Simultaneous incremental neuroevolution of motor control, navigation and object manipulation in 3D virtual creatures

Adam James Stanton

Submitted for the degree of Doctor of Philosophy

August 2017

Keele University

Abstract

There have been numerous attempts to develop 3D virtual agents by applying evolutionary processes to populations that exist in a realistic physical simulation. Whilst often contributing useful knowledge, no previous work has demonstrated the capacity to evolve a sequence of increasingly complex behaviours in a single, unified system. This thesis has this demonstration as its primary aim. A rigorous exploration of one aspect of incremental artificial evolution was carried out to understand how subtask presentations affect the whole-task generalisation performance of evolved, fixed-morphology 3D agents. Results from this work led to the design of an environment–body–control architecture that can be used as a base for evolving multiple behaviours incrementally. A simulation based on this architecture with a more complex environment was then developed and explored. This system was then adapted to include elements of physical manipulation as a first step toward a fully physical virtual creature environment demonstrating advanced evolved behaviours.

The thesis demonstrates that incremental evolutionary systems can be subject to problems of forgetting and loss of gradient, and that different complexification strategies have a strong bearing on the management of these issues. Presenting successive generations of the population to a full range of objective functions (covering and revisiting the range of complexity) outperforms straightforward linear or direct presentations, establishing a more robust approach to the evolution of naturalistic embodied agents. When combining this approach with a bespoke control architecture in a problem requiring reactive and deliberative behaviours, we see results that not only demonstrate success at the tasks, but also show a variety of intricate behaviours being used. This is the first ever example of the simultaneous incremental evolution in 3D of composite behaviours more complex than simple locomotion. Finally, the architecture demonstrably supports extension to manipulation in a feedback control task. Given the problem-agnostic controller architecture, these results indicate a system with potential for discovering yet more advanced behaviours in yet more complex environments.

Keywords: “artificial evolution”; “neural networks”; “incremental evolution”; “virtual creatures”; “3d agents”.

Acknowledgements

Words can't express the debt of gratitude I feel towards the people who have helped me on this journey. This section would be many pages long were I to include everyone who has played a part, but you know who you are, and I thank you sincerely for being there and for everything you've done.

I must above all acknowledge the generosity, kindness, character, and intelligence of Dr Alastair Channon. His supervision and guidance has been a shining light, and his friendship and light touch has made the whole process exploratory and fun. Thank you, Alastair.

Thanks also to all my friends and colleagues at Keele: I will treasure the memories of being a part of this family. Mad props are due to my PhD brother James Borg, who has been an intellectual foil par excellence: our discussions and arguments have helped me form a vision of science, art, politics, and life itself that I couldn't previously have even imagined possible. Thank you for being there for the last eight years.

Ellie Buckley, your strength and your confidence in me has been instrumental. Thanks a million for putting up with the craziness of the final stages of a PhD without even blinking an eye.

Finally, I owe it all to my parents, Roy and Paula Stanton. None of this would have been possible without your unconditional love and support. *Thank you with all my heart.*

Table of Contents

Page

1 Introduction **1**
1.1 Overview . . . 1
1.2 Motivations . . . 1
1.3 Structure of the Thesis . . . 2
1.4 Attributions . . . 3
1.5 Notes On Terminology . . . 3

2 Intelligence in Art and Nature **5**
2.1 Artificial Intelligence . . . 5
2.2 Situated AI . . . 9
2.3 Embodiment . . . 11
2.4 Enaction, Cognition and Life . . . 13
2.5 Artificial Life . . . 14
2.6 Simulation and Modelling . . . 18
2.7 Computational Intelligence . . . 20
2.7.1 Neural networks . . . 20
2.7.2 Evolutionary computation . . . 21
2.7.3 Evolutionary design . . . 24
2.7.4 Neuroevolution . . . 25
2.7.5 Incremental evolution . . . 26
2.8 Open Problems . . . 28
2.9 Thesis Aims . . . 30
2.9.1 Improving environmental complexification in incremental evolution . . . 30
2.9.2 Simultaneous incremental neuroevolution of intelligence . . . 30

3 Complex Behaviours in Embodied Artifacts **31**
3.1 Evolution of Behaviours . . . 31
3.1.1 Virtual environments for intelligent behaviour . . . 35
3.2 The Evolution of Embodied Intelligence . . . 37
3.2.1 Evolutionary robotics . . . 38
3.2.2 3D virtual creatures . . . 41
3.2.3 Agents that exhibit a single behaviour . . . 46

3.2.4 Agents with multiple behaviours 52
3.3 Chapter Summary 53

4 Guiding Incremental Evolution **55**
4.1 Introduction to Chapter 55
4.2 Complexification Strategies 56
4.3 Hypotheses 57
4.4 Methods 58
4.4.1 Physical model 58
4.4.2 Control system 59
4.4.3 Evolutionary algorithm 60
4.4.4 Experimental set-up 61
4.5 Results 62
4.6 Discussion 66
4.7 Conclusions 68

5 Reactive, Deliberative 3D Virtual Creatures **70**
5.1 Introduction to Chapter 70
5.2 Hypothesis 71
5.3 Methods 72
5.3.1 Environment and physical model 72
5.3.2 Neural architecture 74
5.3.3 Evolutionary system 79
5.4 Results 79
5.4.1 Qualitative results 79
5.4.2 Quantitative results 81
5.5 Conclusions 83

6 Object Manipulation in 3D Virtual Creatures **85**
6.1 Introduction to Chapter 85
6.2 Hypotheses 86
6.3 Methods 86
6.3.1 The physical 3D RC problem 87
6.3.2 Agent control 88
6.3.3 Evolutionary algorithm 90
6.3.4 Data collection 91
6.4 Results 91
6.5 Conclusions and Discussion 98

7 Conclusions **100**
7.1 Summary 100
7.2 Contributions 101
7.3 Key Findings, Technologies and Implications 102
7.3.1 Heterogeneous environmental complexification 102

7.3.2 Hybrid control architectures for 3D virtual creatures 104
7.3.3 Adapting hybrid architectures to physical-manipulation tasks 105
7.4 Limitations and Further Work 105

Bibliography and References **107**

CHAPTER 1

Introduction

1.1 Overview

This dissertation examines various aspects of the incremental evolution of 3D virtual creatures (3D VCs). The objective is to explore the cutting edge in evolved 3D VCs and present the problems and some solutions to the problems that are encountered when engineering such systems. Historically, there have been numerous attempts to develop 3D VCs by applying evolutionary processes to populations that exist in a realistic physical simulation. Whilst often contributing useful knowledge, no previous work has demonstrated the capacity to evolve a sequence of increasingly complex behaviours in a single, unified system. This thesis has this demonstration as its primary aim. The work builds on work previously published in the ECAL and ALIFE conference series.

1.2 Motivations

One of the earliest and most striking research areas within Artificial Life is the attempt to generate simulated organisms on an animal scale, whose behaviours are recognisable analogues of real-world animal activity. Whilst these *virtual creatures* are a distinctly technological endeavour, the research programme can nonetheless offer a great deal to the scientific community beyond a mere intellectual curiosity. Research in virtual creatures provides new hypotheses to explore in order to unpick the mechanisms of evolution in the natural world and explain the origins of intelligent, adaptive behaviour in nature. Additionally, this synthetic approach to science has tangible practical benefits, facilitating the construction of a methodology for building advanced intelligent technology. Finally, in terms of broader, longer-term objectives, virtual creatures are, as argued by Channon (2001), crucial to addressing the problem of understanding artificial open-ended evolutionary (OEE) systems. The potential complexity of these systems requires an instinctive validation by observation even when metrics may indicate that ongoing evolutionary activity is taking place. The closer the medium of expression of this activity is to regular human experience, the more we can have subjective confidence in our objective measures. A physically-realistic, evolutionary-expressive environment of the sort ubiquitious in virtual creatures studies is also then an objective for work in the OEE domain.

The research presented in this thesis is thus concerned with the generation of varied behaviour in simulated organisms using nature-inspired principles in order to understand how to build this technology, to explore hypotheses about how this came to be in nature and to provide a step towards flexible and expressive substrate that would allow a future open-ended evolutionary system to be validated against intuitive experience.

1.3 Structure of the Thesis

- Chapters 2 and 3 comprise a review of relevant literature, split into two parts. The first part discusses the general, historical research trends that have led to the current research topics of evolutionary robotics and evolved 3D virtual creatures, and from this presents the main objectives of this dissertation in terms of those disciplines. The second part examines research relevant to addressing these research objectives, identifying gaps in current knowledge and looking at technologies and techniques that can be employed to advance the current state-of-the-art.

- Chapter 4 presents the first contribution of the research, a detailed investigation of a technique to guide evolutionary generalisation in a principled way that has hitherto only been assumed by research in this field. The output of this chapter is a principled and theoretically-informed understanding of environmental complexification, explored and articulated with reference to a 3D agent-based baseline task. The work in this chapter has been presented at ECAL 2013 (Taormina, Sicily) and published in Stanton and Channon (2013).

- Chapter 5 uses this technique as part of a novel synthesis of a 3D virtual creature system that demonstrates for the first time multiple evolved behaviours that go beyond simple locomotion. This work shows the utility of the findings from chapter 4 and the potential that such systems have for further complex, and perhaps unbounded development of behaviours. The output of this chapter is a novel synthesis of existing and new technology to produce a system capable of reactive and deliberative behaviours in 3D virtual creatures and representing a step change in the complexity of such systems. The work in chapter four has been presented at ECAL 2015 (York, UK) and published in Stanton and Channon (2015).

- Chapter 6 takes the neuroevolutionary system developed in chapter 5 and applies it with an environment requiring physical manipulation of an object in the agent's world. Object manipulation is a key component of intelligent behaviour in 3D environments, whether in nature or in silico. The results presented here demonstrate the first ever success at such a task using an evolved architecture capable also of solving deliberative planning problems. The work described in chapter five has been presented at ALIFE 2016 (Cancun, Mexico) and published in Stanton and Channon (2016).

- Finally in chapter 7, an overview of the whole dissertation is presented, the findings and implications from each chapter discussed and directions for further research are examined.

1.4 Attributions

This work was executed under the supervision of Alastair Channon. Work completed explicitly by Adam Stanton is as follows:

- Identification of the question of task presentation strategy which constitutes the research question addressed in chapter 4;

- Development of technology (programming, cluster algorithms, visualisations and design and integration of other technologies and scientific ideas) and experimental procedure (automated data aggregation, statistical analysis, and experiment design);

- Conception of the 3D river-crossing problem as an environment in which 3D agent behaviour could be incrementally evolved past simple locomotion;

- Design and implementation of the hybrid neural network described in chapter 5; and

- Adaptation of this technology to the physical 3D environment and subsequent experimental design and analysis in chapter 6.

1.5 Notes On Terminology

Throughout this thesis, I refer at times variously to 3D agents, virtual creatures and evolved agents, and combinations of these terms to suit the context. Unless specified otherwise, at all times when using this terminology I am referring to *fixed-morphology evolved virtual creatures (EVCs)*, where fixed-morphology is in constrast to evolved morphology and means the agent's physical body plan remains unchanged over evolutionary time.

In addition, I often refer to the general concept of intelligence and intelligent behaviours. The MIT Encyclopedia of the Cognitive Sciences defines intelligence as “the ability to adapt to, shape, and select environments” (Wilson and Keil, 2001, p.409) and it is this broad sense I wish the reader to bear in mind when considering the text.

Finally, I often use the terms *difficult* and *complex*. The OED defines difficult as ‘needing much effort or skill’ and complex as synonymous with *complicated*, ‘intricate’, or something that has been made difficult. These meanings, rather than any more technical senses, are the ones intended in this text.

CHAPTER 2

Intelligence in Art and Nature

> "If you wish to make an apple pie from scratch, you must first invent the universe."
> **–Carl Sagan**

Artificial Intelligence research led to a new conceptualisation of intelligent behaviour, cast not as an abstract algorithm but as a situated, dynamic relationship between agents and environments. Artificial Life seeks to understand the theoretical object of biology through simulation and synthesis of natural processes. At the interface between Artificial Intelligence and Artificial Life lies evolved intelligent behaviour, a phenomenon with tangible technological value as well as the potential to investigate through synthetic means the structures found in biology, and to demonstrate the products of evolution "as they could be" rather then as they are. The following literature review first traces a path through these concepts to justify research in advanced evolved behaviours and explain how it can make a contribution to science and technology. Then, I define the boundaries of the research project and discuss the techniques and technologies that can be used to push beyond the state-of-the-art.

2.1 Artificial Intelligence

> "The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."
> **–Proposal for the 1956 Dartmouth College meeting**

The scientific and technological research programme called *Artificial Intelligence* (AI) was born in the middle of the 20th century and seen from its inception as a discipline concerned with building models of the mind and constructing intelligent machines. At a time when general-purpose computing machinery had begun to deliver real results and philosophy began to grapple in earnest with the possibility of considering “mind” as a physical property of the world amenable to investigation by scientific methods, the analogy of mind-as-computer (and therefore the algorithmicity of behaviour in general) took hold and refused to give way. The decomposition of high-level intelligent behaviours into algorithmic form, using abstract symbols to represent entities in the world, quickly became the dominant approach to building AI, despite some early efforts in the late 1950s to use networks of processing elements to simulate intelligent behaviour: “[human-like intelligence] was imagined as a kind of logical reasoning device coupled with a store of explicit data” (Clark, 1998). This knowledge-based approach was driven by the advances made in computing theory and hardware, and the emergence of cognitivism in psychology as a discipline that legitimised the scientific study of mind in itself (as opposed to study only of its behavioural outcomes) and whose prevailing explanatory framework was symbolic (i.e. representative) and systematic. The many early successes of other branches of informatics also contributed to the sense that computationalism was the best chance of explaining intelligence and building intelligent machines (Harnad, 1990).

This abstracted, analytical and above all *symbolic* mode of thinking about intelligence entailed many successful applications: chess-playing, theorem-proving and expert deduction computer programs all began to exceed the best human performance. Problems whose constituent components and relationships could be clearly formalised and represented in symbolic form amenable to algorithmic manipulation became clear targets for the increasingly powerful computational hardware of the 1960s and 1970s. It was an optimistic time, and many researchers believed that an increase in computational capacity was the only outstanding requirement necessary to achieving the construction of an artifact exhibiting generalised intelligent behaviour. This attitude was encapsulated in the *physical symbol system hypothesis* of Newell and Simon (1976) that made this strong claim about the potential of symbolic AI explicit. However during the same time, theoretical and practical difficulties became evident in the characterisation of some putative “intelligent” activities.

As attempts were made to expand symbolic AI to more generalised problems, even those solved by the simplest of living forms such as controlling a body or navigating a cluttered environment, it became clear that flexible, adaptable intelligent behaviour was difficult to neatly conceptualise, difficult to decompose into logical relationships and difficult to formulate inside tractable algorithms. In more recent history, research has continued to apply formal, symbolic reasoning methodologies to difficult problems, automating previously manual processes and often also improving their efficiency and accuracy. However, high-profile, expensive projects requiring more generalised intelligent behaviour often failed to live up to expectations. Some more theoretical thinkers who were interested in general intelligent behaviour were openly critical of the symbolic, representational methodology. John Searle's famous *Chinese room* argument posits the absurdity of claiming that a rule-based symbol manipulation machine has any intrinsic understanding of its inputs and outputs. To illustrate briefly, a non-Chinese speaker in the room receives Chinese text through a letterbox, looks up a matching answer in a rule book based on the input symbols and posts the corresponding response back to the world. By contrasting this to the action of a fluent Chinese speaker in the room who is working without the book, Searle claims a material difference between the two systems in that the first does not connect the syntax of the symbols to the semantics in the real world–the symbols are not *grounded*. Searle's conclusion is that this problem precludes the possiblilty of achieving human- or animal-level intelligence in a symbolic machine because it has no real knowledge of the world and will forever lack intentionality[1] (Searle, 1980). Although often criticised as unnecessarily restricted to the person's experience in the room, the most common reply being that the *system* as a whole behaves as if the symbols are grounded as much as any person does (Russell and Norvig, 2003), the argument still highlights the gulf between naturalistic intelligence and symbolic AI and hints that the computational paradigm may be lacking in its fundamental assumptions about the abstract nature of intelligent behaviour. A more convincing argument from a logical perspective also obtains in the physical symbol system hypothesis. The *frame problem* asks how a symbol-manipulating machine, starting from a set of axioms about the world, can maintain an appropriately fast and germane behavioural response in the face of a requirement to perform logical inference about the entire world for an indefinite amount of time (Dennett, 1981; McCarthy and Hayes, 1969). The conclusion that the combinatorial explosion entailed by the proposition renders any inferential computation intractable for real-world problems led to growing support for the alternative hypothesis that intelligent behaviour in the natural world operates according to some other, non-computational, *sub-symbolic* paradigm.

Given these intellectual and practical difficulties, a clear and increasing separation was observed between the logical, knowledge-based intelligence that is well fitted to highly constrained problems, and *Artificial General Intelligence* capable of acting intelligently in an animalistic fashion. The need to design a machine able to cope gracefully in an uncertain world with multiple competing goals and partial information drew researchers away from the mainstream, orthodox approach to AI and back toward the ancient underlying aspiration to explain the history and mechanisms of biological intelligent behaviour. This bifurcation was anticipated to some extent by Turing in his seminal 1950 work, *Computing Machinery and Intelligence*. In this paper, Turing mused on whether machine intelligence would be best tested by "a very abstract activity, like the playing of chess", or whether providing the machine with "the best sense organs money can buy, and teaching it to understand and speak English" might be better, concluding

[1] Intentionality: that feature of certain mental states by which they are directed at or about objects and states of affairs in the world. (Searle, 1980)

that "both approaches should be tried" (Turing, 1950). The fundamental difference in perspective between the orthodox and alternative paradigms comes from a disjoint view of the nature of intelligence itself. The traditional approach offers a view of intelligence as a high-level algorithmic process that is limited only by access to a sufficiently detailed world model and sufficiently powerful computer. Intelligent behaviour results from an exploratory search through the representation, so is limited by the amount of available information, the design of the model and the computational resources on hand. In contrast, the alternative view sees complex intelligent behaviour as an intrinsic dynamical relationship between agents and their environments. As noted by Brooks, "The key observation is that the world is its own best model. It is always exactly up to date. It always contains every detail there is to be known. The trick is to sense it appropriately and often enough." (Brooks, 1990).

Support for this point of view comes from an argument from evolutionary history. In slightly later work, Brooks (1991) points out that the bulk of natural history was spent finding out how to do relatively simple survival-related tasks in a dynamic environment:

> "This suggests that problem solving behavior, language, expert knowledge and application, and reason, are all pretty simple once the essence of being and reacting are available. That essence is the ability to move around in a dynamic environment, sensing the surroundings to a degree sufficient to achieve the necessary maintenance of life and reproduction. This part of intelligence is where evolution has concentrated its time–it is much harder."

Ultimately, this paradigm casts adaptive behaviour not as a complex yet deterministic algorithm that by operating on abstract representations of knowledge dictates the functioning of organisms in the world. Instead, intelligence is seen as a tightly-woven set of often probabalistic relationships between living organisms and their environments: a model of the interactions between entities acquired through the ongoing experiences of ecosystems, species and individuals (Van Gelder, 1998). Individuals are understood in relation to their worlds and their interactions, a perspective with an emphasis on the mutual interconnectedness of environment, body and mind (Tschacher and Bergomi, 2011). Harvey captures the essence thus: "cognition, as ascribed to animals or potentially to machines, is something that can only be attributed to the conjunction of an organism and the world it inhabits." (Harvey, 1992b). In the next sections, this alternative view is briefly explored with reference to the contexts of the physical environment *(situatedness)* and the body and body movement *(embodiment)*, and brought together with wider studies of cognition in living systems under Varela's banner of *enaction*. Then I examine practical techniques that can be used to achieve a level of enactive cognitive behaviour in real autonomous systems.

2.2 Situated AI

> "It was assumed that what the conjurer needed was not the correct wiring diagram but the correct recipe, according to which the ingredients would organize *themselves* into an intelligence." (Hillis, 1998)
> **–W. Daniel Hillis, The Pattern on the Stone**

The framing of intelligent behaviour as dynamical relationships between agents and their environments has entailed a great deal of philosophical speculation as well as many concrete efforts to produce machines capable of exhibiting such behavioural properties. The most fundamental points are that intelligence in this view is both *bottom-up*, i.e. reliant on the combination of elementary units that together produce more complex behaviours, and *behaviour-based*, that is, constituting a model of interactions rather than knowledge representation.

As noted by many scholars who have discussed the history of AI and the emergence of this post-computational approach to understanding cognition, parallels can be drawn with the work of 20th century phenomenological philosophers from the continental school which held that human experience is a contingent creation founded on material existence in the world. However following the tradition of Brooks (1991) I will leave this discussion to others more qualified than me, and concentrate on the engineering considerations and their implications for constructing intelligent technology. This being so, I would like to offer a technological example to illustrate the necessity and sufficiency of being in the world for effective control.

In 1788, James Watt popularised an ingenious device that maintained the power output of steam engines at a near constant strength. The *centrifugal governor* is a feedback control architecture that, as the engine's speed increases, acts proportionally on a throttle valve that acts to calm the machine. No information about the engine state is explicitly represented in this system and yet stability is maintained in the presence of external forces that disturb its operation. It is interesting to note that this analogy has been used for other feedback relationships in the natural world. Wallace draws the same analogy with evolving systems and the stabilising mechanism of natural selection:

> "The action of this principle is exactly like that of the centrifugal governor of the steam engine, which checks and corrects any irregularities almost before they become evident; and in like manner no unbalanced deficiency in the animal kingdom can ever reach any conspicuous magnitude, because it would make itself felt at the very first step, by rendering existence difficult and extinction almost sure to follow." (Wallace, 1858)

This self-maintaining, self-governing idea is founded on interactions, not on knowledge. The concepts inspired scientists and engineers interested in building life-like artifacts, and the first practical attempts to build devices in this paradigm are due to early cyberneticists such as W. Ross Ashby, whose *homeostat* was able to find a configuration of internal electronic circuits and relays that "leads to a condition of dynamic internal stability" (Walter, 1950) - a physical manifestation of a property devised by Ashby called *ultrastability* (Ashby, 1948). Ashby viewed the brain as a dynamical system that generates and modulates behaviour, and that its operation in the face of external pertubations to bring the organism back to a balance point was the crucial aspect. Even though Ashby's model was embedded in discrete mathematics to the extent that he was invited by Turing to simulate the homeostat on Turing's new computing machine (Turing, 1946), it was rooted fundamentally in a continuous, abstracted and biologically-inspired medium that transcended any particular substrate of implementation.

W Grey Walter's precocious electronic tortoises, his *Machina speculatrix*, provided a more intuitive illustration of these concepts. "Elmer" and "Elsie", ungainly wheeled robots equipped with minimal sensory and computational hardware but nevertheless autonomous in the confines of their environments showed that complex, life-like behaviours could be observed in even extremely simple machines (Walter, 1950). The tortoises were constructed not with algorithmic logic but with behaviours arising from the machines' interaction with the world, mediated by connections between the effectors and sense organs of their bodies.

Booker describes organisms that show this kind of behaviour-based, just-in-time activity in an environment as *adaptively salient* (Booker, 1982). Casting intelligent behaviour as an adaptation to the task evironment in which the agent finds itself, he cites Charlesworth: "it does not make sense to talk about adaptation without something to adapt to. And if one designates intelligence as an important mode of adaptation, then the intelligent behaviour has to be viewed in terms of environmentally posed problems" (Charlesworth, 1976). Booker's *Learning Classifier System* was one of the first implementations of a behavioural controller designed explicitly with a messy paradigm where inputs match one or many rules and a cascade of activation happens to cause a direct behavioural change. Whilst still instantiated in a symbol-manipulating computational machine, Booker's approach offers one of the first situated alternatives to traditional AI systems. This method anticipates Minsky's later formal treatise on *society of mind* that proposes a different understanding of AI in the form of a system comprising a heterogeneous collection of individual, minimally intelligent *agents* that together form the basis for an adaptive, reactive cognitive model (Minsky, 1985).

The theme continues into modern robotics through Valentino Braitenberg's thought experiments (or as he terms it, "fictional science") *Vehicles* that aimed to demonstrate the enormous complexity in behaviour afforded by rich environments to simple organisms that interact with them, and also to uncover the unavoidable observer-centric interpretation of these behaviours (Braitenberg, 1984). These vehicles comprised straightforward environmental sensors like light or chemical detectors, coupled to motors that cause the vehicles to move around. Even with only one or two links between sensors and motors, Braitenberg's vehicles demonstrate that complex, even life-like behaviours could arise when environments are sufficiently rich. Independently, Wilson produced a dynamic model of a simple artificial animal (coining the term *animat* in the process) operating according to simple, sensorimotor principles (Wilson, 1985). Later, Philip E. Agre's work explicitly rejected formal, plan-based behaviour generation models of human activity (and by implication, much of the extant cutting-edge research in autonomous behaviour and automated planning up to this date) in favour of a dynamic model that treats the agent–environment relationship as composed of categories without objective identity (Agre, 1988). Agre attempted to define a kind of postmodern, Heideggerian ontology in an explicity agent-centric schema which, although strongly reminiscient of predicate logic-based approaches to navigating complex behavioural problems, did away with the need for explicit representation of objects in the world. He supported these ideas with a demonstration computer program called *Pengi*. "Rather than maintaining elaborate world models and constructing symbolic plans, Pengi relies heavily on its interactions with the world to organize its activity." (ibid.) The concepts of interaction-based intelligence can be found more starkly applied to behavioural robotics (and much more loosely related to earlier, GOFAI mechanisms) in Brooks' work on subsumption-based models of robot control (Brooks, 1986) where explicit representation and planning are both proscribed and all behaviour is just-in-time and messy (Brooks, 1991). As in Agre's work, the behaviours in this architecture are still specified at design time but are then selected automatically according to dynamically-arising interactions between individual behavioural modules and the wider environment in which the robot is embedded. This model led to robust and life-like operation of Brooks' robots, beginning with *Genghis*, a hexapod walking machine able to locomote effectively over a variety of terrains (Brooks, 1989).

The shift towards interaction-based intelligent behaviour addresses the primary criticism of symbolic AI–the frame problem. Intelligent agents are no longer paralysed by interminable ratiocination; they give it their best shot right away based on the information available to them. There is still a question over representation, and how agents behaviour can be said to be grounded in their sensorimotor experience. Systems like Agre's and Brooks' deal with the immediacy and are prima facie without explicit representation. However, they still rely on external decompositions to operate and so in a sense are still at the risk of being symbolic manipulations encoded by human designers rather than intrinsic relationships between real-life entities. The philosophy of embodiment begins to address this problem.

2.3 Embodiment

In 1921, Jakob von Uexküll invited his readers to consider the sensory worlds of animals in their totality and to extrapolate the emergence of the animal's behaviours from this (von Uexküll et al., 2010), stressing the relativity of all signs and signals in nature. Von Uexküll saw all living beings as embodied subjects

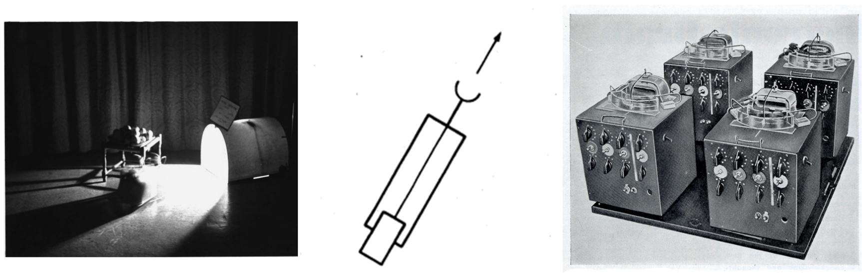

Figure 2.1: Grey Walter's Tortoises, *Elmer* and *Elsie*, near their recharging station; Braitenberg's first *Vehicle*, the simplest demonstration of something that in a complex environment could display life-like behaviours; Ashby's *homeostat*.

that are embedded in an intricate web of intrinsically meaningful relationships, their *Umwelt* (Froese et al., 2010), and thus there is an inseparable and only partially knowable relationship in existence between creatures and the worlds that they inhabit. This relationship is built through the medium common to both–the physical world.

In this sense, living creatures by virtue of their *embodiment* in the world are grounded. Meaningful interactions occur on a fundamental level and the experiences of organisms have a direct bearing on their behaviours. This grounding of knowledge in worldly experience not only deals with the criticisms of GOFAI in the sense of being disembodied and abstracted but more practically makes the problems of existence more apparent and, due to this, causes the builders of artificial intelligent entities to adopt a more naturalistic philosophy towards their construction. The physical aspects of agents' bodies are taken to be both pragmatically and theoretically significant and are located and specific. Context-dependence is a central and enabling feature and activity varies dramatically dependent on contingent facts about particular circumstances (Wilson and Keil, 2001, p.769). The further view that physical embedding is not only a semantic resource for determining reference but a material resource for simplifying thought itself (ibid.) is entailed by the embodied perspective, and naturally implies that "behaviour and cognitive capabilities should emerge out of their interaction with the world" (Tuci, 2004).

It should be noted that some researchers consider symbol grounding a problem that has already been solved early in the history of AI, and claim that a lot of the confusion and debate arising from the question is due to an interdisciplinary misunderstanding concerning the meaning of the term *symbol*. However, they still maintain that a mechanism of *automatically* grounding symbols escaped research until relatively recently. As argued by Steels (2008):

> "So the key question for symbol grounding is not whether a robot can be programmed to engage in some interaction which involves the use of symbols grounded in reality through his sensori-motor embodiment,

that question has been solved. It is actually another question, well formulated by (Harnad, 1990): If someone claims that a robot can deal with grounded symbols, we expect that this robot autonomously establishes the semiotic networks that it is going to use to relate symbols with the world."

2.4 Enaction, Cognition and Life

In a separate and long-standing research direction, biologists Maturana and Varela also put the dynamical relationship between individuals and their environments centre-stage (Maturana and Varela, 1980) in their attempt to reframe understanding the fundamentals of biological organisms, metabolism and living systems. In their view, definitions of life and cognition are inextricably linked through their embedding in the world and there is no sense in which intelligent behaviour can be quantified or demarcated outside of the complete systems viewpoint.

Enaction, first articulated by Varela, Thompson, and Rosch in The Embodied Mind (MIT Press, 1991), breaks from computational cognitive science's formalisms of information processing and symbolic representations to view cognition as grounded in the sensorimotor dynamics of the interactions between a living organism and its environment. A living organism enacts the world it lives in; its embodied action in the world constitutes its perception and thereby grounds its cognition. Enaction offers a range of perspectives on this exciting new approach to embodied cognitive science (Varela et al., 1991). More recently, Solé and Valverde (2013) make clear that physical embedding of evolutionary processes is crucial if macroevolutive processes are to observed. The physical environments promote evolutionary emergence on multiple scales, and facilitate feedback relationships between different layers of the hierarchy.

