

Analysis of a Dynamical Recurrent Neural Network Evolved for Two Qualitatively Different Tasks: Walking and Chemotaxis

Eduardo J. Izquierdo^{1,2} and Thomas Buhrmann^{2,3}

¹Centre for Systems Biology, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK

²Centre for Computational Neuroscience and Robotics, University of Sussex, Brighton BN1 9QH, UK

³Natural Motion Ltd, Oxford OX1 2ET, UK

ej21@sussex.ac.uk, tb30@sussex.ac.uk

Abstract

Living organisms perform a broad range of different behaviours during their lifetime. It is important that these be coordinated such as to perform the appropriate one at the right time. This paper extends previous work on evolving dynamical recurrent neural networks by synthesizing a single circuit that performs two qualitatively different behaviours: orientation to sensory stimuli and legged locomotion. We demonstrate that small fully interconnected networks can solve these two tasks without providing *a priori* structural modules, explicit neural learning mechanisms, or an external signal for when to switch between them. Dynamical systems analysis of the best-adapted circuit explains the agent's ability to switch between the two behaviours from the interactions of the circuit's neural dynamics, its body and environment.

Introduction

All organisms are equipped with a repertoire of distinct behaviours that allow for their survival and reproduction. The nematode worm *Caenorhabditis elegans*, for example, with 'only' 302 neurons shows a remarkable ability to perform a broad range of different behaviours (Hart, 2006; Rankin, 2004). Although our understanding of the neural basis for most of them is still at an early stage (de Bono and Maricq, 2005), it is known that overlap exists between some of the neural circuits responsible for these behaviours (Hobert, 2003). Also, it is well known that the morphology of living organisms is in constant change, both throughout evolution and during the lifetime of the organism. This work investigates how a single neural network that is not structurally divided into separate circuits can produce different behaviours in different bodies.

When modelling adaptive behaviours, assumptions have to be made with regard to the structure of the organism studied in order to simplify the modelling process or the analysis of the model's behaviour. One such assumption that has been made in the past is that functional modularity, the existence of several qualitatively different behaviours in the same organism or agent, should be mirrored by structural modularity in its neural controller. Complex systems are thus often divided into small parts that are synthesized in

isolation. Such a divide-and-conquer approach can be very useful for engineering robots that need to perform multiple complex tasks, not least because it simplifies the understanding of how the robot works. But it is less useful in the context of developing the tools and language to understand biological organisms, as these may not necessarily have evolved to be easily decomposable.

First, we investigate whether a single neurocontroller can exhibit qualitatively different behaviours without imposing constraints on its structure. We use artificial evolution to synthesize a recurrent neural network that when coupled to two different simulated bodies, namely a one legged insect and a two-wheeled robot with a chemical sensor, has to perform legged locomotion in the former and chemotaxis in the latter case¹. A successful agent has to detect which body it inhabits and generate the appropriate behaviour. It must do this in the absence of an external signal and without any on-line changes in the parameters of the controller. We aim to find the smallest network that can solve the task. Although the structure of the network is under evolution, we do not investigate whether the evolved networks exhibit a degree of structural or 'functional'² modularity.

Second, using the mathematical tools of dynamical systems theory, we explain how the circuit in interaction with its body and environment can generate distinct behaviours. We characterize the autonomous dynamics of the best-evolved circuit and how its dynamics vary with inputs. We then study how the observed behavioural patterns are generated through the closed-loop interaction of the neural dynamics with the body and environment, for the two different tasks. Finally, we show how the evolved agent makes use of context-dependent feedback to shape the different transients using the same dynamical landscape. This leads us to suggest a dynamical systems perspective on adaptive behaviour that goes beyond attractors.

¹Both tasks have been studied in some depth in the Evolutionary Robotics literature. See Methods and Related work sections.

²Watson and Pollack (2005) argue that structural descriptions of the network are not sufficient to determine dependence or independence in the dynamics of different subsets of the network.

Methods

Walking task

The walking task employed follows very closely the simple one-legged body described and analysed in (Beer and Gallagher, 1992; Beer et al., 1999). Three variants of this model have been studied in (Beer, 1995), differing in whether sensory feedback is available constantly, only occasionally, or absent. Of the corresponding controllers, namely reflexive pattern generators (RPGs), central pattern generators (CPGs), and mixed pattern generators (MPGs), we focus on the first type only. The leg is controlled by three effectors: one specifies whether the foot has contact with the ground while the other two control clockwise and counter-clockwise torques for the leg's hinge-joint with the body (Figure 1A). The opposing torques model antagonistic muscles, commonly found in animal limbs. Sensory feedback is provided continuously by the leg's joint angle. At the beginning of each walking trial, the state of the leg (i.e. its angle with respect to the body) is initialised at random. The agent is then given 220 units of time to walk. The total distance covered during the trial measures performance.

