



Recent advances in evolutionary and bio-inspired adaptive robotics: Exploiting embodied dynamics

Phil Husbands¹ · Yoonsik Shim^{2,3} · Michael Garvie³ · Alex Dewar³ · Norbert Domcsek³ · Paul Graham⁴ · James Knight³ · Thomas Nowotny³ · Andrew Philippides³

Accepted: 12 February 2021 / Published online: 10 May 2021
© The Author(s) 2021

Abstract

This paper explores current developments in evolutionary and bio-inspired approaches to autonomous robotics, concentrating on research from our group at the University of Sussex. These developments are discussed in the context of advances in the wider fields of adaptive and evolutionary approaches to AI and robotics, focusing on the exploitation of embodied dynamics to create behaviour. Four case studies highlight various aspects of such exploitation. The first exploits the dynamical properties of a physical electronic substrate, demonstrating for the first time how component-level analog electronic circuits can be evolved directly in hardware to act as robot controllers. The second develops novel, effective and highly parsimonious navigation methods inspired by the way insects exploit the embodied dynamics of innate behaviours. Combining biological experiments with robotic modeling, it is shown how rapid route learning can be achieved with the aid of navigation-specific visual information that is provided and exploited by the innate behaviours. The third study focuses on the exploitation of neuromechanical chaos in the generation of robust motor behaviours. It is demonstrated how chaotic dynamics can be exploited to power a goal-driven search for desired motor behaviours in embodied systems using a particular control architecture based around neural oscillators. The dynamics are shown to be chaotic at all levels in the system, from the neural to the embodied mechanical. The final study explores the exploitation of the dynamics of brain-body-environment interactions for efficient, agile flapping winged flight. It is shown how a multi-objective evolutionary algorithm can be used to evolved dynamical neural controllers for a simulated flapping wing robot with feathered wings. Results demonstrate robust, stable, agile flight is achieved in the face of random wind gusts by exploiting complex asymmetric dynamics partly enabled by continually changing wing and tail morphologies.

Keywords Evolutionary robotics · Biorobotics · Visual navigation · Neural dynamics · Chaotic dynamics · Evolvable hardware

1 Introduction

The dynamics inherent in biological systems are varied, multi-levelled, and often complex. They range from the dynamics of individual cells, to those of neural and bodily structures, to individual animal behaviours, to the ecological and evolutionary dynamics of populations. All of these are growing sources of inspiration for the development of new

techniques and approaches in AI and autonomous robotics.

The inherent spatiotemporal dynamics of the nervous system are much richer than those of most artificial systems. Complex electro-chemical processes interact over many different spatial and temporal scales creating multiple layers of adaptive mechanisms that are at the heart of the nervous system's incredible versatility and power [1, 80, 97, 109]. Recent empirical work in neurophysiology provides evidence that the complex information processing that these processes enable is an intrinsic property of various classes of neurons, such that it will arise even in extremely sparse reconstituted networks [100]. The inherent properties of neurons and the structures they build are therefore a very rich source of inspiration for artificial nervous systems intended to generate adaptive behaviour in autonomous mobile robots.

This article belongs to the Topical Collection: *30th Anniversary Special Issue*

✉ Phil Husbands
philh@sussex.ac.uk

Extended author information available on the last page of the article.

At a slightly higher level, the ability of evolution to create, shape and exploit the complex dynamics of neural systems is in itself a powerful inspiration for how to develop control systems for robots. Artificial evolutionary approaches can be employed to design neural controllers [90], or to exploit the inherent properties of some re-configurable physical medium, in which control circuits can be defined.

In nature, embodied behaviours do not arise due to the power of neural dynamics alone. They are generated by the interacting dynamics of the whole brain-body-environment system making up a behaving organism in its environment [33, 93]. The overall morphology of the body, including the layout and properties of the sensors, plays an important role in generating these dynamics. The ways in which these various elements interact, resonate and co-evolve, are what generate, shape – and ultimately enable the exploitation of – the rich dynamics that underpin robust adaptive embodied behaviour.

In this paper we present four case studies of current work from our group where various of these elements are explored in the development of novel biomimetic autonomous robotic systems. The binding theme in these studies is the exploitation of dynamics. Between them, these interrelated strands of work demonstrate a range of ways in which this can be profitably achieved. They cover: exploiting the dynamical properties of a physical electronic substrate in the development of robot controllers; navigation methods inspired by the way insects exploit the embodied dynamics of innate behaviours; the exploitation of neuromechanical chaos in the generation of locomotor behaviours; and the exploitation of the dynamics of brain-body-environment interactions for efficient, agile flapping winged flight.

The very first work on truly autonomous robots, the development of Grey Walter's tortoises [125, 126], was intended as an exploration of how the dynamics of a simple, yet richly interconnected, electronic nervous system could generate adaptive behaviour in an 'artificial creature'. It was meant primarily as a new kind of synthetic method for neurophysiology, but was hugely influential in future developments in robotics and AI. The research described in this paper is motivated by Walter's pioneering approach insofar as it takes inspiration from biology in the development of autonomous robotic systems, but some of the studies also shed new light on aspects of the biological mechanisms that are the source of their inspiration.

The structure of the paper is as follows. The next section reviews recent relevant literature in areas related to the research described in the main body of the paper. Following that, Sections 3–6 present the four case studies outlined above. All make use of bio-inspired adaptive approaches to AI and robotics, focusing on the exploitation of embodied dynamics to create behaviour. Finally, a conclusions section

summarizes and discusses the main findings, pointing out limitations and possible future directions.

2 Literature review

Over the last few years there has been a growing number of examples of using evolutionary and adaptive methods to exploit the dynamics of physical materials to create behaviour generating and/or information processing systems.

A very recent, exciting development was the creation, in 2019, of simple organic robots using a process of artificial evolution [66]. Dubbed xenobots by their creators at the University of Vermont and Tufts University, these tiny biological machines were first designed in computer simulations using the techniques of evolutionary robotics. The xenobots were assemblies of passive and contractile biological cells (the latter can spontaneously contract and relax, that is pulse) which were evolved to perform some simple behaviour (e.g. move forward: the pulsing contractile cells could be exploited to power locomotion). The best designs evolved in simulation were then created in a biology lab from real cells – passive skin cells and contractile heart cells – developed from frog stem cells. The newly assembled organic robots were able to perform the desired behaviours when placed in a petri dish. In the words of the scientists who developed them, the creations are novel living machines, programmable organisms, entirely new lifeforms. It is quite plausible that this kind of technique could be used to develop useful nano robots in the future, maybe to deliver drugs after being injected into the bloodstream. Scaling to larger, more complex creatures will be very challenging, but might be possible in the long-term. As with the evolutionary robotics work described in Section 3, the xenobot simulations were refined using feedback from reality. Even in complex cases such as these, evolved solutions can be made to cross the reality gap. The inherently plastic and degenerate nature of biological material may have been advantageous in that regard.

Howard et al. [52] used a model of a variable resistive memory based novel hardware substrate to evolve spiking neural networks capable of controlling a simulated robot engaged in a T-maze task. While not quite an example of true EHW, this work provides another example of how evolution exploits the dynamics afforded by unconventional media. Recently there has also been renewed interest in analog EHW in other application areas besides robotics. These include the evolution of circuit repair strategies, which require no specific information of the nature or location of faults, in a special very fine-grained FPGA architecture [122]. There have also been more general investigations of other evolvable (continuous) physical substrates, such as liquid crystal, chemical reaction-

diffusion systems and carbon nanotubes, as unconventional mediums for computation and information processing (including for controlling simulated robots) [2, 29, 84, 85]. For such systems, their success was again rooted in the exploitation of rich dynamics.

The research described in the next section demonstrates how another type of physical medium, built from reconfigurable analog electronics, can act as a highly evolvable substrate for developing controllers for a physical robot engaged in visually guided behaviours. The rich dynamics afforded by the medium was exploited to develop very concise, unorthodox, yet highly robust, controllers.

Visually guided navigation is a popular topic within robotics (see [17] for a review), as it holds the possibility of reliable homing over a range of spatial scales, using cameras as cheap and reliable sensors. It is important to note, however, that this topic encompasses a number of problems – e.g. route-following and visual place recognition (VPR; for a review see [77]) – which are subtly distinct. For example, for a robot tasked with following a previously memorised route, VPR may not be much use in the case where it is displaced to an unfamiliar position away from the route, whereas an algorithm such as the visual compass, where memorised views can be used to recall the direction an agent was facing when the view was stored (see Section 4), can be robust to such occurrences: the robot is not required to know where it is in order to know what to do. This is in contrast to map-building-type approaches (e.g. Simultaneous Localisation and Mapping – SLAM [23]) where a metric map is incrementally constructed based on the estimated positions of landmarks: a process which is often computationally expensive, in particular in large-scale outdoor environments. Even among mapless approaches, input images are often preprocessed to extract visual features which are then used as landmarks, e.g. FAB-MAP [28] uses a so-called bag-of-words algorithm, where places are defined based on the presence (or not) of a predefined vocabulary of features. In Section 4 we describe our alternative approach, which eschewing these kinds of abstraction and instead uses direct comparison of raw images – taking inspiration from ants. We show that robust behaviour emerges from this embodied approach, requiring much less computation.

In Section 5 we explore how certain types of complex neural dynamics can be exploited in the generation of embodied behaviour, and how they must work in concert with bodily dynamics. Oscillatory neural dynamics are prevalent in many brain areas and appear to underlie numerous mechanisms involved in information processing and the generation of behaviour [18, 22, 121]. This observation led to the development of various artificial neural network architectures based on coupled oscillators that have been successfully employed as robust control systems for various kinds of

robots [11, 55, 86, 102]. Complex, often chaotic, dynamics are observed in biological nervous systems. Hence some researchers began to explore the properties of chaotic neural oscillators in the generation of behaviour [6]. This work showed the potential adaptive properties of such systems and led to a more active and radical exploitation of chaos as an adaptive force, where the chaotic dynamics arise in a whole embodied neuro-physical system [67, 68]. Such research was generalised and extended by our group to allow the robust development of goal-directed motor behaviours [105]. Section 5 further extends this work, developing a general class of coupled-oscillator neural controllers which are able to generate highly robust, resilient behaviours in a wide range of robots without the need for *a priori* knowledge of the robot or environment. The nature of the chaotic dynamics within the whole brain-body-environment system is explored in some detail.

As we shall see in several of the case studies described in this paper, particularly those covered in Sections 4 and 6, animal body morphologies play an important part in the generation of behaviour, as do the properties of materials and structures making up the body. The development of a ‘PigeonBot’ at Stanford University [26, 82] potentially opens up many interesting directions in the development of hybrid systems where robots exploit the properties of natural biological structures (in this case bird feathers), including the complex dynamics such structures can enable. The primary aim of this research was to understand more about the mechanisms of flapping wing flight in birds, in particular how the wing surfaces dynamically change shape during flight. The team developed a cleverly conceived bio-mechanical hybrid robot — the PigeonBot — by incorporating real pigeon feathers into the wing of the flying machine. The study revealed a lot about how birds fly, particularly how feathers and the ways they are connected enable the powerful mechanisms of continuous wing morphing, and how these dynamics are used. It also points the way to new forms of biomimetic flying robots that could have a lot of useful applications. Wing morphing considerably improves efficiency and manoeuvrability in nature, so further developments of such robots might lead to a new class of agile flying robots that are superior to standard drones. As with much work on the interfaces of robotics and biology, such as that described in Sections 4, 5 and 6, advances were made in AI and engineering while at the same time gaining biological insights. The research described in Section 6 explores the evolution of neural controllers that are able to exploit complex aerodynamics enabled by morphing feathered wings and tail in a simulated flapping wing robot.

The next four sections look in detail at case studies which explore how biomimetic systems can use the various kinds of exploitation of dynamics outlined above.

3 Evolving robot controllers directly in analog electronics: Exploiting the dynamics of an evolvable physical medium

An interesting area that uses evolutionary methods to exploit unconventional dynamics is that of evolvable hardware for robot control. Evolvable hardware (EHW), or evolutionary electronics [117], is the application of evolutionary search algorithms to the design of electronic circuits [24, 25, 41, 75, 76, 96]. This is a field in which our group has made recent advances, namely, the first demonstration of transistor level analog electronic controllers evolved directly in hardware for non-trivial visually guided robot behaviours [40]. This section describes this work and puts it in the context of recent developments in evolutionary robotics (ER) [90, 123]. In particular, it highlights how unconventional dynamics can be exploited for sensorimotor control by evolving controllers in a physical medium, thus tapping directly into its spatiotemporal properties, and how new kinds of fast simulations can be used to evolve behaviours off-line that seamlessly transfer to reality.

EHW has spawned many highly unconventional circuit designs that operated in very different - often superior - ways to conventional hand designed systems [65, 118]. One of the very earliest works in EHW was the evolution of a hardware control system for a physical robot [116]. Evolved digital circuits for controlling similar, or slightly more complex, robot behaviours soon followed [46, 61, 89, 101]. These later systems were all based around bespoke or proprietary (digital) reconfigurable gate-level circuitry (e.g. FPGAs).

However, a crucial aspect of the original work [116] was the fact that it pointed towards the potential power of evolved *analog* processing in the robotic context [118]. The point was that the potentially rich dynamics of such a system was exploited by the evolutionary process to mesh with the dynamics of the robot-environment interactions arising as the robot behaved in the world. Thus tight, efficient sensorimotor loops, running through the environment and the hardware, were evolved, illustrating the power of unconventional electronics created by the EHW approach, as well as feeding into dynamical systems understandings of embodied behaviour [9, 42, 53].

The potential dynamics of unconventional evolved analog circuits is particularly rich. This, combined with the insight from evolutionary robotics [90, 123] that dynamically complex neural networks are highly evolvable [10, 54], suggested that analog EHW might be very well suited to evolving compact controllers operating with small numbers of components, even for visually guided behaviours which traditionally employed high levels of processing [12]. The technical difficulty of evolving component level analog robot controllers has meant that,

until now, this possibility has been left unexplored. An important property of dynamically complex evolved neural networks is their ability to cope with noisy, poor quality sensory data, even when the networks have very few nodes [54]. This suggests that analog EHW might also be a useful approach for low cost realtime hardware applications requiring cheap sensors and simple circuits. In order to explore this hypothesis further, we have conducted a series of experiments where controllers for an autonomous robot were evolved directly in analog hardware. This was done in such a way (mainly through the deliberate use of low grade sensors) that the study has some relevance to the kinds of applications mentioned above.

3.1 Experimental setup

In order to allow the best chance of creating and exploiting rich unconventional dynamics, an attempt was made to evolve circuits at the analog component level. That is at the level of basic electronic components such as transistors and resistors (in our case primarily transistors). Previously there had been success using field programmable analog arrays (FPAAs) as a substrate for the evolution, in hardware, of artificial neural networks (ANNs) for robot control [14]. That work focused on evolving at the higher level of an analog implementation of an ANN, rather than the lower, component level of analog circuits. Because of the inherent technical challenges, there had been no investigation of component-level analog EHW applied to robotic control prior to the work described here.

To allow an exploration of transistor-based analog circuits, all experiments used the Heidelberg field programmable transistor array (FPTA) [71], a non-commercial research chip specifically developed for EHW applications. This FPTA is a 16x16 array of virtual transistors with configurable local routing. Each virtual transistor has a configurable channel width and length. Routing is such that cells can be bypassed. All edges of the FPTA contain I/O blocks (IOBs) which can be configured for buffered/direct input/output. In buffered I/O mode, DACs and ADCs are used to generate and read voltages from the chip and transfer them through a host FPGA board. The aim was to evolve transistor circuits on the chip to control a mobile robot by feeding robot sensor readings into the chip as inputs and using chip outputs as robot actuator signals (both via wireless connections). The configurable chip allows us to evolve the individual transistor properties (width and length, which determine the current-voltage characteristics), and the way in which the transistors are connected together on the array. Sensor inputs and motor outputs were mapped to/from the native chip IO range [0,5V].

An incremental approach to evolution [47] was used. Task difficulty increased from stage to stage; the initial

population was seeded randomly and then each new stage was seeded with the population from the previous stage's final generation. Preliminary experiments indicated that this was the most efficient way to proceed, in keeping with previous explorations of this issue [90, 123]. The following series of behaviours was evolved: obstacle avoidance, visual target approach in an empty environment, visual target approach in a complex cluttered environment. The environment and the fitness function were changed at each new stage to develop the desired new behaviour.