The work of Andy Clark served as a rally point for ideas of embodiment in cognitive science and robotics (Clark, 1998), and ideas of situatedness percolated through to research themes in many other areas, including human psychological research (Bargh and Chartrand, 1999). And so, from a combination of influences spanning an enormous intellectual range from theoretical biology, semiotics and philosophy of mind through to demonstrations of six-legged walking robots, the field of behaviour-oriented AI emerged, concerned directly with the study of relationships between brains, bodies and environments and the application of the ideas to real-world problems (Ziemke, 1998). While the deeper philosophies of mind concern themselves with the nature of enaction as a way to understand cognition, new robotics and new AI have taken a more pragmatic approach and attempted to incorporate these ideas directly into practical, functional artifacts.

As can be seen from the work outlined above, the trend in this behaviour-based *new robotics* was clearly toward a highly integrative view of brains, bodies and environments. The systems perspective elevated the complexities arising from physical embedding in the world to the level of those supposed to exist in natural control architectures like DNA and metabolic networks, and animal brains. As delightfully illustrated in Chiel and Beer (1997):

> "the nervous system is one of a group of players engaged in jazz improvization, and the final result emerges from the continued give and take between them. In other words, adaptive behavior is the result of the continuous interaction between the nervous system, the body and the environment, each of which have rich, complicated, highly structured dynamics."

There is still some debate over whether these principles of situation, embodiment and enaction apply strictly to artifacts in the real world, or whether simulated worlds can reify objects within them in the same way. In this thesis I argue for the latter–that provided that agents' worlds form closed universes in which the ideas of situatedness and embodiment have maximum validity, then the end result is the same. This is aligned with the concept of *functional validity* (Channon, 2001; Channon and Damper, 1998b), whereby the situatedness of agents during development correlates directly with their performance (a question of *reality gaps*). As argued by Channon:

> "if organisms are only ever to inhabit an 'artificial' environment then there should be no concern about them being evolved in that environment. Their 'world' is not a simulation and so the approach suffers none of the problems that occur when trying to use a simulation to evolve robots for the real world."

The myriad interrelations between concepts in new robotics and embodied cognition research show that the development of intelligent machines in the behaviour-oriented paradigm is intimately connected to explanations of the nature of living things in themselves, so it is no surprise that the field shares deep links and a great deal of overlap with the emerging interdisciplinary field called *Artificial Life*, which is the subject of the next section. As described by Steels (1993), in this field the connection between behaviour-oriented AI and biology is almost axiomatic: "Intelligence is seen as a biological characteristic, and the core of intelligence and cognitive abilities is [assumed to be] the same as the capacity of the living."

2.5 Artificial Life

Artificial Life (a-life) is a research programme that aims to explain a variety of biological phenomena, as diverse as metabolic networks and evolutionary dynamics, as well as characterise the abstract properties of life itself: the rules governing "life as it could be" (Wilson and Keil, 2001). This aim is most often approached with a synthetic methodology, where the subtleties of the systems under microscope are teased out by building detailed working models (Brooks, 2000).

This is an interdisciplinary undertaking that seeks to investigate the fundamental properties of living systems and build artificial systems that display recognisable properties of organisms and societies from abiotic components. A-life includes study of biology, chemistry and computer science; the focus in the present work is on computer science, and on how to engineer life-like systems in simulation. This endeavour has a distinguished history and is a founding component of the field.

The earliest examples of a-life systems date back to work by Hungarian mathematician John von Neumann who regarded life not as a mysterious property of matter, but rather as a process that could be decomposed and understood theoretically. Von Neumann conceived the cellular automaton, a mathematical system that simulated life-like interactions (Von Neumann, 1966) and went on to theorise that life-like processes could be modelled using a very simple concept of a *universal constructor*, a coupled system of rules that, when executed, could create copies of itself. This work, conceived in 1948 and published posthumously in 1966, constitutes one of the first attempts to formalise an abstraction of the natural world by breaking down organisms into functional components (for example, neurons, muscles, construction mechanism and structural bodies) and finding logical processes that connect these components in configurations that result in a stable dynamical reproductive process.

Von Neumann's research, although pioneering, was undertaken before a practical method of simulation existed and so was a purely philosophical treatise. However, the arrival of the digital computer allowed researchers to test their ideas empirically. In 1970, mathematician John Conway invented *Life*, a simple, two-dimensional cellular automaton with uniform rules specifying the behaviour of each cell. When executed sequentially, the system produces chaotic patterns resembling the dynamics observed in many real-life processes (Gardner, 1970). The apparently simple simulation environment–much simpler than von Neumann's original conception–produced an unexpected array of stable patterns that were discovered, analysed and documented by researchers of that time. This is a classic example of the exciting and unpredictable nature of a-life: simple systems can produce remarkable results due to the global complexity that arises from local interaction.

Inspired perhaps by Conway's *Life* but certainly by the work of von Neumann, Chris Langton, one of the founders of the a-life field, attempted to marry the local-interaction global-behaviour paradigm of the cellular automaton with von Neumann's ideal functional decomposition of life-like systems (Langton, 1986). In this work, Langton breaks down the biochemical processes in living systems into their functional roles (e.g. catalysis, transport, energy storage) and implements a simulator of this system. It was discovered that virtual automata, that is, emergent functional phenomena, arise frequently in the simulation and were observed to fulfil some of the functional roles identified in real-life systems. Langton argues that this process of starting with a set of behavioural primitives and using them as a set of building blocks to discover more complex behaviours constitutes an abstraction of the essence of life into any virtual medium and is a crucial step between building models of life and building examples of life.

Indeed, the original demarcation of the field as a research programme in itself is due principally to Langton's work of the 1980s (Brooks, 2000). Originally, Langton sought to describe the behaviour of living systems in terms of dynamical systems and cellular automata (Langton, 1986) but then went on to expand his definition of the field to be "the simulation and synthesis of living systems" (Langton, 1989), a working encapsulation still accepted by contemporary researchers.

Philosophical difficulties lurk at the most fundamental levels in a-life; the question of whether a formal system can really be alive–see e.g. Boden (2000), or Lenski (2001) for a more accessible presentation–is still an open, and perhaps unanswerable problem. Philosophers, artists and scientists continue to explore the theoretical, epistemological framework necessary to pin down a firm understanding of living systems (Annunziato and Pierucci, 2002; Shanken, 1998), even as modern thinkers recognise the need for pragmatism in definition-of-life problems when working towards concrete scientific ends (Pennock, 2012). Researchers with a more traditional background in the biological sciences have long recognised the new opportunities for study that arise in a-life (Hokkanen, 1999), whilst critical theorists use a-life as a *limit biology*, posited to exemplify how the ontological category of nature itself has become unstable: no longer is matter necessary to describe life, and the elevation of the concept to the realm of descriptive information, to dynamical systems, invites an understanding of life simultaneously as a category and a construction (Helmreich et al., 2015). This is a perspective on the natural world that is as deconstructed as the enactive, embedded and embodied approach to the theories of meaning and relation in understanding intelligent behaviour that were outlined in the previous sections.

Notwithstanding, this overarching drive toward increased understanding of the theory of self-organisation and emergence, as well as the practical problems of building artifacts to explore the nature of complex systems analogous to biological ones have at the same time advanced our scientific understanding of the world and borne fruitful technological innovation in many areas, including robot control, manufacturing, graphics, games, design, security and telecoms (Kim and Cho, 2006).

Concerning simulation, a-life spans a range of topics. At one end, philosophical discussion concerning the nature of model-building as it applies to scientific endeavour is examined. At the other, simulations of natural phenomena with various simplifying assumptions are played out, with the objective of revealing important, invariant relationships from the turmoil of complexity that comprises the myriad parameters and processes governing their evolution.

The synthesis component of Langton's definition also covers a huge array of activity; at its purest this synthesis aims to produce artifical living systems–a singularity where simulation and synthesis are indistinguishable. More recently and in general more practically, a-life synthesis covers efforts to produce artificial phenomena that do not necessarily map out the natural world but are nonetheless analogous in terms of higher-order emergent properties. This overlaps with research from other fields and includes exploratory optimisation techniques such as Evolutionary Algorithms (EAs) and representational modalities such as Neural Networks (NNs), as well as producing native a-life research like Artificial Chemistries (ACs) and Virtual Creatures (VCs).

The diversity in a-life is due to the early formation of the research programme which was rooted in the technological developments that occurred during the late twentieth century. History shows overlap in ideas and applications between many different emerging cultures, as described by Penny:

> Artificial Life burst onto a cultural context in the early 90s when artists and theorist were struggling with the practical and theoretical implications of computing–that is, it was contemporaneous with virtual reality, bottom-up robotics, autonomous agents, real-time computer graphics, the emergence of the Internet and the web and a general interest in interactivity and human-computer interaction. (Penny, 2009)

Despite this heterogeneity, according to Bedau et al., the focus of a-life research should above all still be on science: “Artificial life is foremost a scientific rather than an engineering endeavor. Given how ignorant we still are about the emergence and evolution of living systems, artificial life should emphasize understanding first and applications second[.]” (Bedau et al., 2000). However, in the same work, which is an overview written to highlight open challenges in the field, it is recognised that technological artifacts are one of the clearest methods for demonstrating this understanding. In particular, the ambition to “Demonstrate the emergence of intelligence and mind in an artificial living system”, a problem directly related to understanding whether mind and intelligence are only meaningful when embodied in living systems. They continue, “The easiest aspect of mind and intelligence to detect is flexible adaptive behavior, i.e., the capacity to act appropriately in a complex dynamic environment”, arguing that this kind of demonstration can help to settle long-standing open problems in a-life and AI. The point concludes by describing the forms that a useful example of this property could take: “having increasingly sophisticated forms of this capacity emerge from increasingly impoverished initial conditions.” (ibid.)

A major open problem in a-life is thus the demonstration of these principles of intelligence in an artificial system. As a technological effort there is clear advantage in such an achievement and the synthetic methodology found throughout a-life lends itself naturally to such demonstrations. However, the epistemological validity of scientific simulations can be brought into question. This is one of the criticisms of the a-life discipline as a research programme, a “fact-free science”, according to Maynard-Smith (Horgan, 1995), so in the next section

I briefly discuss how the study of a-life can contribute scientifically as well as technologically.

2.6 Simulation and Modelling

> "The story so far: In the beginning the Universe was created. This has made a lot of people very angry and has been widely regarded as a bad move."
> **–Douglas Adams**

The scientific method is founded on the observation of nature and the falsification of theories that concern the relationships between the constituent observed phenomena. Until recently, the empirical observations of the world were carried out directly, with instruments designed to focus the world on our senses. Since the advent of the modern day computer, these experiential knowledge-generating processes have been augmented by a new, automated and algorithmic methodology of investigation. The ascendancy of the computational paradigm has had a revolutionary effect on fundamental scientific research. Computerised numerical simulations have come to act as scientific instruments in their own right, affording a temporally-extended view of analytically intractable systems, the long-term histories of which were hitherto opaque (systems that today are known as complex systems). With these new tools we have explored the furthest reaches of time and space and refined our understanding of the cosmos on the largest and smallest scales, an undertaking otherwise impossible due to the difference in temporal scale from our physical vantage point (Ayala and Forero-Romero, 2013), often with tangible technological results. In all these cases, simulation enables the execution of thought experiments accurately and quickly, guiding research directions and suggesting new avenues to explore. Simulation also allows imperfect models of the world, built through incomplete observations or with unavoidable assumptions or simplifications, to be examined for consistency. Sometimes, the ultimate outputs of these simulations are validated empirically to assert the quality of the conceptualisation itself.

The study of life and living systems is also suffused by the epistemic third-way of scientific simulation and some of these research themes have coalesced in the a-life research programme. This interdisciplinary undertaking seeks to investigate and explain the fundamental properties of living systems through formalisms and synthesis, in the realm of material things as well as in simulation (Langton, 1989). A-life studies include biological, chemical and computational perspectives and it was one of the earliest forums for the emerging field of complex systems science. Research in a-life ranges from the smallest scales (artificial chemistries) to the largest (social simulations), and from the most abstract (computational ecosystems of competing program code) to the most grounded (biophysical simulation of whole organisms).

The ongoing incursion of simulation into the natural sciences, and particularly in a-life, is not without drawbacks. As observed by Arnold (2013), “computer simulations are not material in any sense that would liken them to experiments” and “experiments are not intertwined with models to such a degree that the function of models in experiments becomes indistinguishable from the function of models in simulations”. It follows that in many situations, the strength of justifications for theories that arise from simulations is reduced and their function is correspondingly diminished.

These kinds of simulations require external validation of their axioms and outcomes against empirical data if they are to justify knowledge about the world. However, in many cases computer simulations can act as proofs-of-concept or proofs-of-sufficiency. This occurs when the simulation is in a context with an observable, satisfiable objective but without a need to justify theories that also apply in nature. In yet other cases, simulations are truly material and have as much intrinsic empirical validity as real-world experiments.

As Tom Ray writes,

> “I would like to suggest that software syntheses in a-life could be divided into two kinds: simulations and instantiations of life processes. A-life simulations represent an advance in biological modeling, based on a bottom-up approach, which has been made possible by the increase of available computational power.”

he continues:

> “The second approach to software synthesis is what I have called instantiation rather than simulation. In simulation, data structures are created that contain variables that represent the states of the entities being modeled. The important point is that in simulation, the data in the computer is treated as a representation of something else, such as a population of mosquitos or trees. In instantiation, the data in the computer does not represent anything else. The data patterns in an instantiation are considered to be living forms in their own right and are not models of any natural life form.” (Ray, 1993)

Other literature has addressed these epistemological issues more directly. Famously, Webb attempted to characterise simulation models of biological behaviour in a seven-dimensional space (Webb, 2001). By assessing a model’s *relevance*, *level*, *generality*, *abstraction*, *structural accuracy*, *performance accuracy*, and *medium*, Webb argues that a principled approach to biological modelling can be applied to particular problems. From this, Webb discounts models that are based on idealisations and argues for realism, physicality, and strong validation on the assumption that biological behaviour needs to be studied and modelled in context. Vassie and Morlino (2012) take a different view. In this work, the authors propose a taxonomy that elucidates distinct relations between the natural and the artificial:

a *comparative* approach views artificial systems in the same light as natural systems and looks at the similarities and differences; a *modeling* approach uses artificial systems to learn about features of natural ones; and an *engineering* approach, that uses natural systems to draw inspiration for building artefacts.

For the present work, the value of simulation is twofold and is rooted in Vassie and Morlino's *comparative* class. First, simulations can allow us to demonstrate the emergence of intelligence and mind, through observations of behaviours in simulated environments, addressing one of the major open problems outlined above. Within their own closed worlds we can observe how parallels with nature play out, without making any claim as to their correspondence to natural processes, but still benefiting from understanding how the techniques can work to produce adaptive activity. Second, at the point where simulations and the real world become one entity we can observe evolutionary mechanisms in action, with complete access to their histories, and be confident that these process can also inform our understanding of the natural world, and general principles pertaining to evolutionary activity wherever it is found.

2.7 Computational Intelligence

I have now presented the AI ambition, the situated, embodied way of getting there, and the overlap with living systems research in a-life in terms of underlying philosophies and the grand challenges common to both endeavours. I have also justified the use of simulations in the production of knowledge. In this section, I discuss some widely-used technologies that when combined result in a system that can provide the context in which an enactive intelligence can occur. Two main threads are important, from the point of view of an enactive artificial intelligence. First, neural networks, through judicious statistical configuration, can represent arbitrary subsymbolic relationships between sense experiences and actions in the world. Second, evolutionary computation is an algorithmic method that simultaneously implements the statistical training and also ascertains which of these relationships are pertinent to intentional action by agents in that world.

2.7.1 Neural networks

Artificial Neural Networks (ANNs) are abstract computational structures that allow intrinsic behavioural relationships to form naturally between agents and environments. Haykin (1994) defines a neural network as, "a massively parallel distributed processor that has a natural propensity for storing experiential knowledge and making it available for use". These computational devices exhibit a number of desirable properties when used in a control domain: non-linearity, input-output mapping, adaptivity and fault tolerance (i.e. graceful degradation) all contribute to their utility as control architectures. The first research in this field, and perhaps the first work that can be recognised as AI, comprised a model of binary artificial neurons whose activation was dependent upon stimulation by a sufficient number of connected units (McCulloch and Pitts, 1943). McCulloch and Pitts argued for the general computational properties of these neurons, as well as their potential for experiencing learning processes (Russell and Norvig, 2003, p.16). Due to the other factors outlined earlier, as well as criticism from an influential text that underestimated the capabilities of such technology for real problems (Minsky and Papert, 1969), ANNs were not widely explored for AI applications for a number of years, although *connectionism*, an all-encompassing term for AI and cognitive science founded on these principles, remained a consistent alternative idea to the mainstream. It was only with the popularisation of ANN technology after the resurgence of the connectionist paradigm that followed Werbos's invention and Rumelhart's later popularisation of the gradient-descent based error-propagation method of training networks (Rumelhart et al., 1986; Werbos, 1974) that serious attention was paid to the potential of these devices. Their capacity to capture important statistical invariances within inputs, to act as control architectures for autonomous agents and to offer a model of biological neural control, combined with computational architectures able to rapidly optimise their performance secured the position of ANNs as a mainstream technology. In terms of robot control and cognition, they deal with automatically providing symbols–invariant features are learned–and if this learning takes place making reference to an external environment then the symbols are effectively grounded. Moreover, even complex networks allow near real-time updating from input values and as such they are good for robot control.

2.7.2 Evolutionary computation

What is the source of the delicate complexity ubiquitous in natural systems? Many explanations have been proposed to this long-standing question and our understanding of the issue has changed as scientific discoveries have been made. Today we can appeal to our knowledge of the intricate mechanics of the natural world when formulating our answer, but we are required to look back to the work of Charles Darwin and his contemporaries for the first complete and defining explanatory theory of how the awesome diversity of nature has arisen. Darwinism has taken centre stage as a unifying theory in Biology for more than a century. It is agreed, at least in the scientific community, that the process that we now know as evolution, the natural selection of well-adapted organisms through reproductive success and their modification in later generations through an imperfect copying process, is a sufficient explanation for the specialised and limitless diversity of forms found in nature.

This theory has become an a priori assumption for those seeking to explain the natural world; in 1973, evolutionary biologist Theodosius Dobzhansky wrote an essay entitled *Nothing in Biology Makes Sense Except in the Light of Evolution* (Dobzhansky, 1973). This phrase has become somewhat embedded in popular consciousness and stands as a qualitative proof of the explanatory power and intellectual reach of Darwin's hypothesis. Research is now focussed on examining the contributory environmental factors in the struggle for survival that explain our observations of the natural world in terms of evolutionary processes, as well as studying the parameters governing the (non-Darwinian) evolution of the dynamical system over time and the adaptation of the substrate of those processes to new ecological niches (called Evolutionary Dynamics).

These evolutionary ideas have exerted intellectual pressure on other areas of science, and it was only a matter of time before scientists and engineers started to look to nature for inspiration when searching for novel solutions to computational challenges. Certain problems, intractable from an analytical viewpoint, seemed ideal candidates for nature-inspired solutions where stochastic, numerical and once again constructive methods could be used to find acceptable solutions within a reasonable amount of time.

Perhaps the first person to think of evolution as an algorithmic process that could be harnessed and refined in a computational system was Alan Turing. In musing on how one might go about producing a machine to play the Imitation game (the parlour game that inspired the *Turing Test*), he imagined searching for a "child brain [containing] little mechanism and lots of blank sheets" using a trial-and-error approach, allowing the fittest to survive and guiding this man-made evolution with what little domain knowledge is available to the human designer (Turing, 1950).

The first uses of this algorithmic concept to solve real rather than hypothetical problems are due to the early works of Box (1957) who applied the concept to industrial productivity by discovering optimal manufacturing parameters, Friedberg (1958) who used an evolutionary method to search the space of computer programs, and Bremermann (1962), who considered the algorithmic time-complexity benefits of using an evolutionary heuristic to search large parameter spaces. Ultimately three main strands coalesced from the early experiments. *Genetic algorithms*, a general model of adaptive processes (Holland, 1975), *evolutionary programming*, a method to parameterise finite-state automata to achieve an artificial intelligence (Fogel, 1962), and *evolutionsstrategie* to solve general parameter-optimisation problems (Rechenberg, 1973).

All these methods are characterised as stochastic, hill-climbing population-based optimisation strategies. Since the inception of the general idea of evolutionary search the avenue that has shown most positive results and has by far the largest corpus of derivative work is optimisation, due probably to the relatively clear problem specification and evaluation functions found in this area as well as the flexibility and adaptability of evolutionary computation compared to other gradient-descent methods of parameter optimisation (Back et al., 1997). Research has now coalesced into the general field of Evolutionary Algorithms (EAs) whose root analogy is shared by all three initial approaches.

Their basic mode of operation is as follows. A set of strings that represent candidate solutions to a problem is maintained by the algorithm. Each individual solution in the population is given a score, its *fitness* (performance, rather than reproductive fitness), that allows comparison with the other individuals. Typically, an external objective function is used to provide this fitness metric. The algorithm proceeds by comparing individuals, preserving highly-fit individuals and generating new candidates to replace weaker solutions. The new solution can be created in many ways: from random data, from copies of existing individuals or from some kind of default template. In this way the space of solutions is explored and the algorithm improves its ‘best- effort’ guess at a solution as time goes by (Mitchell, 1998). This “black box” functionality makes implementation of this kind of optimisation process feasible for a wide class of problems, especially those where the shape of the solution space is not known. However, there are still constraints that make evolutionary algorithms unsuitable for some problems.

There are many thousands of examples of the EC heuristic being used for practical application since the early examples cited above, and countless more studies examining how the algorithms can be optimised for different problems and algorithmic forms.

As noted earlier, the smoothness of the hypersurface of fitness provided by the objective function as the population moves around in the function’s domain is a good indicator of the efficiency of the algorithm. As mutation events occur on the genotype during the reproduction process, individuals move around on the fitness surface. The mutation rate parameter is the main control for the trade-off between the breadth of search (and therefore the avoidance of local optima) and the fine-grained optimisation around a peak, and translates directly into an approximate distance moved on the fitness surface. A smooth surface is forgiving if mutation is too large–overshoots will be corrected in later generations–but as the number and size of the discontinuities increases, so does the likelihood of a mutation rendering an individual much less viable than its parent. A related problem is particularly apparent in the early stages of the algorithm’s progress through a complex fitness function, when large changes are needed in specific groups of parameters in order to move away from zero-fitness regions. In this case, many individuals will score zero fitness and selection will not make progress. This issue is called the *bootstrapping* problem and becomes more and more evident as the required complexity of any viable solution grows. Numerous methods have been proposed to potentially alleviate this issue, including determining heuristics that constitute sub-parts of the task and modifying the objective function to reward solutions with these attributes (Nolfi and Parisi, 1997). This naïve solution necessarily entails a more complex fitness function and is entirely dependent on the designer’s understanding of the problem domain and the appropriate heuristics, the overall effect being to increase the difficulty of designing the initial fitness function. Nolfi notes in later work (Nolfi, 1998) that any heuristic embedded into the fitness function restricts the available evolutionary trajectories and thus limits the potential of the search process to find optimal solutions.

Other less straightforward methodologies have been proposed to tackle the problem of bootstrapping and jagged fitness landscapes. Incremental evolution uses increasingly complex (more specific) fitness functions to encourage the evolutionary process to find proximal solutions at early stages. Once adequate performance at reduced-complexity tasks has been achieved, the environment is made more difficult (the fitness landscape becomes less smooth) and the process continues until an acceptable solution has been found. This technique has been shown to work very well in some scenarios (e.g. Gomez and Miikkulainen (1997)) but also shown to be less effective than naïve methods in others (Christensen and Dorigo, 2006).

In terms of understanding living systems from a synthetic, a-life perspective, evolutionary methods are fundamental. Evolution finds itself at the root of many first-principle definitions of biology's object of study, e.g. "Life should be defined by the possession of those properties which are needed to ensure evolution by natural selection", (Maynard Smith, 1986), and thus synthetic evolution is a prime target for research in isolation. In addition, evolutionary simulations can be used to address a variety of questions about the behaviour of evolution and the evolution of behaviours observed in the natural world (Todd, 1996), and also find uses, as described above, as a practical technology for exploring a complex, high-dimensional search space to solve engineering goals (Lenski, 2004).

2.7.3 Evolutionary design

Evolutionary Design (ED) is a subset of Evolutionary Computation, seen by some as comprising tools for design optimisation, creative design, evolutionary art, and evolutionary artificial lifeforms (Bentley, 1999). Whilst there is overlap and blurring between these categories, as well as with the wider field of EC, they serve to illustrate that this subfield of EAs is concerned with exploring a large, high-dimensional and highly heterogeneous search space using the evolutionary paradigm. Applications are many and varied although some examples are almost canonical in today's literature. One classic early study looked at the evolution of circuit-board designs to discriminate between 1kHz and 10kHz tones (Thompson, 1996). In this work, evolution was found to discover unusual designs that no human would create, and also to use components in unexpected ways (some designs showed components disconnected from the main circuit but still necessary to its function; the author speculates that the evolutionary process had discovered how to use localised electromagnetic effects or interaction through the power supply wiring to modulate the operation of the FPGA).

ED is not restricted to computer science applications. In Funes and Pollack (1998, 1999), an evolutionary methodology was used to design load-bearing structures. This system produces templates for building the structures from Lego blocks. Saul et al. (2010) applied an evolutionary algorithm to the design of paper-folding meshes, producing objects with various characteristics selected for by the process including height, stability and efficiency. *Interactive* evolution, where a human oracle acts as the objective function by judging aesthetically the results of the evolutionary process, has been used in artistic installations (Mignonneau and Sommerer, 2001; Sommerer and Mignonneau, 1999). Evolutionary approaches to design have also been explored as alternatives to a traditional toolchain in 3D computer graphics (Nishino et al., 2004), allowing users to interactively model 3D objects by using human aesthetic judgements as the objective function being optimised.

Recent applications of ED in an even wider context include a system to generate and explore choreographic sequences when planning dancers' motions through a space (Eisenmann et al., 2011), as well as the inverse task of training virtual dancers to 'move to the music' (Dubbin and Stanley, 2010). Both of these systems were also predicated on interactive evolution but they serve to demonstrate the wide scope of problems that can be addressed using such a methodology.

The use of interactive methodology in these complex design problems highlights a major difficulty in the practical application of evolutionary algorithms in this context: the difficulty of designing an objective function that effectively captures both the defining characteristics of a good solution and the relationship between the independent variables that provide a gradient in solution space toward this optimum. As argued by Zaera et al. in their paper discussing the difficulty of evolving collective behaviours in artificial organisms, "formulating an effective fitness evaluation function for use in evolving controllers can be at least as difficult as hand-crafting an effective controller design." (Zaera et al., 1996). This problem is found throughout the world of ED; it is a good hint at a problem we are going to run into later: creating fitness functions in order to generate interesting behaviours in robots or virtual creatures.

2.7.4 Neuroevolution

The application of evolutionary methods to the problem of optimising a neural network for a given task came swiftly after the resurgence of the connectionist paradigm that followed Werbos' invention and Rumelhart's later popularisation of the error-propagation (back-prop) method of training networks (Rumelhart et al., 1986; Werbos, 1974). Instead of adapting the weights of the network by deterministic gradient descent methods such as back-prop, the network's parameters are encoded in a genetic schema and a population of candidate networks subjected to evolutionary pressure based on fitness in an evaluation problem. The usual rules of EAs apply, but over time and with enough training examples to learn from, the population adapts to the problem and learns the generalisation in the same way. The properties of neural networks outlined above greatly simplify this process: in particular, systems that exhibit a graceful degradation in performance can tolerate small changes in their configuration, meaning that mutation is not catastrophic and thus the fitness landscape being explored is smooth.

There are many ways to represent the network structure in the EA. Many parameters can be under evolutionary control: weights, number of units, transfer functions and other tweaks. The most straightforward is to used a fixed architecture and just evolve the network's weights, an approach first put forward in Montana and Davis (1989). Montana and Davis used the EA as a direct replacement for back-prop as a way to find weights for a set of neural links. Their algorithm outperformed the standard method, although more recent research across a range of tasks and training schemes has highlighted that often the accuracy of the various EA methods is not significantly different from the accuracy reached by backpropagation alone (Cantú-Paz and Kamath, 2005). For complicated reinforcement learning tasks where supervised training is not possible, the application of gradient-descent training becomes difficult. Cantú-Paz and Kamath (2005) also argue for simplicity, showing that for many problems simple EAs perform as well as more nuanced algorithms. For this reason, evolved neural networks are often preferable in problems where reward / error signals are temporally remote.

2.7.5 Incremental evolution

Incremental evolution means two things. On the one hand, it means the gradual progression of an EA through a range of intemediary converged populations towards an overall solution. On the other hand, it can mean the gradual or stepwise changing of an objective function, usually in a direction of increasing problem complexity, in order to guide an evolving population towards an overall solution. In some cases it can mean both of these things at the same time, since they are complementary.

In terms of the second meaning, for the sake of clarity I call this *environmental complexification* (see the section *Categorisation of Complexification Techniques* below for further detail on this term.) Whilst many researchers have used this idea with varying degrees of success, no extant work has examined in detail the mechanisms that underpin this idea, nor have examined a variety of strategies for performing it.

Inman Harvey's *Species Adaptation Genetic Algorithm* (SAGA) paradigm, motivated by evolution in the natural world, set the stage for the computational use of incremental evolution by providing an evolutionary mechanism which allows an evolving species to maintain, at least theoretically, most if not all evolutionary pathways as potential candidates for exploration, no matter how converged the population has become to a single point in genotype space (Harvey, 1992a, 1997). Once a SAGA algorithm is implemented, objective functions can be changed and the population can be expected to adapt to its new circumstances by traversing *neutral networks* - pathways through genotype space that are defined as having equivalent phenotypic characteristics (usually, fitness) (Harvey, 1997, 2001). The requirements for the successful implementation of a SAGA-type incremental process are straightforward: inclusion of mutation as a genetic operator, smooth fitness landscapes and a redundant (high-dimensional) genotype to phenotype mapping which permits neutral networks to percolate through genotype space. In the SAGA literature the term *incremental evolution* is used in a sense that implies continued change, development or acquisition of domain knowledge by the algorithm over time as converged species roam around in genotype space. In addition, where there is a gradual change in the presentation of the objective function as the algorithm progresses, the term *environmental complexification*, as mentioned in (Mouret and Doncieux, 2009), is preferred. Finally, the label *incremental evolution* is also sometimes applied where intermediate solutions are moved to a new objective domain entirely; this case is also considered a flavour of environmental complexification, a point explained in more detail below.