Chemotaxis task

For the chemotaxis task we also follow a methodology similar to that employed in (Beer and Gallagher, 1992). A food patch, placed at arbitrary locations and orientations with respect to the agent, emits a chemical signal (s), whose intensity falls off as a function of the distance from the center of the patch (d): $s = e^{-\lambda d}$, where λ is a constant, -0.0138 . The agent can move freely in an environment without walls and must find and remain in the vicinity of the food patch.

The agent has a circular body and possesses a sensor that can detect the intensity of the chemical signal at its location (Figure 1B). Additionally, it is equipped with two effectors located on opposite sides of its body. These effectors³ can apply forces that move the body forward and rotate it. In the simplified physics of this environment, the velocity of movement is proportional to the force applied.

During a chemotaxis trial, a food patch is placed in a random direction from the agent, anywhere between 10 and 15 units of space apart. This is repeated after 100 units of time. Three food patches are shown in total. Performance is given by: $f_c = (d_i - \bar{d})/d_i$, where d_i and \bar{d} are the initial and average Euclidean distance between the agent and the food patch, respectively.

Neural model

We use continuous-time recurrent neural networks as a model of the agent's internal dynamics. Each component

³Although the mechanics of the body correspond closer to a khepera-like robot, similar physics have been used in idealised models of the nematode worm's movement in (Ferree and Lockery, 1999). Instead of 'wheels', the effectors located on opposite sides model the ventral and dorsal neck muscles of the worm.

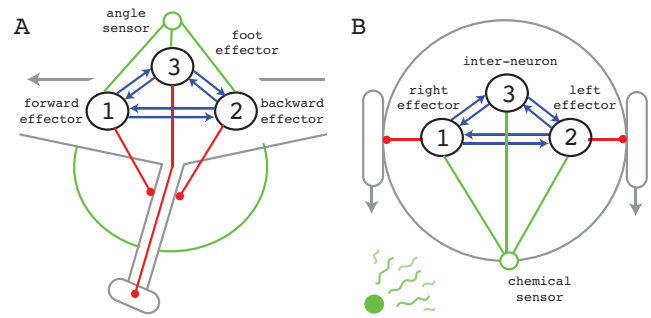


Figure 1: Task set-up. The same neural network circuit is used to control two different bodies: a one-legged insect-like walking agent (A) and a khepera-like chemotaxis agent (B). The circuits are fully inter-connected. The effector neurons control the antagonistic muscles and the foot for walking and the two effectors on the opposite sides of the body for chemotaxis. All of the neurons receive sensory perturbations: from the leg angle during walking, and from the proximity to the food during chemotaxis.

in the network is governed by the following state equation:

$$\tau_i \dot{y}_i = -y_i + \sum_{j=1}^N w_{ji} \sigma(y_j + \theta_j) + s w_i S(t) \quad (1)$$

where y is the activation of each node; τ is its time constant; w_{ji} is the strength of the connection from the j^{th} to the i^{th} node; θ is a bias term; $\sigma(x) = 1/(1 + e^{-x})$ is the standard logistic activation function; and N represents the number of nodes in the network. All nodes have access to the sensory perturbations via a set of connection weights: $s w_i$. The sensory input is normalised to run between 0 and 1 for both tasks. This prevents a solution that switches behaviour reactively as a response to different sensory input ranges. The network is fully connected (including self-connections) and no symmetry is imposed on its weight matrix. In simulation, node activations are calculated forward through time by straightforward time-slicing using Euler integration with a time-step of 0.1.

Evolutionary algorithm

The parameters of each circuit (i.e. biases, time-constants, inter-neuron and sensor-neuron weights for each node) are evolved using a version of the microbial genetic algorithm (Harvey, 2001). There are $N^2 + 3N$ parameters in total. These are encoded in a genotype as a vector of real numbers over the range $[0, 1]$. Offspring of microbial tournaments are generated as a mutation of the winner of the tournament (i.e. no recombination). The mutation is implemented as a random displacement on every gene drawn uniformly from a Gaussian distribution with mean 0 and variance 0.01. Each gene is forced to be in $[0, 1]$: when a mutation takes a gene out of this range it is reflected back. The

offspring replace the loser of the tournament. Genes are linearly mapped to network parameters in the range $[-10, 10]$ for biases, inter-node and sensory weights and to the range $[1, 20]$ for time constants. The size of the population used is 50 and we define a generation as the time it takes to generate the same number of new individuals. A minimal 1D wrap-around geography with demes of size 10 is used, such that only nearby individuals can compete in tournaments. Finally, because the fitness is noisy, agents are re-evaluated every time they participate in a tournament.