All experiments used a K-Team K-Junior wheeled mobile robot [70] with a SmartEvo vision turret having a conical mirror mounted above its upward facing camera to provide 360° vision (Fig. 1). The camera lens, focus and orientation were fixed such that the full panoramic image occupied as much of the camera field of view as possible (Fig. 1). The robot is circular with a 6.5cm radius. Five of its six IR proximity sensors are distributed within a forward facing 180° arc at the front, and the sixth faces backwards at centre rear.

Instead of using the whole camera image as visual input to the controllers, five Haar-like feature detectors (or filters) [124] (whose position, size and type were genetically set) were used to effectively evolve visual sensors which use only part of the image and pre-process it into a low dimensional input for the FPTA. Each of the five filters is genetically set to be one of the eight types illustrated in Fig. 1. This approach, where visual sensors are co-evolved with the controller, builds on [48, 54] but uses more sophisticated filters than in previous work. It is a powerful approach to automatically achieve dimensionality reduction and feature extraction and selection in an integrated way. Haar-like feature detectors calculate the sums and differences of average intensities over adjacent rectangular regions (the black and white regions shown in Fig. 1). They can act as simple edge and line detectors, as well as responding to more complex visual features and picking

out areas of high contrast. A number of these feature detectors acting in concert have been shown to work well in robot navigation, providing an efficient, low dimensional representation of visual data [8]. The output, s (in the range [0, 1]), of a Haar-like feature was calculated by overlaying the filter on the 360° conical mirror view as in Fig. 1 (right) and applying the following equation:

$$s = \sum_{p \in W} v_p + \sum_{p \in B} (255 - v_p) 255 \times (|W| + |B|) + \Gamma \quad (1)$$

where W and B are white and black filter regions respectively, as in Fig. 1, v_p (in the range [0, 255]) is the value for pixel p and Γ is uniform random noise in the range [0, 0.1]. Areas of the visual sensors falling outside of the visual field are ignored.

A generational evolutionary search algorithm was used with a binary genotype, linear rank selection, single point cross-over, mutation and elitism. After preliminary investigations, a population size of 30 was used. With only a single FPTA, which had to be used for each (expensive) evaluation, this population size proved a good compromise. Each member of the population was evaluated in turn by using the FPTA circuit it described to control the robot. Preliminary experiments indicated a cross-over probability of 0.6 and a per bit mutation rate of 0.0016 as the best values to use. Each genotype was a fixed length binary string which encoded a FPTA configuration using 6144 bits describing transistor properties and connections for the array, as defined by the chip configuration protocol [71]. This binary encoded protocol is hardwired into the chip design and *must* be used to configure the FPTA. Hence it made sense to use it directly as the genetic encoding as intended by the chip designers [71]. An extra 120 bits were appended for visual sensor configurations (24 bits each for 5 evolved visual Haar-like filters) determining their sizes, positions and other properties. Numerical values




Fig. 1 Left: the K-Junior robot with the SmartEvo turret having a conical 360° mirror mounted above the upward facing camera; Middle: Haar-like filters available as visual sensors; Right: example of filter ID

6 overlaid on the conical mirror 360° view. White regions are additive and black regions subtractive – see text for details

were represented using a Gray code [44] so that single bit mutations cause incremental changes in values resulting in a smoother fitness landscape. Each filter effectively acted as a separate visual sensor, feeding into its own FPTA input.

All robot behaviours were evolved to take place in a 85×114 cm rectangular arena with bounding walls. Various obstacles were introduced for some experiments and visual features (such as geometric patterns) were stuck on the walls in some places.

Robotics is inherently noisy. Sensor and actuator noise and natural variations in environmental conditions (e.g. lighting) are always present. In our case the physical medium used for the controller (unconventional FPTA circuits) can provide another source of inherent noise (e.g. parasitic capacitance build-up). Because of this, nominally identical fitness evaluations will result in different fitness values. Since we require the controllers to be robust to such variation, as well as to different initial conditions, the evaluation method must be carefully designed. Multiple trials must be used and these should be appropriately weighted in order to produce a selection pressure towards general and robust behaviours.

A set of N fitness evaluations, f_i , were integrated into a final overall fitness value F for each individual in the experiments described below. By giving a heavier weighting to lower fitness values as per [54], robustness is encouraged - the robot must behave well on *all* trials. To achieve this, fitness function F (2) was used in all experiments. The scores, f_i , were ranked (i is the rank, rank 1 is best (highest), rank N worst (lowest)). Thus F weights the evaluations scores inversely proportionally to fitness.

$$F = \frac{2}{N(N+1)} \sum_{i=1}^N i f_i \quad (2)$$

Obstacle avoidance was the basic first level behaviour on which the others were built. The aim of this task was for the robot to cover as much ground as possible without colliding with obstacles. 6 FPTA I/O pins were allocated as inputs from the six robot IR sensors, with values normalized to the [0,5]V range. Values are high when the sensor is close to an obstacle. Vision was not used in this experiment. 2 pins were allocated as outputs to the robot motor. An S shaped obstacle was present in the arena during evolution. Fitness was calculated over 6 trials starting from random positions and orientations. Each trial was scored as in [37, 57], using (3).

$$f_a = \sum_{t=0}^{t=\text{trialEnd}} V(1 - \sqrt{\Delta v})(1 - i) \quad (3)$$

where V is the sum of the instantaneous rotation speed of the wheels (stimulating high speeds), Δv the absolute value of the algebraic difference between the speeds of the

wheels (stimulating straight line forward movement), and i is the normalised value (in range [0,1]) of the IR sensor of highest activation (stimulating obstacle avoidance). This simple behaviour has been achieved many times with various ER/EHW approaches and is not in itself particularly interesting. However, it is a good basic test of the FPTA approach and is necessary as an initial bootstrapping behaviour in our incremental methodology (Fig. 2).

The aim of the second task in the sequence was for the robot to approach a target which it can only recognize by using vision (a red 29cm x 21cm rectangle next to a black 29cm x 21cm rectangle on one of the walls, see Fig. 3). Other visual features (patches of colour or patterns) were stuck to the arena walls and could potentially be used for orientation, but also had to be discriminated from the target. At this stage the evolution of visual sensors was switched on as outlined earlier. A further 5 FPTA pins were allocated to the evolved visual sensors. The robot was to approach the target from a random starting position and orientation in an otherwise empty arena without crashing into the walls (i.e. maintaining the obstacle avoidance behaviour alongside the visual target finding).

Behaviour was evaluated over 6 trials, each starting from a random location and orientation in the arena. The score for an individual trial was calculated as follows:

$$f_g = \frac{\sum_{\forall t, d_t \geq 10.5} (1 - \frac{d_t}{d_1}) + \sum_{\forall t, d_t < 10.5} 2}{t_{end}} \quad (4)$$

where d_t is the robot's distance to the target at time-step t , and t increases from 1 until the trial ends at t_{end} due to robot collision or timeout, whichever comes sooner. Double scores were awarded on every time-step the robot spent near the target without colliding with it. Maximum trial length was 80 time-steps. Versions of this behaviour have been achieved before [48] but here it is used as an incremental step towards the next stage of more complex behaviour. It is also a good validation of the coevolution of visual sensors approach and is the first time an EHW approach has been used for a visual behaviour (of course evolved FPTAs have never previously been used for robotics tasks).

The third task in the incremental sequence was the same as the second but in a complex, cluttered, maze like environment such that from many locations - more than 50% of the arena - the target was hidden (see Fig. 3). This was considerably harder than the empty arena task as a more general visual searching strategy had to be evolved in order to robustly find the target from any starting position in the environment. The evaluation procedure was exactly the same as in the previous task, except the maximum trial length was 200 time-steps. This task is more

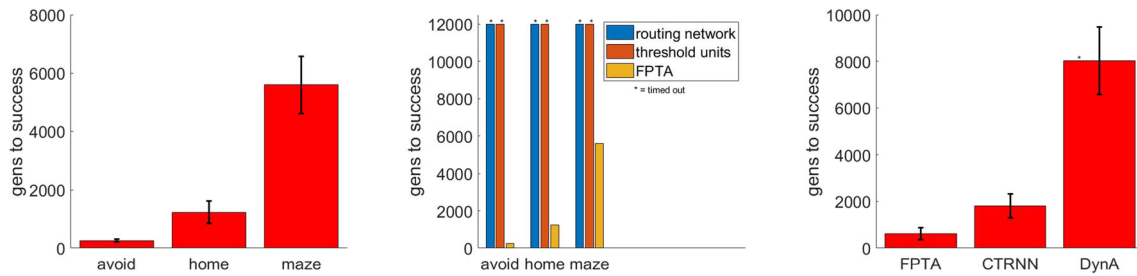


Fig. 2 *Left*: Results of ten incremental evolutionary runs. The bar heights show the mean number of generation to a successful robust solution at each stage, error bars indicate standard error of the means (SEMs). *Middle*: Full FPTA compared with FPTA used as a routing network and with simple threshold units. An asterisk above

a bar indicates all runs under that condition timed-out before success was achieved. *Right*: FPTA compared on avoider behaviour with an array of dynamical units and with a dynamical neural network with a pre-defined architecture

difficult than most previous ER visual behaviours and has not been attempted before [90]. For further details of the experimental setup see [40].

3.2 Simulating the robot and its environment

The evolutionary robotics approach requires large numbers of candidate controllers to be evaluated. In our case this creates the challenge of either evaluating behaviours on the physical robot within a reasonable timeframe, or developing sophisticated enough simulations to allow general visually guided behaviours to evolve that transfer into the real world without loss in performance. While a number of approaches to this latter problem have been developed over the years [56, 90], general vision simulation techniques that cross the reality gap [57] for any but the most low resolution cameras have proven elusive [90, 123]. We tackled the problem by developing new methods for accurately simulating the robot acting in its environment. This allowed us to evolve behaviours using a robot simulation which successfully transferred to the real robot. During evolution the simulated robot was connected to the (real) FPTA in the same way as the physical robot is when evolved controllers are used in the real world.

An efficient physics-based simulator was written in the style of [57] to model the kinematics, assuming a flat floor. Careful empirical measurements were used to model motor responses and sensor readings. Motor and sensor noise was introduced at a level empirically determined to match the real behaviour [40]. The simulator operated at a configurable discrete time step.

The method used to simulate vision employs an empirical sampling technique made feasible by the use of a 360° field of view. The technique builds on a method previously utilized in [8], but with significant extensions and improvements. As far as we know it is the first time such a technique has been used in an evolutionary robotics context. The basic idea was to divide the whole world (the robot arena) into a set of equally sized cells. The image seen by the robot was then sampled in each grid cell to build up a database representing the robot's visual world. Because the robot has 360° vision, the panoramic image at a given location is essentially the same for any orientation of the robot, it has just been rotated. Hence, instead of having to sample at each location for many orientations, a small number of samples is sufficient. The retrieved image can be easily mathematically rotated to match the actual robot orientation. It is this trick – which relies on the rotational

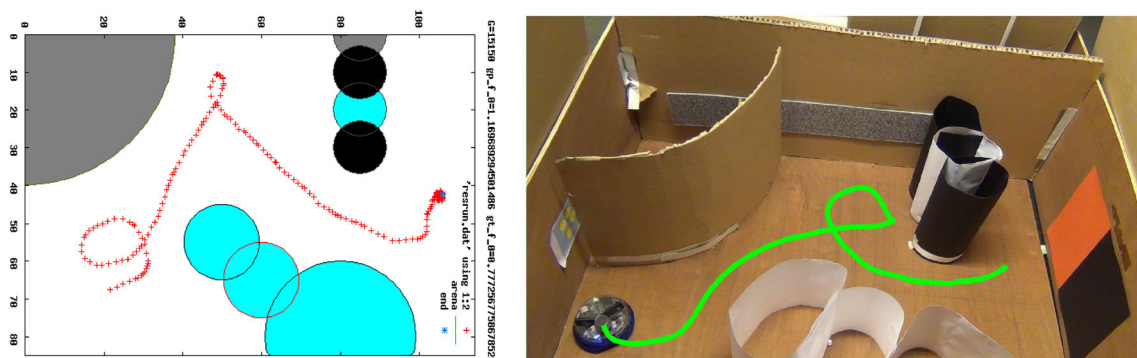


Fig. 3 Successful evolved FPTA controller navigating the maze environment *Left*: in simulation and *Right*: when transferred to the physical robot and started from a similar position

symmetry of the 360° image – that makes the technique feasible, otherwise the number of samples needed would become too large.

The arena was sampled by capturing the world as seen by a north and a south facing robot at every location on a grid of 5 × 5cm cells. At any given moment during simulation, a north or south orientation is chosen at random. The image sampled in that direction is picked from the sampling cell closest to the current position of the simulated camera. The chosen 360° image is then rotated according to the simulated robot orientation. The simulator also added noise to the sampled image (at empirically determined levels). The addition of noise and use of randomly chosen samples (north or south orientation), as well as variations in obstacle positions, forces the evolved controllers to be robust to a range of visual conditions rather than relying on a fixed set of values. Such robustness is essential for transferring to the real world and operating in realistic conditions. The discrete nature of the sampling, and the use of the *nearest* sample to the actual position of the simulated robot, adds further noise and coarseness which increases the pressure to produce general, robust solutions [40]. In some of the experiments using vision, additional features and obstacles were introduced into the modelled arena through Computer Generated Imagery (CGI) injection into the sampled images. Hence the original sampled world can become the basis of new, more complex environments.

3.3 Evolved FPTA results

Figure 2 summarizes the results of a number of incremental evolutionary experiments using the FPTA as an evolvable medium. The bar chart on the left shows the results of ten incremental experiments on the sequence of robot behaviours described in Section 3.1. The height of the bars shows the mean numbers of generation to a robust, successful solution at each stage (that is a high scoring solution that remains the best in the population for 30 generations with re-evaluations on each generation, thus eliminating ‘lucky’ individuals that scored well once; a high score for obstacle avoidance is 500, for both the visual tasks it is 1). The error bars show the standard error of the mean. All runs were successful at each stage, with moderate standard errors, indicating that the methodology is highly robust.

The evolved successful controllers transferred very well to the real robot, producing qualitatively very similar behaviour to the simulation. A trace of the live robot engaged in the end stage behaviour (find target in complex cluttered environment) can be seen in Fig. 3, directly compared with the behaviour in simulation. The robot trace was extracted from an overhead video using motion tracking algorithms. Even though the controllers were evolved

to produce very efficient behaviour in the environment shown in Fig. 3 (with variations during evaluation trials as outlined earlier), and exploited robot-environment dynamics particular to this environment, our methodology still made them general enough to be able to successfully perform the task (find the target and stay at it) in unseen variations of the environment where the target had been moved to a different location or the shape of the environment has been altered [40] (see Fig. 4). This demonstrates that the successful controllers were processing sensory information to generate the behaviour, rather than using some trick to blindly learn the shape of the environment and location of the target.

The results of incremental evolution suggest the FPTA is a suitably evolvable medium for developing robust sensorimotor behaviours, even when the sensors and motors are noisy and unreliable (there was also a 600ms camera latency). The observed, rather subtle, dynamics of the evolved behaviours (especially the visually guided behaviours) suggest the potentially rich dynamics of the FPTA medium are being exploited. Analysis of evolved circuits further supports the importance of the exploitation of analog FPTA dynamics (see [40]).

The middle and rightmost bar charts in Fig. 2 show the results of a series of experiments aimed at probing the hypothesis that the FPTA’s evolvability is at least in part due to the exploitation of rich dynamics.

It is possible to effectively bypass all the transistors in the FPTA by fixing all routing to be ‘pass through’ so that the chip becomes an evolvable routing network. In this mode it is possible to evolve routing networks that can connect sensors and actuators in potentially complex (or relatively simple) ways but which no longer make any use of the transistors and the potential dynamics they can impart. The middle bar chart shows results of ten comparative incremental evolutionary runs, using the same series of behaviours, of the full FPTA, the FPTA used as a routing network (transistors bypassed), and an array of simple threshold units. *All* the full FPTA runs were successful at each stage. It is clear from the plots that *none* of the runs of the FPTA as routing network or of the array of threshold devices were successful at any stage. All of these runs timed out at 12,000 generations. Multiple runs with the threshold devices using a fixed threshold of 0.5, or with evolved weights (range: 0-2) on the outputs were equally unsuccessful, as were runs with an array of different simple linear units (no longer resembling transistors, but still without dynamics) [40].