Some of the earliest work which uses the environmental complexification approach directly is that of Gomez who (in addition to discrete, staged evolution over subtasks) gradually increased the speed of prey in a pursuit- avoidance simulation where neural networks were evolved to control simulated predators (Gomez and Miikkulainen, 1997). This work showed a very large performance gain by using the incremental approach. The work also identified an interesting adaptive approach where complexification is dependent upon agent performance at the current level of complexity. Mouret introduced a more general approach to rewarding sub-task performance than the hand-designed, staged approach common until this point (Mouret and Doncieux, 2009). Complex agent behaviour was evolved incrementally in a two-dimensional task in Robinson et al. (2007) where agents in a discrete world were evolved to navigate a hostile environment by avoiding and building bridges over increasingly challenging obstacles. Environmental complexification was used to evolve swarm robots in (Kadota et al., 2012), although the complexification chosen constituted arbitrary, discontinuous changes to the agents' environment and not a smooth transition over a range of difficulties. Notwithstanding, once again the incremental approach delivered a much higher rate of success in the given task (co-operatively foraging for food in a two- dimensional environment). Oh and Suk (2013) evolved controllers for unmanned aerial vehicles first using a non-incremental strategy. This strategy was found to perform badly as more constraints were added into the objective function so an incremental, task-subdivision strategy was used instead.

Categorisation of Complexification Techniques

Barlow identified two classes of complexifying training schemes: functional incremental evolution and environmental incremental evolution (Barlow et al., 2004). In this definition, functional approaches parameterise fitness functions to increase the apparent difficulty of tasks towards the desired level, whereas environmental approaches modify the environment around the evolving individuals *without* modifying the fitness function, with the same effect. Sub-categories of incremental evolution identified by Mouret in (Mouret and Doncieux, 2009) are *staged evolution*, *environmental complexification*, *fitness shaping* and *behavioural decomposition*. The most striking of these distinctions, common to both Barlow's and Mouret's work is environmental complexification; this category is of particular interest as semantically it can encompass all of the other categories identified and thus becomes synonymous with the sense of incremental evolution where the problem is simplified and made progressively more difficult. Additionally, environmental complexification is the only category that adequately encompasses co-evolutionary systems (which can be seen as auto-complexification), that in turn are the natural precursors to open-ended evolutionary systems, one of the grand challenges facing the field of a-life.

Incremental Learning in Neural Systems

The idea of incremental learning is not confined to evolutionary algorithms; neural network research has also considered this both as a problem (learning invariances is a dataset piecewise) and as a solution (tackling complex problems) for networks generally, outside of any particular training scheme. In the standard approach of using neural networks, training and application are distinct phases: all training data are presented to the network and the system learns the invariances and abstractions in that data using some learning algorithm. Then, this trained network is put to work on unseen data. This method of presentation can make it difficult for the network to adapt to new, unseen data at a later time and cause networks to suffer the phenomenon of *catastrophic forgetting* (McCloskey and Cohen, 1989), where training on new input cases causes previously learned knowledge to be lost from the network. In contrast to this, incremental learning algorithms are designed to allow the neural system to continually adapt to new information whilst maximising the information available in the network from previous training. This is an important concept for real world applications as often data is not available all at once and sometimes learning guides further exploration, meaning that learning is a continuous process rather than a discrete activity (Giraud-Carrier, 2000). One popular solution to the catastrophic interference problem found in these incremental learning schemes is to *rehearse* either already known data or pseudo-data representing the knowledge already in the network, interleaved with training on new information. See (French, 1999) for an overview and (Guajardo et al., 2010) for an example of recent work using this technique.

2.8 Open Problems

This chapter has presented an overview of the embodied, evolved approach to generating AI in terms of the history, scientific and technical hurdles and the links between these topics. The final section on Incremental Evolution has presented some specifics relating to this technology that are explored in a detailed experiment in chapter 4. The following section gives an overview of the two high-level, guiding objectives in a-life and AI that are the motivation for this work: demonstrating advanced behaviours in artificial agents, and providing a mechanism that opens up a path toward the demonstration of open-ended evolutionary processes. I then outline the aims of this thesis. In the next chapter, I review related work by other researchers that has implications for these ambitions and outline the work presented in chapters 5 and 6.

Demonstration of advanced behaviours

As noted above, the grand ambition of AI is to understand the nature of intelligent behaviour, and engineer artifacts that demonstrate it. I also reported that a grand ambition of a-life is to demonstrate the naturalistic emergence of such behaviours from increasingly impoverished starting points–meaning more homogeneous, less decomposed beginnings. This is the major aim and contribution of this thesis: using the principles of situation, embodiment, neural control and incremental evolution to achieve a level of behaviour hitherto unseen in systems specifically designed to automate the entire development process.

Open-ended evolution

One phenomenon that fits neatly into the category of closed computational structures that offer epistemic grounding equivalent to that of the natural world is open-ended evolution. This is another grand challenge in a-life; the much-debated idea that certain evolutionary systems have properties that allow indefinite evolutionary progress in some metric, the Earth's biosphere being the definitive example. Although (as we note in a recent paper) discussion continues in the OEE community about the hallmarks of the phenomenon and potential metrics to determine its presence (Taylor et al., 2016), like evolution more generally OEE is not dependent on observations of the material world, provided the substrate and metrics are sufficiently abstract. The earliest working definitions of the phenomenon were encapsulated in an *evolutionary activity* metric (Bedau and Packard, 1992). Originally, this measure classified the Earth's biosphere as open-ended, and all other extant artificial systems as not so. Channon's work demonstrated the first artifical system capable of passing Bedau and Packard's test (Channon, 2001), and then also passed a more difficult test, designed by Channon to normalise some statistical artifacts. However Channon's system does not offer intuitive confirmation of the OEE result through observation–the evolving entities are too removed from natural human experience to allow this.

Any meaningful metric must have this validation to offer more than an idiosyncractic definition of OEE so a major piece of the required analytical framework is the demonstration of experimental results that help to confirm the validity of the metric by offering a meaningful physical or pseudo-physical exposition of the evolutionary process alongside the statistical measure. The potential of systems that demonstrate advanced evolved behaviours to contribute here is large; behavioural changes in evolved agents are similarly understandable as those in the natural world so the intuitive validation can take place in relatively controlled circumstances.

2.9 Thesis Aims

2.9.1 Improving environmental complexification in incremental evolution

Incremental evolution has been widely applied in a variety of contexts in an effort to improve the efficiency and power of evolutionary algorithms. However, this application has generally been ad-hoc, intuitive and often an afterthought used to improve results on an existing system. The important considerations and effects of different incremental strategies remain unclear. Therefore, the first aim of this thesis is to examine how the different methods of applying incremental evolutionary pressure affect the outcome in a real task. By applying incremental evolution in a principled way, a theoretically-informed conclusion can be drawn about best practice in the use of this technology. The first work presented in this thesis in chapter 4 attempts this examination.

2.9.2 Simultaneous incremental neuroevolution of intelligence

The demonstration of intelligent behaviour is the long-term objective of AI; the demonstration of such behaviour emerging from relatively impoverished initial conditions is a grand aim of a-life. Observing an open-ended evolutionary process in human-scale artificial organisms is an important validation for theoretical research in evolution. Hence, the overall objective of this thesis is to use the principles of simultaneous incremental evolution, neural networks, and enactive cognition to contribute to the cutting edge of this research, by demonstrating a system that takes a step forward in these directions. The following chapter looks in detail at a wide spectrum of research relating to this ambition in order to inform a practicable approach towards achieving this next step in intelligent behaviour. (See chapter 3 section 3.3 for a more developed statement of aims for this part of the thesis.)

CHAPTER 3

Complex Behaviours in Embodied Artifacts

3.1 Evolution of Behaviours

The scope of possibility for the study and generation of animate activity widened with the introduction of new building blocks like classifier systems and neural networks for producing complex behaviour and, in evolutionary algorithms, a principled system for automatically discovering configurations that exhibit complex, emergent phenomena. The late 1980s and early 1990s saw an explosion in applications which combined one or both of these techniques with simulations of self-contained environments composed of specific primitives and rules governing the interaction of these primitives.

Applications and virtual substrates were diverse but all shared the same common objective: to observe organised activity on a systems level emerge from the low-level interactions. These organised behaviours ranged from simple, self-maintaining patterns in cellular automata through the co-ordinated control of legged robots (Brooks, 1989; Lewis et al., 1992), control of general robot behaviours (Beer and Gallagher, 1992; Floreano and Mondada, 1994), simulated insect behaviours (Beer et al., 1990), control of animat behaviours (Wilson, 1985) and emergent population dynamics (Packard, 1989).

These pioneering works, heavily reliant on computational capacity, were severely restricted in scope by the primitive calculating power available at the time. The explosion of computer power during the 1990s made even more ambitious a-life projects possible: one of the most well-known of these was Thomas Ray's Tierra platform.

Tierra is a complete computational ecosystem (Pichler, 2009). The Tierra universe is a simplified von Neumann computing architecture, simulated on a physical machine. Within this universe, populations of computer programs evolve to exploit the architecture based on the natural selection principle that differential reproduction is the only necessary driver of evolutionary progress. Ray saw the emergence of diverse ecological communities, and found organisms with distinctly natural characteristics: parasites, immunity, and hyper-parasites (organisms which co-opted parasitic organisms' reproductive facility for their own ends, driving the original parasites to extinction) (Ray, 1992). Ray's work was one of the first open-ended simulations where reproductive capacity is directly linked to the phenotypic expression of the organism's genotype. No external intervention occurs in the reproductive process: either the organism is born with the capacity to reproduce or it isn't, so after the first generation only those with the capacity to reproduce can reproduce.[1]

This is in contrast to most other work of the time, where evolution works only on a data structure in order to find set of optimal parameters unknown at the start of the simulation, and is in essence an optimisation process. Ray's system was undoubtably ground-breaking and significant even if only the diversity of emergent phenomena are considered: not only does the system exhibit interesting individual behaviours but also shows a general increase in complexity. Ray cites the *unrolled loop* as an example of this complexification in Tierra. Loop unrolling is a clever optimisation technique, invented independently of natural inspiration by humans and yet discovered by the chaotic evolutionary search within Tierra. The emergence of population-level dynamics where organisms discover methods of exploiting one-another as well as their static environment is another. Ray implies that this is suggestive of traditional evolutionary theory where biotic adaptation is the primary driving force behind the diversification of organisms. Tierra is not without criticism however; Channon (2001) argues that the apparent successes of Tierra arise from the simplicity of the underlying mechanics of the simulation environment rather than any specific adaptations on the part of the evolutionary process. Notwithstanding, the system shows a non-trivial emergence of composite behaviours from simple components which is a fundamental principle of descriptions of life-like systems, and has inspired a family of similar models, for example, Computer Zoo (Skipper, 1992), Avida (Adami and Brown, 1994), and Cosmos (Taylor and Hallam, 1997). It can be argued that just because the abstraction is not what was originally intended, there is no implication that the emergent phenomena are less valid as novel solutions to the problem of existence in the Tierra world.

Ray's Tierra was one branch of a three-way diversification in the lineage of large-scale life-like simulations. The second and third branches were pioneered by Larry Yaeger (1993) and Karl Sims (1994). In contrast to Tierra, Yaeger's *PolyWorld* system attempted to abstract at the ethological level rather than at the level of fundamental metabolics, in order to generate emergent phenomena at a higher level still, with the objective of exploring a-life systems as mechanisms

[1] Ray's initial problem of finding a bootstrap to reproductive success was solved by using hand-designed *ancestor* organisms imbued with this capability as an initial population.

for making progress towards artificial intelligence (Yaeger, 1993). This aim was facilitated chiefly by PolyWorld's biologically-inspired fundamental building blocks (i.e. neural processing elements), a naturalistic, 2.5D environment which provides sensory grounding for the organisms in a way that previous systems had not, following the argument made by Cliff in his work simulating the neural visual system of a fly (Cliff, 1991), and an abstraction of base behaviours from the agents' worlds–a calculated step away from autonomous embodiment.

Yaeger's system was inspired by the Braitenberg Vehicles mentioned earlier (Braitenberg, 1984), where emergent behaviours were imagined in two-wheeled vehicles with extremely simple and arbitrary control structures connecting sensors to effectors. Yaeger's extension (aside from creating a real simulation rather than thought experiments) was to substitute Braitenberg's control structure with a neural-network mediated system with a supporting palette of intrinsic interactions based on intuitions about real-world ethological classes (these are the fundamental building blocks) and to incorporate the adaptive mechanism of learning (within the neural network) alongside natural selection. A rudimentary vision system was also incorporated that supplied the controller with data about the organisms' environments. The neural system in PolyWorld is a straightforward multi-layer perceptron customised with evolvable neuronal clustering and topological distortion (in order to facilitate development of the kind of retinatopic maps found in natural organisms). Learning occurs according to the Hebbian principle (Hebb, 1949) where synaptic efficacy is moderated according to the previous activation state of the connected units and the evolution of the learning rate (η) parameter: "neurons wire together if they fire together" (Lowel and Singer, 1992).

Yaeger's results were intriguing; he found that various species evolved successful behaviours which permitted their populations to reproduce and live in the simulation indefinitely but even more interestingly, individual complex emergent behaviours like adapting speed in the presence of certain visual stimuli, grazing (slowing in the presence of food) and following other organisms. These behaviours certainly achieve Yaeger's original aim of finding complex emergent behaviours from a simple suite of primitives, and are further examples of the crossover between real-world emergent phenomena and those of simulated complex systems.

Although Yaeger's groundbreaking approach yielded impressive results, they were marred slightly by the sheer complexity of the underlying simulation. It is difficult to estimate with any accuracy the appropriate attribution of responsibility for high-level emergent behaviours to any particular underlying component of the system, be that the neural architecture, the extensions to this architecture invented by Yaeger, the Hebbian learning process or simply the range of primitive behaviours available to the organisms. It is certainly possible that equivalent behaviours could be found by a simple feed-forward network in an otherwise identical environment, and it is unclear whether any new behaviours would arise beyond those relatively simple emergent activities described above using any control model (see Channon and Damper (1998b) for a full discussion). Yaeger's

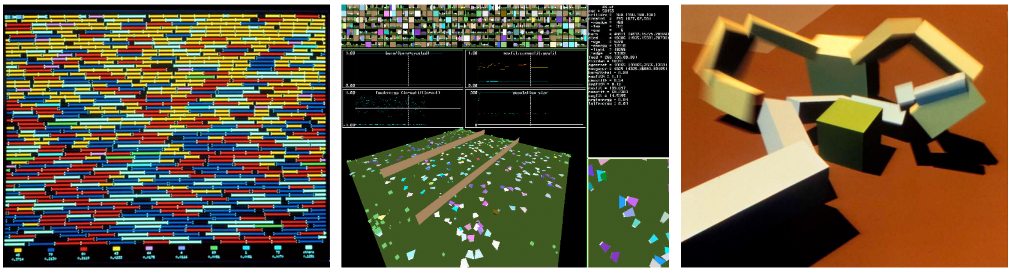

Figure 3.1: Tierra (Ray), Polyworld (Yaeger) and Blockies (Sims)

attempt to ground his organisms within their sensorimotor worlds is admirable, but the simplicity of the underlying model (even with the potential for reflexive internal complexity) may impose a ceiling on emergent ethological complexity due simply to the lack of available space to explore.

Karl Sims' work (1994) is perhaps the best known of the a-life simulations of the early 1990s. Sims created a virtual universe of rigid, 3D structures and constraints (joints) between them and evolved the specification both for the bodies and the neuron-like network structures that controlled forces acting on the constraints (Sims, 1994a). With a focus more on engineering than science, Sims' work is recognised as groundbreaking primarily because of the step-change in complexity and realism introduced with a fully 3D environment. It is the opinion of this author that this work provides the most life-like, engaging and intuitively accessible demonstration of evolved virtual organisms, principally because of the natural affinity the human observer has with the world of the evolving organisms. Sims' demonstrated organisms composed of multiple blocks connected by powered hinges, evolved using artificial evolution to optimise performance at an arbitrary task (locomotion, block-following and competition for resource).

To some extent, it is possible to classify each of the three preceding systems in terms of their scale analogues in the biosphere. Ray's Tierra (as an extension of Langton's work on computational biochemistries) can be seen as focusing on intra-cellular dynamics; the processes which allow single-cell organisms to survive and reproduce. PolyWorld also investigates the single-cell, except abstracting the mechanics of reproduction to the level of interaction between organisms and the environment. Sim's work on block creatures then represents multi-cellular organisms which have to learn to control their own bodies (inter-individual interaction is somewhat an afterthought). Each level of detail reveals higher-level abstractions of life-like processes in these simulations but also each fundamentally solves the same problem: the marshalling of available base resources in the environment in order to interact successfully in that environment and participate in the implicit struggle for for survival.

As noted earlier, results from PolyWorld showing complex behaviours are scant. However, PolyWorld has still served as a useful model for life-like processes. Research using or based on the PolyWorld platform has tended to focus on investigating specific aspects of evolutionary dynamics or neuroevolution, rather then follow the original aims of evolving complex behaviours in order to approach artificial intelligence. In Yaeger and Sporns (2006) the evolution of neurocontrollers was analysed in the PolyWorld context to discover general trends of evolving neural systems. Griffith used PolyWorld as a simplified ecological model in order to investigate the evolution of agents showing an optimal, ideal free distribution of agents to resources (Griffith and Yaeger, 2006). More recently, Yaeger has used the system to investigate the evolution of complexity and passive versus driven trends in the complexification of evolving agents (Yaeger, 2009; Yaeger et al., 2008). It is clear that PolyWorld is a useful tool for simulating evolutionary activity but has not shown recent promise in the search for artificial intelligence.

Another system which bears some superficial similarity to PolyWorld is Geb (Channon, 2001; Channon and Damper, 1998b). Although conceived independently, Geb has a similar level of abstraction in that a palette of behaviour is available to a neural network that expresses these actions in by controlling an agent in a two-dimensional world. The original aim of the Geb system was partly parallel to that of PolyWorld–to build a simulator with the flexibility to produce intelligent behaviour through an evolutionary process. Channon's early work showed success in this area, developing unexpected and interesting behaviours (Channon and Damper, 1998a). Although a spartan model compared to PolyWorld, the main focus of Geb is elsewhere, in the search for the first evidence of open-ended evolutionary systems in artifacts (Channon and Damper, 1998b).

Common to all of these systems is the problem of the complexity ceiling, where highly complex behaviours are not accessible due to either the underlying representation, the evolutionary dynamics or the time available to run the process. As such, research is driven down a narrower track of investigation of specific questions relating to evolutionary processes or autonomous behaviour, and often loses sight of the original objective–the evolution of advanced, general-purpose adaptive behaviours.

3.1.1 Virtual environments for intelligent behaviour

As the range of work in the previous section attests, simulations of living systems have covered a broad spectrum of abstraction but typically aim to exhibit behaviours at the level above that specified in the model's design. When building animat simulations that focus on interactions recognisable at the human scale, such as moving around, fighting and environmental manipulation, one of the key distinctions between designs is the physicality in which agents operate, specifically the choice between 2D and 3D environments. A two-dimensional world such as *PolyWorld* or *Geb* abstracts simulations away from the natural physical domain. Agents in these flat environments generally do not have to solve any complex physical control problems (Channon and Damper, 1998a), as controllers are able simply to signal directions in which to move or turn the agent. Such models can encourage early emergence of more complex composite behaviours but preclude the development of novel motor control which may later allow for a richer interaction between agents and their environments, in the spirit of pure situated, embodied intelligence. Two-dimensional simulations have not tended toward clearly *displaying* the impressive physical interaction observed in nature, whether or not complex (simulated) non-physical interactions have evolved. This can be attributed, at least in part, to the fundamental rigidity and paucity of physical actions in such environments and the lack of situatedness and embodiment that leads to a diminishment of environmental grounding.

The development of an even more low-level system such as *Tierra*, *Avida*, and *Cosmos*, whilst allowing evolution free reign over its own operation on a microevolutive level (also known as the evolution of evolution, or EvoEvo), are even further abstracted from the intuitive naturalistic domain. Results in these worlds may be in a sense situated and embodied, but are not recognisable analogues at common scales and so specialised analysis is required to understand the activity of evolution in these contexts.

By contrast, having three-dimensional articulated bodies in a 3D world provides for much greater intricacy in how agents can interact with their environment and each other. Agents must begin to construct a coordinated motor pattern that results in basic directional motion before richer behaviours can develop as composites of these lower-level patterns. The specific characteristics of the environment are implicitly included in the performance of these motor patterns, and this couples agents to their environment. This coupling is crucial, together with the coupling of brain and body, to two key principals of embodied cognition: “first that cognition depends upon the kinds of experience that come from having a body with various sensorimotor capacities, and second, that these individual sensorimotor capacities are themselves embedded in a more encompassing biological, psychological and cultural context” (Varela et al., 1991). Recent trends reinforce this point of view, highlighting the importance of morphology and soft materials in the embodied loop (Pfeifer et al., 2014), and the effect that this can have on the evolutionary process itself (Corucci et al., 2016).

It is clear that, in terms of the ongoing ambition to evolve advanced life-like behaviour, both 2D and 3D approaches have been fruitful. For example, using 2D non-articulated agent bodies, early work by Yaeger showed (in a 3D environment) the emergence of complex collective behaviour (Yaeger, 1993); Channon demonstrated the first candidate synthetic open-ended evolutionary system using an agent-based (2D) world (Channon and Damper, 1998a); and Robinson et al. (2007) evolved agents capable of reactive and deliberative behaviours in novel and dynamic environments. However, it is in embodied, 3D agents that we begin to see behaviour that is truly recognisable as intelligent and life-like.

3.2 The Evolution of Embodied Intelligence

In thinking about how to achieve Artificial Intelligence in its modern, nature-inspired incarnation we have travelled through the philosophical idea of embodiment, the connection that this has to all living processes, and the practical methods of capturing and expressing a nonlinear relationship between brain, body and world using evolution and connectionist networks. We have seen that intelligent behaviour is an emergent product of a number of components interacting in a complex dynamic and have catalogued some of the important actors in the show. We have observed that in all cases assumptions are made about where to base our abstractions, the possible way we can connect the components together, and mechanisms we can employ to do this so as to move towards recognisably intelligent behaviour. By designing environments, evolutionary algorithms, control architectures and agents with close attention paid to the underlying philosophical principles of embodied cognition, previous research has made substantive leaps forward in the search for artificial general intelligence technology and an understanding of how cognition has developed in nature. The previous section highlighted some successes from the broader field of research in this area, but when the above concerns are attended to closely, we can outline a category of artifacts that clearly stand apart in terms of their relations to real-world intelligent behaviour in natural organisms. Embodied machines with non-linear control architectures, embedded in 3D space analogous to our own, that are sculpted by a wandering evolutionary process have delivered examples of intelligent autonomous behaviour that are truly breathtaking. These devices, whether constructed in the real world or merely reified in their own simulations, demonstrate viscerally animate, autonomous activity in a fashion unmatched by other technology to date.

Of course, problems still attend in such systems. For example, in general the evolutionary process itself is not embodied, in the sense of Watson's *embodied evolution* (Watson et al., 1999) or the previously mentioned work of Channon. In the taxonomy of Schut et al. (2009) which characterised evolutionary systems according to their "embeddedness" in time and space, these systems are of type 4 because the capacity of individuals to reproduce is governed by their spatial location (other examples given include evolving agent societies and on-line robot evolution–ibid.) In contrast, most work in the evolution of embodied intelligence is of type 1; the EA lurks in an abstract space receiving fitness assessments of individuals and remotely orchestrating selection and reproduction with a kind of omnipotent grace.

These constraints do not affect the overall motivation–the evolution of embodied intelligence–only the modes in which it may be carried out. In the short term, we need to understand those conditions that are necessary and sufficient to achieve advanced evolved behaviour. This may entail advanced control architectures, powerful evolutionary algorithms, or complex co-evolutionary environments. It may simply mean a judicious alignment of much more modest components.

The following section looks at work in the area of evolved embodied intelligent behaviour and catalogues the approaches and findings.

3.2.1 Evolutionary robotics

> There are obviously strong ties between behavior-oriented AI and robotics, because the construction of physical agents is seen as a condition sine qua non for applying the method of the artificial properly. But the two fields should not be equated. The goal of robotics is to identify, design, and engineer the most reliable and most cost-effective solution for a sensorimotor task in a particular, usually fixed and known, environment. Behavior-oriented AI uses the tools of roboticists to study biological issues, but very different criteria for success apply. (Steels, 1993)

The broad principle of Evolutionary Robotics (ER) is to use evolution to produce useful values for a parameterised robotics system, be that it's body (morphology), brain (controller), or both (Floreano and Keller, 2010), in order to design machines capable of performing autonomous behaviours. This definition has been expanded to include bioinspired robotics that extends beyond evolutionary design methodologies (Pfeifer et al., 2005). In the words of Bongard (2013), "The long-term goal of evolutionary robotics is to create general, robot-generating algorithms." Practically however, the approach is "useful both for investigating the design space of robotic applications and for testing scientific hypotheses of biological mechanisms and processes" (Eiben, 2014). In general ER is focused on building robotic machines in the real world, as opposed to a purely simulated environment. Early ER research used simulations to accelerate the evolutionary process–time can run much faster in a simulation–but found problems when trying to transfer the species evolved *in silico* to real-world analogues. A body of work was undertaken to solve this problem, including methods of minimal simulation (Jakobi, 1998) that tested species both in simulation and physically, in order to maintain evolutionary trajectories successful in both environments. Macinnes and Di Paolo (2004) explored evolving the morphologies and controllers in simulation for real robots, using Jakobi's ideas of minimal simulation, with some success. Various other versions of this crossover have been tried; Pollack and Lipson (2000) and Lipson and Pollack (2006) demonstrated real robots (morphologies and controllers) that were designed by evolution in simulation by using a rapid prototyping process.

Mahdavi and Bentley (2003) evolved controllers for real robots notionally based on a snake's morphology. They found that although evolution in real robots was difficult, there were advantages in terms of adapting to unexpected constraints or changes in the environmental parameters.

Lund (2003) evolved controllers and morphologies for Lego robots, both automatically and with a user-guided design approach, advocating the approach on the embodiment argument. This work used a *Hox*-gene inspired encoding to lay out the robot's body and controller plans.

Matsushita et al. (2006) investigated pseudo-passive dynamic walking bipeds in simulation. It was found that when evolving fine morphology and controller, pseudo-passive dynamic walkers showed more dynamic stability than actively-controlled examples. The authors argue that this is due to the compliant components in the model functioning as noise filters and passive oscillators.

Samuelsen et al. used a Hox-gene inspired developmental approach, using a two-level description linked to two different axes of physical development to evolve simulated robots with potential for transference into a physical machine (Samuelsen et al., 2013). They then investigated different distance measures in order to reward diversity in populations of these robots, finding that the lengths, branching factor and longest depth resulted in species most adapted to the distance-metric objective function in use (Samuelsen and Glette, 2014).

Bongard (2008) succeeding in producing a robot capable of multiple, chained behaviours by a scaffolding approach that carefully presented the two tasks. This was one of the earliest pieces of work to attempt a chain of behaviours. This work was continued in later work that demonstrated the importance of the order in which the various tasks were presented to the evolving populations (Auerbach and Bongard, 2009).

Bongard (2010) used an evolutionary approach to design the morphology of a robot arm, investigating how the utility of this approach varied according to task complexity. By demonstrating a relationship between task complexity and the utility of an evolutionary design mechanism, Bongard provides yet further evidence that an embodied AI paradigm is important for generalised intelligent behaviour.

This idea is continued in (Auerbach and Bongard, 2010) where a novel representation scheme, *CPPN-NEAT* (based on Stanley's Neuroevolution of Augmenting Topologies (NEAT) algorithm (Stanley and Miikkulainen, 2002) and Compositional Pattern-Producing Networks) is used to evolve controllers and morphologies at multiple spatial scales; the proposed system worked as a proof of concept for this approach.

Glette and Hovin (2010) examined the physics simulator *PhysX* as a platform for evolving artificial muscle-based robotic locomotion, both in simulation and with transfer to real-world machines. The system produced stable locomotion when adapting the cloth feature of PhysX to act as muscle tissue, but the researchers noted that care must be taken when attending to simulator parameters to avoid instability.

Faíña et al. (2014) used a simulation of modular robots to evolve machines which could then be built in the real world. The specific architecture lends itself to rapid physical construction and allows a large number of varied morphologies with only a few different types of module.

Groß et al. (2011) imagined an ambitious system of real-world fundamental evolution, where basic building blocks are able to come together in an agitated medium to produce living non-biological physical organisms. This is perhaps the closest that evolutionary robotics has come to the life-as-it-could-be paradigm of a-life.

Nakamura et al. (2011) used an EA to optimise control parameters for the NN control of a simulated flat fish in an underwater environment.

Jared Moore's 2013 work uses and EA to optimise values in a parameterised sinusoidal controller for a bipedal hopping model (Moore et al., 2013). In Moore and Clark (2014), he designed a system to produce controller for an underwater robot using NEAT. Moore researched a novel method of actuating virtual creature joints (Moore and McKinley, 2014a), and then investigated how these neuromuscular connections affect the performance of quadruped agents, finding that for more complex agents a joint-control coupling was as effective as a simulated musculature (Moore and McKinley, 2014b). The same research group have also worked on optimising the parameters for the caudal fin of a robotic fish, finding through rapid fabrication of the evolved structures that it is possible to search the design space for successful parameters in simulation, and then transfer to a physical machine without losing key aspects of the solution's performance (Clark et al., 2012).

Nogueira et al. (2013) applied a modulatory neural network to the control of a Khepera-type robot in simulation, demonstrating a capacity to evolve a foraging behaviour using simple neuroevolution. The authors comment on the lack of complexity, and the difficulty of increasing the behavioural output of the system.

Very recently, Central Pattern Generators (CPGs) have been used as components of motion primitives in a robot locomotion problem, in order to allow the robot to switch between locomotive modes to cover uneven or unexpected terrain. These primitives are combined by a higher level planner to achieve the robot's goal (Vonasek et al., 2015).

While there have been many novel results and important successes in the study of ER, some researchers still criticise the field from the point of view of lack of generalisation: most ER demonstrations solve toy problems or exist in contrived environments (Nelson, 2014). Notwithstanding, it is clear from this survey that ER has contributed to fundamental technologies that allow us to construct intelligent, autonomous agents in the real world and in simulation. It has also advanced the sciences of intelligence and behaviour, acting as a platform for constructive (synthetic) investigations of hypotheses about how these characteristics of the natural world developed, and has validated the philosophy of situated AI described in the previous chapter. However, these examples show that the focus of ER research has mostly been on either basic locomotor control behaviours, or to investigate the effect that natural principles exert when using artificial evolution to build robot controllers. Only in a few cases have complex and varied behaviours been the focus, and even then the scaffolding required to achieve the results meant that the outcome was more due to design decisions by the research team than the automatic emergence of such behaviours through evolution within the platform.

From the point of view of the ambitions of this thesis, ER research has shown that technologies like evolved neural networks and the undercurrent of enactive AI that they embody are a fruitful line of inquiry when aiming to build general-purpose intelligence artifacts. In addition, Bongard's work has shown that an incremental approach can be used to generate increasing complex behaviours in an evolved AI system. Furthermore, the capacity of an evolutionary robotics system to generate unbounded diversity is only a small conceptual step from current frameworks, with the potential to produce machines with highly heterogeneous behaviours, many of which could be found in nature as convergent evolutionary answers to the common problems of existence in an uncertain world.

