

Adapting to Your Body

Peter Fine, Ezequiel Di Paolo and Eduardo Izquierdo

Centre for Computational Neuroscience and Robotics (CCNR)
University of Sussex, Brighton, BN1 9QH, UK
{p.a.fine,ezequiel,e.j.izquierdo}@sussex.ac.uk

Abstract. This paper investigates the processes used by an evolved, embodied simulated agent to adapt to large disruptive changes in its sensor morphology, whilst maintaining performance in a phototaxis task. By avoiding the imposition of separate mechanisms for the fast sensorimotor dynamics and the relatively slow adaptive processes, we are able to comment on the forms of adaptivity which emerge within our Evolutionary Robotics framework. This brings about interesting notions regarding the relationship between different timescales. We examine the dynamics of the network and find different reactive behaviours depending on the agent’s current sensor configuration, but are only able to begin to explain the dynamics of the transitions between these states with reference to variables which exist in the agent’s environment, as well as within its neural network ‘brain’.

1 Introduction

The notion that perception and action are tightly intertwined concepts is a common theme throughout Artificial Life and Evolutionary Robotics research. O’Regan & Noë [8] outline the idea that perception is not a static, feed-forward means to achieve the representation of the world in the brain, but rather an ongoing process in which action plays an equally important part as perception. Additionally, adaptivity is seen as a crucial component of behaviour, since resilience to changing circumstances can be regarded as one of the basic tenets of life. However, a degree of tension can arise between paradigms aligned towards rapid reactivity through tight sensorimotor coupling, and the relatively slow-timescaled dynamics attributed to the reconfiguration of a system in response to the changing environment.

It is tempting to take a divide-and-conquer strategy to account for the differing requirements of these fast and slow aspects of behaviour. This can occur at a mechanistic level, with the addition of separate ‘adaptive’ processes which sit on top of a sensorimotor system (e.g. [5], where plastic synapses are used to enable an agent to switch between different lights during phototaxis). It may also occur via the imposition of a structural modularity (e.g. [12], where separately evolved networks account for the reactive and the adaptive components of a memory based task).

These approaches can certainly form useful models of adaptive behaviour, since biological organisms exhibit distinct phenomena such as spike time dependent plasticity [2] or slowly diffusing neurotransmitters [6] which warrant investigation. In [4] synaptic plasticity was exploited in an Evolutionary Robotics context to produce an agent robust to sensory inversion. These mechanisms appear to be somewhat dedicated to the slower timescale dynamics found in adaptive behaviour. The study of different, abstracted classes of adaptive mechanism can inform our understanding of what to look for in nature and how to build better engineered systems. However it is also beneficial to investigate the emergence of systems with interacting timescales. Evolutionary Robotics operating on vanilla, dynamical recurrent neural networks (with no additional mechanisms which are designed to account for the relatively slow, adaptive properties of behaviour) can provide a useful, relatively assumption-free paradigm within which to do this (e.g. [3], in which a robot’s locomotion exploits a model of its body which is automatically constructed and refined through a predictive algorithm).

Previous work in Evolutionary Robotics has investigated the potential for fixed-weight networks to display the adaptive functionality more commonly attributed to dedicated weight-changing mechanisms (e.g. [11]). [9] provides a recent discussion and example of this approach. *Our work here differs from this and other attempts in that the environment itself does not change, but rather the agent’s own morphology is what must be adapted to.* Instead of adapting a sensorimotor law to suit a disrupted environment, it is the sensorimotor laws used to control the agent which become disrupted. This occurs when the agent’s body is altered such that the relationships and invariances between the sensory and motor activity change.

[10] offers a mathematical framework based on the sensorimotor contingency principles outlined in [8]. They provide an algorithm which is able to extract ‘laws’ allowing the coordinated control of unknown, arbitrary and high-dimensional simulated robotic systems. It provides a useful proof of concept, demonstrating the potential for sensorimotor systems to discover the structure of their environment and morphology despite knowing nothing about these *a priori*. This paper intends to investigate dynamical approaches to a related self-discovery task, albeit in minimally complex way.