To explore the dynamics issue further a new set of comparative runs was executed on the avoider behaviour. This time the FPTA was compared with units with explicit dynamics. Results of these experiments (ten runs of each) are shown in the rightmost bar chart. An array of dynamical units (DynA in the figure) was created by replacing the




Fig. 4 The evolved controller shown in Fig. 3 generalising to environments unseen during evolution. The black cross in the behaviour plots indicates the location of the target on one of the walls/barriers (dark filled areas). *Left*: target moved to lower barrier (during evolution it had always been on the top wall), *Middle*:

target moved to horizontal barrier, *Right*: target in original position but part of the diagonal barrier removed to create variation in the shape of environment

threshold units in the simulated analog array, with units described by the following (leaky integrator) differential equation:

$$\tau \frac{dy}{dt} = -y + I \quad (5)$$

where τ is an (evolvable) time constant (range:[0.3, 10]), y is the unit's 'activation' and I its total input. A unit's output, O , is given by:

$$O = \frac{1}{1 + e^{-(y+\theta)}} \quad (6)$$

where θ is an evolvable bias (similar to a threshold, range: $[-5, 5]$). This form of equation is widely used in neural modelling and in dynamical neural networks employed in robotics [9, 31]. The connectivity, biases and time constants for all units in the array were evolved using the same machinery and encoding as in the previous experiments. Out of 10 runs with the DynA setup, 5 were successful, the rest timed out after 12000 generations. The DynA bar in the figure is a little misleading as it uses 12000 generations for the unsuccessful runs to calculate the overall averages and standard errors. This is overgenerous but allows some kind of visual comparison with the other runs. Of the 5 successful runs, the mean number of generations to success was 4056 ± 1220 . A fresh set of ten runs with the FPTA were all successful. The mean and standard error, shown in the plot, were similar to those found on the other earlier sets of runs, illustrating the reliability of the FPTA as an evolvable medium for noisy sensorimotor tasks. The partial success of the DynA experiments further suggested

that dynamical capabilities were an important factor in success at this task, which is made relatively tricky by the noisy nature of the simulation (reflecting the noisy nature of the sensors and motors on the physical robot, and used to enable robust transfer to reality as explained earlier).

To examine this in more detail, a set of ten runs were performed which evolved a widely used kind of dynamical neural network [10] with a pre-defined architecture suitable for the task (labelled CTRNN in the bar plot). These runs moved away from the 16x16 array setup, which is directly comparable to the FPTA and in which control architectures have to be 'carved out' by evolving suitable circuits, to the more constrained problem of setting the variables on a fixed neural network. For this case the genotype was an array of reals encoding the network variable as explained below. For comparability, the same form of evolutionary algorithm was used as with all previous experiments, except now the mutation operator was a random variation in a gene based on a normal distribution centred on the current value with std 0.2 (a so-called creep operator). A fixed architecture three layer 6-3-2 network was used. Each layer was fully connected to the previous layer. The six nodes of the input layer each took an input for one of the 6 IR sensors, the two nodes of the output layer provided the motor signals. The 3 nodes in the middle layer also had recurrent connections to themselves and each other. All connections had evolvable weights (range: $[-5, 5]$). Nodes in the network operate according to the following equations:

$$\tau_i \frac{dy_i}{dt} = -y_i + \sum_{j \in \Phi} w_{ji} \sigma(y_j + \theta_j) + I_i \quad (7)$$

$$\sigma(x) = \frac{1}{1 + e^{-x}} \quad (8)$$

where τ_i is a genetically set time constant for the i th node (range: [0.3, 10]), y_i is the node's activation, Φ is the set of all nodes connected to node i , w_{ji} is the genetically set weight on the connection from node j to node i , θ_j is the genetically set bias for node j (range: [-5, 5]), and I_i is any external sensory input to node i . The summation is the total weighted input to node i from all other nodes connected to it (the output of a node is $\sigma(y + \theta)$). All ten CTRNN runs were successful. The high performance of this more elaborate and targeted architecture adds weight to our theory that the ability of evolution to manipulate internal controller dynamics is very useful in finding good solutions in this context.

Although all of both the FPTA and CTRNN runs were successful (as defined at the start of Section 3.3), it can be seen from Fig. 2 that both the average and standard error of the FPTA runs (617 ± 264) is lower than that of the CTRNN runs (1806 ± 513). A non-parametric Mann-Whitney U test revealed that the FPTA is significantly better than the CTRNN at the 95% confidence level ($p=0.049$). Clearly the DynA runs were much worse, half of them not completing successfully.

4 Insect-inspired visual route navigation: Exploiting the embodied dynamics of innate behaviours

The previous section shows how navigating to a goal can be achieved using visual information, which is referred to as visual homing. The use of visual information for navigation is a universal strategy for sighted animals, amongst whom desert ants are experts. Despite having brains of under 1 million neurons (100,000 times fewer than in the human brain) and low-resolution vision equivalent to a 0.001 MPixel camera, desert ants learn long paths through complex terrain after only a single exposure to the training data [62]. These features make ants a model species for both biologists attempting to understand the minimal cognitive requirements of spatial learning as well as engineers seeking to emulate their feats in autonomous robots. But how is such rapid learning achieved with such small brains? At Sussex, we have used an interdisciplinary methodology to answer this question, combining biological experiments with computational and robotic modelling [95]. We have in particular shown that learning is an active process scaffolded by specialised innate behaviours which have co-evolved alongside bodies and brains to allow ants to directly acquire and use task-specific information [131, 132]. These behaviours serve to structure the dynamic flow of information during both

learning and recall, outsourcing computation to the body, and thus enabling complex behaviour to emerge without complex processing. In this section we describe this ongoing research. We show that morphological body constraints that may at first seem limiting, have been exploited by evolution to produce minimal yet highly effective mechanisms for visual navigation in ants – mechanisms that we demonstrated on a mobile robot operating in dynamic outdoor environments.

4.1 Innate behaviours scaffold route navigation in ants

Ants have a suite of innate behaviours which allow robust navigation to emerge. Principal among these is path integration (PI), a mechanism by which an ant keeps a running tally of the distance and direction travelled, allowing it to subsequently home directly back to the start point of a journey. As path integration is subject to cumulative errors, ants learn the visual information needed to guide later routes on the first PI-mediated route. This innate behaviour thus provides a scaffold for learning, constraining the information to be learnt, turning a route into a one-dimensional manifold through a two-dimensional space. However, the ant's embodiment further constrains the incoming information making it even easier to learn and later recall. To understand how embodiment simplifies visual navigation, we first note that if an agent stores a view when facing a given direction, the difference between this view and views from nearby locations will be minimised when the agent is facing the same direction as when the original view was stored [135]. This means a remembered view can be used as a so-called 'visual compass' to recall the direction the ant was facing when the view was stored. For ants, and many wheeled robots, their embodiment constrains them to move in the direction that they are facing and thus a view stored when travelling along a homeward route implicitly defines the direction of movement and thus learning these views as they appear means learning the directions home. This process is illustrated in Fig. 5a where we see a stored view (top panel) and the views perceived by an ant while scanning (middle panels). The closest match occurs when the ant faces in a similar direction to when the view was stored (bottom panel). If the ant was sensitive to this difference, it could use this process to recall the direction it was facing when last near to that position. What's more, when the ant finds the best matching direction, it is already facing the same way and so can simply move forward. To navigate, the ant therefore needs to determine at what heading the current view best matches its memory. However, mentally rotating the images is a computationally demanding task so the ant has outsourced this search to behaviour. The most obvious




Fig. 5 Visual navigation is scaffolded by innate behaviours. A: Innate scanning behaviour embodies a visual compass. If an ant remembers a panoramic view (top panel) it can recall its facing direction by comparing the stored view to views it perceives while rotating or scanning (middle panels; images of an ant scanning are to the right of the panoramic views). The minimum difference

of these is a scanning behaviour, in which the ant will periodically rotate on the spot (images of an ant during scanning are shown on the right of the middle panels of Fig. 5a). The frequency of these scans is higher in unfamiliar environments, suggesting that they are active processes driving navigation and essentially implementing an embodied visual compass [132]. However, the kind of active sampling that scanning brings is also evident in the modulation of the basic sinuosity of ant's paths, where wiggling is upregulated when ants are uncertain [20, 133] and hence would benefit from increasing the sampling rate of visual scenes. These innate behaviours allow the problem of learning and navigating a route to be simplified. During learning, PI constrains the ant's experience and actions so that they are implicitly correct in the sense of leading to the goal (black path and arrow heads in Fig. 5b). These actions are memorised as stored views which can be recalled by physically scanning the world (cartoons and blue arrows in Fig. 5b). In this way, navigation is reframed in terms of an embodied search for familiar views, as, when an agent is facing in a familiar direction, it is likely facing a similar direction to when it was previously at that location so should move in the direction it is facing.

between the images will occur when it is facing a similar direction as when it stored the first view. B: If views are stored as they appear to an agent travelling a route (black line) the correct directions can be recalled when near the route by an agent performing a visual compass style scan (blue arrows)

4.2 Route navigation in robots

Based on the above observations, we have developed a parsimonious insect-inspired navigation algorithm in which a route, or routes, are learnt holistically and route recapitulation is driven by a search for familiar views [7]. The algorithm proceeds as follows: an agent equipped with a low-resolution 360° panoramic visual sensor first travels a route. The views it experiences along this route are used to train an artificial neural network (ANN) which learns a holistic representation of the views encountered. Subsequently, the network is used to estimate the likelihood of whether a given view – and thus a pose – has been experienced before. When trying to repeat the route, the agent derives a direction of movement at a position by visually scanning the environment (either by physically rotating or rotating the view *in silico*). Each rotated version of the current view is applied as an input to the network which outputs an estimate of its familiarity. The agent then moves in the direction corresponding to the view most similar to those encountered during learning.

To estimate view familiarity we follow [7] and use a neural network model that was specifically designed to

perform this task [78]: Infomax. We chose to use this approach mainly because it only requires a single pass through the data, meaning that each view is experienced just once and then discarded lending it biological plausibility as the training data does not need to be memorised. Further as views need only be judged familiar or not, rather than individually recognised as in place recognition, the representation is sparse. The network consists of an input layer and a novelty layer with $\tanh()$ activation functions. The number of input units is equal to the dimensionality of the input which in our case is $[120 \times 25] = 3000$, the number of pixels in a down-sampled view of the world. The number of novelty units is arbitrary and here we use the same number of novelty units as inputs, although using fewer novelty units has worked in simulation and will be tested in future work. The network is fully connected by feedforward connections w_{ij} . Weights are initialised randomly from a normal distribution, normalised so that the mean of the weights feeding into each novelty unit is 0 and the standard deviation is 1. The network is then trained using the Infomax learning rule [13], adjusting the weights so as to maximise the information that the novelty units provide about the input, by following the gradient of the mutual information using (11) which performs gradient ascent using the natural gradient [4] of the mutual information over the weights [73]. During learning the activation of each of the M novelty units h_i is computed as:

$$h_i = \sum_{j=1}^N w_{ij} x_j \quad (9)$$

where x_j is a row vector assembled by concatenating the rows of $C(a, \theta)$ and $N = p \times q$ (the number of input units). The output y_i of the novelty units is then:

$$y_i = \tanh(h_i) \quad (10)$$

and the weights are adjusted using:

$$\Delta w_{ij} = \frac{\eta}{N} \left(w_{ij} - (y_i + h_i) \sum_{k=1}^N h_k w_{kj} \right) \quad (11)$$

where η is the learning rate which is set as 0.0001 for this paper. Finally, the response of the network to the presentation of an unseen N -dimensional input x is computed as

$$d(C(a, \theta)) = d(x) = \sum_{i=1}^N |h_i|, \quad (12)$$

where $|\cdot|$ denotes the absolute value. By applying $C(a, \theta)$ to the ANN for a range of θ , a Rotational Familiarity Function (RFF) can be calculated from $d(x)$ and hence the most familiar direction can be found.

As a control condition, dubbed the Perfect Memory algorithm, we store all memories seen during training rather

than using them to train an ANN. This provides a baseline performance for how well algorithms which navigate via a visual-compass-style matching can perform if storage and computation are not constrained. In the Perfect Memory algorithm, each rotated version of the current view is sequentially compared to every one of the training views. The best matching heading is then defined as the one corresponding to the smallest image difference, across all training views and rotations, with image difference (IDF) calculated as the mean absolute difference between each of the image pixels:

$$\text{IDF}(C(a, \theta), S(b, \phi)) = \frac{1}{p \times q} \sum_{i=1}^p \sum_{j=1}^q |C_{i,j} - S_{i,j}| \quad (13)$$

where $C(a, \theta)$ is a $p \times q$ pixel view captured at location a with heading θ , $S(b, \phi)$ is a $p \times q$ pixel snapshot stored in memory and $C_{i,j}$ and $S_{i,j}$ refers to the intensity of pixels in row i and column j of the captured view and stored snapshot respectively.

We first tested our algorithms using a training route of approx. 60 m through Stanmer Park (50.8634175 deg, -0.093938 deg), a grassy area with some trees shown in Fig. 5b. As a training set we gave the robot two passes through the route. Training images were gathered every 200 msec providing a dataset of 778 images (across the two passes). To assess performance, we drove the robot once more across the route and extracted test positions every 200 msec. The view from each test position was rotated through 360° and the best matching heading recovered from the Infomax and Perfect Memory algorithms. The resultant headings are shown by the arrows in Fig. 6a. Both algorithms perform comparably with mean errors across the test route of 19.4° and 21.4° for the perfect memory and Infomax algorithms, respectively. The Infomax network is however faster: the mean calculation time per image was 8.8 ms compared to 14.4 ms for perfect memory. Note also that whereas for perfect memory, computation time scales linearly with the number of training images, it is constant for Infomax as the network size is fixed. Hence, for routes with thousands of images (i.e. with a longer route or recording at a greater frame rate), Infomax would have a considerable efficiency advantage.

To assess how compact the route memory can be and, as a corollary, how quickly our algorithm can run, we repeated our analysis using ANNs with decreasing numbers of units in the hidden layer. The results are shown in Fig. 7. Surprisingly, the performance is somewhat insensitive to the number of hidden units with similar performances being seen for as few as 10 units. Going much lower than this seems to result in somewhat erratic performance which can be seen by the increasing spread in the data. The higher spread is significant, as we are more concerned with the

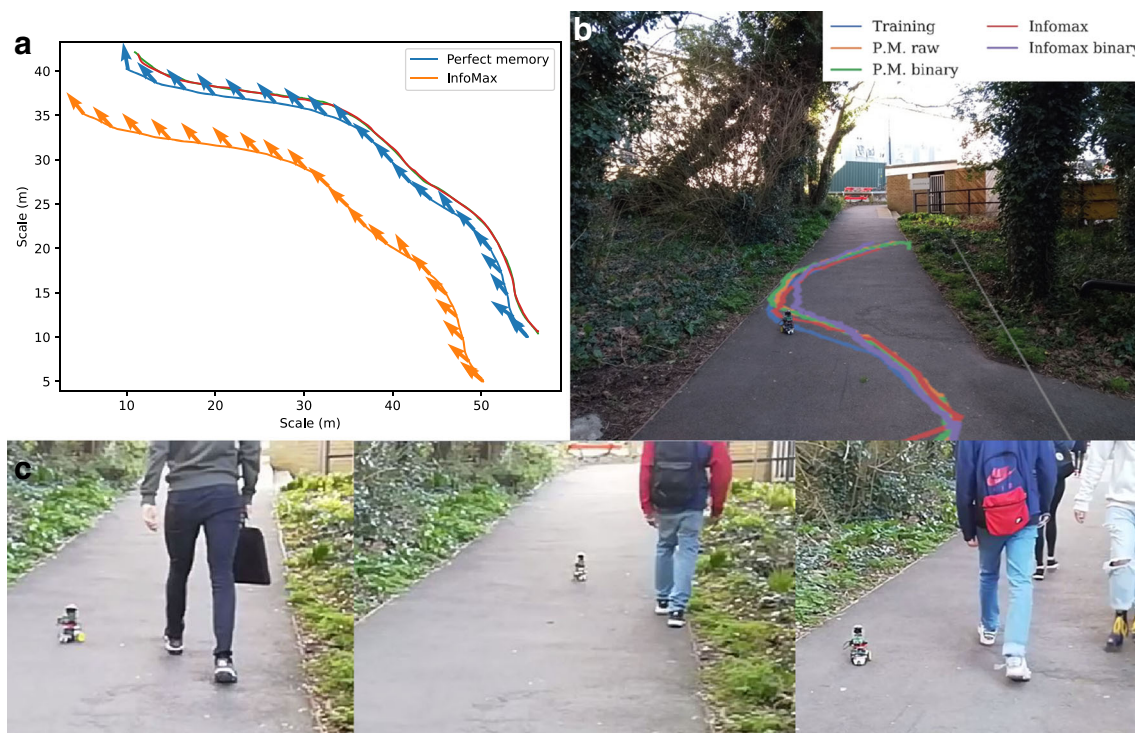


Fig. 6 Route navigation on an autonomous robot. **a:** Both perfect memory and Infomax provide accurate directions over the ca. 60 m route illustrated in Fig. 5b. The red and green lines (the latter is largely obscured by the red) show two passes over the route used for training, the blue shows performance for perfect memory and orange for Infomax. Note that blue and orange routes start at the same starting position as red and green training routes, but have been displaced as they would otherwise obscure each other; the axes are purely for scale. The arrows show the estimated heading direction given at each test point for the two algorithms. **b:** Closed-

loop trials show that route following also works on a real robot platform in a real-world environment. The robot has a panoramic camera and finds its heading by performing periodic scans of the visual scene *in silico*, then heading in the most familiar directions, after first traversing the blue training route. To assess robustness to contrast changes, we used raw images as well as images automatically converted to binary as sky/not-sky (see legend). **c:** Performance is not seriously affected by conspicuous changes to the visual scene, such as passing pedestrians

number of large errors rather than small ones, as large errors disproportionately lead to failure. It may be that the performance of the larger networks would improve with more training data or in different environments, which we will investigate in future work. However, the performance for small networks is very promising as the network calculations scale roughly linearly with the number of hidden units, with small networks providing very rapid heading estimates (approx 0.3 ms per image for the network with 10 hidden units). Having many heading estimates in a short space of time can be very useful in avoiding catastrophic errors as the image data can be very noisy between timesteps due to uneven ground and sudden lighting changes such as lens flare, as it allows the agent to poll across multiple samples and discard noisy estimates.