A successful circuit must maximize: (a) the distance walked when embodied in the insect-like body and (b) the time spent around the chemical-emitting food patch when in the khepera-like body. A fitness evaluation consists of 2 trials of the walking task and 15 trials of the chemotaxis task. At the start of each trial the state of the neural controller and the body is randomised. The performance on each task is averaged over all trials and normalised in the range $[0, 1]$. The fitness of an individual is calculated by multiplying the performance on both tasks.

Results

Evolutionary performance

Evolutionary searches with 3-, 4-, and 5-node circuits were performed. We examined the ability of populations to evolve for both tasks, and conducted control experiments in which either task was evolved on its own. For each condition, 20 evolutionary experiments with different initial random seeds were carried out.

First we compare networks of different size. Figure 2 shows the performance of the set of best circuits grouped according to size on the walking (2A) and chemotaxis task (2B). For each, two whisker plots show the performance of circuits on the task at hand. The grey whisker boxes correspond to populations evolved for only that task. White whisker boxes correspond to populations evolved on both tasks. Circuits of the same size perform better at generating the appropriate behaviour when evolved for only one task than those required to do both. This is true in all conditions and is expected. Also, the difference in performance is smaller for walking than for chemotaxis, which suggests it is the ‘easier’ task of the two. This is also as expected.

Nevertheless, sufficiently successful circuits that performed both tasks did evolve. However, is there a trade-off between walking performance and chemotaxis in the successful circuits? In other words, are some individuals good at walking but poor at chemotaxis and vice versa? In Figure 2C we show chemotaxis versus walking fitness for each of the best individuals from all evolutionary runs on the two tasks. No obvious tradeoffs are noticeable. Instead, most of the successful individuals at one task are also relatively successful at the other task.

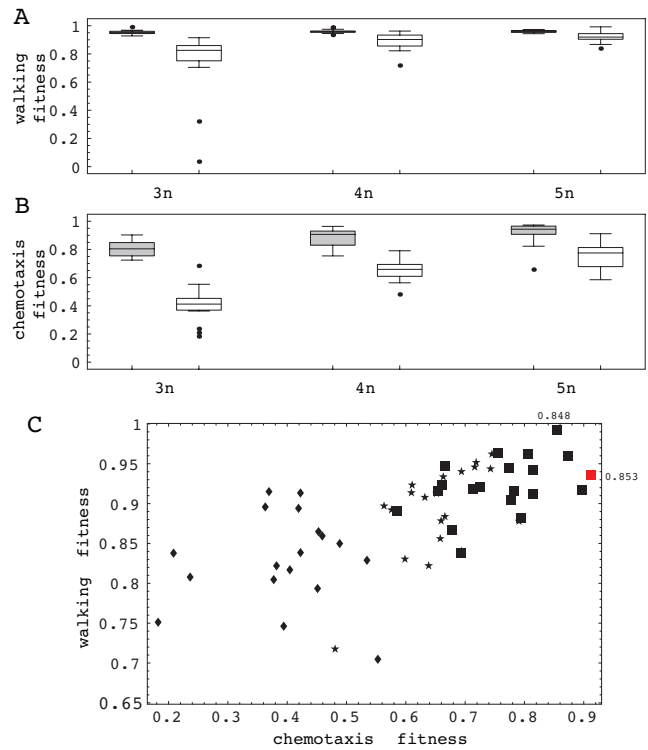


Figure 2: Evolutionary performance. Whisker plots (25% to 75% quantiles and outliers as points) comparing the fitness achieved by 3-, 4- and 5-node networks after having evolved for only one task (grey) and both tasks (white) in the case of walking [A] and chemotaxis [B]. [C] Chemotaxis versus walking fitness on the best 3- (diamonds), 4- (stars), and 5- node (squares) circuits evolved for both tasks.

Performance and behaviours

The best evolved agent achieved a walking fitness of 93.59% and a chemotaxis fitness of 91.18%. This individual corresponds to the red square from Figure 2C. In Figure 3 we show an example trial with this circuit performing both tasks. It is relevant to note that all neurons are active during both tasks.

Walking The optimal walking pattern for this one-legged model has been studied in (Beer et al., 1999). As can be seen from Figure 3B, the evolved pattern is almost perfectly aligned with the optimal pattern, at least geometrically. The different sections in this pattern correspond to particular stages of the walking cycle (labelled in grey): (1) foot up and swing, (2) foot down, (3) stance power, and (4) stance coast. This agrees with results in (Beer et al., 1999). Yet, we know the performance is only 93.59% of the optimal asymptotic velocity that the walking agent can achieve (0.627). The difference is in the timing. The best evolved circuit has 2 units of time delay between the mo-




Figure 3: Example trial for best evolved circuit. [A] Sensory and neural pattern during walking and chemotaxis. Transition denoted by horizontal dashed line. [B] Walking behaviour in black. Angle (θ) versus angular velocity ($\dot{\theta}$). Optimal walking in thick grey. [C] Chemotaxis behaviour over six presentations of food.

ment it sets its foot down (phase 2) and the moment it starts to move the leg backwards (phase 3). This unnecessary foot ‘rest’ causes the degradation in performance, nevertheless maintaining the optimal geometry of the walking pattern. Thus, the best-evolved circuit completes the full cycle at a slightly slower rate than the optimal. If we turn the sensory feedback off, the circuit cannot walk. This is expected from the RPG.