Below, we present an evolved model agent, which is able to perform a sensorimotor task despite a radical perturbation to its morphology. This is analysed in section 4, where we see that the agent exploits its environment to overcome failure imposed by disruption to its sensor configuration.

2 Methods

We carried out experiments using a simulated agent required to perform phototaxis towards a sequence of lights. To achieve this, it was fitted with a single, simple light sensor. By switching the location of the light sensor, however, a single set of sensorimotor relationships is unlikely to be sufficient to always complete the task (as evidenced below).




Fig. 1. A: The agent’s circular body, showing motors and range of angles at which light impinges on the sensors. B: The neural network configuration, fully-connected with one input neuron (no. 1), an interneuron (2) and two motor neurons (3 & 4). The parameters for the best evolved network (see section 4) are depicted. The neuron size reflects time constants, neuron shade depicts the bias (dark is high), weight thickness is its magnitude, and weight colour is dark for excitatory, grey for inhibitory.

The experiment consists of a circular agent of radius 1, which exists in a 2D plane (see Figure 1A). Distance and time units are of an arbitrary scale. It is controlled by a motor on each side, which differentially steer the robot with their output (in range $[0, 1]$). A light sensor is attached to either the front or the back of the robot at any one time. It outputs a value in $[0, 1]$ indicating closeness to a point light source, by linearly scaling the distance to that light (γ), thresholded such that $\gamma > 22.5$ becomes 0 (22.5 is 1.5x the maximum γ when each light is first presented). The sensor only receives input if the light is within the 180° in front (or behind) the robot, centred on the light sensor, since its circular body occludes the light.

Lights are presented sequentially, located at a random angle and distance in $[10, 15]$ from the agent. After a variable (in $[40, 60]$) time period, a new light is generated. Less frequently, the sensor is switched from the front to the back of the agent (or vice-versa) at the start of a light presentation. The agent is controlled by a continuous time recurrent neural network (CTRNN), shown in Figure 1B and governed by equation (1) (based on [1]).

$$\tau_i \dot{y}_i = -y_i + \sum_{j=i}^N w_{ij} \sigma(y_j + \theta_j) + sI(t) \quad , \quad \sigma(x) = 1/(1 + e^{-x}). \quad (1)$$

The state (y) of the i^{th} neuron (out of N) is a function of the time constant τ (scaled exponentially in $[1, e^5]$), the weight w_{ij} in $[-10, 10]$ from incoming neurons, their state y_j and bias θ_j (subject to a sigmoid function). The sensor weight s in $[-10, 10]$ is applied to the input I from either the front or back light sensor (whichever is enabled at a given time). The network (and other simulation variables) are integrated with an Euler timestep of 0.1 during optimisation of the agent’s controller, and 0.01 during analysis to ensure stability.

These network parameters were optimised using a genetic algorithm with tournament selection, with a population of 50. The ‘losing’ genotype in a tournament is replaced with a copy of the winner, subject to mutation via the addition of a random number to each gene selected from a Gaussian distribution (mean 0, variance scaled between 0.01 and 0.05 as the best performing genotype of the previous generation’s fitness varies between 0.4 and 0.8). Each run progressed for 6000 generations (each consisting of 50 tournaments), with fitness calculated as the normalised average distance of the agent from the light during the last 25 time units of each light presentation. The first light after the sensor was switched does not contribute towards fitness, allowing for a possible adaptation phase to occur without punishing the agent.

The evolution progressed according to a shaping scheme consisting of several phases, since it would be unlikely that both phototactic and adaptive behaviour would evolve simultaneously. When the best agent of each of the 10 prior generations had attained a fitness greater than 0.8, the phase was advanced. Initially, the sensor was solely located on the front of the agent, so standard phototaxis was all that was required. For each trial, the agent was subject to 6 sequential light presentations, and the average of 8 of these trials taken as the overall fitness score for a given genotype. Phase 2 consisted of an additional set of 8 trials, with the sensor located on the back of the agent (which was not changed for the duration of each run). The overall fitness score was the mean of all 16 trials. Finally, a third phase proceeded as per phase 2, but with double the number of lights per trial (12). The sensor was switched to the opposite side of the agent after a random number of light presentations in the range [4, 8].