While this shows that our algorithm is feasible for long routes, the robot is not autonomously navigating. To assess navigational success in this case, we used a shorter paved path of approximately 10 m through a wooded area [63]. The robot was manually driven along the blue path in

Fig. 6b recording training images every 100 msec, which resulted in a dataset of 455 images as training data. To subsequently navigate, we ran the following algorithm every 500 msec:

1. Capture and unwrap a panoramic image and perform any image processing
2. Using either the Infomax or Perfect Memory algorithm, calculate the familiarity with the processed image ‘in-silico’ when rotated through 90° .
3. Find the orientation with the highest familiarity and, if it is within 4° , start driving forwards. Otherwise, start turning in the correct direction to align with the image.

The robot successfully recapitulated the training path using each of the navigation and image processing algorithms with little difference in performance apparent (Fig. 6b; mean distance between training and recapitulated paths was 9 cm with standard deviation of 8 cm, within the margin of error for extracting robot location from video). This is perhaps surprising given that the robot was




Fig. 7 Distribution of heading errors for the Infomax algorithm when using different numbers of hidden units. Red lines show medians, boxes show 25th and 75th percentiles, whiskers extend to 5th and 95th percentiles respectively. Circles show outliers which are points outside the whiskers. Note that performance is comparable even with very small networks, with performance declining only for five hidden units

navigating while people walked up and down the path (Fig. 6c) but is a result of the algorithm being able to use wide-field, low-resolution information as it does not need to extract or match features. Overall this shows that the algorithm is both robust to the precise implementation used and conspicuous disruptions.

5 Control from chaos: Exploiting complex embodied neuromechanical dynamics

While the FPTA evolutionary robotics work described in Section 3 used evolutionary search to find ways of exploiting the rich dynamics of a physical medium to develop robust robot controllers, it is also possible to produce embodied robotic system with complex exploitable dynamics through a combination of careful design and continuous adaptation. If a robot has suitable adaptive mechanisms it can learn to exploit its own embodied dynamics to produce a desired behaviour. But the potential dynamics of the system need to be rich and exploitable for this to work effectively. Chaotic dynamics turn out to be very useful in this context. In this section we discuss very recent work from our group which shows how, for a class of embodied neuromechanical systems, it is possible to efficiently exploit chaos in the development and learning of motor behaviours for bodies of arbitrary morphology. Detailed analyses reveal chaos at all levels of the systems; the entire brain-body-environment system exhibits chaotic dynamics which can be exploited to power an exploration of possible motor behaviours.

Intrinsic chaotic dynamics in the nervous system have long been recognized in neuroscience and have been shown to be integral to the operation of the brain [3, 38, 45, 64, 130]. The existence of such dynamics in both normal and pathological brain states across a variety of species, at both global and microscopic scales [130], supports the idea that chaos plays a fundamental role in many neural mechanisms [110]. Chaotic dynamics are known to operate in brain regions – such as the cortex – that are associated with higher-level information processing [98, 110], and also in neural circuitry responsible for motor behaviours [114]. In many motor behaviours chaos seems to occur not just at the neural level but also within the dynamics of the body [99]. For instance, chaotic movement appears to play a crucial role in the development and learning of limb coordination [91].

Following the seminal work of W. Freeman and colleagues [110], various models were developed to explain the existence and possible role of brain chaos [69, 79], and to show how chaos can enhance learning performance (e.g. of complex rhythmic patterns [51, 113]). The latter case inspired work on the adaptive control of robot motor behaviours by utilizing chaotic attractors as a controllable source of information (i.e. pattern reservoirs) for generating desired behaviours as well as enabling flexible transitions between them. One broad approach that emerged in this area involved the control (stabilization) of chaos by sensory input. Such chaos control has been realized both in wheeled mobile [134] and legged robots [112].

While such models have demonstrated how to harness neural chaos for various sensorimotor, perceptual and learning tasks, they assume that the neural system generates chaotic dynamics in the absence of sensory input, with sensory feedback mainly acting as a stabilizer for chaos. However, the identification of chaotic dynamics in natural motor behaviours from multiple *in vivo* studies [43, 91, 103] suggests that continuous sensory stimuli actively participate in the generation of chaotic dynamics. In particular, the sensory input received while engaged in motor behaviours contains information about the physical body and its environment, emphasizing the *embodied* nature of such chaotic behaviours.

This has been reflected in a strand of research which has proposed a more active and radical exploitation of chaos, where the chaotic dynamics arise in a whole neuro-physical system [67, 68]. These models implemented a ‘bodily-coupled’ neuro-musculo-skeletal system inspired by cortico-medullo-spinal circuits. The neural system consisted of a group of identical electrically decoupled neuromuscular units – each implementing an individual reflex loop, with sensory input, driven by a central pattern generator (CPG, modelled by a neural oscillator). Although each CPG communicated only locally with the correspon-

ding muscle, information was indirectly channeled between CPGs through the inertial and reactive forces from the physical body and its environment, giving rise to a variety of sustained or transient coordinated rhythmic movements which could be spontaneously explored and discovered. Generating chaotic dynamics in these systems is crucial for the exploration of self-organized motor coordination. It requires a proper set of tonic (slowly changing) descending signals (which act as parameters) for each CPG. These tonic ‘command’ signals descend from the brain in most spinal animals and are usually related to sensory input. However, maintaining and controlling chaos in these early models was challenging and it often rapidly dissipated.

Significantly extended models introduced by our group [105, 106] addressed this issue by incorporating an adaptive local neural mechanism to achieve the controllable chaotification of a similar embodied model for use with an arbitrary physical system, where the sensory inputs for CPGs were homeostatically regulated (i.e. maintained within appropriate ranges) while identical descending signals were fed to all CPGs as a bifurcation parameter. This model enabled performance-driven exploration and goal-directed learning of sustained locomotor behaviours in robotic systems of arbitrary morphology. The bifurcation parameter was dynamically adjusted between chaotic and synchronized regimes, in response to a performance feedback signal, to allow the system to escape from low performing behaviours and be entrained in high performing ones (Fig 8). The system performs a kind of ‘chaotic exploration’, where chaotic dynamics power a form of search through the space of possible system dynamics, settling on a high performing configuration. This scheme defines a very general class of models in which there is one adaptive neuromuscular unit for each muscle (actuator) degree of freedom (DoF) in the body, hence it can be applied to many different bodies.




Fig. 8 The overall chaotic exploration concept for motor behaviours. *Left:* Performance feedback is transformed into a descending input to all CPGs which acts as a bifurcation parameter. The level of chaotification of the system is inversely proportional to its performance. Actuator sensory signals pass through a homeostatic adaptive process (SA). *Right:* Overall search dynamics of the method. Chaotic exploration samples the population of attractors

While the behaviour of these performance-driven systems is impressive, until very recently analysis of their dynamics has been limited. In this section, as well as demonstrating the efficacy of the overall approach, we outline detailed analysis of the dynamics, showing that the *whole* neuro-physical system is exhibiting chaotic dynamics that are exploited to generate goal-directed behaviour.

5.1 CPGs, adjustable chaoticity and homeostatic sensory adaptation

A central element of the chaotification of the embodied neuromechanical system stems from Asai and colleague’s version of the two coupled Fitzhugh-Nagumo (FHN) neuron model [5, 6]. FHN neural models [36, 88] have become important tools in theoretical studies of chaotic neural systems [30, 108]. They are widely used two-dimensional simplifications of the biophysically realistic Hodgkin–Huxley (HH) model of neural spike initiation and propagation [50].

The equations describing two reciprocally coupled FHN neurons are:

$$\dot{u}_1 = c(u_1 - \frac{u_1^3}{3} - w_1 + z_1) + \delta(u_2 - u_1) \quad (14)$$

$$\dot{w}_1 = \frac{1}{c}(u_1 - bw_1 + a) + \epsilon u_2 \quad (15)$$

$$\dot{u}_2 = c(u_2 - \frac{u_2^3}{3} - w_2 + z_2) + \delta(u_1 - u_2) \quad (16)$$

$$\dot{w}_2 = \frac{1}{c}(u_2 - bw_2 + a) + \epsilon u_1 \quad (17)$$

where u describes a neuron’s output and w is its refractory, or ‘recovery’, variable, $a = 0.7$, $b = 0.675$, $c = 1.75$ are constants and $\delta = 0.013$, $\epsilon = 0.022$ are coupling strengths. The constants and coupling strengths were empirically determined [6, 105], such that the neurons

(motor patterns) that describe the intrinsic dynamics of the embodied system by driving the system orbit through the state space. The orbit is entrained in a high performing basin of attraction. This process warps (mutates) the attractor landscape producing a new landscape that inherits major parts of the structure of the previous state space. The process repeats as fitter and more stable behaviours emerge

exhibit biologically plausible dynamics. z_1 and z_2 are the external stimuli acting as the control parameters for the coupled system. While a single isolated FHN (with $\delta = \varepsilon = 0$) exhibits subcritical Hopf bifurcation at $z = z_h \approx 0.38247$, the coupled system can generate autonomous oscillations in a narrow range below z_h .

An interesting characteristic of this coupled FHN system is that it can generate a rich variety of dynamics ranging from multiple synchronised and quasiperiodic oscillations to chaotic orbits, depending on the two control inputs z_1 and z_2 [5, 6, 107] (Figs 9, 10, 11 and 12). In particular, it has been shown that the system exhibits chaos in a certain region of the parameter space defined by the values of z_1 and z_2 and the degree of their asymmetry. Taking the equations' left-right symmetry into account, the Largest Lyapunov Exponent (LLE) map of two coupled FHNs (Fig 13a) on the $z_2 - dz$ space ($dz = z_2 - z_1$) confirms the existence of chaotic dynamics within a diagonal belt-shaped area that was identified in a previous more qualitative study [5] (a strictly positive LLE indicates chaotic dynamics).

In the embodied neuromechanical systems illustrated in Fig 8 each CPG, i , uses a single FHN oscillator described by (18)–(19), communicating locally with its dedicated muscle/actuator by giving motor signal u_i and receiving sensory signal I_i from the muscle/actuator proprioceptors.

$$\dot{u}_i = c(\tilde{u}_i - \frac{\tilde{u}_i^3}{3} - w_i + z) + \delta(I_i - \tilde{u}_i) \quad (18)$$

$$\dot{w}_i = \frac{1}{c}(\tilde{u}_i - bw_i + a) + \varepsilon I_i \quad (19)$$

where \tilde{u}_i is the CPG output translated by $\tilde{u}_i = u_i + A_{\text{ref}}$, A_{ref} is a reference offset in order to bring the output of the CPG into a zero-centered oscillation. Other symbols and constants are as in (14)–(17).

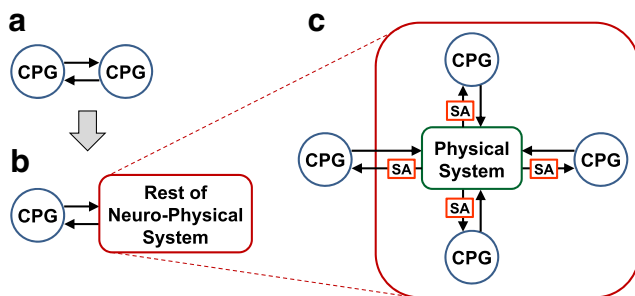


Fig. 9 Scale invariant interaction between each oscillator and the rest of a neuro-physical system. Every neural oscillator communicates with all the other subsystems only through the local coupling to its corresponding muscle. An oscillator interacting with the rest of the neuro-physical system (b) is analogous to the interaction between two coupled oscillators (a), in that any oscillator sees its incoming information from the entire rest of the system as if from another oscillator (boxes with dashed lines in (b) and (c)) via homeostatic sensor adaptation (SAs in (c))

The input I_i is a realtime modulated signal from the muscle/actuator proprioceptors which is passed through a homeostatic sensory adaptation process (SA in Fig 8). Here we use a systemic model of dynamic sensory modulation inspired by the adaptive fusimotor action in muscle spindles [81], which attempts to maintain the amplitude and offset of a rhythmic sensory signal close to those of a reference CPG. The raw proprioceptive sensor signal, s_i , is transformed into the adapted sensory signal I_i , which is fed to the corresponding CPG, according to the following equation:

$$I_i = (s_i - \hat{s}_i) \log(1 + e^{\alpha_i}) + (\hat{s}_i + \beta_i) \quad (20)$$

where α_i and β_i are dynamic variables that control the homeostatic process. α_i determines the amplitude scaling and β_i the offset bias. \hat{s} is the moving average of s , as determined by a simple leaky integrator which is used to smooth the adaptation. The dynamics of α_i and β_i are:

$$\tau_h \dot{\alpha}_i = P_{\text{ref}} - p_i \quad (21)$$

$$\tau_h \dot{\beta}_i = -\hat{I}_i \quad (22)$$

Where \hat{I}_i is the smoothed moving average of I_i and p_i is the smoothed moving average of the log power of the signal I_i , which is a measure of the signal strength. P_{ref} is the target value for p from the reference CPG. Hence α_i is dynamically changed to keep the signal strength close to that of the reference CPG output. The offset variable, β_i , changes to maintain I_i as zero-centered, like the reference CPG. See [108] for further details of the system dynamics and the exact properties of the reference CPG.

This homeostatic adaptation, keeping the sensory signal close to that of a reference CPG, has a crucial and powerful effect: since the sensory signal is its only connection with the rest of the neuromechanical system, each CPG acts as if the entire rest of the system is its pair in the two coupled FHN oscillator system described by (14)–(17). Every CPG views the rest of the system (i.e. the physical system and the other CPGs) as a single autonomous system (Fig. 9). This allows a scale-invariant two coupled CPG scheme for an arbitrary physical system that is able to generalize, amplify and exploit the chaotic regimes of the coupled FHN system.

5.2 Embodied chaotic exploration in action

Figures 10, 11 and 12 illustrate the chaotic exploration system in action for two different simulated robots, highlighting some particularly effective aspects of its operation.