3.2.2 3D virtual creatures

As noted above, results of evolution that are generated from a pure simulation are often difficult to transplant into physical reality: imperfect sensorimotor signals and enormously variable environments can expose the limitations of the assumptions necessary to build artificial analogues. However, the requirement for evolved robots to exist in the real world can be relaxed and this opens a swathe of opportunity for exploration of evolved behaviours not possible in physical machines. These wholly artificial *evolved virtual creatures* (EVCs) are born as pure simulations, entirely shaped by the assumptions built into their worlds by their designers. This research direction has occurred in parallel to real-world evolutionary robotics, but its initial ambitions were more diverse. Castro and Gudwin describe virtual creatures succinctly in work that sets out to use virtual creatures to examine the role of episodic memory in decision-making animats:

> "An artificial creature is a special kind of autonomous agent, which is particularly embodied in a given environment (there may be autonomous agents which are not embodied). A virtual creature, for its turn, is a special kind of artificial creature, where the environment is a virtual environment, so the creature's body is not a concrete one, like in a robot, but just an avatar in a virtual environment."(de Castro and Gudwin, 2010)

In fact, virtual creature ecosystems have a long history at the interface between science, technology and art. Some of most forward-thinking work in this area had a focus on educational outreach, creativity and imagination as well as driving forward technological development.

At the turn of the 21st century, some of the best known examples of digital biota from the a-life research programme were Tom Ray's Tierra, Karl Sims' 3D virtual creatures and Biota's Nerve Garden (Jensen, 1999). As noted earlier, Ray's work is certainly the most pure from an a-life perspective in that its only assumption is the computational medium where the coding for reproduction can take place, but Ray himself sees EVCs as sitting on the same conceptual continuum of evolutionary processes that have, to varying degrees, been liberated from a pure objective optimisation (Ray, 1997). In Nerve Garden the objective was to synthesise an immersive, networked virtual environment using nature-inspired principles in which people could interact and explore the world (Damer et al., 1998, 2003). Ray makes the point that evolved virtual organisms have an intrinsic beauty and can affect human observers in the same way as natural organisms: “the creatures may appear beautiful, elegant, sensuous, nervous, bizarre, strange.” (Ray, 2001). The intuitive correspondence between these artificial organisms and our experience of the natural world is made clear by Hayles:

> “when we attribute to Sims’s virtual creatures motives and intentions, we interpolate their behaviors into narratives in which events are causally related to one another and beings respond to their environments in purposeful ways.” (Hayles, 1999).

One of the earliest and most prevalent examples of EVC applications in the literature is as a procedural animation technique, either to achieve realistic low-level behaviours automatically or to provide high-level behaviour and the ability for entities to adapt to and anticipate situations dynamically in unknown environments (Duthen et al., 2010). In these systems, only the appearance of autonomous behaviour is important; in general any means to this end is acceptable.

Perhaps the first such example to be widely published is that of Miller (1988), whose models of legless figures (snakes and worms) were for the time an extremely biologically plausible generation of the motions of these animals.

Contemporary to Sims’ 1994 work which is looked at in more detail in a later section, Terzopoulos produced seminal work on the simulation of fishes. This graphically-focused but biologically-rooted system had hand-crafted and yet complex behaviours, and the situated nature of the system led to a much increased level of graphical realism (Terzopoulos, 1999; Terzopoulos et al., 1997; Terzopoulos and Rabie, 1995; Terzopoulos et al., 1994), and also began to achieve a degree of automation in the search for realistic behaviours (Grzeszczuk and Terzopoulos, 1995).

McKenna and Zeltzer (1990) used the simulation of legged locomotion to demonstrate the utility of a fast dynamics simulation. Similarly, Raibert and Hodgins (1991) used a dynamics simulator to produce locomotion in animats with various numbers of legs. Nikovski (1995) combined an explicit plan-based scheduling with an evolutionary algorithm to find parameters for legged locomotion in a hexapod robot, in simulation. In these cases, although the authors had a goal of animation, the techniques employed were beginning to touch on the situated, embodied approach since this had the potential to produce the most naturalistic appearance. Indeed, Thalmann et al. (1996) discussed how a complex ‘virtual life’ could be created in simulation, mentioning the problems of perception and action that arise when trying to automate interactions in such worlds. This philosophy continued with other groups who wished to automate to varying degrees the interactions in virtual environments to achieve a realistic and yet interestingly novel environment, although the a-life philosophy is explicitly criticised as not having the aesthetic potential to satisfy the goals of the VL endeavour (Wang and Mckenzie, 1998).

This approach has continued up to the present, employing increasingly sophisticated simulation and graphical techniques to incrementally increase the realism of the animations. Liu et al. (1997) used an evolutionary strategy to search for parameters to achieve predefined gaits in bipedal and hexapod creatures. Heleno and dos Santos (1998) modelled a river ecosystem using hand-crafted controllers for various organisms residing within it. Eccles et al. (2000) engineered a graphical ecosystem by pre-evolution of control parameters for constituent organisms (*bugs* and *fish*) and then a subsequent stage of manual integration into the ecosystem. The bugs are stuck on the 2D lake surface; only the fish has 3D motion capability and this is limited to abstract behavioural primitives of move-up, move-down, etc. However, the creatures in this world were controlled by neural networks allowing an emergent complexity to unfold when the fully-integrated ecosystem simulation is run.

Many researchers in this area also found EA-like optimisation techniques (often combined with a structured search space based on domain knowledge) to be useful to produce realistic animations. Buendia and Heudin (2000) used artificial evolution to produce behaviour in video game characters. Their approach was founded on the principle of the subsumption architecture (Brooks, 1991), evolving useful values for parameterising an hierarchical behaviour space. Boeing (2008) used an EA combined with a spline-based representation of skeleton animation to produce working models which could then be tuned by animators. In Wampler and Popović (2009), preset morphologies and gaits were adapted using a novel optimisation technique to produce realistic gaits in a variety of 3D animal forms. Furukawa et al. (2010) presented a computationally efficient air-drag simulation and evolved flapping creatures within it to demonstrate its realism, citing the difficulty of solving Navier-Stokes for the wing parameters as the reason to use the evolutionary method. Tan et al. (2011) used Covariance Matrix Adaptation (CMA) to optimise control parameters for simulated fish locomotion and Nakamura (2015) used evolved neural networks to control a fish. Allen and Faloutsos (2009) generated high-quality bipedal locomotion using evolved neural networks created with the NEAT framework, in a realistic physical simulation. This work included haptic sensors in each foot that detected contact with the ground and bilateral symmetry in the controller by using two copies of the evolved controller, one for each leg. Most recently, Geijtenbeek et al. produced extremely realistic locomotive behaviours for graphical animation using a flexible, muscle-based approach combined with CMA (Geijtenbeek et al., 2013).

As the potential of larger virtual worlds began to be realised, some research also took the EVCs in a more generalised direction, either to harmonise research in the field, engage more of the general populace in the idea of virtual environments, or to explore very large-scale simulations. Prophet (2001) constructed a large, open world and invited users to create creatures in the environment that would go on to interact and reproduce. Miranda et al. (2001) built a generalised platform for virtual creatures research called *Arena/WoxBox*, with the aim of using the system for a variety of research; behaviour modelling, examination of learning algorithms, social behaviour, and population dynamics amongst others.

The work discussed above focused on graphics, and in many cases achieved impressive results. EVCs have been an inspiration across science; in the Black Shoals arts project, artificial evolution was employed to design articulated creatures which interact with a world of real time financial data. Their phenotypes are composed of multiple interacting elements in a discrete time simulation of Newtonian physics, and a key element of the artwork was to create a world of artificial 'speculators' able to 'feed' on the shifting real-time patterns of the world's stock markets. (Hoile, 2014). However, in any of these examples it can not be claimed that these artifacts are exhibiting the properties of intelligent, living organisms any more than a marionette. The similarities are at a surface level only; grounding for these systems is provided by clever decompositions on the part of the designers. To generate life-like behaviour more autonomously, we must refer back to the grand ambitions of a-life, that is, to present intelligent behaviour that has emerged naturally from increasingly impoverished starting conditions. The philosophies of enactive AI–chiefly situation and embodiment–are also key: for true realism we must build systems that automatically ground their own symbols and and whose behavioural interactions occur as a result of a natural development between agents and their environments. The system developed by Karl Sims remains the exemplar for this approach to this day.

3.2.2.1 Karl Sims's Blockies

In 1994 Sims published the seminal SIGgraph paper in which he presented his first forays into artificial evolution (Sims, 1994b). As mentioned throughout this thesis so far, this and the related co-evolutionary work has had a resounding impact on research across the a-life, computer graphics and artificial evolution programmes. In these papers and the accompanying video, Sims describes how, using a bespoke physics engine based on Featherstone's reduced co-ordinate approach to rigid-body physics simulation (Featherstone, 1987), virtual creatures with arbitrary body shapes and behaviours are evolved. Evolution takes place against a variety of objective challenges - walking, swimming, phototaxis and (in later work) competition between two creatures for a resource. Sims's evolutionary mechanism follows a developmental model, using a recursive graph-generating grammar to encode body shapes and controller configuration, inspired by the graph grammar presented in Kitano (1990). Maintaining a population of 300 individuals, Sims's algorithm follows the normal procedure for implementing an evolutionary process. Virtual creatures are born in the world according to the genetic specification, and their reproductive success is determined by their performance as measured against the external task. A percentage of creatures

Figure 3.2: Megan Daalder imitates the block creatures of Karl Sims in her 2010 performance piece, “Tribute To Karl Sims”.

survive to reproduce (around 1 in 5) and as time goes by, the species gradually improves in quality. Sims ran the experiments for 100 generations, with a creature lifetime of 10 seconds. From Sims’s many runs of the system, a large variety of diverse body plans and behaviours were observed, although in some cases a degree of interactive evolution was present as well. Nonetheless, Sims’s work remains to this day an example of the visceral realism that a-life simulations can invoke. Sims’s blockies have inspired not only a-life practitioners and computer graphics researchers, but have also been used to illustrate the power of evolution as a general, universal principle and as artistic pieces that juxtapose the natural and the artificial, a noteworthy case being the performance art piece by the artist Megan Daalder (figure 3.2) whose physical imitation of Sims’s creatures reverses the arrow of artificiality and imbues the creatures with a realism that escapes their digital confines.

Sims’s virtual creatures left a lasting legacy and remain an inspiration for a-life researchers today. The goal of generating ever more interesting behaviours in virtual creatures is a worthy ambition for graphics and robotics in itself but, as with evolutionary robotics, the objective also bears directly on two of the grand challenges of a-life that were outlined earlier: the observation of complex intelligent behaviours emerging from systems with increasingly impoverished starting conditions, and the demonstration of ongoing evolutionary activity in an artificial system. In the following section I examine work that is, in one way or another, at least in part directly deriviate of Sims’s initial ideas. The first part looks at systems designed to exhibit only simple behaviours–these systems often explore one aspect of the generation of behaviour in EVCs. In the second part I survey systems whose objective is the production of more complex, composite behaviours. These systems are more aligned with the wider objectives of a-life research.

3.2.3 Agents that exhibit a single behaviour

In fact, nearly all recent research is focused on the evolution of a single gross behaviour in virtual creatures–usually locomotion. This is for a variety of reasons; often, locomotion is an acceptable objective against which various hypotheses about sufficiency of conditions or efficiency of algorithms can be tested. As documented below, much research has been carried out to understand evolutionary systems operating in these contexts, as well as how best to improve our application of the technologies involved.

Fontijne (2000) implemented Featherstone's rigid body physics solver and attempted to follow Sims's strategy of evolving virtual creatures for various tasks to test the simulator. Whilst enjoying success with some tasks, those requiring feedback or oscillation (including locomotion) were out of reach for this system, potentially due to short run-times.

In (Komosiński and Rotaru-Varga, 2001), an investigation into the benefits of using indirect encodings when evolving morphology and control of EVCs in the Framsticks (Komosinski, 2000; Komosinski and Ulatowski, 1998, 1999) environment. It was found that indirect (recurrent and developmental) encodings improve the performance of the evolutionary search by restricting the search space, although a consequence was a clear difference in the appearance of the behaviours between the indirect methods. This research was continued in a system that aimed to present gene-phene mappings intuitively, also using EVCs as the underlying evolutionary objective (Komosinski and Ulatowski, 2004). See also Komosinski (2003) for an accessible overview of work with this platform. Stredwick (2005) designed a system of virtual creatures in 3D based on the SodaRace-type 2D systems (McOwan and Burton, 2005) to explore the maintenance of convergence and stagnation in populations evolving morphology and behaviour. This study noted the difficulties involved in the analysis of behaviours resulting from the evolutionary process.

Hornby and Pollack (2001) also examined encodings, using a parametric Lindenmeyer system (P0L-system) to encode the morphology of virtual creatures with simple oscillating joints and a distance covered metric for the objective function. This work found that the indirect encoding produced creatures comprising a larger number of parts and possessed of a higher degree of regularity. This work has also been extended to incorporate a neurocontroller rather than a simple oscillator circuit (Hornby et al., 2001), and subsequently to a more generic generative encoding (Hornby and Pollack, 2002).

Taylor and Massey (2001) reimplemented Sims's original work using off-the-shelf physics components.

Teo and Abbass (2002) investigated the use of a Pareto-frontier Differential Evolution algorithm to parameterise a neural network in order to control quadruped locomotion whilst minimising the size of the resultant network. The same system is then extended to compare various optimisation methods (Teo and Abbass, 2004).

Shim et al. produced a body of work based, unusually, on simulated *flying* creatures. Initially the objective was simple locomotion through the air (Shim and Kim, 2003), but in this work the authors discovered some of the problems inherent in evolving multiple behaviours simultaneously: learning to turn was incompatible with learning stable flight. The neural model of this work was not generic; it comprised numerous functional (sine, cosine, etc.) units rather than just a generalised McCulloch-Pitts model. The authors extended this system to a two-step, incremental evolutionary system to solve the problem of turning (Shim et al., 2004a,b) and finally analyse their model's validity in terms of the morphologies produced for different body masses of creature compared to real-world examples (Shim and Kim, 2006).

Druhan (2004) used a tree-structured genetic program to represent a construction recipe for virtual creatures, evolving a population of such programs in an EA with distance moved as the objective.

Marks et al. (2006) apply a hybrid CPG approach to generating gaits in creatures with varying numbers of limbs, although encountered difficulties in evolving a realistic gait. They concluded that fine-tuning the parameters of the EA used to explore the behaviour space (a modified Evolution Strategies approach) entailed as much effort as a hand-designed controller, and proposed several routes out of the problem including increasing the complexity of the controlling networks, restricting the fitness function yet further and applying a form of online "human-in-the-loop" evolution to guide and support the process. Ohono et al. (2007) used an evolutionary technique to optimise values for a parameterised sinusoidal controller operating on the intersegmental joints of an evolved 3D agent with cuboid body parts and presented results showing that by combining these segments, evolved creatures were able to accomplish tasks otherwise impossible to achieve. Komosinski and Polak (2009) evolved morphology and neural control architectures of simulated 3D ski-jumping agents.

Chaumont et al. (2007) presented a system that used artificial evolution in the *Breve* environment (Klein, 2003) to find morphology and control parameters for virtual creatures of two types, walkers and block-throwers (catapults), demonstrating the potential for comparatively simple architectures to produce impressive results. This system used simple distance-based objective measures of fitness. Chaumont et al. found difficulty in producing oscillating motion:

> "It appears to be easier for the genetic algorithm to evolve single-impulse controllers-giving rise to jerks-than controllers producing periodic repetitive motions." (ibid.)

Krčah (2007) used off-the-shelf physics components to recreate Sims's original work, achieving walking and swimming behaviours. In subsequent work, Stanley's NEAT algorithm is used and the evolutionary search time with this mechanism is compared to a standard EA. The results show that in all cases the evolutionary time is significantly reduced, whilst the quality of the solutions found is improved (Krčah, 2008).

Lassabe et al. (2007) used the evolution of virtual creatures' morphologies and control architectures to explore the use of classifier systems as controllers, rather than the (now more traditional) neural-type systems, and to increase the level of complexity of EVC environments. In this case, the classifiers manipulate a corpus of motor patterns in order to produce complex joint activations in the virtual creatures. The rich environment of these creatures was a key part of this work; its complexity allowed the authors to explore the performance in many tasks with real-world analogues (i.e. not just walking, but skating, pushing objects and so on) although in each case only a single behaviour was optimised.

Jones et al. (2008) explored the importance of bilateral symmetry in an eel-like virtual creature, concluding that this kind of body plan is important for directed locomotion and energy minimisation and that a correspondingly symmetrical control architecture is preferred by an evolutionary optimisation process when constrained by this kind of morphology. Jones's Ph.D. thesis looked more generally into how body plan affects neural development, again using virtual creatures as the platform (Jones, 2010).

Mazzapioda et al. (2009) investigated the use of a generative process that involved development modulated by regulatory substances (known as *artificial ontogeny*) in an EA to produce virtual creatures able to move around in a 3D environment. The authors noted that this method has a high replicative stability across experiments, and also that exposure to a more complex environment produced agents better adapted to simple environments as well. The work also investigated the use of task-independent fitness measures and concluded that these measures help to increase the robustness of the evolved solutions, if not the absolute performance.

Heinen and Osório (2009) looked at the evolution of gait controllers in a quadruped, specifically, a model of a dog. This work compared finite-state machine controllers with neural networks, and it was found that the evolved neurocontroller performed better than the FSM model. The research also looked at the evolution of morphology as well as control, and adapted the model to bring the limb segment lengths under evolutionary control. It was found that morphology evolution achieved significantly higher fitness, but other factors may have been at play (the authors did not control for absolute size, for example) that confounded the results.

Valsalam (2010) developed a 3D, fixed-morphology legged robot system to explore how symmetry can play a role in the development of controllers for such a robot and concluded that symmetry is an extremely important factor. In this work, robots evolved with a symmetrical bias showed naturalistic and fluid motion in contrast to the asymmetric controllers which were less effective, did not scale and seemed 'crippled'.

Pilat and Jacob (2008b, 2010) evolved light-following behaviours in virtual creatures, training a neural network using an incremental evolutionary approach. This two step method added rudimentary vision at a later stage in the evolutionary process, after the species had found a stable locomotive behaviour. Through this process the authors were able to train the controllers to produce phototaxis even under modestly varying environmental conditions.

Bainbridge (2010) examined whether digital control networks could be used as an alternative to the more common floating-point type controllers in virtual creatures, finding that such controllers offered similar performance to the traditional type in varied EVC problems - pole balancing and locomotive control of evolved morphologies.

Hiller and Lipson (2010) evolved multi-material freeform shapes for locomotion, instead of the more usual rigid-body creatures. These animats were actuated by periodic volumetric expansion and contraction and the results showed some success in achieving locomotion.

The soft body approach is taken further in Rieffel et al. (2014) where a virtual creature platform is used to demonstrate methods of co-discovering soft-robot morphology and control, including muscle placement on a fixed body shape and material properties, parameterised material properties, and developmental encodings. In all cases, the maximum behavioural complexity sought is at the level of locomotion.

Lehman and Stanley (2011) created a virtual creatures platform, again with locomotion as the gross objective, to examine novel methods to generate diversity within evolving populations. In this work, evolution is provided with a novelty objective that encourages diverse morphologies and a local competition objective that rewards the winners in local groupings of similar individuals. It was found that this method discovers more functional morphological diversity than models with global competition, and thus it is claimed the technique constitutes a principled approach to combining novelty search with pressure to achieve.

Azarbadegan et al. (2011) attempted to recreate Sims's work with the added constraint of understanding how environmental pressures can cause a bipedal gait. The work hypothesises that pressure to carry and move at the same time promotes bipedalism and the authors contrived a fitness function to restrict the design space to solutions exhibiting such characteristics.

Krčah (2012) examined the novelty search concept in more detail in an evolutionary simulation designed to produce animats with the capability either to move toward a goal position around an interposed object (a so-called *deceptive* task at which greedy optimisers tend to fail) or to move around in an underwater-type unconstrained environment. These agents used the NEAT algorithm, and the research found that novelty search was only of benefit where the task was clearly deceptive.

Bongard (2011) examined how different types of *scaffolding* interact to assist the evolutionary process in progressing toward a complex solution. In this work, legged robots are evolved to perform phototaxis and a synergy is demonstrated between morphological and environmental scaffolding, provided they are combined in a particular way (morphological, followed by environmental). This thinking was further explored in Auerbach and Bongard (2012) (again using EVCs), where the relationship between environmental complexity and morphological complexity was tested and a correlation discovered. Most recently, the same platform has been used to investigate passive versus driven trends in the development of complexity using a formal metric of morphological complexity and the authors conclude that different niches exert different selective pressures toward complex body plans (Auerbach and Bongard, 2014).

Rada and Aguilar (2012) used a particle-swarm optimisation (PSO) algorithm instead of an EA to find body plans and controllers for virtual creatures in order to achieve a high-speed locomotion behaviour.

Kou and Kawaguchi (2012) used an evolved sensor-actuator style neural network (Van de Panne and Fiume, 1993) to control each leg of a quadruped and concluded that this is a viable method of generating gaits. Similar techniques have been used to find controllers for other simulated movement behaviours: Ooe et al. (2013) used an evolved ANN to control a flying creature, finding that the diversity of input signals about target points increased the generalisation capability of the evolved controllers.

Ouannes et al. (2012) evolved virtual creatures with neural networks to produce animats able to move around and search for food in their environments, as a first step toward the construction of a virtual ecosystem.

In a similar vein, Pilat et al. (2012) evolved foraging agents able to move towards the closest source of food in their environments. This work claims to be an effective demonstration of the utility of physical simulation environments for studying biological phenomena, although this is a preliminary study and is essentially a proof-of-concept.

Ito et al. (2013) use populations of predator and prey EVCs to examine the relationship between the evolution of morphology and the evolution of behaviour, in order to find out which precedes the other. It was found in this model that morphology tends to precede behaviour when new strategies emerge independently, and behaviour precedes morphology when responding to the innovations of opponents.

Cheney et al. (2013) applied Stanley's novel Compositional Pattern-Producing Networks (CPPNs) (Stanley, 2007) as a generative encoding for soft-bodied creatures in order to achieve locomotion. These creatures are composed of voxels with various properties, including actuation according to a sinusoidal input and structural attributes like material stiffness. They found that the CPPN approach produced faster and more natural gaits in the soft robots compared to direct encodings.

Tibermacine and Djedi (2014) used Stanley's NEAT algorithm to generate locomotive controllers in 3D virtual creatures with a selection of pre-defined morphologies. This was found to be a successful strategy for locomotion generation in these agents.

Krčah (2014) used a similarity metric to implement speciation in evolving populations of virtual creatures, claiming an increase in efficiency compared to older methods although other differences between systems do make this comparison harder to use to draw firm conclusions.

Lessin et al. (2014b) demonstrated that much of the control architecture required for single-task objectives (locomotion, jumping and so on) can be moved from the traditional NN-like controller to the musculature of the creatures' bodies, relieving the monolithic controller of the difficulty of coordinating various parts of the body to achieve a task. Lessin's argument is that this paves the way for more complex processing to occur in the central controller, with physical tasks delegated to the dynamics of the creature's physical form.

Ito et al. (2014) co-evolved predators and prey in a 3D VC simulation designed to examine the interaction between population dynamics and trait evolution. This work found cyclical dynamics in the short term corresponding to the Lotka-Volterra type interactions. In the long term, correlations were found between the population sizes and volume of the prey that led to the emergence of defensive and non-defensive morphological adaptations.

Moreno et al. (2015) supports the results found in Stanton and Channon (2013). A complex environment was decomposed into simpler sub-environments and evolving species trained on these evironments, sequentially. This was compared to a 'whole task at once' approach and it was found that the sequential approach is more reliable and consistent.

O'Kelly and Hsiao (2004) successfully evolved fighting creatures in a similar manner to Sims's co-evolutionary research, leveraging the fitness gradient provided by adaptive opponents to drive the search. Creatures with a variety of strategies were found, but again these creatures were limited to a single behaviour governed by the morphology, even if that behaviour performed multiple functions simultaneously (typically, moving, attack and defence behaviours).

The work of Thomas Miconi from 2005 to 2008 focused on using evolved 3D virtual creatures to understand various aspects of evolutionary dynamics. Miconi's first published work in this area re-created the work of Sims but with an aim of well-adapted general behaviour rather than realistic appearance. To this end, Miconi replicated the locomotion results from Sims's work by using general McCulloch-Pitts neurons in the creatures' controllers rather than Sims's tailored signal-generating units, and also avoided the human-in-the-loop elements of selection present in Sim's demonstration (Miconi and Channon, 2005). This system was subsequently improved but it was noted that a difficulty arises when requiring oscillation to occur due to the sensorimotor loop rather then endogeneously in the controller (Miconi and Channon, 2006b). In later work, Miconi implemented a co-evolutionary system (Miconi and Channon, 2006a), presenting a new analysis for co-evolutionary activity based on *Master Tournament matrices* (Floreano and Nolfi, 1997; Nolfi and Floreano, 1998). He then went on to analyse the co-evolutionary dynamics occuring in a relatively large, open world in detail, concluding that the approach showed promise for generating a more diverse, if not open-ended, evolutionary dynamic (Miconi, 2008a). This work also included the first co-evolutionary *fighting creatures* (Miconi, 2008b).

As in O'Kelly's earlier work, even though Miconi's later EVCs solve multiple problems in their environments due to the implied nature of fitness in the system, it is not clear that evolution's solution has produced distinct behaviours in any single organism. Rather, the animats adopt a general behaviour pattern that works equally well for locomotion or fighting–there is no modal change between behaviours due to environmental stimuli.

Similarly in very recent work, Arita et al. (2016) used eco-evo-devo approaches to evolve predator-prey virtual creatures.

In terms of virtual creatures *in themselves* however (as opposed to virtual creatures whose function is to provide a methodologically-agnostic problem space in which one can test EA or dynamic control methods), it is less obvious why behavioural development is so limited. Taylor identifies the issue early on in the history of EVCs, from a perspective of using evolution to create interesting behaviour in artificial characters: "the complexity of behaviours that have been generated so far has been somewhat limited; most of the work has produced characters that can walk, crawl, swim, jump, or, at best, follow a moving target." (Taylor, 2000). This is attributed in this work to the need to codify the goal of such systems in terms of an objective function to be optimised, and the fact that complicated objectives frustrate the search when starting from random initial locations in the search space. Taylor suggests several avenues out of this problem, including behavioural/task level decomposition of the problem ("there are no general guidelines to suggest the most appropriate way to do this"– ibid.), co-evolution, virtual ecologies, lifetime learning, behavioural primitives, and user-guided evolution. Furthermore, there is a more practical issue in that the computational overhead required to achieve more generalised evolutionary dynamics in a 3D, physical simulation is still prohibitive.

3.2.4 Agents with multiple behaviours

To this day there are notably few exceptions to the general approach in EVCs of producing machines that perform only a single behaviour. Yonekura and Kawaguchi (2008) demonstrate one such system, where neural feedback through bodily dynamics allows a dynamical controller to switch between limit cycles and therefore exhibit qualitatively different behaviours. Although limited in range and complexity, this work does show at least in principle that this is possible in EVCs.

Pilat and Jacob (2008a) present a general system for performing a-life experiments with EVCs, *Creature Academy*. This work specifically addresses the problem of evolving creatures in order to tackle a combined walking–jumping behaviour environment by using a two-tiered training regime. Although this involves two separate behaviours in the early stages of evolution, the behaviours are fused in the final part of training as the individual component gains are consolidated.

Spector et al. (2007) presented an ambitious project that involved an energetically-conserved 3D simulation of interacting, partially connected and actuated, neurally-controlled blocks, called *Division Blocks*. Spector describes this work as "an attempt to extend Sims's idea to a considerably more open-ended evolutionary and ecological context", in order to investigate interactions between development, form and behaviour. Fitness in this world is implicit; good energy harvesters survive longer and good reproducers, reproduce more. The research is interesting visually, and certainly is the most 'open-ended' work to date in the EVC programme but the complexity and lack of observable structure in the resulting morphologies and behaviours is testament once again to the need to carefully construct such a system from the ground-up rather than throw many components together into a milieu and observe the results.

Some of Dan Lessin's recent research focus, culminating in his Ph.D. thesis, is one of the few efforts that tackles the narrow behavioural repertoire of EVCs directly. In this work, Lessin's *encapsulation*, *syllabus*, and *pandemonium* (ESP) system uses a human-designed syllabus to decompose complex behaviours into a sequence of smaller learning tasks. These skills are encapsulated (protected somewhat from evolutionary change) and finally a mechanism is used to resolved disputes between competing skills or drives within the brain (Lessin et al., 2013). This system was extended to include full morphological adaptation beyond the first skill learned, increasing the variety and quality of evolved creature results (Lessin et al., 2014a).

Rossi and Eiben (2014) compared two different strategies for the evolutionary learning of multiple tasks in simulated robots, using a $(\mu + \lambda)$ *Evolution Strategies* algorithm, where $\mu = 1$ and $\lambda = 1$. The work assessed whether learning one task at a time (building over previously trained individuals) is a better strategy than trying to learn all the tasks at once. It was found that the incremental strategy (known in this literature as the *Robot School*) reduces the learning times and shows less variance in the quality of the resulting gaits.

3.3 Chapter Summary

ER and EVCs both contribute to the science and technology that together allow us to construct intelligent, autonomous agents. In both cases, the endeavours have acted as platforms for constructive (synthetic) investigations of hypotheses about how the defining characteristics of natural behaviour have developed, and in both cases have validated the philosophy of situated AI described in the previous chapter. Whilst differing sometimes in focus–ER's pull toward real-world robotics puts emphasis on practicability and engineering, whereas EVCs tend toward the hypothetical–the underlying problems and solutions overlap to a large extent.

Techniques like neural networks and incremental evolution have been shown to be practical solutions in both fields, supporting the more theoretical arguments around enactive AI and the subsymbolic new robotics approaches championed by Brooks. Furthermore, the intuitive, subjective understanding the observer experiences when observing artificial organisms in physically-realistic environments lends itself to the notion that validation of OEE theory could naturally be undertaken in these ‘life-as-it-could-be’ domains. The capacity of these systems to generate unbounded diversity is only a small conceptual step from current frameworks. There is potential to produce machines with highly heterogeneous behaviours, many of which could be found in nature as convergent evolutionary answers to the common problems of existence in an uncertain world.

It is, however, also clear from the research surveyed that despite this potential, little research has actually made progress toward complex behaviours in artificial organisms. Most published work has used ER or EVCs to demonstrate other principles–those of evolutionary efficiency or compositional heterogeneity for example–avoiding the larger goal of automatically generating complex behaviours from increasingly impoverished initial conditions.

Therefore, the further objective of this thesis, (to build upon that described in section 2.9.2) is to demonstrate a sequence of behaviours evolving in a physically-realistic, 3D virtual creature environment. This aim is explored in chapters 5 and 6, making use of the techniques and technologies discussed so far and remaining theoretically aligned with the principles of situated, embodied intelligence.