3 Results

16 runs (with different, independent random seeds) were carried out for each of the 0, 1 and 2 interneuron cases (giving a total of 3, 4 and 5 neurons for each). Of these, none of the 0 interneuron experiments succeeded in generating a successful phase 3 (re-adaptation to sensor switching) controllers. Two networks with 1 interneuron and one with 2 interneurons did however achieve a high level of fitness.

It is clear that the success was limited to only a small number of the evolutionary runs, and discovering why this task appears to be so difficult with the current experimental configuration could prove an interesting direction of study. However we concentrate here on the analysis of one of the best performing runs, since we are primarily concerned with the kind of dynamics which can satisfy the task, rather than evolutionary performance. The best agent of the most successful 1-interneuron run is thus considered from this point on. The progress of that evolutionary run can be seen in Figure 2. In 2000 trials performed on this agent, in which 25 lights were presented with a switch in sensor location after every 5 lights, the mean fitness score was 0.806 ($\sigma = 0.043$), indicating a reliable and adept controller.




Fig. 2. Evolutionary performance. Black is the best agent of each generation, grey is the average. Only the first 2000 generations (of 6000) are shown.

4 Analysis of a Successful Agent

In order to understand how the high fitness scores were achieved, an analysis of the agent from the run shown in Figure 2 was carried out. Firstly, the behaviour of the agent was examined, and then its internal neural dynamics studied.

4.1 Behavioural Performance

Figure 3 demonstrates the behaviour of the agent in a series of trials separately showing the four main demands which the task requires of this agent. The top two plots of Figure 3A show the trajectories generated after the agent has been allowed sufficient time to adapt to the two possible sensor configurations. It is clear that the agent is able to navigate towards the light reliably from different relative headings, and along reasonably un-convoluted paths. However it is also evident that it follows quite different routes depending on whether the sensor is on the front or back. This is clarified in Figure 3B, where the difference (left–right) of the two motor outputs (shown on the x axis) reflects the turning behaviour of the agent. A positive difference produces a left turn, near 0 causes forward motion, and negative values a right turn. It demonstrates that when the sensor is on the front, the agent only ever turns left or moves forward, whereas when on the back, turns in both directions are made.

The trajectories during re-adaptation (the lower two plots of Figure 3A) show more complex behaviour. In both cases (front→back adaptation and the converse), the light is not reached by the end of the first light presentation, suggesting that there is not (at least that evolution could find) a trivial, reactive behaviour permitting both front and back sensor phototaxis. In both cases, the agent at first moves away from the light, which can be considered a temporary maladaptive strategy, from which it recovers. In the front→back case, the agent is sometimes successful by the second light, but in others, requires three.