Chemotaxis Even though it has often been used as an example in the artificial life literature, chemotaxis has not been studied in as much depth as the one-legged insect walker. Our agent has only one non-directional sensor. Hence the only way to detect the chemical gradient is by moving about. Organisms that are too small to sense the gradient along the length of their own bodies are in a similar situation. The only strategy available is to use subsequent sensory signals to estimate the chemical gradient in time rather than space.

We can identify four relatively different phases in the best-evolved agent during a chemotactic run: (1) circling search, (2) decreased turning in direction towards the gradient, (3) straight run, (4) circling around food patch. However, this is only a simplified observer-perspective heuristic

and the phases are not always clear-cut. The full story is provided only by a geometrical account.

For the initial problem of finding the gradient the agent employs a heuristic that involves circling on the spot while the distance from the food is either constant or decreasing, and moving straight otherwise (i.e. while it is increasing). To confirm this hypothesis, we performed a series of experiments in which we allowed the agent to move while controlling the sensory information arbitrarily. We considered the initial transient behaviour under two conditions: when the sensor value is fixed or decreasing, and when it is increasing at a constant rate. During the former, the agent circles around a small region of about 2 units of space. During the latter, the agent reduces the turning behaviour as a function of the rate of increase of the sensor activity.

Interestingly, the chemotaxis behaviour of the best-evolved circuit employs a strategy similar to that observed in very simple organisms. In *E. coli*, for example, chemotaxis is achieved by modifying the frequency of ‘tumbling’ (Macnab and Koshland, 1972). In *C. elegans*, the turning behaviour is referred to as a ‘pirouette’ but the heuristic is similar (Pierce-Shimomura et al., 1999).

Switching behaviours We know the agent can perform well doing each task independently. In order to test whether it can also switch between them during its lifetime we change the circuit’s body without resetting the state of the neurons and evaluate the circuit’s performance. Although populations were not evolved to cope with this transition, most of the successful circuits managed to switch between tasks in both directions, including the best one analysed here. The example shown in Figure 3A is for a successful transition in one of the directions: from walking to chemotaxis. We will answer why this is possible in the last section of the results.

Dynamics of the decoupled circuit

As a first step towards understanding the evolved behaviours, we consider the dynamics of the circuit when decoupled from the environment. We do this by examining the asymptotic behaviour of the circuit after replacing the time-varying sensory input with a fixed parameter, thereby reducing it to an autonomous system.

Bifurcation diagram In Figure 4 we show the asymptotic behaviour of the circuit as a function of the possible sensory perturbations that it can receive. Solid black trajectories represent attractors, dashed black trajectories represent saddle-nodes and grey dots correspond to limit cycles. Three bifurcations can be observed and are shown in the figure as colored disks. From left to right, the first bifurcation is a saddle-node bifurcation (red disk), from which a fixed point (a_2) and the saddle node (sn) arise. Fixed point a_1 is a stable spiral point for $s < 0.38$. This spiral is weakened and a

stable limit cycle (lc) arises near the origin in what is likely to be a Hopf bifurcation (green disk). The size of the cycle first increases slowly and then comes crashing inwards until it reverts to a stable spiral point for $s > 0.77$.