Figure 11 focuses on chaotic exploration for the 3-arm swimmer shown in Fig 10a, in which limb coordination had to be learnt. The swimmer was constructed using a 2D mass-spring-damper system, where the stiffnesses of the springs were set differently to represent three distinct types of body part: rigid structures, compliant edges, and




Fig. 10 Simulated robots. (a) 3-arm 2D mass-spring-damper swimmer. (b) 8 DoF quadruped. (c) Antagonistic torsional muscles for a joint of quadruped

actuating muscles, as indicated in Fig. 10a. All point masses were set to 1kg, and the spring rest lengths were set to those at the neutral pose of the robot as shown in Fig 10a, except $r_m = 0.075$ for the three muscle edges. For each outer edge of the robot, a fluid force acting in the normal direction was calculated [108]. The required behaviour was to move straight ahead in an efficient manner. The evaluation measure for the robot was thus based on its forward speed. Since the system has no prior knowledge of the body morphology of the robot, it does not have direct access to the direction of movement or information on body orientation. In order to facilitate steady movement in one direction without gyrating in a small radius, the center of mass velocity of the robot was continuously averaged by leaky integration, and its magnitude was used as the performance value [105, 106], defined as:

$$E = \|\bar{\mathbf{v}}\| \tag{23}$$

$$\tau_E \dot{\bar{\mathbf{v}}} = -\bar{\mathbf{v}} + \mathbf{v} \tag{24}$$

where $\tau_E = T$, the CPG period, and \mathbf{v} is the center of mass velocity.

With three CPG-muscle units acting in combination with the partially compliant body, the swimmer has a total of 25 degree of freedom (DoF) described by the system variables. The behavioural stability landscape for the swimmer robot, shown in Fig. 11a, was obtained empirically by repeatedly running the system for 3000 sec starting from 50×50 phase difference points on the grid. Then all the movement vectors in the same grid cell were averaged to generate the ‘flow field’ of phase differences between the three CPGs. The permanently stable behaviours were also found numerically by long term observations of the system running from many different initial phase differences. The performance landscape (Fig. 11b) was also empirically generated on the phase space in the same way, except sensory feedback was disabled in order to maintain the initial phase differences of the CPGs. Considering the radial symmetry of the robot body, the stable behaviours that emerged define three qualitatively different modes: high performance propulsion

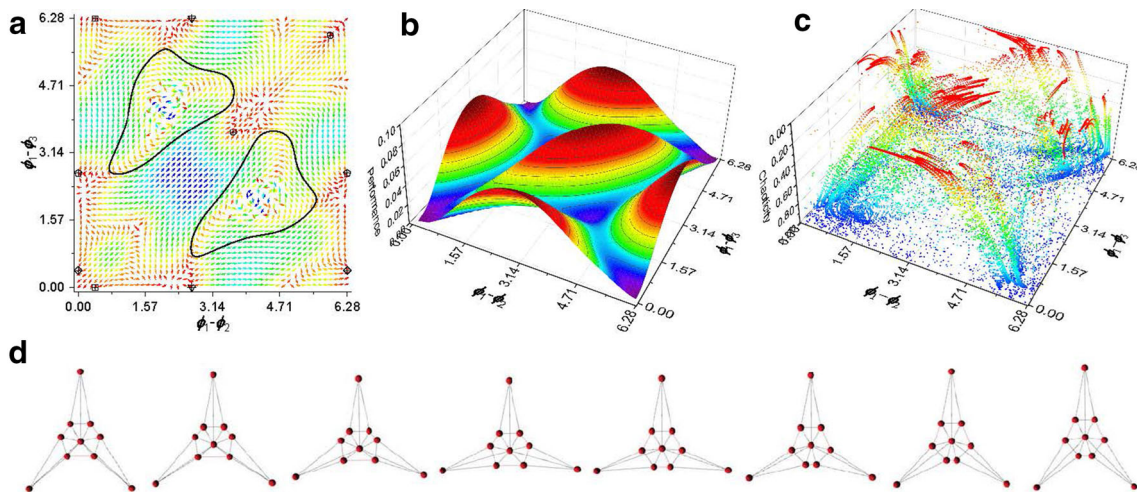


Fig. 11 Chaotic exploration of 2D swimmer. **a** behavioural stability landscape. 6 stable points (on a 2-torus surface) and 2 periodic transitions emerged symmetrically due to the radial symmetry of the robot morphology. **b** Performance landscape. 3 stable points have the highest performances (i.e. locomotion speed), whereas the other 3 points

(nearly all-in-phase motions) and the 2 transient behaviours show low performances. **c** Long-term visits during chaotic exploration. **d** Snapshots of high-performing locomotion (the behaviour point at the middle of (A); (3.62,3.62))

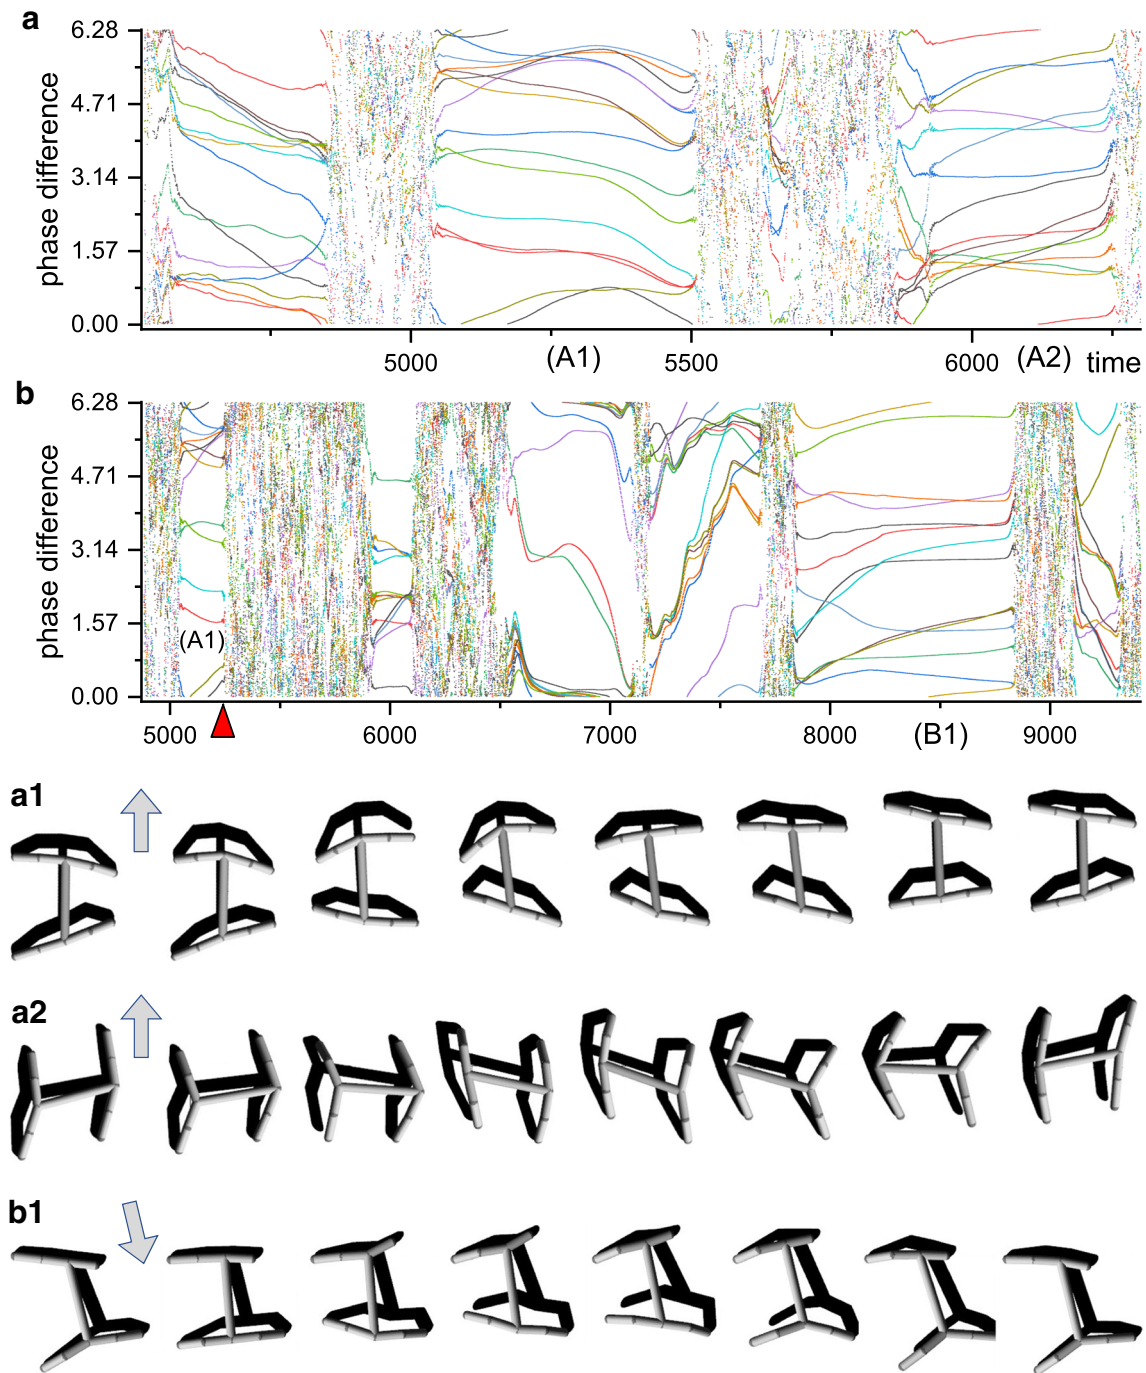


Fig. 12 Chaotic exploration of Quadruped. **a** An example of the time courses of phase differences between CPG-1 and the other 15 CPGs during exploration. Two high performing locomotor behaviours are shown as (a1; quadruped walking gait) and (a2; side-walk like gait) with corresponding snapshots. **b** A scenario for the realtime recovery

from damage where the one of lower limbs was removed during the course of (a1) behaviour (the moment of damage is indicated by the red arrowhead), a new high performing behaviour (b1; hobbling walking gait) was quickly found. The gray arrows in (a1), (a2), and (b1) indicates the directions of movement

using two arms (Fig. 11d), small arm movements by nearly in-phase action, and periodic transition of phase differences which result in circling movement with no forward locomotion. Since the patterns for the three high

performing peaks are also highly stable, we artificially forced the system to eventually escape from those states by gradually increasing the chaoticity whenever the system is stabilized to any of the discovered patterns, which

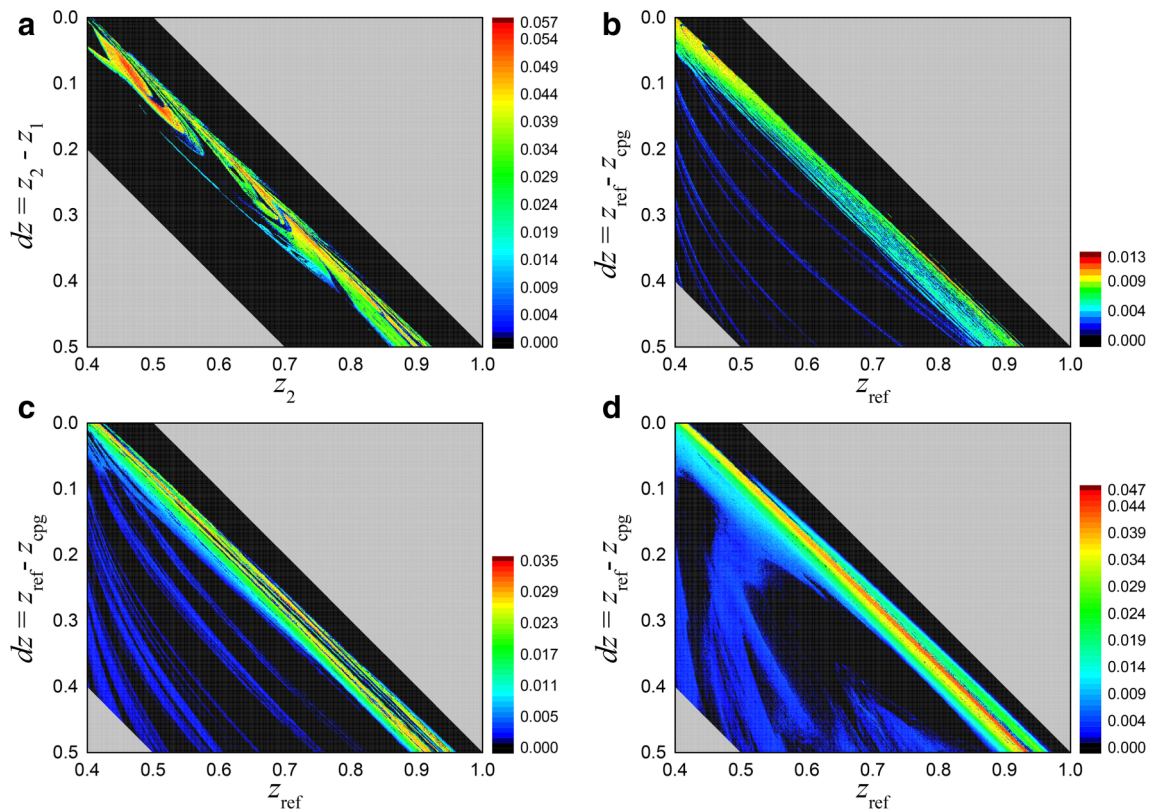


Fig. 13 LLE maps for (a) Two coupled FHN oscillators; (b) 9DoF embodied neuromechanical system; (c) 16DoF system; (d) 25DoF system. Due to the finite computation time, LLEs less than 0.0005 are discarded and rendered as black. Positive non-zero LLEs indicate chaos

allowed us to illustrate the resulting long term exploration dynamics (Fig. 11c). The exploration statistics show the highest performance peaks are most visited, demonstrating the efficacy of the method.

Further generality of the method was demonstrated using a simulated quadruped with 8 degrees of joint freedom, where each joint is driven by a pair of antagonistic torsional muscles, resulting in 16 CPGs (2 per joint DoF, Fig 10b, c). See [105, 106] for the full details of the physical parameters for this type of actuator. The time plots of 15 phase differences between all 16 CPGs during exploration are shown in Fig 12. The environment was deliberately challenging, complete with non-smooth reaction forces due to ground friction with slip, to encourage transients so that the long term exploration dynamics could be illustrated. Although transient, most of the high performing patterns found were surprisingly stable, typically lasting for hundreds of walking cycles. Without such challenging slip, very stable long-term locomotion patterns were found, similar to those in [105]. Figure 12a, a1, and b1 show two different locomotor behaviours (forward and side walking) that were discovered by the exploration process, which exemplifies how the system is able to find completely different modes of locomotion for a given physical system. Because of the nature of the environment, many of the

discovered legged motions included some foot slippage, which is energy-inefficient if too great. However, an interesting and unexpected discovery was that the method found particular combinations of different foot slips and asymmetric limb movements resulting in relatively efficient close to straight locomotion of the whole body (as an alternative to bilaterally symmetric gaits). The realtime and online operation of the exploration process allows practical and challenging scenarios such as re-adaptation after damage. This is illustrated in (Fig 12b and b1), where the robot simply resumes the exploration of new locomotor behaviours for the new (i.e. damaged) body. In this case one leg was chopped off at the knee; after a period of exploration triggered by a drop in performance, leading to an increase in chaoticity, a new stable, relatively efficient ‘hobbling’ gait was quickly discovered where the phase difference patterns, and hence limb coordination mechanisms, were completely different from those used pre-damage.

Figure 13 shows Largest Lyapunov Exponent (LLE) maps for the two coupled FHN system and three representative embodied neuromechanical systems (with 9,16 and 25 DoFs). LLEs were calculated at a fine resolution over the parameter space defined by z_{cpg} and z_{ref} which are the representative factors corresponding to z_1 and z_2 in two coupled FHN CPGs (14)–(17), where z_{cpg} is the descending

input for all CPGs and z_{ref} determines the reference values for homeostatic adaptation described in Section 5.1). The resolution of the LE calculations on the parameter region under investigation (i.e the non-gray pentagonal area) was 0.001 on both axes, resulting in total of 170,801 data points for each neuromechanical system. Strictly positive LLEs indicate chaos.

In order to mitigate the massive computational load of calculating full LE spectra over the whole space, the analysis was divided into two stages by first filtering out non-chaotic points by calculating LLEs (λ_1) using Wolf's method [129] (with a time step of 0.001s for 20000 seconds) which is about an order of magnitude faster than calculating full LE spectra. Next the points with $\lambda_1 > 0.0005$ were identified as chaotic and were processed using a standard QR decomposition method [32] for computing LE spectra by numerically updating both the model and its variational equations using Runge-Kutta 4th order integration with a time step of 0.001s for 10000 seconds. Accurate fine resolution LLEs were thus calculated. The whole procedure was processed for two weeks on a parallel computing platform using 160 virtual CPUs.

