, 2000). The design and study of novel integrated systems of this sort may well be one way for evolutionary robotics to contribute useful information back to biology in the proximate future.

The above results show that significant and consistent robustness is obtained by introducing certain changes in neural controllers. These changes, at first sight, appear unintuitive and yet they find their motivation in biological data both on neural mechanisms and the role of central patterns of nervous activity. The combination of oscillatory and compensatory dynamics at the level of the individual neuron means that a controller will not be able to rely on 'informational traces' at this level, and so, much less on computation as a design strategy. In contrast, a dynamical perspective come into focus as the natural operational and explanatory option for such systems. The impossibility of functional specificity at the neuronal level seems the most likely source of robustness at the performance level. When functional specificity is allowed, either by slowing down the permitted periods of oscillations, or allowing differences in long-term average activation, robustness disappears, as is also absent in the non-oscillatory non-homeostatic control runs.

The implications of these results are not just conceptual. On the practical side of things, it is of extreme interest to understand the origins of adaptivity, robustness, and stability in robot controllers and other complex systems. The results presented here suggest that when dealing with multi-component systems, the less the causal and explanatory burden is concentrated at the micro-level, the more the whole system is forced into globalised and de-centralised solutions. Such solutions, by definition, can be expected to be more robust than solutions that depend on causal specificity at the level of single components. Removing the capability of artificial neurons to assume specific functional roles according to their state is one way of doing this. Evolution is constrained to find solutions that integrate central patterns with sensorimotor couplings in ways that the desired behaviour is achieved. This should not be taken as implying that real nervous systems achieve robustness in the same way.

In view of this conclusion it is possible to advance the following hypothesis: *In complex multi-component systems, robustness will be likely to be obtained if functional specification at the level of individual components is minimized.* The only course left for such systems is to achieve the desired performance by means of large-scale coordination between components. Functionality is obtained at a level which is different from that of the lower mechanisms. Such systems are difficult to design. Constraints on performance and robustness do not translate cleanly into obvious mesoscopic synergies. The most promising design strategy for such systems remains, therefore, a synthetic approach such as evolutionary robotics.

It is important to emphasize that this hypothesis presents a one way link. Robust performance may well result from different design principles. Further investigation into similarly constrained architectures is needed to test the usefulness and limitations of this idea and explore its relation to other not-so-distant issues such as plasticity and adaptivity.

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

References

Ashby, W. R. (1960). *Design for a Brain: The Origin of Adaptive Behaviour.* London: Chapman and Hall, second edition.

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

Beer, R. D. (1995). Computational and dynamical languages for autonomous agents. In Port, R. and van Gelder, T., (Eds.), *Mind as Motion: Explorations in the Dynamics of Cognition*, pages 121–147. MIT Press, Cambridge, Mass.

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.

Beer, R. D., Chiel, H. J., Quinn, R. D., and Espenschied, K. S. (1992). A distributed neural network architecture for hexapod robot locomotion. *Neural Computation*, **4**:229–230.

Braitenberg, V. (1984). *Vehicles: experiments in synthetic psychology.* Cambridge, MA: MIT Press.

Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE, J. Rob. Autom.*, **2**:14–23.

Dewey, J. (1896). The reflex-arc concept in psychology. *Psychol. Review*, **3**:357 – 370.

Di Paolo, E. A. (2000). Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions. In Meyer, J.-A., Berthoz, A., Floreano, D., Roitblat, H., and Wilson, S., (Eds.), *From Animals to Animats 6: Proc. 6th Int'l Conf. on the Simulation of Adaptive Behavior*, Paris, France. Cambridge MA: MIT Press.

Di Paolo, E. A. (2001). Rhythmic and non-rhythmic attractors in asynchronous random Boolean networks. *BioSystems*, **59**:185–195.

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.

Floreano, D. and Urzelai, J. (2000). Evolutionary robots with on-line self-organization and behavioral fitness. *Neural Networks*, **13**:431 – 443.

Fujii, A., Ishiguro, A., Aoki, T., and Eggenberger, P. (2001). Evolving bipedal locomotion with dynamically-rearranging neural networks. In Kelemen, J. and Sosík, P., (Eds.), *Advances in Artificial Life: Proceedings of the Sixth European Conference on Artificial Life*, pages 290 – 298. Springer Verlag.

Harvey, I. (1997). Cognition is not computation; evolution is not optimisation. In Gerstner, W., Germond, A., Hasler, M., and Nicoud, J.-D., (Eds.), *Proceedings of the 7th International Conference on Artificial Neural Networks, ICANN'97.* 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.

Ijspeert, A. J., Hallam, J., and Willshaw, D. (1998). From lampreys to salamanders: Evolving neural controllers for swimming and walking. In Pfeifer, R., Blumberg, B., Meyer, J.-A., and Wilson, S., (Eds.), *From Animals to Animats 5: Proc. 5th Int'l Conf. on the Simulation of Adaptive Behavior*, Zürich, Switzerland. Cambridge MA: MIT Press.

Kanamaru, T. and Okabe, Y. (2000). Associative memory retrieval induced by fluctuations in a pulsed neural network. *Phys. Rev. E*, **62**:2629 – 2635.

Marder, E. and Bucher, D. (2001). Central pattern generators and the control of rhythmic movements. *Current Biology*, **11**:R986–R996.

Mathayomchan, B. and Beer, R. D. (submitted). Center-crossing recurrent neural networks for the evolution of rhythmic behavior. *Neural Computation.*

Maturana, H. and Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living.* Dordrecht, Holland: D. Reidel Publishing.

Merleau-Ponty, M. (1942/1963). *The structure of behaviour.* London: Methuen. Translated by A. L. Fisher.

Rodriguez, E., George, N., Lachaux, J.-P., Matinerie, J. Reanault, B., and Varela, F. J. (2001). Perception's shadow: long-distance synchronization of human brain activity. *Nature*, **397**:430–3.

Skarda, C. E. and Freeman, W. J. (1987). How brains make chaos in order to makes sense of the world. *Behavioral and Brain Sciences*, **10**:161 – 195.

Thompson, A. (1996). Silicon evolution. In Koza, J. R. and et al., (Eds.), *Proceedings of Genetic Programming 1996*, pages 444 – 452. MIT Press, Cambridge, Mass.

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

Varela, F. J., Lachaux, J.-P., Rodriguez, E., and Matinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nature Reviews Neuroscience*, **2**:229–230.

Varela, F. J., Thompson, E., and Rosch, E. (1991). *The embodied mind: Cognitive science and human experience.* Cambridge, MA: MIT Press.

Williamson, M. M. (1998). Neural control of rhythmic arm movement. *Neural Networks*, **11**:1379–1394.

Yamauchi, B. and Beer, R. D. (1994). Sequential behavior and learning in evolved dynamical neural networks. *Adaptive Behavior*, **2**:219–246.