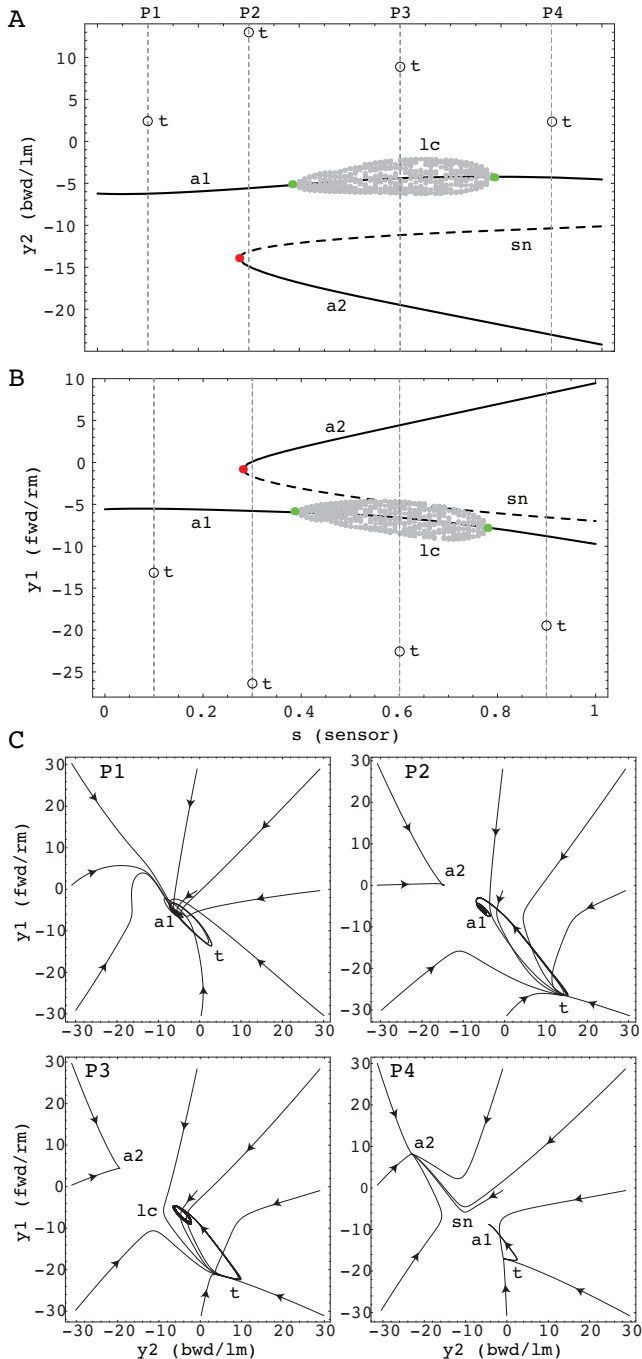


Figure 4: Bifurcation diagram and example phase portraits. [A and B] Two-dimensional slices of the 6D bifurcation space (5 neurons + sensor). [C] Two-dimensional slices of phase-portraits $P1$ through $P4$ for the two effector neurons, y_1 and y_2 . See main text for an explanation of the labels.

Phase portraits The bifurcations divide the space of qualitatively different dynamics available to this circuit in four. The dashed vertical lines in Figures 4A and 4B represent the slices of parameter space that are studied in Figure 4C, which shows two-dimensional slices of phase-portraits $P1$ through $P4$. The portraits are shown only for neurons 4 and 5, which control forward and backward movement in both insect and khepera bodies. $P1$ corresponds to the family of phase portraits available when $s < 0.28$ before the first bifurcation occurs. As there is only attractor (a_1) in the system all trajectories are drawn to it. The paths taken to get to it are not direct, but follow a spiral in its vicinity. Prior to this, however, a subset of the trajectories follow a much longer transient involving a loop near the region labelled t .

$P2$ corresponds to the family of phase portraits available between the first and second bifurcations ($0.28 < s < 0.38$). It comprises two stable fixed points a_1 and a_2 , and a saddle-node (not shown). Trajectories starting in the top-left corner approach the newly created stable point, a_2 , whose basin of attraction (not shown) is smaller than a_1 's. Hence, most trajectories approach a_1 . The transient towards it is, again, not direct. In fact, as we will see, this is always the case for this attractor. What varies is the spatial extent of the spiral loop. When looking through the perspective of neurons 4 and 5, any trajectory bound for a_1 will first navigate towards t . In $P3$ the spiral attractor becomes a stable cycle. The transient remains similar.

In $P4$ the cycle disappears and gives way to a stable fixed point. Also, a_2 's basin of attraction becomes larger, with certain initial configurations ending up in a_2 that previously ended in a_1 . Also, the effect of the saddle-node (sn) becomes more obvious in this portrait. The transient loop (t) still exists, but it is relatively closer to a_1 .

Finally, approximations of the turning point of the transient loop are incorporated into our bifurcation diagram as disks labelled t in Figures 4A and 4B. As these are not real limit sets of the system, they do not show up in our bifurcation analysis. They will play, however, a fundamental role in the agent's autonomous behaviour. If the phase-portrait of the system is changing sufficiently fast (due to rapidly varying input), and if the neural state falls in the basin of attraction of a_1 , then we can predict that it will most likely be seen around t and never actually reach a_1 .