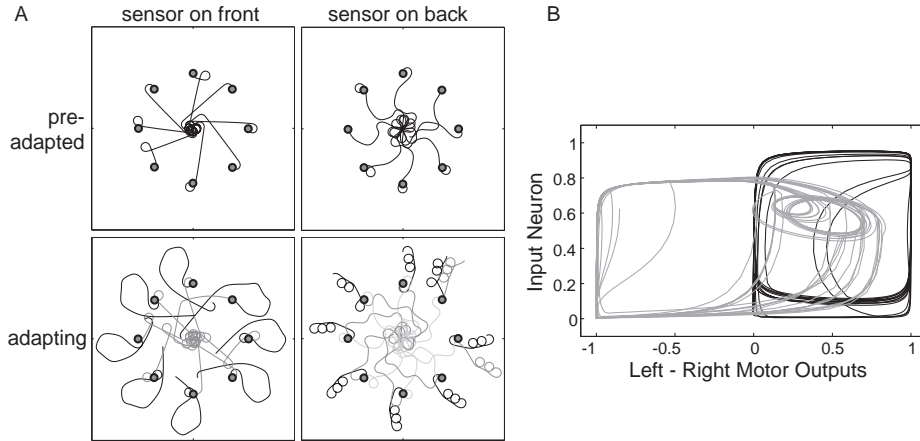


Fig. 3. A: The behaviour of the agent under different circumstances. In each, the agent was started at 8 different locations (marked with a small circle), with the light remaining in the same place (the centre of each figure). Each trajectory in the top two plots of A show the behaviour during a single light presentation of duration 50, with the agent previously being allowed to adapt to its sensor location for 3 successive random light presentations. The lower pair of plots demonstrate the behaviour after the sensor has been switched to the opposite side of the robot. The black line shows the first light presentation, immediately after the switch, whilst the lighter lines show the second and third (with the agent replaced at its original location), where these extra lights were required to complete adaptation. The agent’s initial heading at the start of a light presentation alternates between 0 and π on each subsequent presentation (during the above analysis only). B: shows the motor outputs against input neuron output for the same trajectories as the top two plots of A above, with the black line representing data taken with the sensor on the front, and the grey with the sensor on the back.

4.2 Neural Dynamics

The system under consideration is fully described by four internal continuous variables (node activities), in addition to one continuous and one discrete environmental variable (the distance to the light, and whether it is in sight at a given time, respectively). Since we cannot observe the full structure of this six dimensional hybrid dynamical system, we must carefully choose appropriate lower dimensional projections which adequately capture the coupled agent-environment dynamics. The trajectory plots above adequately demonstrate what the agent does, but not how it achieves its adaptivity. We therefore analyse the neural dynamics to further understand the behaviour of the agent.

Figure 4 shows the behaviour of the interneuron against that of the input neuron during the same four scenarios as Figure 3A. There are quite clearly two separate stable modes in the interneuron’s state space, with mean values of

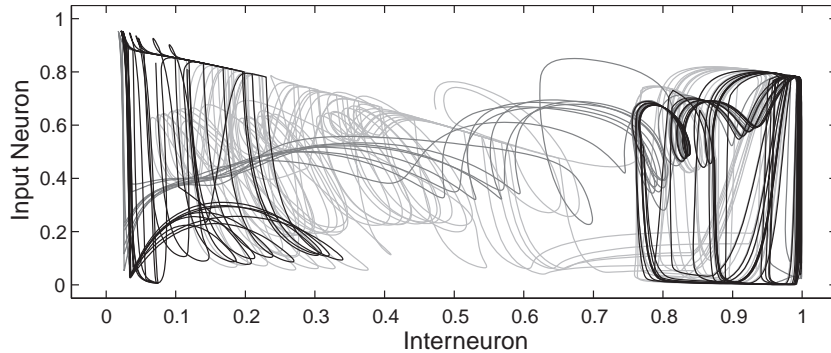


Fig. 4. Phase plot showing the output of the interneuron against that of the input neuron. Neural data from 8 different trials are superimposed. Black shows the behaviour after the agent has adapted to its sensor configuration, with the black trajectory on the left of the plot taken from a period with the sensor on the front, and the right-hand black path with the sensor on the back. Dark grey shows presentations of the light which immediately follow the switching of the sensor from back→front, and light grey represents switching from front→back (in the latter case two presentations are shown since adaptation takes multiple lights to complete). Note that the sensor weight is negative (-0.8), and therefore a high input neuron value is an indication of a low sensory input.

0.11 and 0.94 for the left and right black-lined areas in the diagram, respectively (see caption for details). Since these two bundles of black lines were separately recorded from the agent under each sensor configuration (whilst the agent undergoes successful phototaxis), we can be confident that they correspond to the two different stable sensorimotor behaviours exhibited (depending on that sensor location). In fact, if we fix the value of this interneuron to its mean value in each of these states, the performance of the agent is strongly biased. On setting to 0.11, the mean fitness (over 1000 ten light trials) with the sensor on the front is 0.8258 ($\sigma = 0.05$), whereas the mean fitness is 0 with the sensor on the back. Conversely, on setting the interneuron to 0.94, mean fitness is 0.8093 ($\sigma = 0.02$) with the sensor on the back, and 0 with the sensor on the front. However if we threshold the interneuron’s output so that it is binary (if greater than 0.5, the output is clamped to 0.94, otherwise it is clamped to 0.11), we find a drastic loss in the agent’s capability to adapt.