CHAPTER 4

Guiding Incremental Evolution

"Nothing is a waste of time if you use the experience wisely."
–Auguste Rodin

The research documented in chapter four can be summarised with the question, "Within limited computational constraints, is there a specific strategy of environmental complexification that maximises the performance of an incremental evolutionary system whose objective is to produce a controller for 3D virtual creatures able to solve a general task?". General principles relating to the problem in 3D VCs and wider incremental evolutionary systems are proposed and investigated empirically in a controlled task. The novel result that a strategy that takes into account the need to revisit and progressively increase levels of complexity is presented. This chapter is based on work previously published in the proceedings of the 12th European Conference on Artificial Life (ECAL), Taormina, Sicily.

4.1 Introduction to Chapter

This dissertation has the overarching aim of producing evolved virtual creatures that display complex adaptive behaviour across a range of different tasks, and across a range of variation within a single task. As such, it is important to build evolutionary systems that find generalised behaviours, rather than those that succeed in only a specific or narrow range of parameters.

The same principle is writ large in biological systems: there are many examples in nature of problems that, whilst congruent in general terms, differ in their parameterisation in specific instances. Differences in morphology due to growth, differences in strength due to available food and differences in the abiotic environment due to locality are examples of this principle that are common to many cognitive entities in the natural world. Evolution's approximation of this fuzziness is embedded in the structures and processes of all species, whose constituent organisms are able to operate effectively in a many-dimensional cloud of uncertainty.

Generalisation across a relatively sparse set of training examples is also important even in relatively simplistic evolved 3D VC models. In the simplest case, a single task whose parameters in specific instances are variable, the degree to which agents can accommodate the variability is a function of the degree to which they are exposed to the whole range of task complexity. Presenting all combinations of parameters during each individual evaluation and incorporating the agent's performance measured in each trial into the agent's overall fitness is the most straightforward method to address this problem. However, as the number of parameters and thus combinations increases, computational constraints quickly render this approach infeasible. Thus, the target of this work is to find evolutionary approaches in which each agent is evaluated on a small subset of parameters, in this chapter on a single value for a single behaviour, and yet result in agents able to perform over the full range of parameters by having evolved generalised behaviours.

One approach to tackling this problem is to make use of incremental evolution (see section 2.7.5). Incremental evolution has been used extensively to improve the quality of evolutionary search in many complex, non-linear problem spaces. The work in this chapter first disambiguates the lexicon around incremental evolution, advocating the term *environmental complexification* to represent the progressive complexification of the problem domain by incremental exposure to a range of component complexity. Then, this complexification is divided into two types of strategy, homogeneous and heterogeneous, and instances of each type identified. The strengths and weaknesses of strategies are objectively compared. To summarise, in homogeneous complexification strategies, for any short sequence of successive generations the population is exposed to a single or tightly clustered range of objective functions, while heterogeneous strategies present many, covering a range of complexity.

The specific example explored in this chapter is the problem of finding a controller for a fixed-morphology 3D VC simulated with rigid-body physics, that allows it to climb over arbitrarily-tall obstacles rather than just those of a specific (maximal or other) height.

4.2 Complexification Strategies

As noted in section 2.7.5, previous work has successfully leveraged the power of incremental evolution through successive increases in environmental complexity. However, attention in extant work has been focused solely on the practical outcome of using this incremental modification to the canonical strategy as compared to a non-incremental system, and the particular strategies used to complexify the environment have not been examined in detail. A careful investigation of the complexification approach is necessary both for the reliable practical application of incremental evolution and the further elucidation of the interplay of agent and environment in co-evolutionary settings which can ultimately lead to unbounded evolutionary activity.

The naive strategy presents one single task (often, the most difficult or complex) to the evolving species at every opportunity. The failings of this approach spurred the development of alternative, progressive strategies, such as monotontically changing tasks as time passes, often in a direction of increasing difficulty as defined by the experimenter. This *linear* approach has often been used to circumvent the bootstrapping problem, one of the first attempts occurring in (Gomez and Miikkulainen, 1997). The appeal of this approach is strengthened by the simplicity of its implementation and the broadness of its potential application. Many task decomposition strategies can also be considered as implementations of this strategy (with linear or at least monotonic increase in complexity), albeit in discrete units rather than tweaking a continuous variable. Gomez also proposed an extension to the linear increase in task complexity where the task is only changed when the evolving species achieves a certain level of performance against the current objective function. This interesting strategy has not been developed in detail by others but it is a good candidate for analysis as it enforces gradient at every point of decomposition, potentially solving some or all of the issues described in the introduction to this chapter.

Although not often described in previous work, random presentation of different task components may also be useful and finally, drawing upon ideas from incremental learning in neural systems, a strategy of repeated presentation of tasks to which the evolving species has already been exposed is proposed. These strategies may have something to offer beyond linear or adaptive monotonic changes in task complexity.

4.3 Hypotheses

It is anticipated that homogeneous complexification strategies, for example direct presentation of difficult tasks or linearly-increased complexity, will perform poorly due to either loss-of-gradient or temporally-local over-fitting (analogous to *catastrophic forgetting* in neural systems), as argued in section 2.7.5. Heterogeneous strategies are the approach proposed in this chapter to overcoming forgetting, as an analogue of rehearsal, with smoothly changing heterogeneous strategies, such as oscillatory strategies, also overcoming the loss-of-gradient problem. For oscillatory strategies, an amplitude parameter determines the extent to which the whole range of task complexity is exposed to the evolving species. A gradual increase of this range may be expected to show improved performance. At very low frequencies, such a strategy would degenerate to the homogeneous linear strategy, and at very high frequencies to the random strategy. Thus, the following hypotheses are investigated:

H1: Homogeneous strategies will fail to achieve good coverage on the evaluation task.

H2: Heterogeneous strategies (with the possible exception of random) will achieve better coverage than homogeneous strategies.

H3: Heterogeneous strategies with a range of difficulties increasing over time will outperform heterogeneous strategies with constant range (other hyperparameters remaining constant).

H4: A heterogeneous strategy using an oscillatory approach, as an analogue of *rehearsal*, will exhibit an optimal frequency for any particular problem.

4.4 Methods

The general setup of our experiment is designed to test the above hypotheses in a task which provides a smooth fitness landscape and neutrality in genotype space. In this work, the evolution of controllers for three-dimensional agents is chosen as the base platform, tasked with learning how to walk and climb over an obstacle. The height of the obstacle represents the complexification parameter of the system; task difficulty varies somewhat as obstacle height varies but the ultimate objective for the agents is to deal with every possible obstacle - this is the most complex case. Thus, the many possible complexification strategies (that is presentation of tasks of various difficulties) can be assessed, in order to determine which provide the strongest gradient for the evolutionary system to climb and the most robust evolved agents.

4.4.1 Physical model

In the tradition founded by Sims (1994b) and continued by many others, all experiments are performed on agents in a three-dimensional virtual world consisting of collidable rigid bodies connected by powered constraints. Unlike Sims, the morphology used here is a fixed quadruped which is controlled by a feed-forward three-layer perceptron augmented by sinusoidal input. The cuboid quadruped torso (length 0.2m) is supported by four limbs, each comprising an upper and lower portion (length 0.075m each). Constraints with two degrees of freedom limit the motion of torso and upper limb at the hips; constraints with one degree of freedom limit the motions of lower limb and upper limb at the knee. See figure 4.1 for a visual representation. The range of motion of each knee joint is limited from 0 to $\frac{\pi}{2}$ radians, so knees cannot bend outward. The maximum torque that can be applied at any constraint is 0.5 N m. The obstacle is situated 1m from the agent's origin and extends to infinity in x and for 0.05m in y. The height of the obstacle is varied as described elsewhere. The physical simulator used was *Open Dynamics Engine* (ODE) version 0.12, using double-precision arithmetic, the standard big-matrix step function and a step size of 0.02s. Coulomb friction

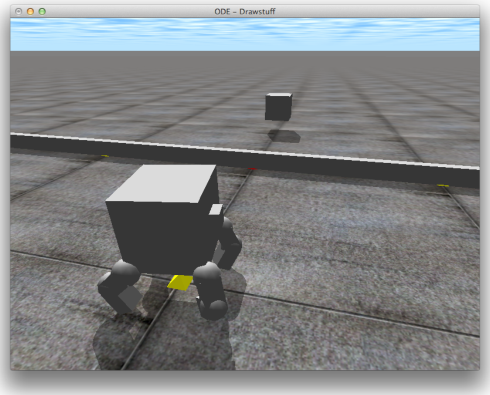

Figure 4.1: Visualisation of physical environment. Agent, obstacle and target location are shown.

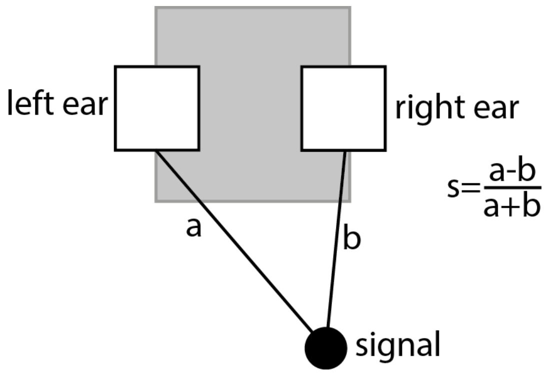

Figure 4.2: Target sensor scheme.

was applied at contacts between the agent, the obstacle and the ground plane with $\mu = 2$.

4.4.2 Control system

The agent controller is modelled by a standard three-layer feed-forward neural network with 12 hidden nodes. Networks receive 4 real-valued inputs in addition to 12 joint-angle sensors. Inputs comprise two sinusoidal oscillators (sine and cosine, period 1 second), an input describing the target location in relation to agent position and orientation (as in Reil and Husbands (2002), the difference between the distances from target to each “ear”, divided by the distance between the ears across the agent’s torso) and an up-sensor which describes the orientation of the agent’s head relative to the ground plane. The target sensor scheme is shown in figure 4.2.

Neural network updates are made synchronously with physics integration. Each hidden node activation is a weighted sum of its inputs with a hyperbolic

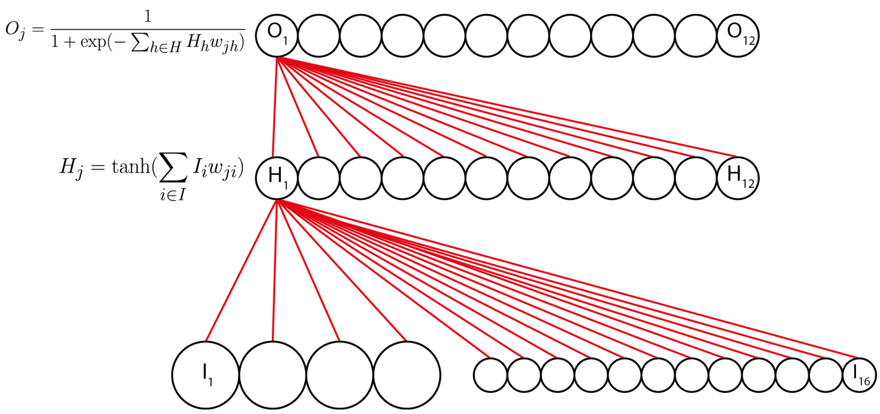

Figure 4.3: Feed-forward neural network showing transfer functions at each layer (hyperbolic tangent in the hidden layer; logistic in the output layer). The four larger input nodes indicate the two sinusoidal inputs, the target sensor and the orientation sensor; the twelve smaller, the joint-angle sensors. Not all weights are shown.

tangent activation. Each output node activation is a weighted sum of hidden nodes with a logistic activation function (see figure 4.3 for an illustration). The output node's value propagates to the joint motors of the agent through a proportional-derivative (PD) controller, in the same manner as described in Reil and Husbands (2002). The network's output for a joint is treated as a target angle in the range of action for that joint. The PD controller applies a torque to the joint relative to the correction required to achieve the desired position, also applying a damping factor proportional to the angular velocity of the joint in the previous timestep:

$$T = k_s(\theta_d - \theta) - k_d\dot{\theta}$$

where T is the torque applied, k_s is the spring constant, k_d is the damping constant, θ_d is the target angle and θ is the current angle. Preliminary experiments were conducted to ascertain appropriate spring and damper values. The agents body was positioned in a standing pose and the PD controller used to maintain that pose in opposition to gravity. The values $k_s = 0.5$ and $k_d = 0.5$ were found to be acceptable, preventing the agent from collapsing but also allowing the pose to stabilise. As noted in Reil's work, the PD approach applies simple control at the mechanical level, relieving the controller of basic problems such as standing up without the agent's legs buckling under the mass of the torso.

4.4.3 Evolutionary algorithm

Individual genotypes specify floating-point weights for the neural control system. Initial values for the first generation are drawn from a uniform distribution $x \in [-1, 1]$. In each run, the evolutionary simulation is progressed for 5000 generations using a population of 50 individuals. Individuals are evaluated for 20 simulated seconds and the objective function is defined as the reduction in distance in the x-y plane to a target position (x_*, y_*) situated on the other side of the obstacle:

$$F(x,y) = -\sqrt{(x_* - x)^2 + (y_* - y)^2}$$

At each new generation, individuals are scored according to the objective function and ranked in order of fitness. The lower half of the population is replaced with mutated, crossed-over variants of the upper half. Single-point crossover is implemented at a random point on the genotype and crosses the current parent individual with another random individual from the best half of the population (possibly itself). It is necessary to consider the *competing conventions* problem, also known as the *permutation problem*, when implementing recombination operators in neuroevolution. This problem occurs when two networks compute the same function by differing mechanisms, meaning that crossover of their components risks incompatibility. The problem was first described in Montana and Davis (1989), and adequately addressing it was one of the motivations of Stanley's NEAT architecture (Stanley and Miikkulainen, 2002). It has been found in recent doctoral work that for realistic problems, competing conventions are extremely rare. Further, crossover operators do not significantly impede the progress of evolutionary search of weight space, at worst acting as macromutation operations rather than recombination mechanisms (Haflidason, 2010). Mutation occurs on average twice per genotype and consists of adding a value drawn from a Gaussian distribution with $\sigma = 1$ and $\mu = 0$.

4.4.4 Experimental set-up

Sixteen possible strategies for environmental complexification have been identified and tested. Each of these strategies modifies the height of the obstacle in the environment for the current generation of the species. In every case the maximum height of the obstacle, tau, is 0.1m. The height function h for generation G and wavelength λ is defined for each strategy as follows:

1. Direct presentation of an environment with complexity τ at every generation (Strategy 1):
$$h(G) = \tau$$
2. Presentation of a randomly complex environment at each generation, with complexity drawn from a uniform distribution between 0 and τ (Strategy 2):
$$h(G) = x\tau, x \in \mathcal{U}(0,1)^{\mathbb{R}}$$
3. Gradual complexification of the environment, with complexity interpolated linearly between 0 and τ from generation 0 to generation 4000 and fixed at τ from generation 4001 to 5000 (Strategy 7):
$$h(G) = \begin{cases} \frac{\tau G}{4000}, & G < 4001, \\ \tau & \text{otherwise} \end{cases}$$

4. Oscillating complexification of the environment (λ = 50, 100, 200, 400 generations), with complexity following a sinusoidal increase and decrease over wavelength λ with maximum amplitude τ (Strategies 3, 4, 5 and 6):

$$h(G,\lambda) = \tau\,\frac{1+\sin(\frac{2\pi G}{\lambda}-\frac{\pi}{2})}{2}$$

5. Oscillating complexification of the environment as above, with maximum amplitude interpolated linearly between 0 and τ from generation 0 to generation 4000 and fixed at τ from generation 4001 to 5000 (Strategies 8, 9, 10 and 11):

$$h(G,\lambda) = \begin{cases} \frac{\tau G}{4000}\,\frac{1+\sin(\frac{2\pi G}{\lambda}-\frac{\lambda}{2})}{2}, & G < 4001, \\ \tau\,\frac{1+\sin(\frac{2\pi G}{\lambda}-\frac{\pi}{2})}{2} & \text{otherwise} \end{cases}$$

6. Adaptive modification of 3, where h is instead increased by 1% of τ when the population's average fitness has increased or remained the same and otherwise decreased by 1% of τ, while being kept in the range $[0,\tau]$. (Strategy 12):

7. Adaptive modification of 5 where the value multiplying the oscillator is instead increased by 1% of τ when the population's average fitness has increased or remained the same and otherwise decreased by 1% of τ, while being kept in the range $[0,\tau]$. (Strategies 13, 14, 15 and 16).

4.5 Results

Table 4.1 shows that *no homogeneous complexification strategy (direct, linear or adaptive) was able to achieve success on all task difficulties*, in any experimental run. In contrast, *all heterogeneous strategies did*. The adaptive oscillating (λ=50) strategy achieved 100% success in 20% of runs and 95% success in 48% of runs.

Figure 4.4 shows a complete view for each strategy, with λ=50 selected for each oscillating strategy and each strategy's 100 runs sorted along the horizontal axis by proportion of successful evaluations (shown on the vertical axis). Note that we are primarily interested in the upper portion of this graph, that is in those populations able to complete the task at most obstacle heights. The adaptive strategy generated fewer populations than the linear strategy, successful on fewer than 50% of evaluations (over the full range of obstacle heights) but of greater interest is that it generated only a comparable number of populations successful on more than 90% of evaluations. The random strategy, whilst better than all homogeneous strategies, is by far the worst method of the heterogeneous strategies. In turn, the simple oscillating strategy is outperformed by the increasing oscillating and adaptive oscillating strategies.

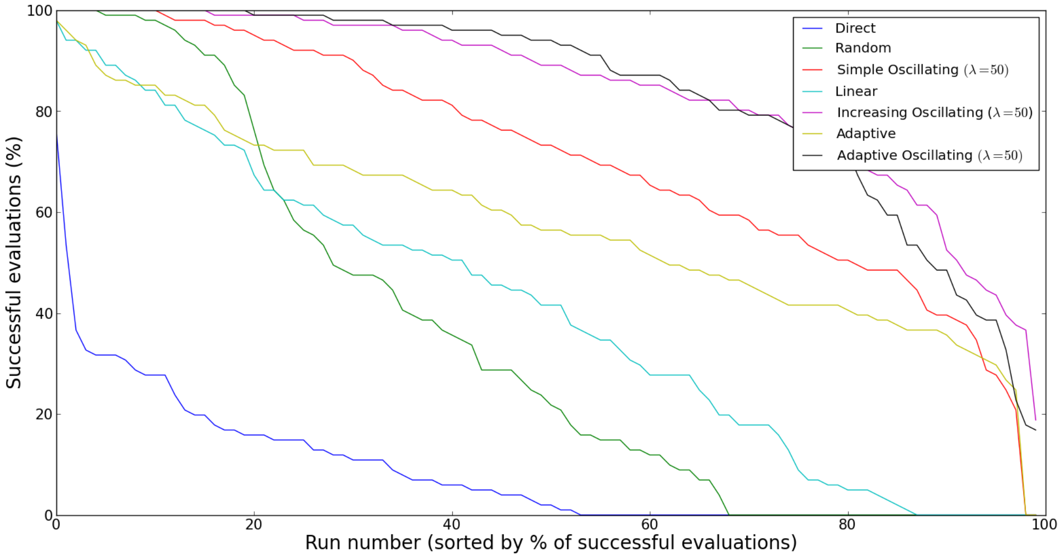

Figure 4.4: Performance of various strategies, 100 runs per strategy sorted best to worst. 1–direct; 2–random; 3–simple oscillating (50); 7–linear; 8–increasing oscillating (50); 12–adaptive; 13–adaptive oscillating (50).

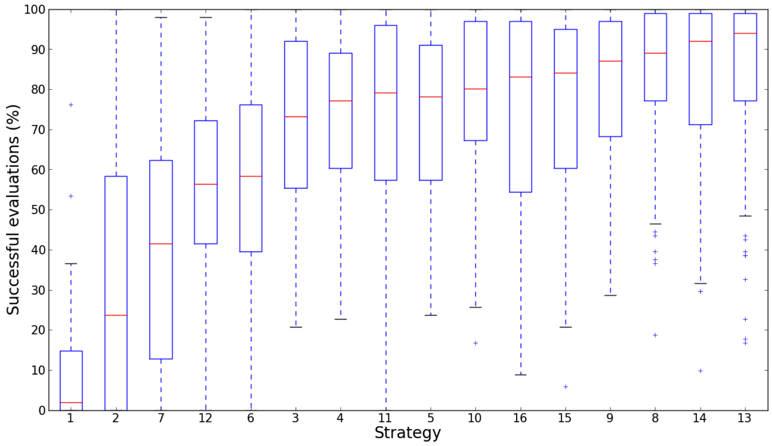

Figure 4.5: Aggregate success rate over all obstacle heights for various strategies, sorted by median success rate. Each evolved population was evaluated on the task at heights 0%, 1%, ... 100%. (See Table 4.1 for description of the numerical labels.)

Number	Strategy	% success: 95%	100%
Homogeneous Strategies			
1	Direct	0%	0%
7	Linear	1%	0%
12	Adaptive	2%	0%
Heterogeneous Strategies			
2	Random	13%	5%
3	Simple Oscillating (λ=50)	21%	11%
4	Simple Oscillating (λ=100)	16%	7%
5	Simple Oscillating (λ=200)	17%	8%
6	Simple Oscillating (λ=400)	10%	2%
8	Increasing Oscillating (λ=50)	39%	16%
9	Increasing Oscillating (λ=100)	30%	14%
10	Increasing Oscillating (λ=200)	30%	12%
11	Increasing Oscillating (λ=400)	29%	10%
13	Adaptive Oscillating (λ=50)	48%	20%
14	Adaptive Oscillating (λ=100)	44%	16%
15	Adaptive Oscillating (λ=200)	26%	9%
16	Adaptive Oscillating (λ=400)	31%	9%

Table 4.1: Number of runs achieving success on 95% and 100% of obstacle heights.

Figure 4.5 shows a box plot of successful evaluations (%) for each strategy (with whiskers to 1.5 interquartile ranges below and above the lower and upper quartiles), complete with a range of wavelengths (50, 100, 200 and 400 generations). Mann-Whitney U tests were performed to examine significant differences in median number of evaluative successes between strategies and within strategies (by varying wavelength). In table 4.2, a left arrow indicates that the strategy corresponding to the row number has a significantly higher ($p < 0.05$) median success rate than the strategy corresponding to the column number (and an up arrow vice versa), shown particularly clearly by strategies 8 and 13. Within each of the increasing and adaptive oscillating strategies, the median number of successful evaluations was found to be significantly higher ($p < 0.05$) for strategies with wavelengths of 50 to 100 generations when compared to the same strategy with four times the wavelength or higher. Within the simple oscillating strategy, the long wavelength (400 generations) produced a significantly lower median ($p < 0.05$) than shorter wavelengths (50, 100, 200 generations).

Strategies which oscillate showed the best performance. No significant difference was found in median between the increasing and adaptive oscillating

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1		↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑
2	←		↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑
3	←	←				←	←	↑	↑	↑		←	↑	↑		
4	←	←				←	←	↑	↑	↑		←	↑	↑		
5	←	←				←	←	↑	↑	↑		←	↑	↑		
6	←	←	↑	↑	↑		←	↑	↑	↑	↑		↑	↑	↑	↑
7	←	←	↑	↑	↑	↑		↑	↑	↑	↑	↑	↑	↑	↑	↑
8	←	←	←	←	←	←	←			←	←	←			←	←
9	←	←	←	←	←	←	←				←	←	↑			
10	←	←	←	←	←	←	←	↑				←	↑	↑		
11	←	←				←	←	↑	↑			←	↑	↑		
12	←	←	↑	↑	↑		←	↑	↑	↑	↑		↑	↑	↑	↑
13	←	←	←	←	←	←	←		←	←	←	←			←	←
14	←	←	←	←	←	←	←			←	←	←			←	←
15	←	←				←	←	↑				←	↑	↑		
16	←	←				←	←	↑				←	↑	↑		

Table 4.2: Significance table showing one-tailed statistical relationship between strategies, $p < 0.05$. Between any two strategies x and y, the arrow in the box at (x, y) points to the statistically dominant strategy (significantly higher median value). An empty cell indicates no statistical difference between strategies. Statistical test used was the Mann-Whitney U-test.

strategies at equal wavelengths. It was found that either a linear or an adaptive increase in maximum amplitude over the training time performed significantly better than simple oscillation. For both increasing oscillating and adaptive oscillating the two lower wavelengths (50 and 100 generations) showed a significantly higher ($p < 0.05$) median number of successful evaluations than the simple oscillating strategies at all wavelengths.

On average, the adaptive strategy performed significantly better than the direct, linear and random strategies, and significantly worse than every oscillatory strategy (except the simple oscillating strategy at wavelength 400 for which there was no significant difference).

The linear strategy resulted in a significantly higher median number of successful evaluations than the direct and random strategies (even though the random strategy produced more highly fit populations from many more runs) and a significantly lower median than all other strategies.

On average, the random strategy performed significantly worse than all other strategies except for the direct method, which was significantly worse than all other strategies.

In order to determine whether the poor results of the linear strategy is due to either evolutionary loss or failure to gain, the proportion of successful evaluations at each obstacle height was determined throughout the evolutionary progress, for each strategy. All strategies achieved 8% success at all obstacle heights, with the exceptions of direct (for which obstacle height is always 100%) and adaptive (low coverage at high obstacle height). The linear strategy achieved more successful evaluations than the simple oscillating strategy at all wavelengths during the evolutionary phase, indicating that its ultimate failure is due to evolutionary loss rather than a failure to gain. Only 10% of the final populations from linear runs contained an individual able to walk to the target with no obstacle, compared to at least 69% for the increasing and adaptive oscillating strategies. Figure 4.6 shows the performance of the linear strategy and the three oscillatory strategies of wavelength 50 generations, against obstacle height, during evolution.

As in figure 4.5, figure 4.7 shows the number of successful evaluations for each strategy, but drawn only from those runs able to reach the target with no obstacle (that is eliminating those runs which experienced the greatest evolutionary loss). It shows that in these cases, linear performance has a range comparable to the simple oscillatory strategies and a median comparable to the increasing and adaptive oscillating strategies.

To investigate the dependency of success rate on oscillatory frequency the simple, increasing and adaptive oscillating strategies were evaluated across as range of wavelengths from 2 to 10000 generations; figure 4.8 demonstrates this relationship. As wavelength approaches zero, the proportion of successful evaluations approaches that of random. As wavelength approaches total evolutionary time (number of generations), the proportion of successful evaluations approaches that of linear. Between these points, it can be seen that for each strategy there is an optimal wavelength (for the current algorithm, around 50-100 generations).

4.6 Discussion

It is clear from the results presented above that there is a strong distinction between the homogeneous and heterogeneous strategies. No homogeneous strategy achieved 100% coverage of the evaluation task in any run (Table 4.1) whereas all heterogeneous strategies did. Within the homogeneous category, the trivial, direct method of presentation was by far the least successful (Figures 4.4 and 4.5). The linear strategy was more successful but the best strategy in this category was the adaptive strategy. The poor performance of the homogeneous category can be explained by *evolutionary forgetting*: these strategies have either lost evolutionary gradient and drifted away from any early successes (linear) or over-specialised on later parts of the problem (adaptive).

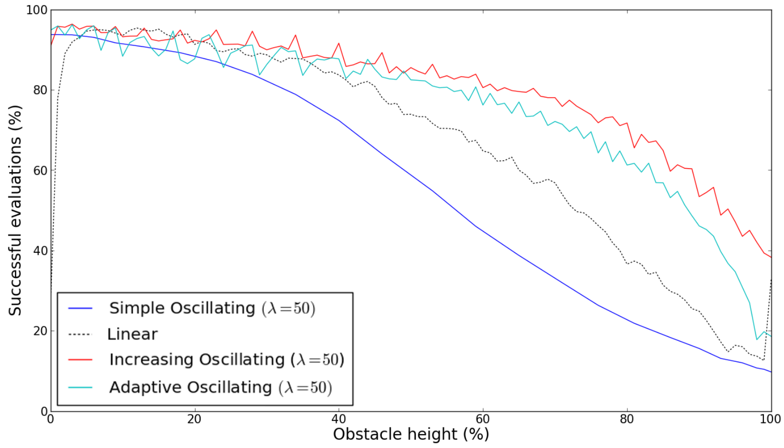

Figure 4.6: Strategy performance against obstacle height during evolution.

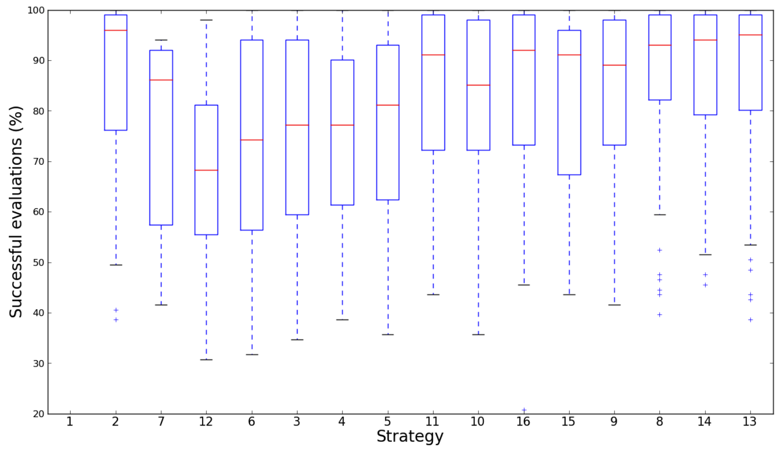

Figure 4.7: Success rate over all obstacle heights for various strategies (only aggregates runs which solved the task at zero-height). (See Table 1 for description of numerical labels; order preserved from Figure 4.5.)

The heterogeneous strategies performed better than the homogeneous group: all the most successful strategies examined made multiple presentations of easier tasks at later stages of the evolutionary run, at the expense of fewer consecutive presentations of very similar tasks. These strategies all performed well at the hardest task and had the best generalisation performance over the whole range of tasks, suggesting that our hypothesis has merit.