Brain-body-environment coupled dynamics

Let us now consider the behaviour of the agent when coupled to the environment and how it relates to the underlying dynamical landscape described in the previous section. Figure 5 depicts the trajectories of the controller when driven by the agent's sensor, which is itself influenced by the circuit's effectors and the corresponding changes to how the agent perceives the environment. Red lines correspond to the walking task, blue lines to chemotaxis. The trajectories are imposed over a simplified version of the circuit's

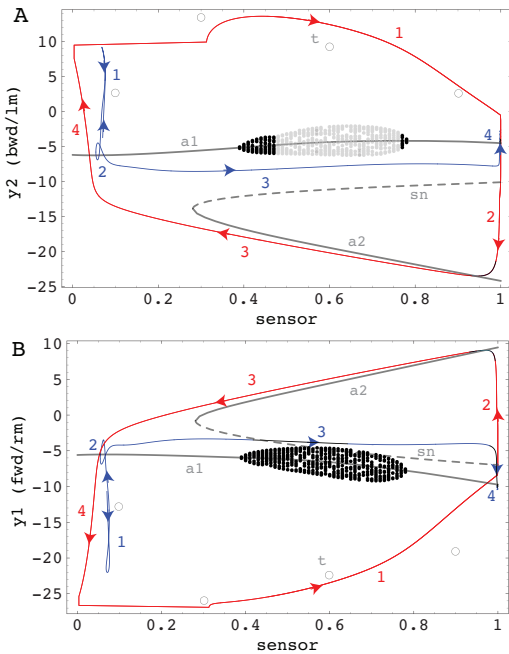


Figure 5: Brain-body-environment coupled dynamics for the two different tasks: walking (red) and chemotaxis (blue). Trajectories are imposed over the bifurcation diagram of the non-autonomous nervous system (grey). Two slices of the 6-dimensional space are shown. The same two slices shown for Figure 4A and 4B, respectively.

autonomous dynamics from Figures 4A and 4B, using the same projections. While the dynamics of the two tasks are significantly different, they share the same underlying dynamical landscape.

During the walking task, the dynamics of the circuit are constantly switching between approaching attractor $a1$ when swinging the leg forward, and approaching attractor $a2$ during the stance power and coast. However, while the system gets close enough to $a2$, it ends up relatively far from $a1$. In fact, the cycle that arises from the coupled system is observed to switch between $a2$ and t (the longer transient towards $a1$). This agrees with our prediction from the circuit's autonomous dynamics, which suggested that it is being driven at a relatively fast rate. We test this hypothesis in the next section. We also note that the cycles in y_2 and y_1 follow opposite directions, clockwise and counterclockwise, respectively. This reflects the antagonistic muscle cooperation necessary to produce the swinging of the leg.

The trajectory during chemotaxis is more subtle and is produced solely within the basin of attraction of $a1$. The circling search behaviour is produced by the longer transient towards t , in combination with the spiral shape in the vicinity of $a1$. However, the state of the neural controller doesn't really reach $a1$ until the agent gets close to the food

patch, at which point the sensor gets maximally activated. As soon as the gradient towards the food is found, the sensory value increases and the phase-portrait shifts, leaving the state of the effector neurons in a region of space where the power of the opposing effectors are balanced, which corresponds to moving straight. Interestingly, the spiral attractor and indeed the limit cycle around $a1$ ensure that if the gradient ceases to increase, the agent will circle on the spot until it increases again. This agrees with our observation of the agent's chemotactic heuristic. Finally, once the agent reaches the top of the gradient, the dynamics come cycling in towards $a1$, which ultimately leaves the agent turning on the spot near the food patch.

Behaviour coordination: driven circuit

How does the neural controller perform the appropriate behaviour at the appropriate time? What our analysis shows is that it is *not* the neural controller itself that coordinates the change of behaviours. Instead, different patterns of feedback are created when the neural system is coupled to a different body, and it is these patterns that ultimately produce the distinct transient behaviours. It is worth emphasizing that brain, body and environment form a closed loop such that no single part is the sole cause for the difference in the dynamical patterns. The shape of the feedback is as much the result of neural output and body dynamics as the neural activity itself is the result of environmental feedback.

Is a particular feature of the feedback signal associated with the change of behaviours? From the previous section, we observed that the walking task is generated by the system's movement between two basins of attraction, in such a way that it never actually settled on any one of them. This suggests that fast switching between the two basins generates walking. During chemotaxis, on the other hand, the dynamics stay within the basin of $a1$ and movement is sufficiently slow to allow for it to draw close to the attractor, settling only when sufficiently close to the food patch.

In summary, the essential feature of the feedback is its time-scale. While the sensory feedback from the insect-like body is relatively fast, the sensory feedback from the khepera-like body is much slower. We test this hypothesis by driving the neural system with fast and slow sine waves, and compare the observed dynamics in internal space (Figure 6A) to the dynamics during walking and chemotaxis (Figures 6B and 6C). We find that, depending on the frequency, it will either: (i) jump from one attractor to the other, which is relevant to the walking behaviour, or (ii) stay on the central attractor, which is relevant to chemotaxis. Finally, this provides an explanation for why the evolved agent can switch between behaviours during its lifetime. The behaviours don't depend on where in neural space the state of the system is, but on the rate at which it is being driven by the feedback from its interactions with the environment, as a product of the mechanics of its body.