Another point to note is that the time constant parameters have evolved to within 0.05 of their fastest permitted value of 1 for all of the neurons except the interneuron, which took a larger value (6.6), which corresponds to a slower integrating neuron.

By outputting different values depending on the sensor’s location, the interneuron is able to push the network into a different region of motor dynamics, which produces the different behaviours demonstrated above. But this does not explain how the dynamics during the grey phases shown in Figure 4 make the




Fig. 5. The trajectory of the interneuron, against the agent’s distance from the light. The left case shows the back→front transition in grey, with the black line showing the dynamics prior to the sensor switch. The right-hand figure shows the front→back transient in a similar fashion. Each show 10 different simulations, superimposed. The arrows show the general trend that the dynamics follow, their direction indicating the passage of time since the start of each light presentation.

transition to the alternate state. In fact, for parts of the grey trajectories (after the sensor has switched), the input neuron appears to follow a similar regime to its activity prior to the switch. However if we consider what is happening to the agent’s relation with its environment (in terms of phototactic performance) during these ambiguous parts of the transient dynamics, the picture becomes clearer. In Figure 5 we consider an additional, environmental variable: the distance from the agent to the light.

In contrast to Figure 4 (which only shows neural outputs), Figure 5 provides a more adequate disambiguation of the attractor and transient dynamics. In each case, the black lines show the agent’s performance after it has been given time to adapt. They commence each light presentation at a distance in the range of 10 to 15 away from the light, and moves downwards (in these plots) towards a high-fitness, low distance location. The grey lines start in the same region of light distance at the start of a new light presentation, but now the sensor configuration has been changed. From here, the light distance grows instead of diminishing. The interneuron is at first still within the same region as it was prior to the switch, however following the same sensorimotor strategy as before the switch now moves the agent away from the light. Eventually, this maladaptive behaviour will lead to diminished sensory inputs, since the sensor is distance-sensitive. These changing input patterns likely cause the interneuron to transition to its opposite state, pushing the network’s dynamics into the alternate, now successful, region of sensorimotor behaviour. Explaining these patterns in more detail, however, has not yet been achieved.

Finally, the agent was tested with the sensor mounted at different angles — not just the ones which it experienced during evolution. The average fitness during these tests were low when the angle deviated more than a small distance from either 0° or 180° .

5 Discussion

A model agent has been presented, which is capable of consistently performing phototaxis, regardless of whether its light sensor is mounted on its front or back. By design, this performance cannot be achieved by a fixed reactive control strategy. We have demonstrated how this is realised through a time-extended adaptation process with no explicit plasticity. After a switch of sensor location, the agent’s now maladaptive sensorimotor behaviour produces a change in the results of its actions. It is then able to exploit the new situation to ensure its dynamics enter a different, task-satisfying phase. We have analysed the neural activity in the agent’s controller, pointing to two separate strategies for the two different sensory morphologies. At first glance these two strategies correspond to the low and high firing regimes of the interneuron. However a useful explanation of how the dynamics bifurcate between these states required the inclusion of an environmental variable in the analysis.

We did not *a priori* provide a mechanism designed to produce the relatively slow dynamics of the adaptive process (when compared to the fast, reactive sensorimotor dynamics). However neural activity of multiple timescales were present within the system — this is particularly demonstrated by the fast time constant values of the sensor and motor neurons, in contrast to the slower dynamics of the interneuron which correlated with the adapted state of the agent.

It is interesting to compare this situation to what we would expect if we had to design such an agent by hand. On a first approximation, a rule based design would look at establishing sensorimotor correlations (such as light sensor changes with respect to motor activity). These correlations may seem sufficient to determine the body configuration of the agent. This approach would be very much in line with the abstract framework proposed by Philipona et al. [10], whereby invariance in the structure of the sensorimotor correlations can be exploited to generate appropriate control laws. Whilst this approach remains generic, one issue we must confront when dealing with a time-extended, behaving agent is how to combine the fast timescales of behaviour and the slow timescales of adaptation into an integrated functional unit. A hand designed, algorithmic approach would tend to separate these two requirements into specialised sub-modules. We find that our evolved agent handles these two requirements with a single mechanism. Our analysis has shown how fast dynamics corresponds to a stable behavioural regime and how, upon changing the sensor configuration, longer term transients in the interneuron corresponding to phases of distancing from and approaching a new light are successfully combined without being functionally separated.