The resulting maps clearly show that the main chaotic regions of all systems, including the initial two-CPG model, are very similarly spread on the same area, thus demonstrating that the homeostatic adaptation does indeed result in systems that support the scale-invariant two-CPG scheme illustrated in Fig. 9. Examining the diagonally stretched bands of the chaotic area suggests that chaotic dynamics mainly take place around the Hopf bifurcation point of a CPG ($z=z_{\text{cpg}}=z_h \approx 0.3812$ in (18)-(19)). More detailed examinations of this region [108], outside the scope of this paper, demonstrates that dynamic Hopf bifurcation underlies the dynamics that power this successful chaotic exploration method. Because the scale invariant two-CPG scheme encompasses the entire bodily system's indirect interactions with each CPG, the maps for the embodied neuromechanical systems reveal chaotic dynamics at all levels of the system, from the neural oscillators to the bodily movements. The dynamics of the whole brain-body-environment systems had areas rich with complex and chaotic regimes: all these systems exhibited chaos and hyperchaos. It is this that is exploited in learning locomotion behaviours.

6 Efficient flying: Evolved controllers that exploit aerodynamics

Our final example highlights ongoing research into evolving dynamical neural net controllers for bird inspired flapping wing flight. In this case the dynamics of the neural networks must mesh with the physics of the body and wings as the

(simulated) winged robot moves through the air. Evolution must shape the overall brain-body-environment dynamics so they can be exploited to create efficient flight in which the robot is able to manoeuvre in an agile way while being subjected to external disruption through sporadic gusts of wind.

One of the most impressive features of the behaviour of ornithopters, such as birds, is their stable and rapid aerial manoeuvring [127]. Generating such behaviours in artificial systems is challenging and remains an open problem [74, 87, 104]. In order to achieve a successful compromise between the contradictory properties of stability and manoeuvrability in flapping wing flight, the work described here employs two bio-inspired concepts: the important role of morphology in generating the overall embodied behavioural dynamics [59, 93, 94, 128], and mechanosensory reflexes which are embodied as Reflexive Pattern Generators [39]. These are effectively merged into a simple and tractable robot model using a flexible wing composed of a series of partially independent sub-panels, acting like feathers, which have mechanosensors connected to the control system. This system allows us to study the under-explored role of asymmetric wing and tail movements (much work in the area uses symmetric wing beats) which is widely used by birds and underlies their superb manoeuvrability.

The robot comprises two single-armed wings which each have three degrees of freedom (dihedral, sweep, and twist) and a tail with 3 degrees of freedom (bend, twist, spread) (see Fig. 14). A wing is composed of four feathers which are attached to its skeleton using hinge joints with nonlinear angular springs. The Open Dynamics Engine was used to simulate the articulated rigid body dynamics.

Just as the morphological constraints exploited by ants to produce extremely efficient navigation strategies were the core inspiration for the work described in Section 4, and the evolved morphology of the visual sensors were an important aspect of the EHW robot controllers (Section 3.1), the feathered wing introduces a degree of 'morphological computation' [94] to the flying robot. The shape of the wing changes, enabling aerodynamics that can be actively transformed and exploited by the evolved controllers for efficient flight. The flexible feathers act as an 'aerodynamic cushion' in that they reduce the stiffness of motor control. In completely rigid wings, a small difference in stroke force between two wings will result in a drastic change of net aerodynamic force on the body. Conversely, in a flexible wing, the change will be small due to the passive bending of each feather. This property also confers robustness against external perturbations. Also, as in real bird flight, our feathers can hold laminar air flow through a large range of angles of incidence of a wing arm without stalling [21]. This results in a wide range of effective stroke angles in which




Fig. 14 The flapping wing flying robot (left), the feathered wing (centre), and the tail with asymmetrically controlled spread

the feathered wings are able to produce more lift than rigid wings.

Another important advantage of using feathered wings is that a robot, like a bird, can sense the aerodynamic forces and distributions on each wing through the degree of feather bending [19]. This work utilises feather sensing in a way that is analogous to the use of touch or pressure sensors on legged robots to deal with uneven terrain and external perturbation. Supplementary to the gradual descending command from the optic and vestibular systems, in a bird agile sensory reflexes from oscillatory feathers can be effectively entrained to the pattern generation of wingbeats and play a crucial role in active stabilisation [127]. We observed similar phenomena in the evolved flying robot.

The wing panels receive different aerodynamic forces through the application of a realistic elliptical lift distribution, whereby the innermost panels receive the highest forces and the tip panels the least. Such a distribution is known to hold for finite wingspans. The aerodynamic forces on the tail are calculated using a model of a thin triangular wing with low aspect ratio [115], which results in realistic tail spread dependent forces and distinct lift distributions.

A nonlinear angular spring for feather bending was simulated using a first order differential equation so that the bend angle smoothly decays toward the equilibrium position between aerodynamic torque and the spring torque. At each time step, the bend angle rate of the i th feather receiving aerodynamic torque T ($=|\mathbf{T}|$ in Fig. 14) is described by:

$$\dot{\theta}_i(t) = P(T_i(t-1) - k\theta_i(t)) \quad (25)$$

where P is a proportional factor and k is the spring constant. We set $P = 100$ and $k = 0.1$. The bending torque is calculated from the net aerodynamic force exerted on the center of mass of the feather.

A pair of bilaterally symmetric fully connected continuous-time recurrent neural networks (CTRNNs) were at the heart of the controller, with weights, biases and time constants set by an evolutionary algorithm. The CTRNNs were modelled by (7)–(8) (Section 3.3), except that in this case $\sigma(x) = \tanh(x)$. Figure 15 shows the structure of the control system. In each of the CTRNNs there were 6 neurons for wing control (2 for each DoF), 2 for the tail and

the rest are interneurons (just one was used for the results shown here).

The tail has four joints controlled by the four tail neurons. A pair of tail bend neurons controls both the tail pitch and roll, and a pair of spread neurons control asymmetric spread actions (left and right halves of tail fan independently spread, see Fig. 14, which is more avian-realistic than previous work). Tail bend/pitch, twist/roll and spread angles are controlled by the following equations:

$$\theta_{bend} = g_b(B_L + B_R) \quad (26)$$

$$\theta_{roll} = g_r(B_L - B_R) \quad (27)$$

$$\theta_{spread}^L = g_s S_L \quad (28)$$

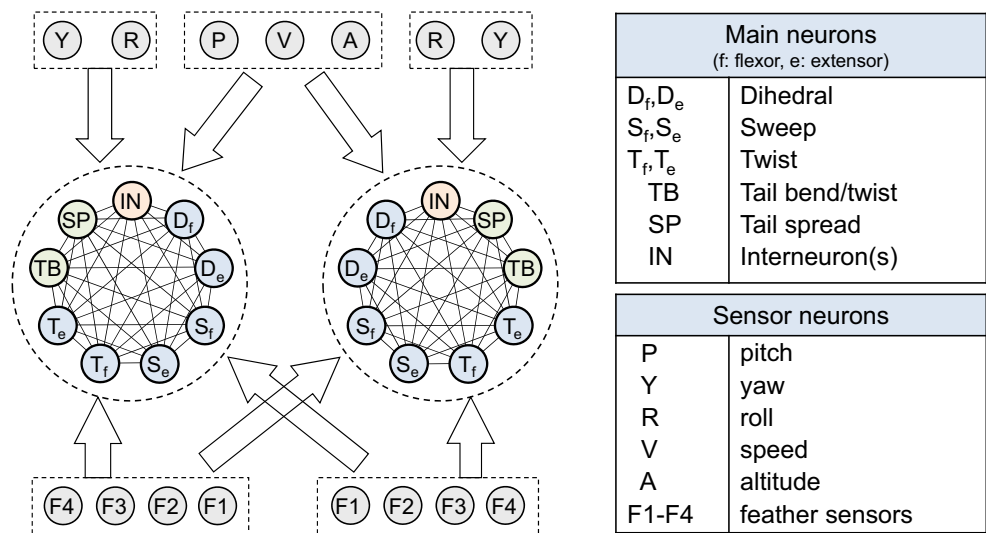
$$\theta_{spread}^R = g_s S_R \quad (29)$$

where g are gains and B , S are signals from the tail neurons (left and right). This arrangement models two pairs of tail pitch muscles bilaterally connected on each side, such that the flexion of one side results in both tail pitch and twist. If the muscles on both sides contract equally together, only tail pitch is affected. This avian-inspired setup enables subtle control during flight.

The neural circuits were integrated using the forward Euler method with a step size of 0.01. The output signals from motor neurons are fed to the simulated servomotors as desired angular positions.

In order to achieve robust, efficient flight a multi-objective evolutionary algorithm (MOEA) was used, based on the SPEA2 algorithm [136]. The behaviour required was for the flyer to follow a straight target path at a given altitude (10m), a ‘target point’ is located 5m ahead of the robot, moving along the straight path (Fig. 16). Periodically during flight, the robot receives perturbations from wind gust that last for a few seconds. Before the perturbations begin, the robot only has to reach the target altitude after takeoff, but once the perturbations start, the robot should return to the flight path line as soon as possible and continuously maintain it. These evaluation conditions require the evolution of sophisticated, robust flight control. Because of the highly nonlinear nature of the robot-environment interactions, the robot is forced to capture a variety of fairly complex manoeuvres. The robot effectively ‘senses’ the target point on the linear trajectory

Fig. 15 Neural control structure for the flying robot. Neurons in dashed-lined circles are the main neurons. Large arrows represent all-to-all connection. All neuron parameters and connection weights are bilaterally symmetric (mirrored)



to follow because the azimuth of the target point and the direction of gravity relative to the robot’s orientation were used to calculate the body pitch, yaw, and roll sensor values, as well as the flight speed and altitude (Fig. 17).

In order to achieve this behaviour in an efficient way, three objectives (to be minimized) were used in the MOEA approach: the average distance to the target path (D), mechanical power consumption (P), and a measure of wing flapping (W). The MOEA selection pressure is guided by the Pareto fronts formed using these objectives (Fig. 18).

$$D = \frac{1}{20T} \sum_t d(t) \tag{30}$$

$$P = \frac{t_a}{T} \sum_t \sum_i |\tau_i(t) \omega_i(t)| \tag{31}$$

$$W = \frac{t_a}{3T} \left(\frac{A}{A_{min}} + \frac{F}{F_{min}} + \frac{S}{S_{min}} \right) \tag{32}$$

where t_a is the time until the robot falls to the ground, $\tau_i(t)$ and $\omega_i(t)$ are the torque and angular speed of wing motor i . A , F are the time averaged amplitude and frequency of the two wing motions over all axes and S is the flight speed during the period the robot is airborne.

The evolutionary progression of the population during a typical run can be seen in Fig. 18. Very good solutions were reliably found after a few hundred generations. Some solutions had very low values for all objectives. Because the minimum wingbeat objective (W) takes its best possible values when the flyer can successfully remain airborne, the population always saturated towards low W in the later generations as can be seen in the figure. Hence the population always moved toward robust, stable flight regardless of the target following and the power consumption measures, although many members had low values for these too, as can be seen. Videos of typical good solutions in action can be found at <https://youtu.be/-JBBwaw0x8Q> and https://youtu.be/N9A_Z4hYZ3s. These show the natural and smooth transitions between flapping wing flight and gliding as the robot manoeuvres.

The evolved controllers produce robust efficient flight by exploiting complex asymmetric dynamics partly enabled by continually changing wing and tale morphologies. Continuous changes in wing morphology can be seen in Fig. 17 which shows flapping wing motion as the robot tracks the moving target. See the videos mentioned above for the full sequences.




Fig. 16 Flapping wing flyer evaluation strategy




Fig. 17 Sequential frames of the simulated flying robot performing wing flaps. Changes in the feathered wing morphology are visible. The coloured square indicates the ‘target’ and the arrow shows the direction of motion of the target

7 Conclusions

This paper has highlighted a number of different kinds of evolutionary and adaptive techniques used to exploit dynamics in the creation of embodied behaviours.

The experiments on evolving FPTA controllers for a visually guided robot, described in Section 3 demonstrated

three things. First, it is possible to evolve, directly in hardware, component level analog electronic circuits to generate non-trivial visually guided sensorimotor behaviour in a mobile robot. More generally, it was established that concise evolved transistor-based circuits could successfully coordinate sensory input and actuator output to produce robust behaviour even when the sensors and




Fig. 18 Evolving flapping wing flyer population. The red individuals represent the Pareto frontiers, green represents archived members from earlier Pareto fronts, blue represents the rest of the population

actuators were low-grade, noisy and unreliable. By integrating visual feature extraction and selection into the evolutionary approach, highly robust embodied sensorimotor dynamics emerged which were readily exploited by evolution.

Second, and perhaps most interesting, controller circuit analysis and comparative experiments established that the successful evolved circuits exploited the rich dynamics of the FPTA hardware medium. The evolved solutions to non-trivial visual navigation tasks can be viewed as dynamical systems with (behavioural) attractors that result in completion of the task regardless of start conditions [42, 53]. The continuous analog medium of the FPTA seems a particularly good substrate to enable the evolution of such attractors. This possibility of rich unconventional dynamics to be exploited is a large part of what makes the FPTA a highly evolvable medium for this kind of application. Naively it might be thought that the large search space defined by the FPTA genetic encoding used would make it much more difficult to find solutions than for the more constrained, smaller search space of the fixed-architecture CTRNN controllers, which were also rich with dynamics. Comparative experiments showed this was not the case, with the FPTA being significantly more evolvable. The unconventional, potentially complex, dynamics afforded by the physical properties of the hardware medium, increases the degeneracy of the FPTA as an evolvable substrate. We use degeneracy as it is applied to biological systems [34, 120]: multiple, often interacting, ways of achieving an outcome (in this case implying many different, easily accessible, routes through the fitness landscape towards high fitness areas). Degeneracy has been shown to greatly boost evolvability [54].

Third, with a carefully constructed, special kind of simulation, it was possible to evolve robot controllers that transferred seamlessly to the real world. Our methodology involved refining the simulation in light of problems with early transfers (making sure the visual latency matched that of the physical robot, and the severe limitations introduced by errors in the vision turret, and so on, were replicated in the simulation). The robustness and generality of evolved behaviours was such that the robot controllers could handle unseen variations in the environment and continued to perform well when used in completely different environments [40], thus exhibiting behavioural resilience.

The ant-inspired navigation algorithms described in Section 4 provide insights into the way insects navigate with low resolution vision and modest neural resources. Their successful demonstration on robots operating in dynamic outdoor environments also shows that they can form the basis of autonomous robot navigation systems in applications where processing is at a premium

(e.g. planetary exploration), or GPS and mapping data is unavailable or infeasible (and hence makes SLAM approaches [119] difficult). In summary, this research shows that taking inspiration from insects and taking advantage of the embodied dynamics of innate behaviours allows a parsimonious approach and simple algorithm which does not require precise place recognition. By removing the need for localisation, we can use a simple ANN trained for familiarity and not image recognition, meaning that all computation can be performed on board a small autonomous robot. The next stages are to take further advantage of the bio-inspired embodied dynamics and to 1) embody the image search process using sinuous paths modulated by the visual input [111]; 2) optimise the visual processing by copying insect eyes [83] and 3) add temporal dynamics into the ANN [60].

The modular coupled oscillator control architecture described in Section 5 has been shown to be highly effective in a number of ways. First, adjustable chaos is generated at all levels of the embodied neuromechanical systems being controlled, and is exploited to power a performance-directed exploration of the space of possible motor behaviours. Second, this exploration and discovery of high performing behaviours does not require any prior or built-in knowledge of the robot body morphology or the properties of the environment. The control architecture automatically adjusts itself to whatever body-environment system it is connected to, as long as the basic setup is as illustrated in Figs. 8 and 9; homeostatic sensory adaptation being a crucial part of the system. Third, neuromechanical systems in this class are resilient: they are able to compensate in realtime to bodily damage, failures and changes to the environment, rapidly finding new control patterns that produce the desired behaviour.

Most other approaches to the development of resilient machines, for instance from the field of evolutionary robotics, make use of some form of self-model [16, 27]. Evaluations based on the model guide an explicit search process where possible new behaviours are tried out in an internal simulation. As the size of the system increases this approach can become computationally expensive and time consuming and requires significant amounts of *a priori* knowledge of the robot and its environment. In contrast, the chaotic exploration method is model-free, it requires no expensive internal self-simulations or *a priori* knowledge, and occurs in realtime.