Figure 6: System driven by sinusoidal waves of different frequencies. [A] Trajectories of the neural state of the system are imposed over the circuit's bifurcation diagram (grey, limit cycles not shown). When the system is driven slowly ($f = 0.272$, thin black trajectory), the trajectory remains near the $a1$ attractor. As the system is driven faster ($f = 0.136$ and $f = 0.068$, thicker trajectories), the state of the system starts oscillating between attractors $a1$ and $a2$, but because of the longer transient towards $a1$, the oscillation is effectively between t and $a2$. [B] Two-dimensional slice through the space of neural activity during walking (red) and chemotaxis (blue) imposed over one of the phase-portraits ($P3$). [C] Neural activity for the system when driven by a fast ($f = 0.272$) and slow ($f = 0.068$) sinusoidal wave.

Related work

Synthesizing neural controllers to generate multiple qualitatively different behaviours is a challenge that has been posed by many. However, the focus has been on the role of modularity. Togelius (2004) showed how subsumption architecture models could be merged with an evolutionary robotics approach for a simulated robot on a learning task. Nolfi (1997) investigated modularity for evolution of a garbage collecting robot that had to cope with subtasks such as recognizing, picking up and disposing of desired objects. Although the networks had a hard-wired modular architecture, evolution was free to choose how these modules were used. Calabretta et al. (1999) addressed the same task, but used a system in which neural modules could evolve from a population of non-modular networks through gene duplication. Although both reported improved performance relative to monolithic networks, Ziemke et al. (1999) showed that in

a more difficult version of the task, a monolithic recurrent network outperformed all modular architectures.

Many animals can rapidly change between different modes of locomotion. In (Ijspeert et al., 2007), the problem of designing the neural controller for switching between swimming and walking in a salamander-like robot is presented. In (Buckley et al., 2008), an agent is evolved to do phototaxis with the sensor in two different positions (front and back of the body) while constraining the dynamical system controller to use a single basin of attraction. In both papers, however, the two behaviours share a large range of qualities. Our work is different from theirs in that the two tasks (chemotaxis and walking) were chosen to be as different as possible, while sharing sensor and effectors.

Yamauchi and Beer (1994) evolved a simulated robot that had to learn which of two environments it was placed in and take an appropriate action such as to approach a desired position. They only succeeded after dividing the network into separate modules with explicitly assigned roles that were evolved separately. They then evolved a classifier network to determine which of the modules is to control the agent. Tuci et al. (2002) later successfully evolved a controller for a very similar task without dedicated modules. Finally, Beer and Gallagher (1992) evolved agents for chemotaxis and walking, but not for the same dynamical system controller.

Discussion

We have shown that small dynamical neural networks are able to implement qualitatively different behaviours as distinct transients on a single dynamical landscape. Specifically, we evolved an agent that could perform locomotion when coupled to a one-legged body and chemotaxis when controlling a khepera-like robot. We demonstrate this is possible without imposing structural modules on the controller, and without employing complicated fitness functions or evolutionary shaping protocols. Neither was it necessary to introduce parameter changes in the controllers or to provide a signal for when the swap of bodies and corresponding behaviour was to occur. The interactions of neural controller, body and environment alone are sufficient to create distinct transient dynamics appropriate for solving both tasks.

The divide-and-conquer approach championed in engineering would suggest that separate modules should be evolved to produce the two tasks independently. This however wouldn't necessarily simplify the problem, as the main challenge would then be to design a mechanism of coordination and a sophisticated sensing machinery to detect when to switch between the modules. More importantly however, while modular structures and synaptic-plasticity exist in living organisms, they were selected based on the adaptiveness of their behaviour and not on how apprehensible their internal mechanisms are. We therefore argue that understanding networks whose structure is not imposed from the top down will help us develop the tools to understand how multiple

behaviours are generated in living organisms.

Our dynamical systems account of the evolved agent indicates that it is misleading to associate a behaviour with an attractor or a basin of attraction in the decoupled internal dynamics of the controller. In fact, there can be many behaviours in the same basin of attraction, as shown in (Buckley et al., 2008), or single behaviours that require several different attractors, as in the RPG shown here and in (Beer, 1995). Furthermore, as this paper demonstrates, multiple behaviours may use an overlapping set of diverse attractors and their basins. This provides an example of the importance of understanding behaviours as a result of the interactions between brains, bodies and environments, where transients play an equal, if not more important, role than attractors.

A possible objection that could be raised about this work is that the behaviours presented here are either too simple or not sufficiently different from each other and that modular and hierarchical architectures and additional learning rules will be required for more ‘complex’ scenarios. We believe this is an important, but mostly open question.