The random strategy is the least successful strategy in this category. This may be due to the same problem of gradient loss as in the homogeneous group. As found in the homogeneous group, the linear and adaptive modifications of the oscillating strategy showed the best performance of all; the slow increase in task difficulty maintains a strong evolutionary gradient and the cyclical nature of task presentation consolidates earlier gains and causes the evolving population to

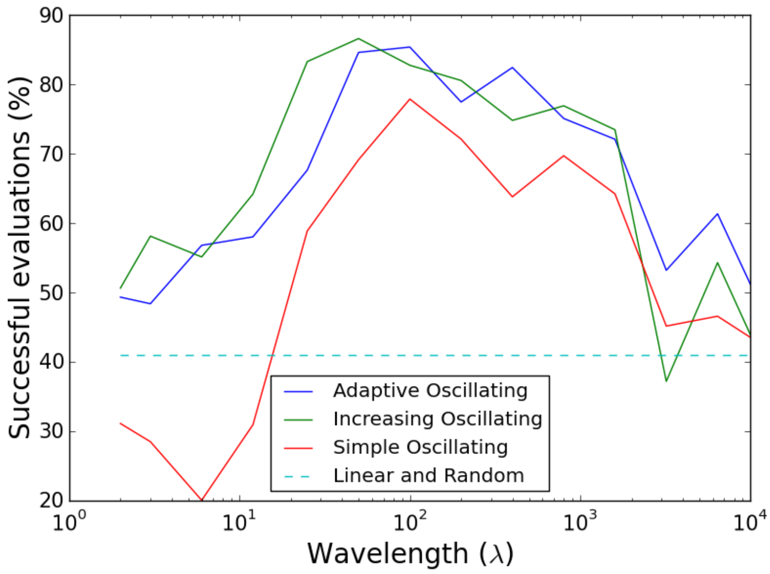

Figure 4.8: Strategy performance (% success) against wavelength for oscillating strategies.

prefer generalised solutions abstracted over the whole problem domain.

This consolidation is dependent on the frequency of re-presentation of earlier, or easier, parts of the task. When investigating this frequency, it can be seen that a clear optimum exists in the frequency domain where cyclical strategies are able to maximise this consolidation without losing gradient. This optimum is likely to be problem-specific and a range of values should be explored for any given task. However, in the limit of wavelength, that is at very low and very high frequencies, it can be seen that the performance of the evolving populations begins to approximate, for low and high frequencies respectively, the linear and random strategies. This offers an abstract insight into the underlying mechanism at work - the maintenance of selective pressure and whole-task capability. As these components reduce in effectiveness due to the change in wavelength, so the oscillating strategies degenerate into the simpler strategies described above. The successful cases are those where environmental change is fast enough to induce a generalisation in the agent's approach to the task but slow enough to prevent catastrophic loss of gradient when evaluating partial solutions.

4.7 Conclusions

The points made in the discussion section support the hypotheses outlined earlier in this chapter. The homogeneous strategies showed weak performance on the evaluation task, with no strategy achieving full coverage in any run. Conversely the heterogeneous strategies, including surprisingly the random strategy, all achieved full coverage in some runs. Those heterogeneous strategies with a range of difficulties increasing over time (increasing and adaptive oscillating) outperformed the simple (constant range) oscillating strategies, showing a much higher proportion of successful runs. Finally, we demonstrated that oscillating strategies do exhibit an optimal frequency.

Complexification strategies for incremental evolution offer a powerful mechanism for adaptive problem solving. However, this power comes at a price: it is easy to lose information learned earlier in the process. In order to fully exploit this power appropriate complexification strategies have to be realised in order to drive populations along desirable adaptive pathways. There are many options for formulating these strategies: much previous work has involved, in one manner or another, a simplification of the objective function and then a progressive complexification as time passes. In this work it was found that many strategies encounter loss-of-gradient or over-fitting problems. A solution was presented in the form of heterogeneous complexification strategies which combine solutions to those problems to deliver robust populations. Our approach can be translated to many scenarios where progressive complexification is used to guide an incremental evolutionary process; further exploration of the limitations and advantages of heterogeneous complexification within different problem domains would be useful in order to generalise these conclusions. Additionally, the oscillating strategies exhibited an optimal wavelength for re-presentation. It is unclear whether this optimum is task-dependent or whether there is an underlying principle and optimal wavelength for this type of training; this question also merits further work.

Finally, on the strength of these results, it is advisable that in general while a random presentation of subtasks or objective difficulty levels is preferable to a linear increase, as a minimum guideline an increasing heterogeneous complexification strategy should be used. This rehearsive, cyclical approach to presentation not only maintains evolutionary gradients but also promotes generalisation amongst the evolving populations from subtask-specific adaptation to performance across the super-task.

CHAPTER 5

Reactive, Deliberative 3D Virtual Creatures

"We first make our habits, then our habits make us."
–John Dryden

In chapter five, a 3D virtual creature system is presented which unifies several earlier ideas, leveraging the technical discoveries made in chapter four. The research question, "Can an evolving 3D virtual creature environment produce lifelike reactive and deliberative behaviours?" is addressed. The motivation for this chapter is to take the understanding from chapter four and move away from the abstract task and a single behaviour, towards navigation and object manipulation, as well as locomotion. A complex task is chosen that requires these kinds of behaviours, but also encapsulates the key ideas from chapter four. A novel hybrid neural network controller is also invented to this end. Results are presented which show that the simulation produces reactive and deliberative behaviours hitherto unseen in 3D virtual creatures.

5.1 Introduction to Chapter

Living systems exhibit a large variety of coordinated activities at many different scales. We find, for example, homeostasis, locomotion, learning, group and social behaviours throughout the natural world. Since the earliest days of Artificial Life, a defining ambition has been to understand how to engineer systems that exhibit some of these complex behaviours, either to solve problems or to understand the underlying principles that gave rise to them in nature (Langton, 1989).

The specification of a model requires assumptions to be made concerning the degree to which its most basic units and the rules governing their behaviour are able to act as reliable proxies for their natural analogues. The granularity of a system has a direct impact on both its speed and its potential to accurately mimic nature, and on the strength of conclusions about the natural world based on phenomena observed to emerge from interactions within it.

The work in this chapter constitutes a first attempt to combine the incremental neuroevolution of reactive and deliberative behaviours with the neuroevolution of a 3D agent's motor control. The overarching aim is the incremental evolution of sophisticated behaviours, for the population to overcome increasingly complex challenges in the agents' environment over evolutionary time.

The challenge is difficult because deliberative behaviour will be limited by necessary performance in motor control. An incremental approach can take this subtask-interdependency into account and prevent loss or lack of evolutionary gradient early in evolution. However, as elucidated in detail in chapter 4, care is required when designing such incremental steps. Changing selection pressures too rapidly or too slowly can, respectively, cause evolution to lose gradient or over-fit to the current challenge. That work also demonstrated that it is necessary to revisit earlier incremental steps in order to prevent the loss of evolved abilities and therefore to find general solutions.

There is then a question of how to implement deliberative processing alongside physical control in a single controller. Deliberative planning systems generally learn a state-based action policy in order to select the best next state given a set of available actions. In contrast, flexible control of 3D motion requires a continuous-time closed-loop control system to keep physical variables within operational parameters. Also, for locomotive behaviours, an endogeneous oscillation within the controller or body–controller action loop is necessary to achieve a reliable gait.

The requirements of each of these control systems is fundamentally different; it is difficult to design an architecture that can effectively learn the two different problems. The choice is between either an architecture that is general enough to be capable of both episodic categorisation and time-based close-coupled motor control, or a combination of the two architectures each tailored to a specific part of the problem and integrated elsewhere. The work in this chapter opts for the latter, as a pragmatic step toward a more general architecture.

5.2 Hypothesis

The present work examines the following hypothesis: *that it is possible to produce a sequence of reactive, deliberative behaviours in three-dimensional virtual creatures using a simultaneous incremental evolutionary paradigm to optimise an implementation of the hybrid neural architecture detailed below*. The "River Crossing" (RC) task devised by Robinson et al. (2007) is used as the baseline reactive–deliberative problem. This task is adapted by the addition of a requirement of physical motor control in 3D, and the complete problem against which agents are tested is hereafter referred to as the 3D River Crossing or 3D RC task.

The remainder of this chapter presents details of the 3D RC task, the agent and its hybrid neural architecture, and the evolutionary system, before reporting qualitative and quantitative results and conclusions. It provides an existence proof that demonstrates the sufficiency and overall success of the design.

5.3 Methods

The main contribution of this chapter is the novel fusion of multiple neural architectures, each addressing different aspects of the 3D RC task, in order to enable the incremental evolution of agents that achieve the full task. This section of the chapter introduces the environment and physical model and then describes the hybrid neurocontroller in detail, making reference to the inputs and outputs defined by the agent–environment relationship. Finally, the evolutionary algorithm is described in terms of the parameters of the neural architecture, and the experimental set-up is outlined.

5.3.1 Environment and physical model

The environment for the evolutionary problem is a modified version of the RC task first used in Robinson et al. (2007). In this task, agents exist and move around in a discrete, 20×20 bounded grid world. Each grid cell has attributes which can affect the agent: *traps* kill it, as does *water* (drowning); *grass* is neutral and *stones* can be picked up and put down. Stones can be placed on water, enabling bridges to be built. The final attribute, *resource*, is the agent's goal. The RC task is an incrementally difficult challenge, with a staged introduction of difficulties. By collecting the resource, agents progress through more complicated environments, eventually arriving at a 20×n-cell river, where n is the increasing width of the river and thus the difficulty of the bridge-building task.

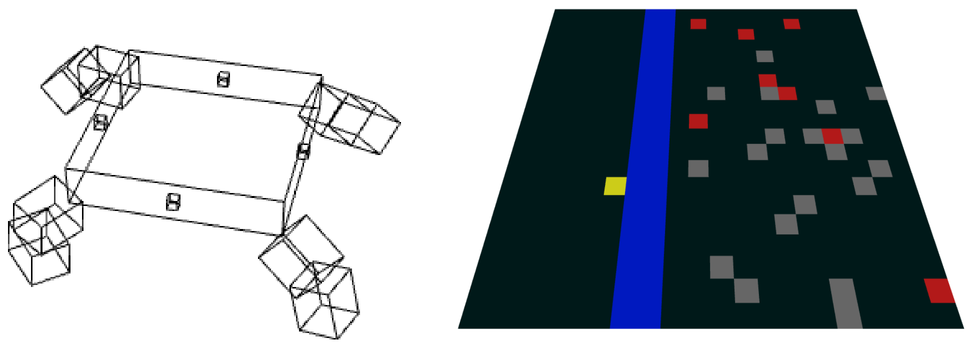

Figure 5.1: Agent morphology and environment, showing resource in yellow, river in blue, traps in red and stones in grey.

The 3D RC environment used in this work extends the 2D RC environment. Agents have a symmetrical quadruped body plan (figure 5.1) comprised of a torso (dimensions 1.0×1.0×0.2 cell-widths), four upper limbs (0.5×0.2×0.2), four lower limbs (0.5×0.2×0.2) and four small sensors (0.05×0.05×0.05). The upper limbs are attached to the torso at each lower corner with a 2-axis constraint. The constraint limits the range of motion of the upper limb relative to the torso, to $\frac{\pi}{2}$ radians around the vertical axis, and π radians around the line lying tangent to the agent's torso in the plane of the torso. Lower limbs are connected to upper limbs via a knee constraint which limits the range of motion between the two parts to $\frac{\pi}{2}$ radians around the y-axis. The sensors are attached with fixed constraints to the centre of each of the four faces of the agent's torso perpendicular to the ground plane. The physical simulator used was Open Dynamics Engine (ODE) version 0.13.1, with friction pyramid approximation for contact response ($\mu = 10.0$) between agent and the ground plane, universal ERP of 0.2 and CFM of 5×10^{-5}.

In order to bootstrap the evolution of locomotive behaviour, two additional levels were added at the start of the incremental RC task. The first level distributes "food" around the RC world. This confers additional fitness on agents once collected. The second level ("dash") has only one occupied cell, containing the resource. These levels together promote locomotive behaviour, and ultimately optimise the behaviour for speed of movement.

The difficulty of the RC environment is increased incrementally across six progressively more challenging levels. An agent's fitness is incremented from zero by 100 each time it successfully finds the resource, a requirement to progresses to the next level.

- Level 1: *Food*. The RC environment contains only cells with the resource (one cell) and food (probability 1/20 per cell). Interaction with a food cell removes the food from the environment and increments the agent's fitness by 1.
- Level 2: *Dash*. This level contains only a single resource cell which agents must discover.
- Level 3: *Stones and Traps*. This level contains eight traps and twenty stones, as well as the target resource.
- Level 4: *Easy bridge*. This level is as level three but with a river of width 1 crossing the terrain.
- Level 5: *Medium bridge*. As level four, but width 2.
- Level 6: *Hard bridge*. As level five, but width 4.

On completion of level 6, agents are returned to level 1 and can continue to accumulate fitness until the time limit of 10 simulated minutes is reached, when evaluation is terminated. These challenges constitute a strategy of revisiting and gradually changing task components, as in the heterogeneous strategies of chapter four. Agents are exposed to previously-seen examples both through the revisiting of earlier levels when successful in level 6, but also through the encapsulation of earlier challenges within later ones. For example, the “easy bridge” level encapsulates the “stones and traps” level. Whilst the breakdown of the whole task (success at any level) is less straightforward in this presentation than in that of chapter four, the underlying principles are exploited in the same way.

5.3.2 Neural architecture

A neural architecture capable of solving the 2D RC task was a major contribution of Robinson et al. (2007) and is extended in the present work. In the 3D RC task, an agent’s neurocontroller transforms sensory inputs into torque values for motor control, which gives rise to behaviour in the physically simulated environment. The control system must produce directed locomotive behaviour in the quadruped, and change locomotive behaviour over the stages and sub-stages of the RC task, according to external (sensory) and internal (neural) state.

The hybrid neural architecture (figure 5.2) integrates the outputs of the RC world *decision network* (DN) and the diffusive *shunting model* (SM) with the inputs of the *physical network* (PN), and then use this information to pilot the agent through the world by affecting the operation of the agents’ *pattern generator* (PG) neurons. This combined architecture is a novel integration of several technologies from neural network, robot control, and biological modelling literature.

The Decision Network

The DN architecture follows the design laid out in Robinson et al. (2007). The DN is a standard feedforward neural network which takes inputs representing the attributes of the agent’s current location in the RC world, and an input indicating whether or not the agent is currently carrying a stone. The hidden layer contains four neurons which sum over the inputs and apply a hyperbolic tangent activation function. The output layer sums over the hidden layer, applies a hyperbolic tangent activation function and tests at the thresholds -0.3 and 0.3; output neurons have three possible values: -1, 0 or 1, and determine the *iota values* used in the SM. These iota values indicate the saliency of the attributes in the environment, so the DN outputs iota values for each attribute (resource, stone, water and trap) except grass (which has an iota value of zero).

The Shunting Model

The SM was first used as a novel approach to motion planning by Meng and Yang (1998). The approach uses the homomorphism between the varying external environment and the intrinsic dynamics of the architecture to achieve route

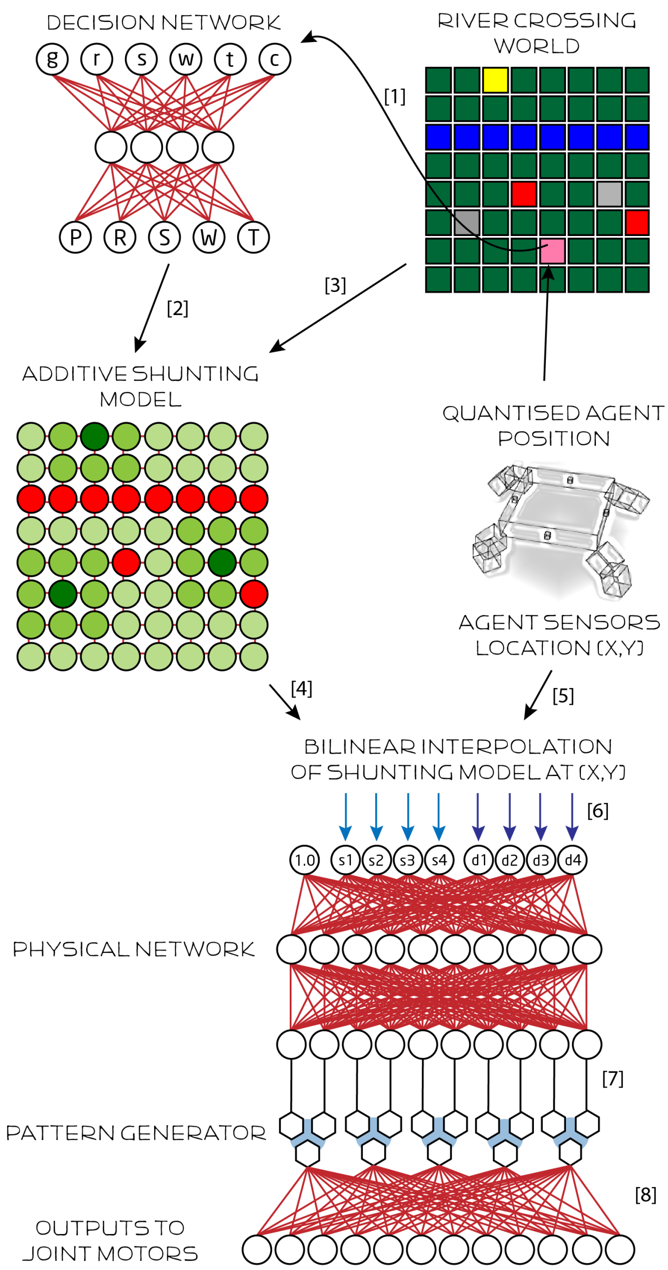

Figure 5.2: Neural architecture. Attributes at the agent's position (g=grass, r=resource, s=stone, w=water, t=trap, c=carrying flag) determine inputs to the Decision Network [1]. The Shunting Model constructs a landscape using iota values output by the DN [2] (P=pickup action, R=resource, S=stone, W=water, T=trap) and the locations of objects [3]. The SM activity landscape is interpolated [4] at the positions of the animat's four sensors [5], and these values fed to the Physical Network [6]. PN outputs are fed to the Pattern Generator Network [7], which outputs to neuromotor controllers. Links in red are genetically specified.

generation (planning) without explicitly searching over possible paths. It is a generalisation of the potential field approach of Glasius et al. (1995), historically an evolution of the model of neural connectivity first proposed in Hodgkin and Huxley (1952). The SM uses a locally-connected, topologically-organised network of neurons to propagate desirable states across the entire network of transitions in the space. This produces an *activity landscape* with peaks at target states and valleys at configurations to avoid. One of the most common implementations of the SM is the *additive model* (Grossberg, 1988), which sacrifices gain control (and thus, stability) for simplicity. This model defines the following differential equation to model the diffusion of input values across the state landscape:

$$\frac{dx_i}{dt} = -Ax_i + \sum_{j \in N_i} w_{ij}[x_j]^+ + I_i \tag{5.1}$$

where each neuron in the SM corresponds to one discrete cell in the environment; x_i is the activation of neuron i, taken to be zero outside of the environment; A is a passive decay rate; N_i is the receptive field of i; w_{ij} is the connection strength or weight from neuron j to neuron i, specified to be set by a monotonically decreasing function of the Euclidean distance between cells i and j (zero outside of the neighbourhood); the function $[x]^+$ is $max(0, x)$; and I_i is the external input to neuron i.

This technique was used in Robinson et al. (2007) to model the state space of the RC problem by directly representing the discrete RC world in the configuration of the SM, with each cell's receptive field set to be the eight cells in its Moore neighbourhood, within which all $w_{ij} = w$, and external input I_i determined by the attributes present in cell i and the saliency (*iota value*) for those attributes as computed by the DN. Neural activations propagate from external input I according to the local connectivity of the neurons, and the entire network can be considered a diffusive model that produces landscapes in which following positive gradients leads to target states. With well-chosen constant multipliers, this method exhibits no undesirable dynamics and has been found to be considerably versatile in a variety of subsequent works, including those of Borg and Channon (2011) and Luo et al. (2014).

In this work, we simplify and clarify the setting of of decay rate and scales for distance (or weights) and iota values. A stable solution ($x_i^{new} = x_i$ for all i) to equation 5.2 is a stable solution ($\dot{x} = \mathbf{0}$) to equation 5.1. We absorb the constant A into the scales for iota values and distances, and set and limit weights and activation according to neighbourhood size (8) and maximum iota value ($maxI$=15), resulting in equation 5.3.

$$x_i^{new} = \frac{1}{A}\left(\sum_{j \in N_i} w_{ij}[x_j]^+ + I_i\right) \tag{5.2}$$

Following the computation of external inputs I by the DN, we zero SM activations and then iterate equation 5.3 fifty times to allow activity to propagate and stabilise across the 20×20 array of SM neurons.

$$x_i^{new} = min\left(\frac{1}{8}\sum_{j\in N_i}[x_j]^+ + I_i,\ maxI\right) \tag{5.3}$$

5.3.2.1 The Physical Network

The PN controls the agent's behaviour in the world. It receives as inputs the SM activations (interpolated) at the positions of the four sensors located on the four sides of the agent's torso. Since the SM represents a neural quantisation of the continuous landscape in which the sensors move, a single value is calculated for each sensor using a bilinear interpolation of the SM's activity values at the four points around the relevant sensor:

$$\begin{aligned} a(x,y) \quad = \quad & f[\lfloor x\rfloor, \lfloor y\rfloor](1-\{x\})(1-\{y\}) + \\ & f[\lceil x\rceil, \lfloor y\rfloor]\{x\}(1-\{y\}) + \\ & f[\lfloor x\rfloor, \lceil y\rceil](1-\{x\})\{y\} + \\ & f[\lceil x\rceil, \lceil y\rceil]\{x\}\{y\} \end{aligned} \tag{5.4}$$

where $a(x,y)$ is the interpolated activity at $(x,y)\in\mathbb{R}^2$, $f[i,j]$ is the SM activation at the discrete point $(i,j)\in\mathbb{Z}^2$ and {x} denotes the fractional part of x.

These four sensor values are normalised (divided by $maxI$) and then fed into the PN, together with four values that indicate which sensor has the maximum value. The PN operates as a standard feedforward neural network where hidden nodes receive a weighted sum of the inputs. The hidden layer uses a hyperbolic tangent activation function in order to maintain negative values. The output layer uses a sigmoid activation function.

5.3.2.2 The Pattern Generator Network

The PG is a set of pre-evolved oscillatory neural circuits which are modelled on the networks of leaky integrators presented in Beer and Gallagher (1992) and used for locomotor pattern generation in many subsequent works, including Reil and Husbands (2002) and Stanton and Channon (2013). The circuits themselves are three-neuron motifs evolved to produce 1Hz sinusoidal oscillations from an output node in the presence of an input signal, and to be quiescent otherwise. Each complete PG network has a set of five identical motifs, initially isolated, which receive input from the PN via a set of weights and send their outputs to the final stage of the agent's controller. The neurons comprising these motifs are simple continuous-time leaky integrators, with behaviour governed by the following equations:

$$\tau_i \frac{dA_i}{dt} = -A_i + \sum_{j=0}^{n} w_{ij} O_j \tag{5.5}$$

$$O_i = \tanh \frac{\alpha_i - A_i}{2} \tag{5.6}$$

where A_i is the activation of a neuron i, O_i is the output of neuron i, w_{ij} is the weight from neuron j to neuron i, α_i is the bias of neuron i and τ_i is the time-constant of neuron i. At each iteration of the update algorithm ($dt = 0.01s$), equation 5.5 computes the change in the activity of the *i*th neuron for all neurons, and then equation 5.6 computes the output value for all neurons. It is this output value that is used by the neuromotor controllers.

To generate the original motif, a population of 1000 randomly initialised three-neuron networks was created with weights, time-constants and biases defined by a real-valued genotype. These networks were evaluated against a fitness function which measured the match between the desired frequency and the output response by summation of the undesirable (non-target) frequencies found in the frequency domain after application of Fourier transform. Networks were simulated for 10 seconds, twice. Once with a high input and a target frequency of 1Hz, and once with no input and a target quiescent state. Through three-genome tournament selection, strong candidates were used to generate new, mutated members of the population using the same evolutionary parameters as the general system described below.

5.3.2.3 Neuromotor Controllers

In the final stage, 12 motor controllers (one for each degree of freedom in the agent's morphology) receive the outputs of the PG network via a weighted sum and sigmoid activation function. These motor controllers implement a proportional-derivative (PD) controller, as used by Reil and Husbands (2002), which takes network outputs to be target angles within each joint's range of motion and applies a torque to the joint according to the following formula:

$$T = k_s(\theta_d - \theta) - k_d\dot{\theta} \tag{5.7}$$

where T is the torque applied to the joint, k_s is the spring constant, k_d is the damping constant, θ_d is the target angle and θ is the current angle. In this work, $k_s = 0.25$ and $k_d = 0.175$ were found to produce stable action at joints using the same preliminary studies as chapter four. The difference in these values between this chapter and the previous one is due to minor differences in simulator version, parameters, and agent morphology. This method has the advantage of relieving the neurocontroller of the problem of balancing an agent's weight against the force of gravity.

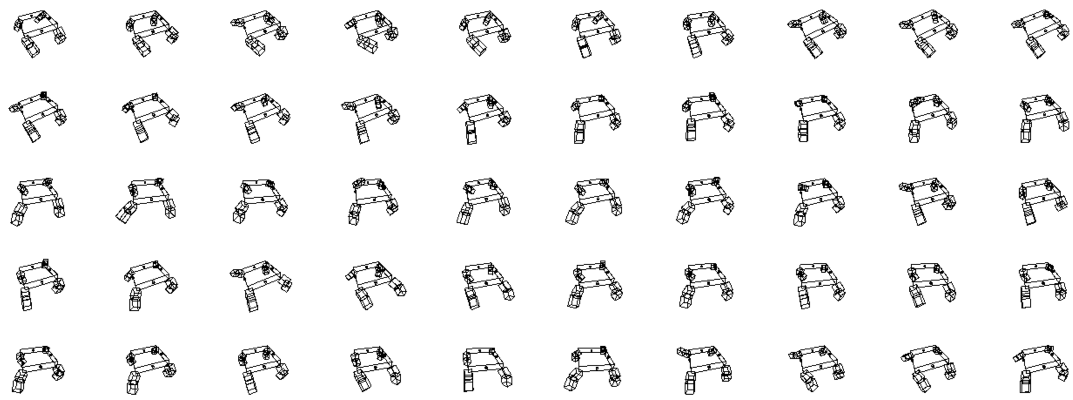

Figure 5.3: Example of a “galloping” locomotive behaviour. Time axis is left to right, top to bottom.

5.3.3 Evolutionary system

A steady-state evolutionary algorithm was used, in which a population of 150 agents is evaluated in randomly-chosen groups of three and the least-fit individual from the group replaced by a mutated single-point crossover progeny of the fitter two.

5.3.3.1 Genetic Representation

Individuals’ neurocontrollers are represented as an array of floating-point values. The sections are laid out as arrays of weights for each network stage as outlined above: the DN input–hidden and hidden–output weights, the PN input–hidden and hidden–output weights, the PG interneuron weights and the PG–motor weights.

5.4 Results

Twenty runs were carried out, each for 10^6 tournaments.

5.4.1 Qualitative results

In those runs scoring highly on the final level of the task, intricate and diverse behaviours can be observed as the agents progress through their environmental challenges. In any single species, several different locomotive strategies can be observed depending on whether the agent is near or far from its target, and whether there are obstacles in the way. In the case of a “clear run”, agents often gallop (figure 5.3) toward the target, whereas if more careful movement is required agents will progress more slowly, making time to avoid unexpected sensory conditions (i.e. traps and water). In both cases, directed control is observed as agents update their heading whilst engaging in locomotion to remain

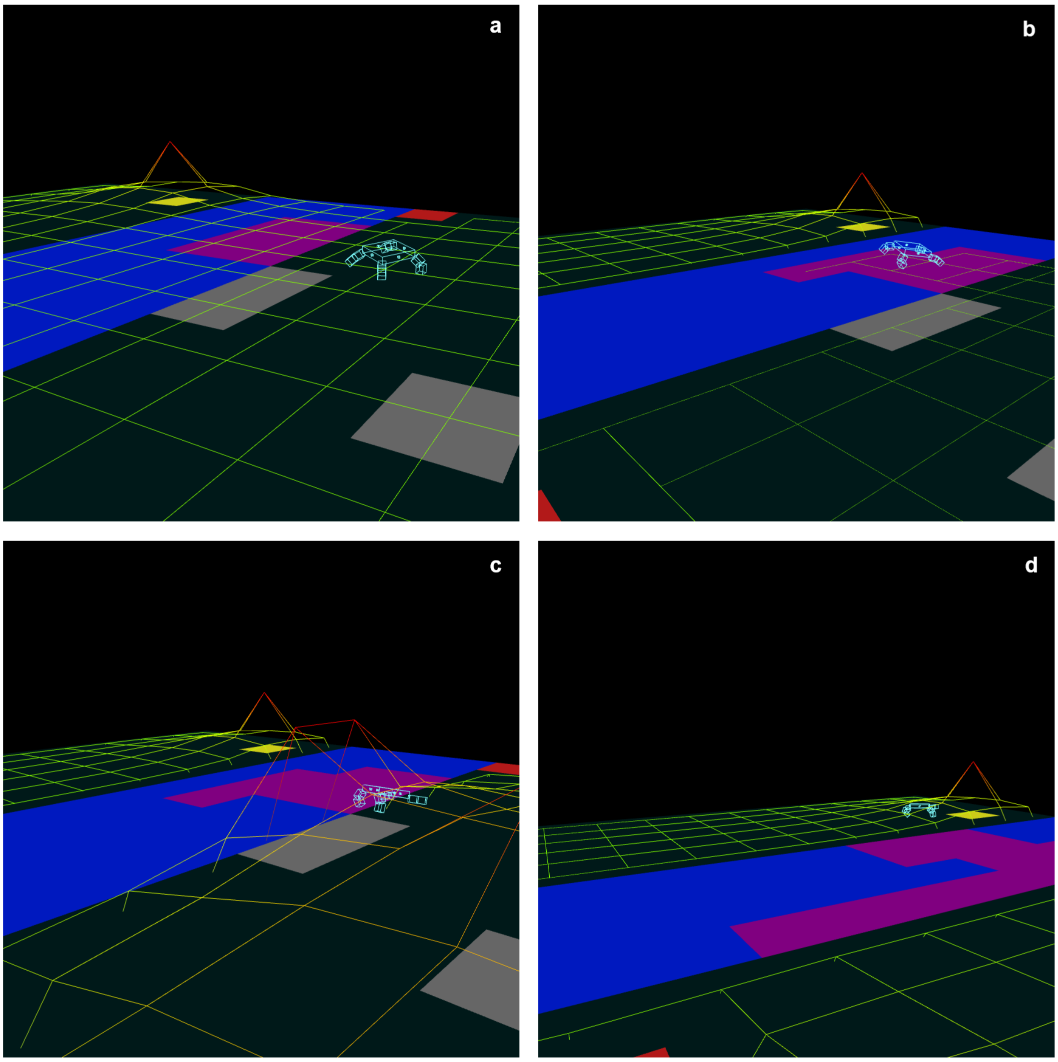

Figure 5.4: Bridge building in action. In (a) the agent has already started to build a bridge and is returning to collect another stone. In (b) the agent has just dropped a stone and is beginning to turn around. In (c) the agent is carrying a stone to drop on the water. In (d) the agent has completed the bridge and is about to reach the resource. The figure also illustrates the SM activity landscape superimposed on the 3D RC world and shows the changes to this landscape due to the updated iota values that occur as the agent's state, and thus DN inputs, vary.