The importance of the interneuron’s transient dynamics is demonstrated by the experiment where its activity is made binary. While the high and low mean activity states are sufficient for producing their corresponding behaviours, the interneuron is also involved in generating the transition between the two, and its function cannot be reduced to that of representing either of the two conditions.

Whilst this experiment does demonstrate an adaptive agent which must explore its relation with the world in order to progress, it is limited to switching between just two discrete states corresponding to two distinct strategies. A model

which extends deeper into notions of flexibly adaptive behaviour would be able to produce viable sensorimotor strategies when encountering a continuous range of morphological configurations. Although it had not been evolved to handle sensor mountings other than on its front or back, we tested the agent's ability to deal with other sensor angles, which it largely failed to adapt to. [7] investigates learning on a continuum in an embodied scenario using recurrent neural networks (CTRNNs) without the need for a separate, dedicated adaptive apparatus, and could provide a basis for the extension of our model into the continuous domain.

Acknowledgements. We would like to thank Hiroyuki Iizuka, Marieke Rohde and other members of the CCNR for valuable discussion. This work was funded by the EPSRC grant EP/C51632X/1.

References

1. Beer, R.D., (1995), "On the dynamics of small continuous-time recurrent neural networks.", *Adaptive Behavior* 3(4), pp. 471–511.
2. Bi, G.Q. & Poo, M.M., (1998), "Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type.", *Journal of Neuroscience* 18, pp. 10464–72.
3. Bongard J., Zykov V. & Lipson H., (2006), "Resilient Machines Through Continuous Self-Modeling", *Science* 314(5802), pp. 1118–1121.
4. Di Paolo, E. A., (2000). "Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions.", In J-A. Meyer et al (Eds.) *Proc. of the Sixth Int. Conf. on the Simulation of Adaptive Behavior*, MIT Press, pp. 440–449
5. Floreano, D. & Urzelai, J., (2001), "Neural Morphogenesis, Synaptic Plasticity, and Evolution.", *Theory in Biosciences* 120(3-4), pp. 225–240.
6. Hölscher, C., (1997), "Nitric Oxide, the Enigmatic Neuronal Messenger: its Role in Synaptic Plasticity.", *Trends in Neuroscience* 20, pp. 298–303.
7. Izquierdo-Torres, E. & Harvey, I., (2006). "Learning on a Continuum in Evolved Dynamical Node Networks.", L. Rocha et al. (Eds.): *Proc. of the Tenth Int. Conf. on the Simulation and Synthesis of Living Systems*, pages 507–512, MIT Press.
8. O'Regan, J.K. & Noë, A., (2001), "A sensorimotor account of vision and visual consciousness", *Behavioral and Brain Sciences* 24(5), pp. 939–1011.
9. Phattanasri, P., Chiel, H.J. & Beer, R.D., (submitted)., "The dynamics of associative learning in evolved model circuits." Submitted to *Adaptive Behavior*.
10. Philipona, D., O'Regan, J.K., Nadal, J.P. & Coenen, O.J.-M.D., (2004), "Perception of the structure of the physical world using unknown sensors and effectors, *Advances in Neural Information Processing Systems* 15.
11. Tuci, E., Quinn, M., & Harvey, I., (2003), "An evolutionary ecological approach to the study of learning behaviour using a robot based model.", *Adaptive Behavior*, 10(3), pp. 201–221.
12. Yamauchi, B. & Beer, R.D. (1994), "Integrating reactive, sequential and learning behavior using dynamical neural networks.", D. Cliff et al. (Eds.), *Proc. of the Third Int. Conf. on Simulation of Adaptive Behavior*, pp. 382–391, MIT Press.