Finally, an important feature of recurrent neural networks is that their history of activations allows them to respond to otherwise identical stimuli in a context-dependent fashion. In other words, a system with internal state, when embodied and situated, is not constrained to a single sensori-motor mapping (as was shown in our example). In von Uexkull (1957)’s terms, such systems could be said to “bring forth their own *Umwelt*”. But while the act of interpreting sensory input contextually is usually attributed wholly to the agent, the example presented here shows that “meaningful” behaviour is the result of interactions in the brain-body-environment system as a whole.

Acknowledgements

We thank Peter Fine for feedback on an earlier draft of this paper. Mathematica tools developed by Randall Beer were used for parts of the dynamical analysis. This work was partly supported by the ESIGNET 6th Framework EU grant.

References

- Beer, R., Chiel, H., and Gallagher, J. (1999). Evolution and analysis of model CPGs for walking: II. General principles and individual variability. *Journal of Computational Neuroscience*, 7:119–147.
- Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72:173–215.
- Beer, R. D. and Gallagher, J. (1992). Evolving dynamical neural networks for adaptive behaviour. *Adaptive Behavior*, 1(1):91–122.
- Buckley, C., Fine, P., Bullock, S., and Di Paolo, E. (2008). Monostable controllers for adaptive behaviour. In Asada, M., Tani, J., Hallam, J., and Meyer, J.-A., editors, *Proc. of the Tenth Int. Conf. on the Simulation of Adaptive Behavior*.
- Calabretta, R., Nolfi, S., Parisi, D., and Wagner, G. (1999). Emergence of functional modularity in robots. In *Proc. of the Fifth Int. Conf. on Simulation of Adaptive Behavior*, pages 497–504. MIT Press.
- de Bono, M. and Maricq, A. (2005). Neuronal substrates of complex behaviors in *C. elegans*. *Annual Review of Neuroscience*, 28:451–501.
- Ferree, T. and Lockery, S. (1999). Computational rules for chemotaxis in the nematode *C. elegans*. *Journal of Computational Neuroscience*, 6:263–277.
- Hart, A. (2006). Behavior. In Ambros, V., editor, *WormBook*.
- Harvey, I. (2001). Artificial evolution: a continuing SAGA. In Gomi, T., editor, *Evolutionary Robotics: From Intelligent Robots to Artificial Life*. Springer-Verlag LNCS 2217.
- Hobert, O. (2003). Behavioral plasticity in the *C. elegans*: Paradigms, circuits, genes. *Journal of Neurobiology*, 54:203–223.
- Ijspeert, A. J., Crespi, A., Ryczko, D., and Cabelguen, J.-M. (2007). From Swimming to Walking with a Salamander Robot Driven by a Spinal Cord Model. *Science*, 315(5817):1416–1420.
- Macnab, R. and Koshland, D. (1972). The gradient sensing mechanism in bacterial chemotaxis. *Proc. Natl Acad. Sci.*, 69(9):2509–2512.
- Nolfi, S. (1997). Using emergent modularity to develop control systems for mobile robots. *Adaptive Behavior*, 5(3-4):343–363.
- Pierce-Shimomura, J., Morse, T., and Lockery, S. (1999). The fundamental role of pirouettes in *Caenorhabditis elegans* chemotaxis. *Journal of Neuroscience*, 19(21):9557–9569.
- Rankin, C. (2004). Invertebrate learning: what can’t a worm learn? *Current Biology*, 14:617–618.
- Togelius, J. (2004). Evolution of a subsumption architecture neurocontroller. *Journal of Intelligent and Fuzzy Systems*, 15(1):15–20.
- Tuci, E., Quinn, M., and Harvey, I. (2002). An evolutionary ecological approach to the study of learning behavior using a robot-based model. *Adaptive Behavior*, 10(3-4):201–221.
- von Uexkull, J. (1957). A stroll through the worlds of animals and men: A picture book of invisible worlds. In Lashley, K. and Schiller, C., editors, *Instinctive behavior: the development of a modern concept*. International University Press.
- Watson, R. and Pollack, J. (2005). Modular interdependency in complex dynamical systems. *Artificial Life*, 11(4):445–457.
- Yamauchi, B. and Beer, R. (1994). Integrating reactive, sequential and learning behavior using dynamical neural networks. In Cliff, D., Husbands, P., Meyer, J., and Wilson, S., editors, *Proc. of the Third Int. Conf. on Simulation of Adaptive Behavior*, pages 382–391. MIT Press.
- Ziemke, T., Carlsson, J., and Bodn, M. (1999). An experimental comparison of weight evolution in neural control architectures for a ‘garbage-collecting’ khepera robot. In Loffler, A., Mondada, F., and Ruckert, U., editors, *Proc. of the First Int. Khepera Workshop*, pages 31–40.