The evolutionary robotics approach, as exemplified by the work described in Sections 3 and 6, can discover highly unconventional systems when it is unconstrained. For instance the evolved FPTA control circuits did not have to conform to some pre-specified architecture. When properties, including the morphology, of the sensors and/or the robot body are coevolved with the controller (as in

Section 3.1), a powerful shaping and harmonising of all levels of dynamics underlying behaviour can be achieved. But this comes at a cost – very large numbers of robot evaluations must be undertaken. In contrast the chaotic exploration approach is much more efficient, occurring in realtime through the intrinsic dynamics of the control architecture, without the need for costly offline evaluations. In this case the architecture is constrained and pre-specified (if very general), but its potential applications are nonetheless widespread (any system that can conform to the scheme shown in Fig. 8).

Interestingly, there is a closer relationship between evolutionary search and chaotic exploration than may at first be obvious [106]. In fact the overall chaotic exploration process has a number of parallels with evolutionary dynamics. The whole system (literally) embodies a population of (motor behaviour) attractors which is sampled by chaotic exploration. The proprioceptor-driven homeostatic adaptation process warps (mutates) the state space such that a new landscape of attractors is created, but one that inherits the major properties of the previous (ancestor) landscape (replication with variation). The process repeats with the new population being sampled by chaotic exploration (Fig. 8). Since the population of attractors is effectively implicit – the intrinsic dynamics of the system drive it to sample the space of attractors – our embodied system can be thought of as a kind of generative search process. The overall brain-body-environment system (literally) embodies a population of motor pattern attractors through its dynamics; it cannot help but sample them during the exploration phases. This is loosely analogous to the generative statistical models used by Estimation of Distribution Algorithms (EDAs) [72, 92], which are well established as part of the evolutionary computing canon. Instead of using an explicit population of solutions and the traditional machinery of evolutionary algorithms, EDAs employ a (often Bayesian) probabilistic model of the distribution of solutions which can be sampled by generating possible solutions from it. Search proceeds through a series of incremental updates of the probabilistic model guided by feedback from sampled fitness. In an analogous way our generative system (the overall system dynamics) is incrementally updated in relation to evaluation based feedback. The overall system dynamics is the generative model, the exploration phase is the sampling step, with the performance evaluation, E , controlling a selection pressure, and the homeostatic adaptation process provides a kind of mutation which facilitates the replication (with variation) of the whole phase space, now containing a slightly different population of attractors but with a bias towards preserving more stable and fitter areas. This work thus points towards the possibility of intrinsic mechanisms, based entirely on neuro-body-environment interaction dynamics, that might

be involved in creating Darwinian processes that could continually run within the nervous systems of future robots [35].

To some extent morphology is important to all of the case studies used in this paper. Be it the way the sensor morphology co-evolves with the robot controllers in the FPTA study, or the behavioural strategy relies on morphological constraints in the ant-inspired navigation work, or the way in which continually changing wing and tail morphologies generate beneficial airflow patterns in the flapping wing flyer example. In common with the use of chaotic exploration for the development of locomotion behaviours, in all these cases information processing is not located solely in the nervous system of the machine; it is spread out over the brain-body-environment system. Some strands of work in soft robotics seek to push this idea further, to blur the line between body and nervous system even more, with greater amounts of processing offloaded onto the body [15, 49, 58, 59].

The behaviours described in this paper, although robust and resilient, are still fairly simple. The dynamics exploited, although often complex, are limited when compared with those used in many animal behaviours. There is still much to explore in the two-way exchange between biology and robotics.

Acknowledgements Thanks to members of the CCNR for helpful discussions, also to Ittai Flascher, as well as members of the INSIGHT consortium, and to the anonymous reviewers for comments on an earlier draft of this paper.

Funding Funding for the various strands of work described in this paper came from: EPSRC grants EP/P006094/1 and EP/S030964/1, EU ICT FET FP7 project INSIGHT, and Intel Corp. ND was supported by an EPSRC studentship.

Declarations

Conflict of Interests The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abeles M (1991) *Corticonics: Neural circuits of the cerebral cortex*. Cambridge University Press, Cambridge
- Adamatzky A (2013) *Reaction diffusion automata: Phenomenology, Localisations, Computation*. Springer, New York
- Aihara K, Matsumoto G (1982) Temporally coherent organization and instabilities in squid giant axons. *J Theor Biol* 95(4):697–720
- Amari SI (1998) Natural gradient works efficiently in learning. *Neural Comput* 10(2):251–276. <https://doi.org/10.1162/089976698300017746>
- Asai Y, Nomura T, Abe K, Sato S (2003) Classification of dynamics of a model of motor coordination and comparison with Parkinson's disease data. *Biosystems* 71:11–21
- Asai Y, Nomura T, Sato S, Tamaki A, Matsuo Y, Mizukura I, Abe K (2003) A coupled oscillator model of disordered interlimb coordination in patients with Parkinson's disease. *Biol Cybern* 88:152–162
- Baddeley B, Graham P, Husbands P, Philippides A (2012) A model of ant route navigation driven by scene familiarity. *PLoS Computational Biology* 8(1). <https://doi.org/10.1371/journal.pcbi.1002336>
- Baddeley B, Graham P, Philippides A, Husbands P (2011) Holistic visual encoding of ant-like routes: Navigation without waypoints. *Adapt Behav* 19(1):3–15
- Beer R (1997) The dynamics of adaptive behaviour: A research program. *Robot Auton Syst* 20:257–289
- Beer R, Williams P (2015) Information processing and dynamics in minimally cognitive agents. *Cogn Sci* 39:1–38
- Beer RD, Quinn HJCRD, Ritzmann RE (1998) Biorobotic approaches to the study of motor systems. *Curr Opin Neurobiol* 8:777–782
- Bekey G (2005) *Autonomous robots: From biological inspiration to implementation and control*. MIT Press, Cambridge
- Bell AJ, Sejnowski TJ (1995) An Information-Maximization approach to blind separation and blind deconvolution. *Neural Comput* 7(6):1129–1159. <https://doi.org/10.1162/neco.1995.7.6.1129>
- Berenson D, Estevez N, Lipson H (2005) Hardware evolution of analog circuits for in-situ robotic fault-recovery. In: 2005 NASA/Dod conference on evolvable hardware (EH'05). IEEE comp. Soc. Press, pp 12–19. <https://doi.org/10.1109/EH.2005.30>
- Bongard J (2014) Why morphology matters. In: Vargas P, DiPaolo E, Harvey I, Husbands P (eds) *The horizons of evolutionary robotics*. MIT Press, pp 125–152
- Bongard J, Zykov V, Lipson H (2006) Resilient machines through continuous self-modeling. *Science* 314:1118–1121
- Bonin-Font F, Ortiz A, Oliver G (2008) Visual navigation for mobile robots: A survey. *J Intell Robot Syst Theory Appl* 53(3):263–296. <https://doi.org/10.1007/s10846-008-9235-4>
- Bressler SL, Kelso JAS (2001) Cortical coordination dynamics and cognition. *Trends Cognitive Sci* 5(1):26–36
- Brown R, Fedde M (1993) Airflow sensors in the avian wing. *J Exp Biol* 179:13–30
- Buehlmann C, Fernandes ASD, Graham P (2018) The interaction of path integration and terrestrial visual cues in navigating desert ants: what can we learn from path characteristics? *J Exp Biol* 221(1)
- Burt Jr. E., Ichida J (2006) Selection for feather structure. *Acta Zool Sin* 52:131–135
- Buzsaki G (2006) *Rhythms of the brain*. Oxford University Press, Oxford
- Cadena C, Carlone L, Carrillo H, Latif Y, Scaramuzza D, Neira J, Reid I, Leonard JJ (2016) Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Trans Robot* 32(6):1309–1332
- Cagnoni S (ed) (2009) *Evolutionary image analysis and signal processing*. Springer, New York
- Campos P, Lawson D, Bale S, Walker J, Trefzer M, Tyrrell A (2013) Overcoming faults using evolution on the panda architecture. In: *Proceedings of IEEE Congress on Evolutionary Computation (CEC'13)*. IEEE Press, pp 613–620
- Chang E, Matloff LY, Stowers AK, Lentink D (2020) Soft biohybrid morphing wings with feathers underactuated by wrist and finger motion. *Sci Robot* 5(38). <https://doi.org/10.1126/scirobotics.aay1246>. <https://robotics.sciencemag.org/content/5/38/eaay1246>
- Chatzilygeroudis K, Vassiliades V, Mouret JB (2018) Reset-free trial-and-error learning for robot damage recovery. *Robot Auton Syst* 100:236–250
- Cummins M, Newman P (2011) Appearance-only SLAM at large scale with FAB-MAP 2.0. *Int J Robot Res* 30(9):1100–1123
- Dale K, Husbands P (2010) The evolution of reaction-diffusion controllers for minimally cognitive agents. *Artificial Life* 16(1):1–19. <https://doi.org/10.1162/artl.2009.16.1.16100>. PMID: 19857145
- Davison E, Aminzare Z, Dey B, Ehrich Leonard N (2019) Mixed mode oscillations and phase locking in coupled fitzhugh-nagumo model neurons. *Chaos An Interdiscipl J Nonlinear Sci* 29(3):033105. <https://doi.org/10.1063/1.5050178>
- Dayan P, Abbot L (2005) *Theoretical neuroscience*. MIT Press, Cambridge
- Dieci L, Russell R, Van Vleck E (1997) On the computation of Lyapunov exponents for continuous dynamical systems. *SIAM J Numer Anal* 34(1):402–423
- DiPaolo E, Buhmann T, Barandiaran E (2017) *Sensorimotor life: An enactive proposal*. Oxford University Press, Oxford
- Edelman G, Gally J (2001) Degeneracy and complexity in biological systems. *Proc Natl Acad Sci USA* 98(24):13763–13768
- Fernando C, Szathmáry E, Husbands P (2012) Selectionist and evolutionary approaches to brain function: A critical appraisal. *Front Comput Neurosci* 6(24). <https://doi.org/10.3389/fncom.2012.00024>
- Fitzhugh R (1961) Impulses and physiological states in theoretical models of nerve membrane. *Biophys J* 1:445–466
- Floreano D, Mondada F (1994) Automatic creation of an autonomous agent: Genetic evolution of a neural-network driven robot. In: Cliff D, Husbands P, Meyer J, Wilson SW (eds) *From Animals to Animats III: Proceedings of the third international conference on simulation of adaptive behavior*. MIT Press, Cambridge, pp 402–410
- Freeman W, Viana Di Prisco G (1986) EEG Spatial pattern differences with discriminated odors manifest chaotic and limit cycle attractors in olfactory bulb of rabbits. In: Palm G., Aertsen A. (eds) *Brain theory*. Springer, London, pp 97–119
- Gallagher J (2001) Evolution and analysis of non-autonomous neural networks for walking: Reflexive pattern generators. In: *Proceedings of 2001 IEEE Congr. evol. comput. (CEC 2001)*, vol 1. IEEE, pp 245–250
- Garvie M, Flascher I, Philippides A, Thompson A, Husbands P (2020) Evolved transistor array robot controllers. *Evolution Comput* 28:1–32. <https://doi.org/10.1162/evco.a.00272>
- Garvie M, Thompson A (2003) Evolution of self-diagnosing hardware. In: Tyrrell A, Haddow P, Torresen J (eds) *Proc. 5th Int. Conf. on Evolvable Systems (ICES2003): From biology to hardware, LNCS, vol 2606*, Springer, pp 238–248
- van Gelder T (1995) What might cognition be if not computation? *J Philos* 92(7):345–381

43. Glass L (2009) Introduction to controversial topics in nonlinear science: Is the normal heart rate chaotic? *Chaos* 19(028501)
44. Gray F (1953) Pulse code communication. <https://www.google.com/patents/US2632058>. US Patent 2,632,058
45. Guevara MR, Glass L, Mackey MC, Shrier A (1983) Chaos in neurobiology. *IEEE Trans Syst Man Cybern SMC-13*:790–798
46. Haddow P, Tuft G (2000) An evolvable hardware FPGA for adaptive hardware. In: *Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2000)*. IEEE Press, pp 553–560
47. Harvey I (1992) Species Adaptation Genetic algorithms: A basis for a continuing SAGA. In: Varela FJ, Bourgine P (eds) *Towards a practice of autonomous systems: Proc. 1st Eur. Conf. on artificial life*. MIT Press, pp 346–354
48. Harvey I, Husbands P, Cliff D (1994) Seeing the light : Artificial evolution, real vision. In: Cliff D et al (eds) *From Animals to Animats 3: Proc. 3rd Int. Conf. on simulation of adaptive behaviour*. MIT Press, pp 392–401
49. Hauser H, Ijspeert A, Füchslin R, Pfeifer R, Maass W (2012) The role of feedback in morphological computation with compliant bodies. *Biol Cybern* 106(10):595–613
50. Hodgkin A, Huxley A (1952) A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal Physiol* 117:500–544
51. Hoerzer GM, Legenstein R, Maass W (2014) Emergence of complex computational structures from chaotic neural networks through reward-modulated Hebbian learning. *Cereb Cortex* 24:677–690
52. Howard G, Bull L, de Lacy Costello B, Gale E, Adamatzky A (2014) Evolving spiking networks with variable resistive memories. *Evol Comput* 22(1):79–103
53. Husbands P, Harvey I, Cliff D (1995) Circle in the round: State space attractors for evolved sighted robots. *Robot Auton Syst* 15(1-2):83–106
54. Husbands P, Philippides A, Vargas P, Buckley C, DiPaolo E, O'Shea M (2010) Spatial, temporal and modulatory factors affecting gasnet evolvability in a visually guided robotics task. *Complexity* 16(2):35–44
55. Ijspeert A, Crespi A, Cabelguen J (2005) Simulation and robotics studies of salamander locomotion. *Neuroinformatics* 3:171–195
56. Jakobi N (1998) Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behaviour* 6(2):326–368
57. Jakobi N, Husbands P, Harvey I (1995) Noise and the reality gap: The use of simulation in evolutionary robotics. In: Morán F et al (eds) *Advances in artificial life: Proc. 3rd eur. Conf. on artificial life (ECAL95)*, LNAI, vol 929. Springer, pp 704–720
58. Johnson C, Philippides A, Husbands P (2016) Active shape discrimination with compliant bodies as reservoir computers. *Artificial Life* 22(2):241–268
59. Johnson C, Philippides A, Husbands P (2019) Simulating soft-bodied swimmers with particle-based physics. *Soft Robotics* 6(2):263–275
60. Kagioulis E, Philippides A, Graham P, Knight JC, Nowotny T (2020) Insect inspired view based navigation exploiting temporal information. In: *Living machines 2020, lecture notes in artificial intelligence*, vol 12413
61. Keymeulen D, Durantez M, Konaka K, Kuniyoshi Y, Higuchi T (1996) An evolutionary robot navigation system using a gate-level evolvable hardware. In: Iwata M., Liu W. (eds) *Proceeding of the first international conference on evolvable systems: From biology to hardware (ICES'96)*. Springer, Berlin, pp 195–210
62. Knaden M, Graham P (2016) The sensory ecology of ant navigation: From natural environments to neural mechanisms. *Annual Rev Entomol* 61:63–76
63. Knight JC, Sakhapov D, Domcsek N, Dewar ADM, Graham P, Nowotny T, Philippides A (2019) Insect-inspired visual navigation on-board an autonomous robot: Real-world routes encoded in a single layer network. In: *Artificial life conference proceedings*. MIT Press, pp 60–67
64. Korn H, Faure P (2003) Is there chaos in the brain? II. Experimental evidence and related models. *Comptes Rendus Biologies* 326:787–840
65. Koza J, Keane M, Streeter M, Mydlowec W, Yu J, Lanza G (2003) *Genetic programming IV: Routine Human-Competitive machine intelligence*. Kluwer Academic Publishers, New York
66. Kriegman S, Blackiston D, Levin M, Bongard J (2020) A scalable pipeline for designing reconfigurable organisms. *PNAS* 117(4):1853–1859
67. Kuniyoshi Y, Sangawa S (2006) Early motor development from partially ordered neural-body dynamics: Experiments with a cortico-spinal-musculo-skeletal model. *Biol Cybern* 95:589–605
68. Kuniyoshi Y, Suzuki S (2004) Dynamic emergence and adaptation of behavior through embodiment as coupled chaotic field. In: *Proceedings of IEEE international conference on intelligent robots and systems*, pp 2042–2049
69. Laje R, Buonomano DV (2013) Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat Neurosci* 16:925–933
70. Lambercy F (2011) K-junior user manual v1.2. Tech. rep., K-Team, Yverdon-les-Bains, Switzerland
71. Langeheine J (2005) *Intrinsic hardware evolution on the transistor level*. Ph.D. thesis, Kirchhoff Institute for Physics
72. Larrañaga P, Lozano J (2002) *Estimation of distribution algorithms: A new tool for evolutionary computation*. Kluwer Academic Publishers, Boston
73. Lee TW, Sejnowski TJ (1997) Independent component analysis for Mixed Sub-Gaussian and Super-Gaussian sources. *Joint Symposium Neural Comput* 441:6–13. <http://ukpmc.ac.uk/abstract/CIT/106463>
74. Leutenegger S, Hurzeler C, Stowers A, Alexis K, Achteik M, Lentink D, Oh P, Siegwart R (2016) Flying robots. In: Siciliano B, Khatib O (eds) *Springer handbook of robotics*. 2nd edn. Springer, pp 623–670
75. Lohn JD, Hornby GS (2006) Evolvable hardware: Using evolutionary computation to design and optimize hardware systems. *IEEE Comput Intell Mag* 1(1):19–27
76. Lopez B, Valverde J, de La Torre E, Riesgo T (2014) Power-aware multi-objective evolvable hardware system on an fpga. In: Vladimirova T (ed) *Proceedings of the 2014 NASA/ESA conference on adaptive hardware and systems*. IEEE Press, pp 61–68
77. Lowry S, Sunderhauf N, Newman P, Leonard JJ, Cox D, Corke P, Milford MJ (2016) Visual place recognition: A survey. *IEEE Trans Robot* 32(1):1–19. <https://doi.org/10.1109/TRO.2015.2496823>
78. Lulham A, Bogacz R, Vogt S, Brown MW (2011) An infomax algorithm can perform both familiarity discrimination and feature extraction in a single network. *Neural Computation* 23(4):909–926. https://doi.org/10.1162/NECO_a.00097
79. Magdoo KN, Subramanian D, Chakravarthy VS, Ravindran B, Amari S, Meenakshisundaram N (2011) Modeling basal ganglia for understanding Parkinsonian reaching movements. *Neural Comput* 23(2):477–516
80. Markram H, Lubke J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic apss and epsps. *Science* 275:213–215
81. Marques HG, Bharadwaj A, Iida F (2014) From spontaneous motor activity to coordinated behaviour: A developmental model. *PLoS Computat Biol* 10(7):e1003653