Level-Cover	>0%	20%	40%	60%	80%
1 (Food)	100%	100%	100%	100%	100%
2 (Dash)	100%	100%	100%	100%	100%
3 (Traps)	100%	100%	100%	100%	95%
4 (River 1)	85%	85%	85%	30%	10%
5 (River 2)	85%	65%	50%	20%	-
6 (River 4)	65%	20%	-	-	-

Table 5.1: Proportion of runs with >0%/20%/40%/60%/80% of their final 1000 tournaments successful at level 1/2/3/4/5/6 of the 3D RC task.

aligned with the target. Agents also often display a distinct “turning” behaviour which will engage if the agent is beyond some angular threshold away from facing its target. Figure 5.4 shows an example evolved agent solving 3D RC task.

One of the most lifelike behaviours to be observed is avoidance: due to the non-spreading negative values in the activity landscape agents can unexpectedly encounter a highly negative region. In this case, agents will often crouch and spring back from the hazard, minimising the chance of falling on it due to imprecise control or previous momentum. Finally, in the case where no activation is present on the landscape around the agent, i.e. all directions are of equal saliency, agents engage in a form of random walk reminiscent of similar exploratory behaviour that can be seen in many simple animals. The temptation to interpret these actions in a human or animal context is ever present–agents can seem to exhibit surprise on encountering an unexpected danger, confusion if trapped in a mediocre part of the landscape and even happiness as they gallop toward the resource.

5.4.2 Quantitative results

The fitness scores of the three agents in each tournament were collected. Figure 5.5 shows the progress of the population from a typical run, in solving each level of the 3D RC task. Table 5.1 shows an overview of the performance of the entire system by aggregating and examining the results of the final 1000 tournaments from each run. From this table, it can be seen that every run was able to complete levels one and two in at least 80% of the final 1000 tournaments, and 95% of runs were able to complete level three to this standard too. Performance fell sharply against the bridge-building challenges, although 10% of runs were still able to complete level four in at least 80% of evaluations. At the hardest level of the task, 65% of runs achieved at least 1 evaluation which was able to complete level 6, and 20% of runs achieved at least 20% evaluations able to complete level 6. Figure 5.6 shows this aggregate data for all runs and levels and makes clear the spread of success across the whole problem in the experiment; a clear divide can be seen between the first half and latter half of the problem.

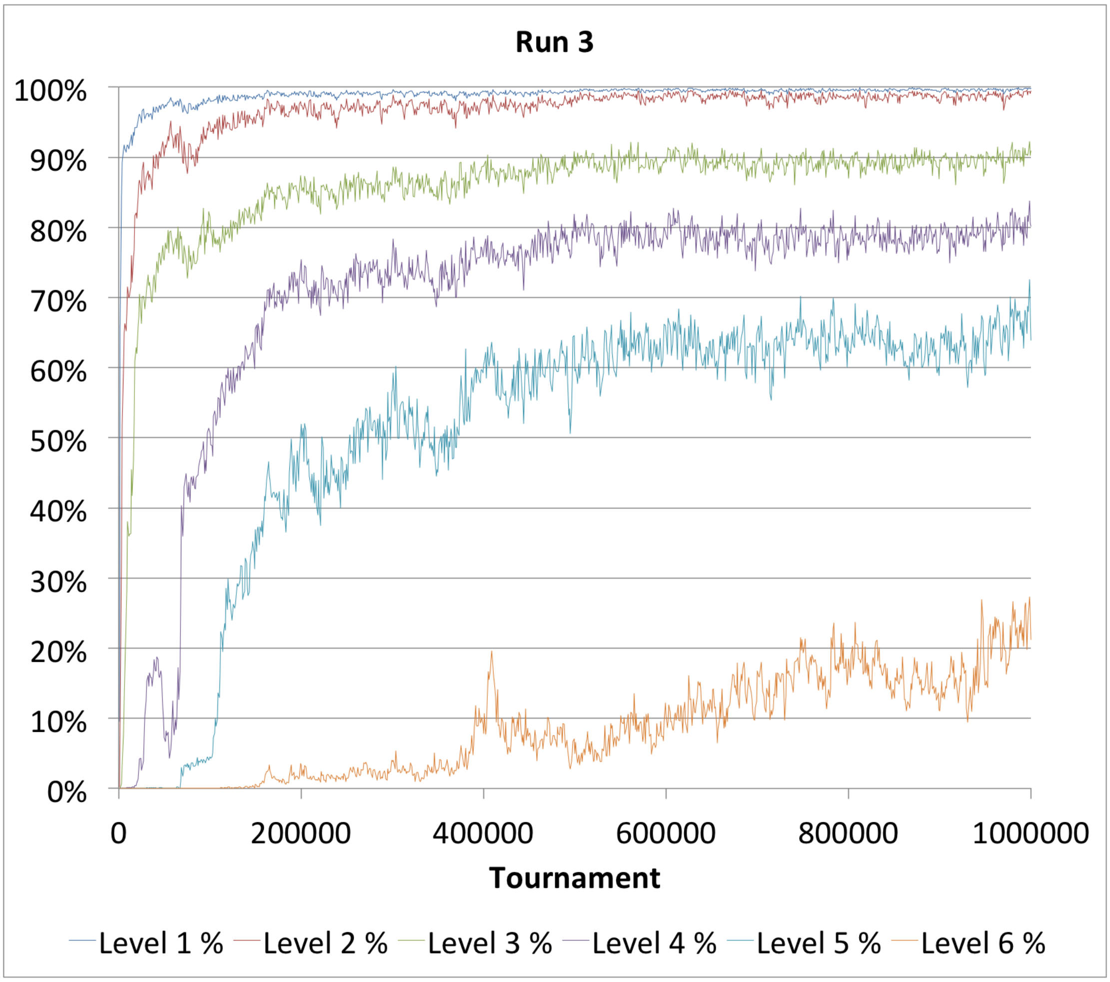

Figure 5.5: Progress of a typical run over one million tournaments. The graph shows the percentage of evaluations successful at completing each level of the 3D RC task, averaged over 1000 tournaments.

When examining the progression of the evolutionary algorithm in individual runs, it can be seen that the first level of the problem is solved early on in the search–typically after only 10000 tournaments. Success at level two soon follows as the problems are similar. Success at the third level (traps and stones, but no river) also occurs early on, in most runs. Levels four, five and six cause a longer delay in the search, and solutions do not appear at all in some runs even though the earlier levels have been solved in similar time to other, successful runs. When solutions do occur, there is often a delay between the solution for level four and later levels.

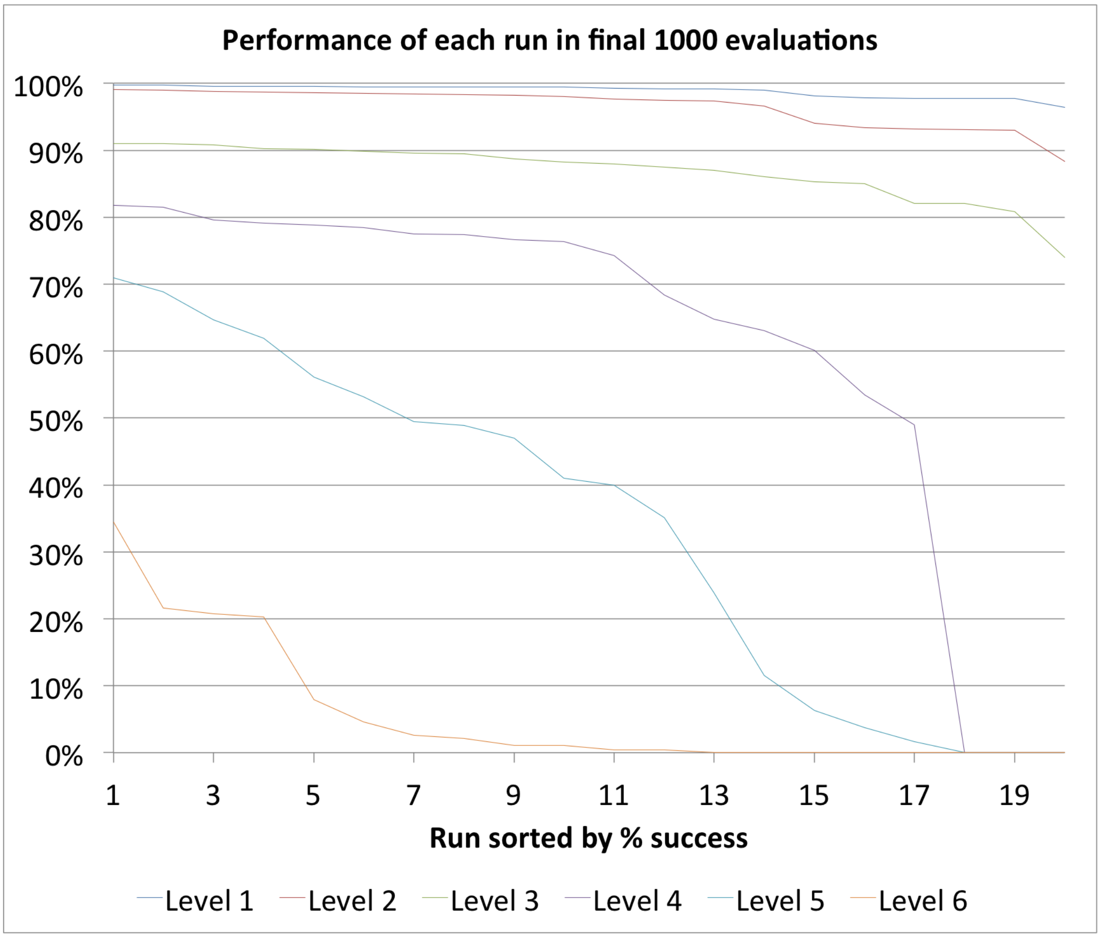

Figure 5.6: Success rates of all runs. The graph shows the performance of each 1000000-tournament run, evaluated from the final 1000 tournaments (3000 evaluations) of each run as the number of these evaluations that successfully completed each level of the 3D RC task. Runs are sorted in descending order for each level of the task.

5.5 Conclusions

This work demonstrates that a standard evolutionary algorithm is sufficient to find parameters for a hybrid neural architecture comprised of loosely-coupled continuous-time and discrete-time neurons to produce reactive and deliberative behaviour in 3D, rigid-body virtual creatures requiring motion control.

By covering the range of task complexity over evolutionary time, species experience an evolutionary pressure (no loss of gradient) whilst still being able to consolidate progress already made. This incremental approach allows species to first develop a locomotive behaviour, and then to use and adapt this ability to explore the space of solutions to the bridge-building river-crossing task.

This work has also shown that a hybrid approach to neurocontroller design that includes a generalised oscillatory component (in this case, an evolved network of leaky integrators) is sufficient to produce agents that exhibit task-dependent behaviours including locomotion, turning and avoidance. The architecture is also able to optimise the strategy for long-term deliberative planning in the 3D RC world at the same time.

The integration of a deliberative decision network and a mechanism to generate reactive behaviour in 3D virtual creatures, via a shunting landscape model, was successful and shows promise for future, more complex work in this area. The limitations of the model are due to the simplicity of the decomposition of the world into the agents' phenomenal space–there is no reason this relationship could not be integrated.

Future work will examine behaviour and neural dynamics to determine what aspects are important. In addition, in order to generalise the applicability of this work to a broad range of tasks, it will be necessary to remove the problem-specific aspects of the neural architecture's design. A first step could be to make the distinction between the DN, SM and PN less explicit. Ultimately a single neural type and architecture, with genetically specified parameters, would be the most general design.

Other possibilities for increasing the coherence in the sensorimotor loop include finer-grained distinctions in the environment, for example iota values for boundary conditions, and the addition of noise to smooth behavioural transitions.

CHAPTER 6

Object Manipulation in 3D Virtual Creatures

"One touch of nature makes the whole world kin."
–William Shakespeare

In chapter six, the work of chapter five is extended to involve an additional abiotic element.

6.1 Introduction to Chapter

The earlier work in this thesis has examined how a sequence of complex behaviours could arise in a 3D virtual creature system. In this chapter, the limits of the previous system are examined in order to move towards the full brain–body–environment, dynamical-systems paradigm of adaptive intelligent behaviour.

The 3D River Crossing (3D RC) task, first presented in Stanton and Channon (2015), provides an ideal base from where the evolution of sensorimotor intelligence and related issues of physical embodiment can be explored. In that work we adapted the *shunting model* of Grossberg (1988) and Yang and Meng (2000), used for the first time in an a-life context in Robinson et al. (2007), to build an evolutionary environment able to evolve control architectures of 3D virtual creatures that exhibit both reactive and deliberative behaviours. However, the problem-solving aspect of the 3D RC task in that work was abstracted from the physicality of the agent's morphology. Although each agent's joint motors were driven by some of the outputs from its neurocontroller, other neural outputs only notionally represented manipulation of physical objects in the agent's world.

An important extension of the earlier work into richer interactions is thus to introduce aspects of the deliberative problem to the agents' physical world, requiring an intricate manipulation of simulated objects to solve the challenge. In this work, we take a first step toward that goal by investigating whether the neural architecture outlined in that work can successfully constitute the control system for a simple manipulation task: displacement of a physically-modelled block in the agent's world, requiring feedback control, hereafter called the *block displacement* (BD) task.

Our general approach is to consider populations of agents in a new environment that provides the physical block challenge. The agents' neural control systems are sensitised to the location of the block by direct interaction with the shunting model, simplifying the adaptive problem. We investigate evolution on the BD task from both random (unevolved) populations and from populations of creatures previously evolved in the 3D RC environment. Hereafter we refer to random populations as *unevolved* populations and populations evolved only in the 3D RC environment as *naive* populations.

6.2 Hypotheses

The objective of this work is the evolution of agents able to successfully complete the BD task, as observed through 3D visualisation. In addition, we developed and tested the following hypotheses in order to further understand the interactions between the various components of the system and explore the limitations of the 3D RC architecture:

H1. The hybrid architecture is sufficient to achieve feedback control that allows agents to successfully manipulate and guide an external object;

H2. There is some overlap between the earlier 3D RC task and the BD task due to the requirement for speedy and accurate movement in both environments;

H3. Species evolved in the 3D RC task show increased performance after evolution in the BD environment, and the final perfomance is not significantly different to species evolved from random in the BD environment.

The remainder of the chapter presents an overview of the method used to generate the agents, and the results of the evolutionary and ablative experiments designed to test the above hypotheses. We then present conclusions and a discussion that relates the design of the base system to the observed results.

6.3 Methods

In this section, we describe how the overall objective of implementing a system capable of using an evolutionary algorithm to produce agents able to manipulate objects in a 3D, realistic physics world was achieved. The solution is split into three parts. The first part is the design of the evolutionary problem that the agent species must evolve to solve; the second part documents the abstractions made in the agent's morphology and control architecture that are under the control of the evolutionary algorithm and the third part describes the evolutionary algorithm itself. Finally we describe the data collection scheme we use to collect outputs from the experiments. In all experiments, simultaneous evaluation of evolutionary scenarios was facilitated by the use of GNU Parallel (Tange, 2011).

6.3.1 The physical 3D RC problem

The general problem used in this work is an adaptation of the 3D RC task described in earlier work (Stanton and Channon, 2015), following the same key ideas described and used to various ends in Robinson et al. (2007) and Borg and Channon (2011). The innovation in this work is the addition of a requirement for agents to physically manipulate objects in the environment; in chapter five, only the body of the agent is physically simulated and all environmental interaction is through a two-dimensional, grid-world abstraction. In the original RC task, 2-dimensional agents are able to move between discrete cells in a 20x20 grid world containing hazards (*traps* and *water*) and resources (*stones* and *resource*. Stones can be carried by the agent and placed into water, enabling bridges to be built. Success in this environment is determined by agents' ability to avoid hazards and reach the single *resource* by learning an appropriate action policy given the current state; this includes capturing an element of deliberative planning in order to build bridges in worlds containing an otherwise impassable stretch of water.

This task was adapted in chapter five work to three dimensions (the 3D RC task), making the problem significantly harder. Agents are embodied in a four-legged fixed-morphology physical form that is simulated using a Newtonian rigid-body mechanics system, meaning that physical control (principally, the locomotive and orienting behaviours required for moving between grid cells) must be part of any solution. The agent's position in 3D is projected and quantised to the 2D RC world and any output from the control architecture translates directly into motor control in the 3D environment.

This work introduces the Physical 3D RC (P3D RC) task where the physical problem is extended beyond the agents' control of their bodies, to the wider environment. Solutions to the P3D RC task involve manipulation: in addition to the agents' bodies, a cube representing a *stone* in the world is also physically simulated. Any solution must use physical motor control to manipulate the cube into a configuration that allows the agent to access the resource objective.

As a step toward the P3D RC task, we first investigate simpler problems where agents must simply move blocks around in the world, without the requirement to solve the deliberative component of the RC challenge. This chapter addresses the first of these challenges, where the problem is to move the environmental block as far as possible.

As in earlier work and summarised here, agents have a symmetrical quadruped body plan comprising a torso (dimension $1.0 \times 1.0 \times 0.2$), four upper limbs and four lower limbs (dimensions $0.5 \times 0.2 \times 0.2$ each). Upper limbs are attached to the torso at each lower corner with a 2-axis constraint, limiting the range of motion relative to the torso. Knees connect upper libs to lower limbs, constraining their relative motion to a hinge. Four small sensors are also modelled in the physical environment as fixed appendages to the agent's torso; this is for convenience of updating sensor values based on their position and the sensors have no effect on the physical operation of the agent. The physical simulator used was Open Dynamics Engine (ODE) version 0.13.1, with friction pyramid approximation for contact response ($\mu = 10.0$) between agent and the ground plane, universal error reduction (ERP) of 0.2 and force-mixing (CFM) of 5×10^{-5}. In addition to the agents, a $1 \times 1 \times 1$ block is simulated at the centre of the environment ($\rho = 0.1$). On initialisation, agents are randomly positioned on a circle with radius 5 units from this point.

6.3.2 Agent control

Given the above problem, a strategy to solve it necessarily requires a control architecture that receives sensor data from the environment and produces appropriate motor stimulation to guide the agent through the challenges of the world. We use a bespoke, hybrid neural network (HNN) to this end. The HNN comprises feed-forward networks for saliency calculation from sensor data, a locally-connected, topgraphically-organised *shunting* neural network (Yang and Meng, 2000) for modelling the agent's world, a feed-forward bridge between this model and the motor control parts of the architecture and a series of recurrent leaky integrator networks in the style of Beer's Continuous-Time RNNs (Beer and Gallagher, 1992) that actually produce motor output from the control system. These components we label the *decision network* (DN), the *shunting model* (SM), the *physical network* (PN) and the *pattern generators* (PG). Together, these components are able to successfully solve the 3D RC task, as demonstrated in the previous chapter, as well as in Stanton and Channon (2015).

Since the details of this hybrid architecture are elucidated in the previous chapter, only a summary of the architecture is presented below, along with notes on aspects that have been modified for the present work. See figure 6.1 for a detailed exposition in graphical form.

In chapter five, the DN and SM follow the ideas presented in Robinson et al. (2007) closely. Together and properly configured, they provide a neural-like encoding of a fixed action policy relating current state (position, local objects and carrying state) to action (preferred movement direction, and a pick-up or put-down action). In the first part of the present work, the focus is on physical performance rather than the species' capacities to learn an appropriate state–action policy. As such, we hard-code ι-values (saliency values) for objects in the agents' worlds rather than learn appropriate weights in the DN.

The PN controls the agent's behaviour in the world. The state transition landscape produced by the agent's SM is sampled at four points physically located on the agent's body and these values are used as input to this network.

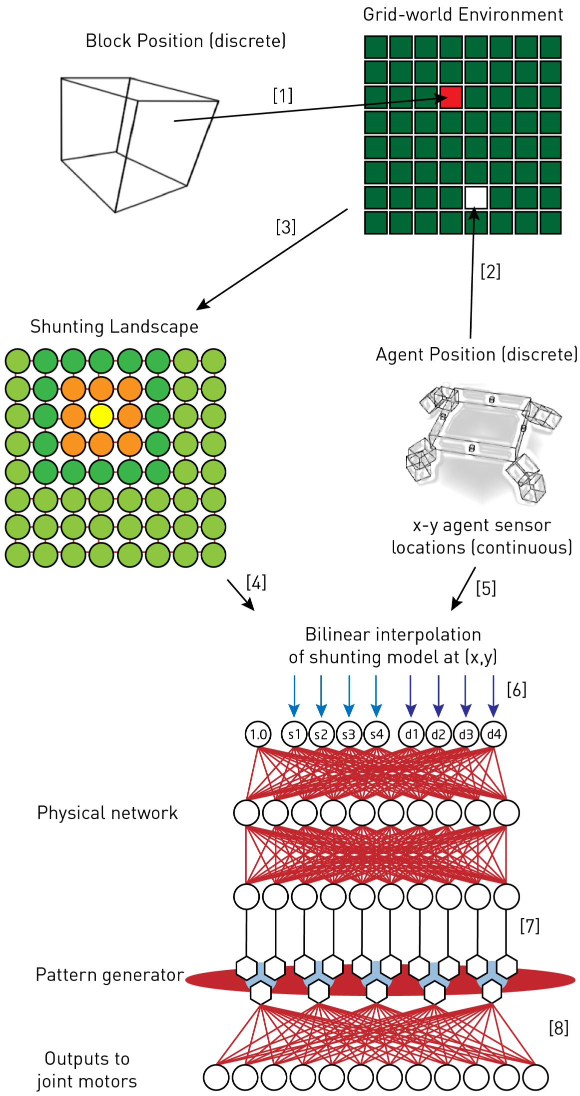

Figure 6.1: Neural architecture. The agent's 3D world, containing the agent and the block, is discretised into a 2D grid (1 and 2); grid locations are given ι-values where salient objects exist and this is used to generate the diffusive shunting landscape (3). Agents sample the landscape (4) at four different continuous positions given by their four sensors (5) by interpolation of values around the sensor location (6). These values pass through a feed-forward network and affect the dynamical trajectories of pattern generators (7) that ultimately output values to effectors via weighted links to joint motors (8). Links shown in red are subject to evolutionary optimisation, both in the pre-evolutionary phase and in the later block task. This includes the red region around the five preset pattern generators whose interneuron weights are also variable: within a single generator preset weights are adapted; across generators weights are initialised at zero but can also move from this value.

Thus, information about desirable state transitions (in this model, directions to move) is available to the PN and can be used by agents to discriminate important features of the preferred state configuration relative to the agent's configuration. The agent's configuration can then be updated to climb the gradient in the state space.

Actual control of the agent's body to achieve this reconfiguration is mediated by the PG network. The network is an array of five three-neuron oscillator circuits, comprising simple leaky-integrator neurons governed by a set of coupled differential equations, modelled after those of Reil and Husbands (2002). The PG network receives input from the PN that perturbs the oscillating cycles which in turn affects the agent's behaviour in the world. The oscillator circuits are a given abstraction in the agents' design, generated by a pre-evolutionary phase that is documented in previous work and summarised in the next section.

Last, outputs from this network are used as target angles for the various joints in the agent's body; actual torques are applied according to a proportional-derivative (PD) equation based on the difference between the current and desired angles at the joint.

6.3.3 Evolutionary algorithm

Pre-evolution

As noted above, populations exploring the block task have been pre-evolved in other environments and also contain specific neural circuits that were produced in an additional, separate environment. These circuits were produced in isolation: three-neuron motifs were evaluated for their capacity to stably generate a 1Hz sinusoidal oscillation in the presence of an input signal and to be quiescent otherwise using an objective function based on the Fourier transform of their output over a 10-second window. The major pre-evolutionary phase involved the simulation of 20 species of agent in the original 3D RC environment. These species progressed through the documented incremental evolutionary phases of food collection, sprinting and hazard avoidance; the evolutionary process was halted before the deliberative part of the incremental challenge. (Specifically, agent populations were allowed 250k 3-individual tournaments; it was found that all 20 species had progressed to the deliberative component by this point.) The 20 species, all capable of tropotactic locomotion, were then installed in the block environment.

Evolutionary Parameters

In all cases the evolutionary algorithm is a three-individual tournament selection-based optimisation process, operating on a population of 150 genomes. Individuals' neuro-controllers are represented as an array of floating-point values. On reproduction, single-point crossover occurs between the two winning individuals in the tournament, and Gaussian mutation is applied to alleles of the resulting child genome with probability $1/l$, where $\mu = 0$ and $\sigma = 1$.

Objective Functions

For the pre-evolution of oscillator circuits, the objective function was the number of non-1Hz frequencies in the frequency domain of a ten-second sample of the output neuron's signal in the input-high state, and the total number of frequencies in the input-low state. During the pre-evolution of gradient-ascending virtual creatures, the objective was as defined in (Stanton and Channon, 2015); agents of high fitness completed many of the incremental stages of the 3D RC task. For the evolved block-pushing task, the objective is to maximise the distance covered by the block in the discrete grid-world.

6.3.4 Data collection

To examine hypothesis H1, we collected observations of agent behaviour, including extracting trajectory data from the highest-scoring individual from the BD task under various deafferent sensory conditions. Deafferentation of control inputs was achieved by systematically disabling sensors, and the agent's progress in a controlled version of the BD task was recorded across a two minute time interval. For each treatment, we examine *approach* and *control*. In both cases the block is positioned at (20,20); for *approach* the agents start far from the block at (5,5), and for *control* they start very close at (18,20). The trajectories followed by agents and block in the two scenarios illuminate the dependence of the gaits on sensory feedback. To examine H2, we used mean evolutionary performance data from the final 1000 tournaments of the 3D RC pre-evolution phase in comparison to the mean score of the same species in the BD task, evaluated for 2 minute and 10 minute periods. To examine H3, the naive BD score before evolution takes place of each individual in each population was measured, in 10 randomly initialised trials. Each trial evaluates the individual for 10 minutes in the BD task. After the evolutionary phase, we repeat the process. We also collected evaluation data for each individual in each of 20 populations over 10 trials of 10 minutes each, after 100k tournament evolution when individuals begin with random genotypes.

6.4 Results

Figures 6.2 and 6.3 show the progress of the two different types of run undertaken over evolutionary time: populations evolved from random starting points and from populations previously successful in the 3D RC task. One extremely high-fitness run was noted in the *random* category; upon visual comparison with other runs, this species is a classic degenerate solution whose strategy is to rapidly vibrate the block to achieve high fitness.

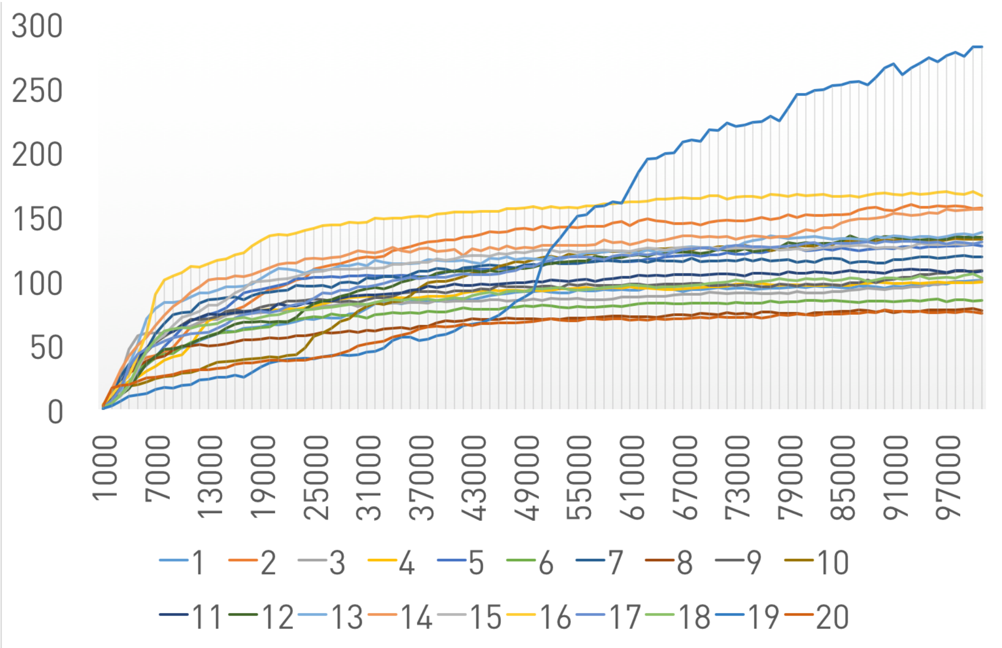

Figure 6.2: Fitness on the BD task (moving average over a 1000-tournament moving window) for evolution from a random (unevolved) population.

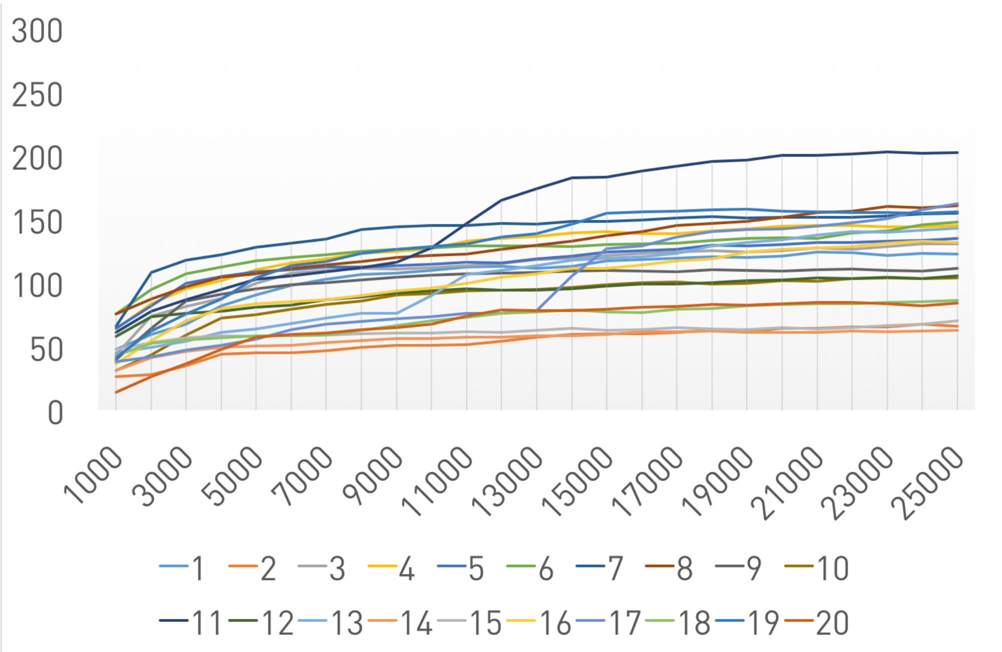

Figure 6.3: Fitness on the BD task (moving average over a 1000-tournament moving window) for evolution from a naive (evolved in 3D RC) population.