82. Matloff LY, Chang E, Feo TJ, Jeffries L, Stowers AK, Thomson C, Lentink D (2020) How flight feathers stick together to form a continuous morphing wing. *Science* 367(6475):293–297. <https://doi.org/10.1126/science.aaz3358>. <https://science.sciencemag.org/content/367/6475/293>
83. Meyer S, Nowotny T, Graham P, Dewar A, Philippides A (2020) Snapshot navigation in the wavelet domain. In: *Living machines 2020, lecture notes in artificial intelligence*, vol 12413
84. Miller J, Harding S, Tufte G (2014) Evolution-in-materio: Evolving computation in materials. *Evol Intel* 7(1):49–67
85. Mohid M, Miller J, Harding S, Tufte G, Massey M, Petty M (2016) Evolution-in-materio: Solving computational problems using carbon nanotube-polymer composites. *Soft Comput* 20(8):3007–3022
86. Molioli RC, Husbands P (2013) Neuronal assembly dynamics in supervised and unsupervised learning scenarios. *Neural Comput* 25(11):2934–2975
87. Mouret JB, Doncieux S, Meyer JA (2006) Incremental evolution of target-following neuro-controllers for flapping-wing animats. In: Nolfi S, Baldassarre G, Calabretta R, Hallam JCT, Marocco D, Meyer JA, Miglino O, Parisi D (eds) *From animals to animats 9*. Springer, Berlin, pp 606–618
88. Nagumo J, Arimoto S, Yoshizawa S (1962) An active pulse transmission line simulating nerve axon. In: *Proceedings of the IRE* 50, pp 2061–2071
89. Naito T, Odagiri R, Matsunaga Y, Tanifuji M, Murase K (1996) Genetic evolution of a logic circuit which controls an autonomous mobile robot. In: Iwata M., Liu W. (eds) *Proceeding of the first international conference on evolvable systems: from biology to hardware (ICES'96)*. Springer, Berlin, pp 210–219
90. Nolfi S, Bongard J, Floreano D, Husbands P (2016) Evolutionary robotics. In: Siciliano B., Khatib O. (eds) *Springer handbook of robotics*. 2nd edn. Springer, Berlin, pp 2035–2067
91. Ohgi S, Morita S, Loo K, Mizuike C (2008) Time series analysis of spontaneous upper-extremity movements of premature infants with brain injuries. *Phys Ther* 88(9):1022–1033
92. Pelikan M, Goldberg D, Cantu-Paz E (2000) Linkage problem, distribution estimation and bayesian networks. *Evol Comput* 8(3):311–340
93. Pfeifer R, Bongard J (2006) *How the body shapes the way we think*. MIT Press, Cambridge
94. Pfeifer R, Iida F (2005) Morphological computation: Connecting body, brain and environment. *Japanese Scientific Monthly* 58:48–54
95. Philippides A, Graham P, Baddeley B, Husbands P (2015) Using neural networks to understand the information that guides behavior: A case study in visual navigation. In: *Artificial neural networks*. Springer, pp 227–244. https://doi.org/10.1007/978-1-4939-2239-0_14
96. Ping Z, Rui Y, Junjie D (2017) Design of self-repairing control circuit for brushless dc motor based on evolvable hardware. In: *Proceedings of the 2017 NASA/ESA conference on adaptive hardware and systems*. IEEE Press, pp 214–220
97. Purves D, Augustine G, Fitzpatrick D (2001) *Neuroscience*, 3rd edn. Sinauer, Sunderland
98. Rapp P, Zimmerman I, Albano A, Deguzman G, Greenbaun N (1985) Dynamics of spontaneous neural activity in the simian motor cortex: The dimension of chaotic neurons. *Phys Lett A* 110(6):335–338
99. Riley M, Turvey M (2002) Variability and determinism in motor behaviour. *J Mot Behav* 34(2):99–125
100. Roberts T, Kern F, Fernando C, Szathmáry E., Husbands P, Philippides A, Staras K (2019) Encoding temporal regularities and information copying in hippocampal circuits. *Scientific Reports* 9:19036. <https://doi.org/10.1038/s41598-019-55395-1>
101. Roggen D, Hofmann S, Thoma Y, Floreano D (2003) Hardware spiking neural network with run-time reconfigurable connectivity in an autonomous robot. In: *Proc. 2003 NASA/DOD conference on evolvable hardware*. IEEE Press, pp 189–198
102. Santos BA, Barandiaran XE, Husbands P (2012) Synchrony and phase relation dynamics underlying sensorimotor coordination. *Adapt Behav* 20(5):321–336
103. Scafetta N, Marchi D, West BJ (2009) Understanding the complexity of human gait dynamics. *Chaos* 19(026108)
104. Shim Y, Husbands P (2007) Feathered flyer: Integrating morphological computation and sensory reflexes into a physically simulated flapping-wing robot for robust flight manoeuvre. In: Almeida e costa F (ed) *Proceeding ECAL 07, LNCS*, vol 4648. Springer, pp 756–765
105. Shim Y, Husbands P (2012) Chaotic exploration and learning of locomotion behaviours. *Neural Comput* 24(8):2185–2222
106. Shim Y, Husbands P (2015) Incremental embodied chaotic exploration of self-organized motor behaviors with proprioceptor adaptation. *Frontiers Robotics and AI* 2:7
107. Shim Y, Husbands P (2018) The chaotic dynamics and multistability of two coupled fitzhugh-nagumo model neurons. *Adapt Behav* 26(4):165–176
108. Shim Y, Husbands P (2019) Embodied neuromechanical chaos through homeostatic regulation. *Chaos An Interdiscipl J Nonlinear Sci* 29(3):033123. <https://doi.org/10.1063/1.5078429>
109. Shim Y, Philippides A, Staras K, Husbands P (2016) Unsupervised learning in an ensemble of spiking neural networks mediated by itdp. *Plos Comput Biol* 12(10):e1005137. <https://doi.org/10.1371/journal.pcbi.1005137>
110. Skarda C, Freeman W (1987) How brains make chaos in order to make sense of the world. *Behav Brain Sci* 10:161–195
111. Steinbeck F, Graham P, Nowotny T, Philippides A (2020) Can small scale search behaviours enhance large-scale navigation? In: *Living machines 2020, lecture notes in artificial intelligence*, vol 12413
112. Steingrube S, Timme M, Worgotter F, Manoonpong P (2010) Self-organized adaptation of a simple neural circuit enables complex robot behaviour. *Nat Phys* 6:224–230
113. Sussillo D, Abbott LF (2009) Generating coherent patterns of activity from chaotic neural networks. *Neuron* 63:544–557
114. Terman D, Rubin J (2007) Neuronal dynamics and the basal ganglia. *SIAM News* 4(2) <https://archive.siam.org/news/news.php?id=1092>
115. Thomas A (1993) On the aerodynamics of birds' tails. *Philosophic Trans R Soc B Biologic Sci* 340:361–380
116. Thompson A (1995) Evolving electronic robot controllers that exploit hardware resources. In: Moran F. (ed) *Proceedings 3rd European Conference on Artificial Life (ECAL'95)*. Springer, Berlin, pp 640–656
117. Thompson A (1998) *Hardware evolution: Automatic design of electronic circuits in reconfigurable hardware by artificial evolution*. Distinguished dissertation series. Springer, New York
118. Thompson A, Layzell P, Zebulum R (1999) Explorations in design space: Unconventional electronics design through artificial evolution. *IEEE Trans Evol Comput* 3(3):167–196
119. Thrun S, Burgard W, Fox D (2005) *Probabilistic robotics*. MIT Press, Cambridge
120. Tononi G, Sporns O, Edelman G (1999) Measures of degeneracy and redundancy in biological networks. *Proc. Natl. Acad. Sci. USA* 96:3257–262
121. Traub RD, Jefferys JGR, Whittington MA (1999) *Fast oscillations in cortical circuits*. MIT Press, Cambridge

122. Trefzer MA, Lawson DMR, Bale SJ, Walker JA, Tyrrell AM (2017) Hierarchical strategies for efficient fault recovery on the reconfigurable panda device. *IEEE Trans Comput* 66(6):930–945. <https://doi.org/10.1109/TC.2016.2632722>
123. Vargas P, DiPaolo E, Harvey I, Husbands P (eds) (2014) *The horizons of evolutionary robotics*. MIT Press, Cambridge
124. Viola P, Jones M (2001) Rapid object detection using a boosted cascade of simple features. In: *Proceedings IEEE computer society conference on computer vision and pattern recognition*, vol 1. IEEE Press, pp 511–518
125. Walter W (1950) An imitation of life. *Sci Am* 182(5):42–45
126. Walter W (1951) A machine that learns. *Sci Am* 185(2):60–63
127. Warrick D, Bundle M, Dial K (2002) Bird maneuvering flight: Blurred bodies, clear heads. *Integr Comp Biol* 42(1):141–148
128. Wheeler M (2005) *Reconstructing the cognitive world: the next step*. MIT Press, Cambridge
129. Wolf A, Swift JB, Swinney HL, Vastano JA (1985) Determining Lyapunov exponents from a time series. *Physica D: Nonlinear Phenomena* 16(3):285–317
130. Wright J, Liley D (1996) Dynamics of the brain at global and microscopic scales: Neural networks and the EEG. *Behav Brain Sci* 19:285–320
131. Wystrach A, Mangan M, Philippides A, Graham P (2013) Snapshots in ants? new interpretations of paradigmatic experiments. *J Exp Biol* 216(10):1766–1770
132. Wystrach A, Philippides A, Aurejac A, Cheng K, Graham P (2014) Visual scanning behaviours and their role in the navigation of the Australian desert ant *Melophorus bagoti*. *J Comp Physiol A* 200(7):615–626
133. Wystrach A, Schwarz S, Graham P, Cheng K (2019) Running paths to nowhere: Repetition of routes shows how navigating ants modulate online the weights accorded to cues. *Anim Cogn* 22(2):213–222
134. Zang X, Iqbal S, Zhu Y, Liu X, Zhao J (2016) Applications of chaotic dynamics in robotics. *Int J Adv Robot Syst* 13(2):60
135. Zeil J, Hofmann MI, Chahl JS (2003) Catchment areas of panoramic snapshots in outdoor scenes. *JOSA A* 20(3):450–469
136. Zitzler E, Laumanns M, Thiele L (2001) *Spea2: Improving the strength pareto evolutionary algorithm*. Tech Rep, vol 103. ETH Zurich, TIK Report

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Phil Husbands is Research Professor of Artificial Intelligence and founding co-director of the Centre for Computational Neuroscience and Robotics (CCNR) at the University of Sussex. His research interests include biologically inspired adaptive robotics, evolutionary systems, computational neuroscience, history and philosophy of AI and cybernetics, machine learning and creative systems.



Yoonsik Shim is a lecturer in the Department of Game and Multimedia Engineering, Pai Chai University. He obtained a Doctoral degree in Informatics from the University of Sussex. He is a visiting member of the CCNR at Sussex. His research interests are focused on bio-inspired adaptive robotics, computational neuroscience, chaotic neurodynamics, self-organisation, and evolutionary robotics.



Michael Garvie read Computer Science at the University of Cambridge, graduating in 2001, and completed a doctorate degree at the University of Sussex in 2005. His research explores the use of evolutionary algorithms for design of analog and digital hardware applied in the fields of concurrent error detection and robotics.



Alex Dewar received his Ph.D. in computational neuroscience from the University of Sussex, where he is now employed as a Research Fellow in Robotics. His main research interests lie in visually guided navigation in robots, taking inspiration from insect species, particularly social insects.



Norbert Domcsek is a PhD Student at the University of Sussex. He is conducting research on insect inspired navigation in autonomous robots.



Paul Graham is Professor of Neuroethology at the University of Sussex where he co-directs the Centre for Computational Neuroscience and Robotics. During a Psychology degree Paul became interested in Artificial Intelligence and robotics, specifically the prospect of capturing aspects of biological intelligence by mimicking insects. This led to a PhD at Sussex with Prof Tom Collett, where they studied spatial cognition in ants. In the subsequent 20 years, Paul

has continued to study insect behaviour and neuroscience, always with a view to how we can learn from nature for the design of algorithms for robot control.



Thomas Nowotny received a PhD in theoretical Physics from the University of Leipzig and after postdoctoral training at the University of California, San Diego joined the University of Sussex in 2007 where he is now a Professor of Informatics and the head of the AI research group. His research interests include computational Neuroscience, in particular olfaction in insects, bio-inspired and embodied AI, GPU accelerated simulation of spiking neural networks

and hybrid brain-computer systems.




James Knight has a MPhil in Advanced Computer Science from the University of Cambridge and a PhD in Computer Science from the University of Manchester. He is a Research fellow in Informatics at the University of Sussex working on problems in computational neuroscience, including neuromorphic architectures and GPU acceleration, and insect-inspired autonomous robotics.



Andrew Philippides is Professor of Biorobotics at the University of Sussex where he co-directs the Centre for Computational Neuroscience and Robotics. His research spans Bio-inspired AI, Evolutionary Computation, Computational Neuroscience and Agent-based Modelling. By considering intelligence as an active process in which adaptive behaviour emerges from the interaction of body, brain and environment, he aims to both better understand intel-

ligence and develop novel AI and biorobotic algorithms. This is exemplified by his work on robotic visual navigation and exploration inspired by the remarkable visual navigation and learning abilities of ants and bees.

Affiliations

Phil Husbands¹  · Yoonsik Shim^{2,3} · Michael Garvie³ · Alex Dewar³ · Norbert Domcsek³ · Paul Graham⁴ · James Knight³ · Thomas Nowotny³ · Andrew Philippides³

¹ Centre for Computational Neuroscience and Robotics, Dept. Informatics, University of Sussex, Brighton, England

² Department of Game and Multimedia Engineering, Pai Chai University, Daejeon, South Korea

³ CCNR/Dept. Informatics, University of Sussex, Brighton, England

⁴ CCNR/ School of Life Sciences, University of Sussex, Brighton, England