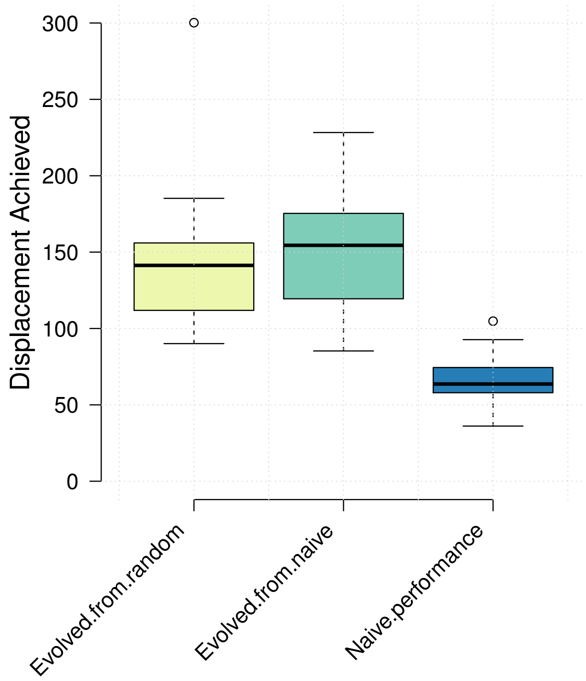

Figure 6.4: Comparison of the best individuals from the naive population, and from populations evolved from the random (unevolved) and naive-evolved populations.

H1: The hybrid architecture is sufficient to achieve feedback control that allows agents to successfully manipulate and guide an external object

Visualisations of agent behaviour can be seen at `https://youtu.be/gZaUvXcdMK8`, and figure 6.5 provides a static view of an agent. The zoopraxiscopic figures, presented in the style of Eadweard Muybridge (Muybridge, 1887), show a time-series of snapshots that illustrate how agents approach the block from a distance (figure 6.6), and manipulate the block in their world (figure 6.7). In order to gauge the importance of the architectural components, agents were observed under impoverished sensory conditions. The behavioural results of sensory deafferentation are presented in figures 6.8 and 6.9. Figure 6.8 shows the planar trajectory followed by agents approaching the block from a distant point under various deafferentation conditions; figure 6.9 shows the response of agents to the same sensory culling in a closer, control scenario.

Figure 6.8 shows an agent that begins at (5,5) and attempts to reach the block at (20,20). The unaltered agent's trajectory is shown in the top left; this agent tends to overshoot its target and then correct by rotating, as the two loops in the

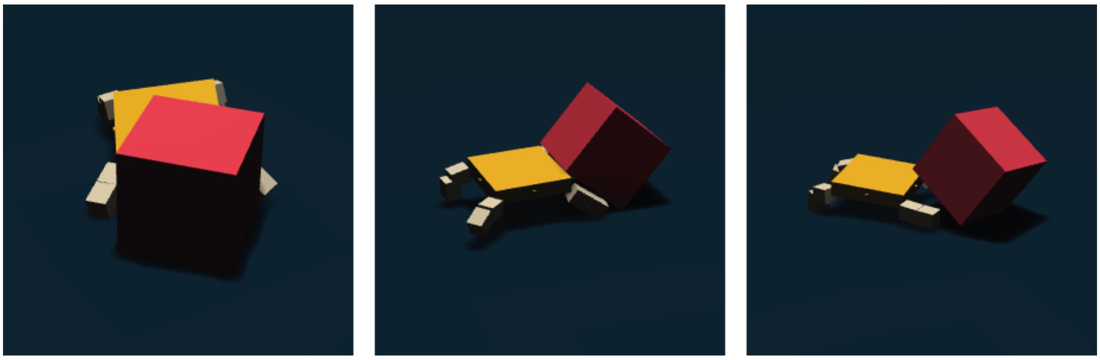

Figure 6.5: Visualisation of a single agent in the block-displacement world. Agent is displaying a low, heavy gait suitable for block pushing.

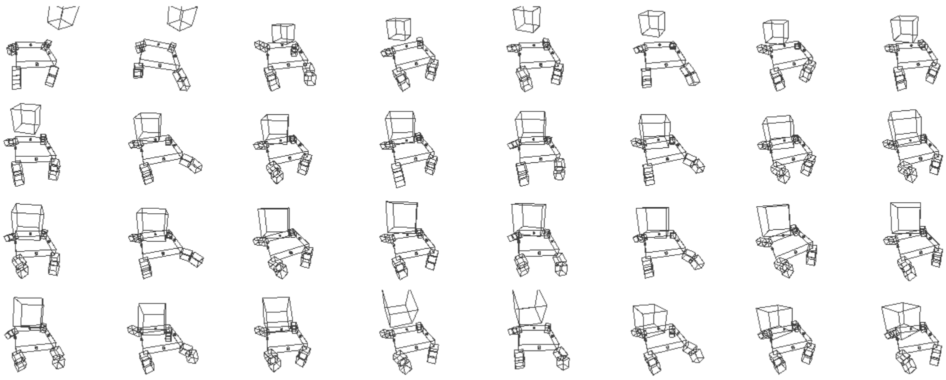

Figure 6.6: Approach gait. The agent is moving toward the block from a distance. All limbs are contributing to the movement.

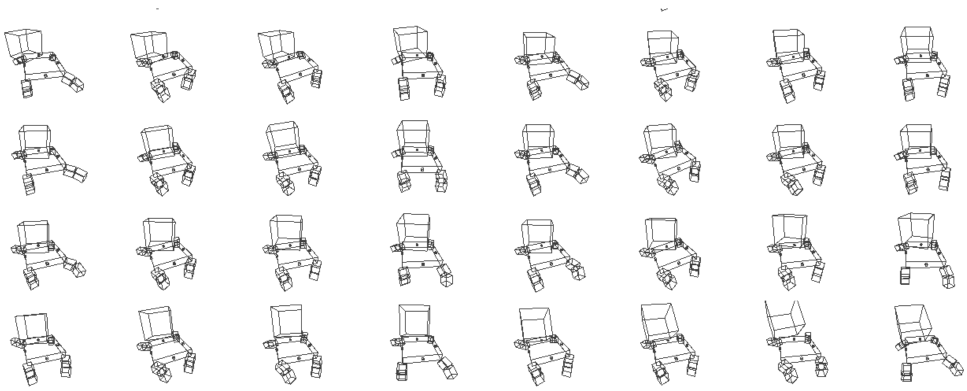

Figure 6.7: Control gait. The agent is pushing forward with its 'back' limbs, maintaining the block between its forelimbs.

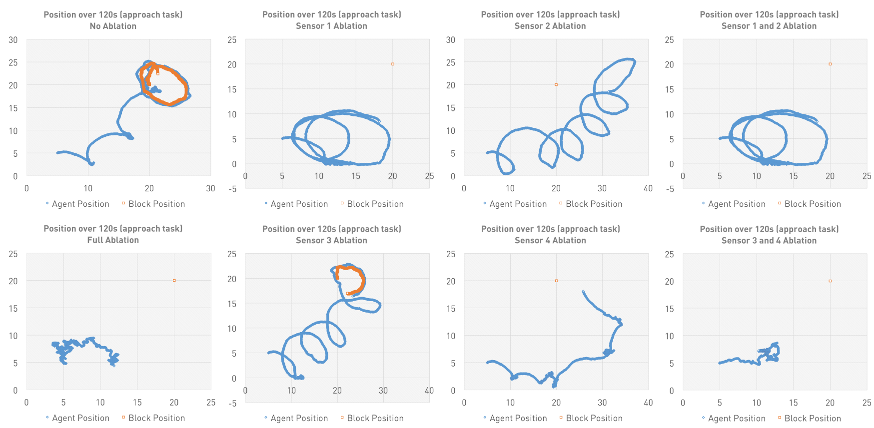

Figure 6.8: Agent-block trajectories of best agent from best overall trained population under various deafferent sensory treatments; approach task.

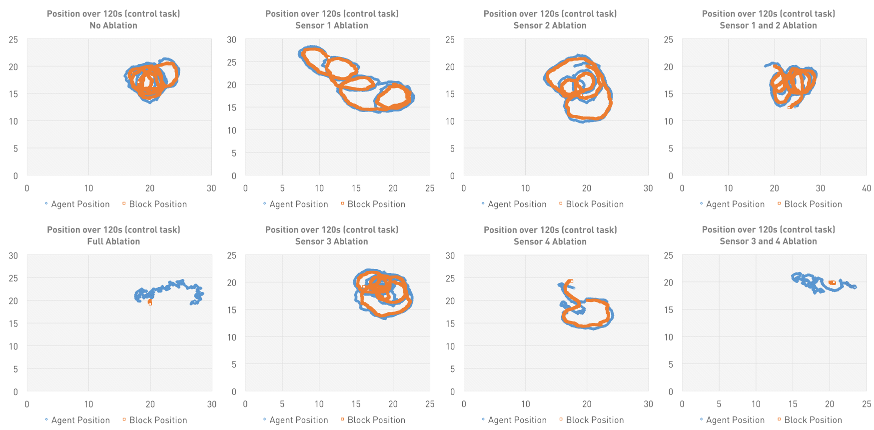

Figure 6.9: Agent-block trajectories of best agent from best overall trained population under various deafferent sensory treatments; control task.

path record. All sensors have some effect on this behaviour although sensor 1 is by far the most pronounced difference in a single cut. In complete deafferentation (bottom left) the agent moves randomly. In contrast, in figure 6.9, the agent begins at (18,20), adjacent to the block at (20,20). The unaltered agent pushes the block in a tight circle to maximise fitness (top left). Sensory deafference does not have a catastrophic effect as in the approach task; all single cuts still maintain block movement although the trajectory is less efficient, as does the dual cut of sensors 1 and 2. Only by cutting sensors 3 and 4 or complete deafferentation was failure to displace the block at all observed. The figures together demonstrate that information from the agents' sensors are being used together to generate reliable gaits for distance approach and block control.

H2: There is some overlap between the 3D RC task and the BD task due to the requirement for speedy and accurate movement in both environments.

A non-parametric correlation analysis was undertaken between the species' relative ranks for mean fitness during the final 1000 tournaments of the 250k-tournament 3D RC pre-evolutionary runs and the mean score on the BD task. Figure 6.10 presents this correlation graphically for both two minute and ten minute evaluation times. In the 10m trial we found a statistically significant although weak correlation ($\rho = 0.38; H_0 \;\; p < 0.05$). The correlation between 3D RC and BD performance in the 2m BD trial is much stronger ($\rho = 0.51; H_0 \;\; p < 0.05$).

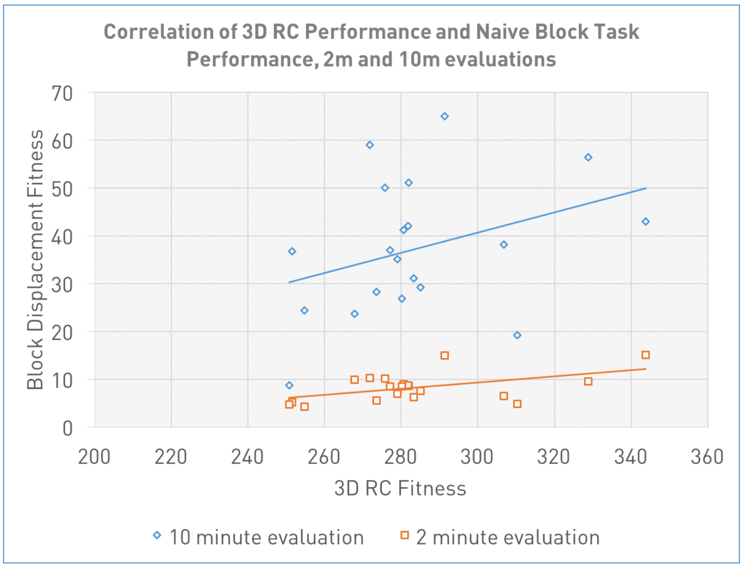

Figure 6.10: Across-species correlation comparing 3D RC performance and BD performance. Outcomes across the two tasks are more correlated when evaluation time is shorter ($\rho = 0.51$), indicating that movement speed is a factor in success in the block task and shared between the two problems. However, a strong gait is required to push the block and this is not selected for in the 3D RC task, hence the lesser correlation in the 10m task ($\rho = 0.38$).

H3: Species evolved in the 3D RC task show increased performance after evolution in the BD environment, and the final perfomance is not significantly different to species evolved from random in the BD environment.

Table 6.1 shows the naive BD score (column 4) compared (column 6) to the evolved score (column 5) for all 20 species. There is a clear improvement in all cases over the 25k tournament evolutionary run: the mean fitness over all naive populations was 37.29, compared to 124.16 in the evolved set ($H_0 \;\; p <$

10^{10}). Figure 6.2 shows progress of runs beginning from random genotypes over evolutionary time (100k tournaments in 1k tournament averages). Figure 6.3 shows the same view of populations beginning from naive genotypes, over 25k tournaments. Both treatments show a levelling off of fitness and there is no significant difference between the best individual evaluation performance of the two starting conditions across all replicates, demonstrated by the proximity of the two treatments' box plots in figure 6.4. A correlation was found between naive score and evolved scores across the 20 species ($\rho = 0.59; H_0 \; p < 0.01$) but no correlation between the naive scores and the magnitude of the change in fitness ($\rho = 0.26; H_0 \; p > 0.1$). Figure 6.11 demonstrates these relationships: δ-fitness is uncorrelated with naive fitness. It was also noted that observed behaviours of the two types of agent were qualitatively indistinguishable, implying that there are potentially few avenues to solve the task, but crucially also showing that the employed decomposition of the task does not hamper the search for this behaviour.

Run	3D RC	Nve2	Nve10	Evl10	δ-f
1	306.78	6.50	38.17	123.07	84.90
2	254.81	4.25	24.42	65.84	41.43
3	283.30	6.22	31.10	130.69	99.58
4	281.91	8.63	51.09	144.80	93.72
5	328.82	9.58	56.36	133.40	77.04
6	271.77	10.29	58.90	147.82	88.92
7	343.80	15.10	42.98	154.68	111.70
8	291.27	14.97	64.90	160.82	95.92
9	284.95	7.58	29.23	110.74	81.52
10	267.79	9.87	23.66	105.04	81.38
11	275.82	10.09	49.98	205.49	155.52
12	251.57	5.21	36.72	103.01	66.29
13	280.65	8.93	41.19	144.42	103.23
14	310.35	4.91	19.16	60.29	41.13
15	281.82	8.71	41.98	69.98	28.00
16	273.59	5.52	28.29	130.22	101.93
71	277.03	8.53	37.01	164.83	127.82
18	278.92	6.99	35.11	87.02	51.92
19	280.14	8.54	26.86	155.60	128.74
20	250.72	4.75	8.71	85.42	76.71

Table 6.1: Results table showing all 20 species' performance in 3D RC, naive block (2m and 10m evaluation) and evolved block tasks, and difference between naive and evolved. There is a relationship between prior and post performance, but not between prior performance and δ-fitness.

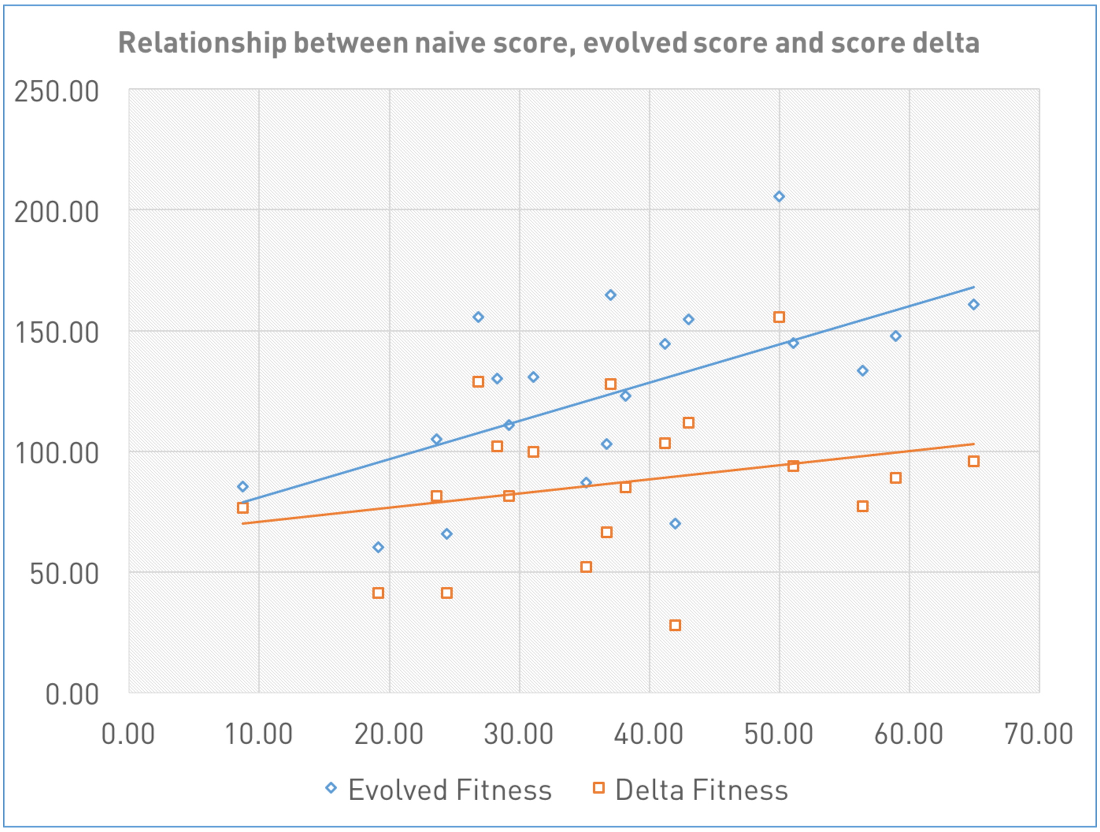

Figure 6.11: Correlation of evolved fitness with naive fitness ($\rho = 0.59$), and delta fitness with naive fitness ($\rho = 0.26$).

6.5 Conclusions and Discussion

We have shown that feedback motor control in evolved agents is possible with the given architecture, and that the architecture is flexible enough to support and adapt to a variety of evolutionary scenarios presented sequentially. This demonstrates that the platform has the potential to support environments that require even more sensorimotor control and is a reasonable starting point from where physical complexities can be added into the 3D RC task, eventually approaching a full physical model of the problem. Observations of the agents' behaviour gained through 3D visualisation have revealed a rich variety of evolved strategies for solving the problem. Different classes of gait for approaching and manipulating the block appear due to the genetic heritage of species, and it is clear that low, heavy gaits work best for pushing the object in the BD task. From the deafferentation studies it can be seen that these gaits are not self-generating, blind gaits that simply aim the agent to the target location, but are more complex aggregates of sensory data that depend on the agents position relative to the block in order to successfully achieve increased displacement.

When we consider the two evolutionary scenarios, 3D RC and BD, we found some overlap between the two problems. A strong correlation was observed between performance in the BD challenge before evolution in a two-minute evaluation, and performance at the end of the 3D RC task, indicating that some components of both challenges contribute similarly to relative agent fitness. This is likely to be the speed and directness of movement in the world which has a greater effect in a smaller evaluation period. As the evaluation period grows larger, this correlation decreases indicating that the block-pushing dimension of fitness in this scenario is not well captured in the 3D RC task and ultimately is the most important component. (It was also observed by measuring the time taken by agents to reach the block that most naive species sacrifice movement speed for block pushing capability during evolution, and this aspect should be investigated more thoroughly to determine whether this is an artifact or a consistent trend.) We showed that performance from either starting point (3D RC or unevolved genotypes) is comparable, demonstrating that an incremental approach incorporating both types of environment is possible in principle. One extremely high-fitness run in the *random* category was noted; upon visual inspection this species is a classic degenerate solution whose strategy is to rapidly vibrate the block to achieve high fitness. It is possible that more complex environments (such as 3D RC) prevent this kind of trivial solution by requiring a richer agent–environment interface. The results comparing BD fitness before and after evolution demonstrate that whilst naive performance is an indicator of final performance, it is not an indicator of how much any particular species will improve. There is a risk that incrementally presenting new environments to only the most successful species could exclude good general solutions, a problem potentially mitigated by heterogeneous presentation of multiple environments.

Further work

The results of this chapter support the conclusion that the P3D RC task is a viable next-step from the 3D RC task presented in chapter 5, towards an embodied virtual creature world where animats display complicated chains of discrete behaviours. It has been demonstrated in this chapter that the block manipulation task does not conflict with the evolutionary trajectories of agents from the 3D RC world, and thus the integration of the BD task into the incremental evolutionary scheme is not precluded from the start.

Ongoing work is toward the P3D RC task: a physically-embodied deliberative river crossing problem. The next step is to consider not just displacement but also positioning of the block using the shunting landscape. This is likely to demand significant revision of the underlying control architecture to incorporate reasoning about relative positioning. Additionally, the question of whether specific types of solutions in the 3D RC world have specific performance profiles in the BD world could be addressed by examining in detail whether some species always slow down and some always speed up. Additionally, it is possible that evolved morphology could significantly contribute to physical manipulation behaviours.

CHAPTER 7

Conclusions

In chapter 7, I conclude the dissertation with an overview of the work undertaken, the discoveries and technical achievements made, and an exposition of the general principles that I believe are necessary for scientific and technical progress to continue in the field of evolved intelligent behaviours in 3D virtual creatures.

7.1 Summary

A-life is a relatively young field of study that brings together under a common banner research in the fundamentals of life, the practicalities of building adaptive technology and the new descriptive and predictive paradigms of complex systems science.

The automatic development of physically-realistic autonomous agents in closed simulated worlds (commonly known as evolved 3D virtual creatures) has been a long-term research ambition in a-life. Since their beginnings with the work of Sims, virtual creatures have inspired and entertained through the visceral feelings of animacy that they generate in observers. In chapters 2 and 3, the review of the literature demonstrated how the research has offered a platform for exploring how to produce intelligent behaviour using artificial evolution, and how the natural world developed the behavioural complexity that it has. However, it was also shown that the rate of development of 3D virtual creatures toward more advanced, more diverse behaviours has been disappointingly slow since early successes in the 1990s.

A comprehensive theory of open-ended evolution (OEE), the continued and ever-expanding production of diversity in an evolving ecology, is another major research program within a-life, part of which is the perhaps closer sub-goal of its practical demonstration in an artificial system. Efforts have been made both to define OEE theoretically and to capture it using statistical metrics applied to artificial systems. However, it is recognised by many authors that an intuitive visualisation of OEE in action in an artificial context is a necessary step toward validating statistical mechanisms that constitute hallmarks of the phenomenon (Taylor et al., 2016). By producing a system in which we can use the innate biases that help us as natural organisms to understand complex behaviours in nature, we can observe open-ended evolution taking place *in silico* and verify that these theories or metrics have meaningful correspondence to the natural principle. To date, all putative OEE demonstrations are opaque and not intuitively valid.

Given these motivations, the primary body of work presented in this dissertation concerns the construction and exploration of 3D virtual creature simulations, in order to contribute both to the development of research in 3D virtual creature technology by advancing the state-of-the-art beyond simple locomotion, as well as to begin to examine its potential for use as a platform for open ended evolutionary systems. These objectives required exploration of evolutionary systems that may begin to deliver complex behaviours unforeseen at design time, as well as the specification of a body-brain-environment system that is compatible with the evolutionary substrate and the design aims of the project. In chapter 4, one aspect of the evolutionary systems necessary to achieve generalised complex behaviours was explored. In chapter 5, a proof-of-concept architecture was presented that takes a step forward in this technology, with the evolutionary system informed by the work in chapter 4. In chapter 6, the ideas were placed in a more physically realistic environment and their potential for further scaling was evaluated.

7.2 Contributions

The literature review established that research in 3D virtual creatures has focused almost exclusively on producing systems that exhibit single behaviours, often focusing either on wider issues in evolutionary design and using EVCs as a test-bed, or else examining how more generalised, bio-inspired approaches can achieve comparable results to the more constrained and directed earlier attempts. The work reported on in the dissertation thus makes a step forward in the 3D VC research programme and contributes to wider knowledge in complex evolutionary design problems, with the following specific contributions (in the order they are presented in the text):

1. An overview of the field of Artificial Life and the grand ambitions of artificial open-ended evolutionary activity and 3D virtual creatures research;

2. a detailed description of the current state-of-the-art in virtual agents and 3D virtual creatures;

3. an investigation into various strategies of environmental complexification applied as a heuristic in evolutionary algorithms, in order to achieve generalisation across different variations of a task, resulting in the first original contribution. It is demonstrated that evolutionary forgetting and over-fitting are real problems in evolutionary adaptation, that much existing research in applied EAs does not consider this problem systematically, and that the problems can be countered with appropriate heterogeneous complexification strategies;

4. the specification of a 3D agent system able to display increasing diverse and complex behaviours resulting from random initialisation of an evolutionary process, offering an original design for controller for such a system and a potential model to use for further development of complex agent-environment interactions in 3D, physically-realistic spaces; and

5. a further analysis of this brain-body-environment system adapted for physical manipulation, as a step toward a full 3D implementation able to both demonstrate increased sensori-motor embedding and situation and also express the complexities of an open-ended evolutionary system in terms that allow intuitive validation of hitherto statistical-only metrics of such activity.

7.3 Key Findings, Technologies and Implications

7.3.1 Heterogeneous environmental complexification

The existing literature in evolutionary robotics and evolved 3D virtual creatures uses a range of techniques to accelerate the exposure of evolving species to multiple facets of their environments or tasks. This includes ideas such as the decomposition of tasks into subtasks, progressively increasing some measure of difficulty, and incrementally adding constraints into the system. These ideas share a general theme of complexification of the environment and have been shown in multiple areas to be beneficial, compared to naive environments which expose agents to the whole task from the outset.

It is not clear how this complexification actually affects the evolutionary system. What factors are important in successfully complexifying environments, and what are the underlying mechanisms that govern the operation?

In chapter 4, an experiment was conducted using 3D virtual creatures and an obstacle traversal task to determine any differences between possible strategies of complexification. The height of the obstacle was used as a proxy for different parametrisations of a more general task, namely successfully climbing the obstacle at all heights. The hypothesis that the order of presentation of different heights during evolutionary time would have a bearing on the success of adapting to the general task was explored. It was found that more traditional, intuitive presentations of the sort that are often strategies in virtual creatures research such as direct presentation of the hardest task, random presentation of wall heights and a linearly-increasing wall height are much less efficient that strategies that adaptively increase based on current performance and, more importantly, regularly revisit parametrisations already explored at earlier stages in the task.

The central argument of this chapter is that these strategies solve the twin problems of overfitting to recent parametrisations (so-called *evolutionary forgetting*) and being unable to find any improvement on current performance through random mutation (*loss-of-gradient*). This result has implications for any further research that attempts problem decomposition and environmental complexification using such strategies: the exploratory potential of a species (in terms of its population size, mutation rate and representation of solution in the problem space) has a direct bearing on the frequency of re-presentation of earlier parts of the problem that is necessary to achieve high generalised performance on the whole task. In addition, when using an extrinsic fitness function, an increase in the rate of change of the task presentations should be proportional to the derivative of fitness over time to avoid loss of gradient whilst still driving the population toward less-frequently presented parametrisations.

It was also demonstrated that there are limit cases in these heterogeneous, oscillating strategies that approach the results seen in the random strategy (for very high-frequency oscillations) and the linear strategy (for very low-frequency oscillations). This result is suggestive of a deeper connection between these presentation strategies and the underlying evolutionary system and thus implies a more general result for evolutionary algorithms beyond their use in 3D virtual creatures.

This generality is hinted at in many other works. The approach would seem suited to the types of problems where either a clear decomposition across a range of parameters can be made, or where the principles of the approach are implicit in the configuration of the evolutionary environment. In the first case, the canonical example of which is the task in chapter 4, a number of different extensions or modifications to the task could be used to demonstrate the generality of the approach. For example, adapting the task to encourage the robot to turn on the spot according to a signal. Training would proceed by varying the turning signals according to the heterogeneous approach, and generalisation assessed in terms of the success over a range of turning signals. In the second case, this is demonstrated by the success of the design of the system presented in chapter 5: the 3D RC task encodes the heterogeneous strategy within the individual levels, which permit the consolidation of progress and yet push towards a proximal zone of development where agents can out-compete each other according to Darwinian principles.

7.3.2 Hybrid control architectures for 3D virtual creatures

The history of research in 3D evolved virtual creatures shows a consistent trend since the first demonstration of such a system by Karl Sims in 1994: increasing generality of control architectures and evolutionary set-ups but with an enduring focus on single behaviours such as locomotion. The central aim of chapter 5 is to develop a 3D evolved virtual creature system that is able to demonstrate multiple behaviours from a single starting point. The technology developed to achieve this takes into account the findings of chapter 4 by structuring the environmental challenges faced by evolving species in a way that respects the need to revisit and refresh early successes whilst maintaining evolutionary pressure. Using this principle, along with current research in robot control and other virtual creatures work, continues a line of thinking that bridges the gap between representational and anti-representational AI. The system is the first such demonstration of multiple behaviours and exhibits an intuitive realism that draws the viewer into the virtual world.

The working implementation comprises a variety of technology and ideas in autonomous systems and evolutionary methods. A hybrid neural network composed of problem-solving (deliberative) and physical control (reactive) components allows an action policy to be built that provides a direction for agent behaviour. The deliberative network takes the current state of the agent in problem-solving space as input and delivers a measure of salience for all objects in the agents' environments. The saliency is propagated across a *shunting network*, a quantised, diffusive model of the agents environment. The resultant landscape is representative of the next step of a plan which agents can follow to complete a longer-term goal. The behaviours that implement this plan in the agents' physical environment (by activating joint motors to move limbs) are built of a collection of oscillatory circuits, implemented as as a network of leaky-integrator units (so-called continuous-time recurrent neurons.) This network responds to changes in the shunting landscape by sampling the landscape at four spatial co-ordinates on the agent's body and propagating values through a feed-forward network. All of the parameters for the model are found by stochastic gradient descent in the evolutionary paradigm, based on an error signal from a fitness function that determines how much progress an agent has made through an incremental evolutionary environment. The environment is an idealisation of a planning-problem, called the 3D River Crossing (3D RC) task. In this problem, agents must learn to avoid environment hazards of traps and water, and collect stones to place on water in order to build a bridge to reach a resource. Informed by the work of chapter 4, this heterogeneous, incremental presentation provides evolving species with evolutionary pressure whilst maintaining performance on other parts of the task. A population of 150 agents in a species compete through 3-individual tournament-selection over an evolutionary history of 1 million tournaments to solve the task. Later individuals can be seen to pass through simple locomotion and avoidance behaviours to discrete behavioural modalities including turning, walking, and galloping and, at the same time, performance on the river-crossing problem increases until agents can solve a river requiring the placement of four stones in a row.

Aside from being the first such demonstration, the technology also has implications for research in the deeper a-life objective of demonstrating open-ended evolution and validating statistical measures of unbounded activity. It can be seen from this work that 3D virtual creatures could provide a rich substrate upon which evolutionary processes can operate and, through the intuitive nature of the 3D space in which they're embedded, these processes can be laid bare for human observers to grasp intuitively.

7.3.3 Adapting hybrid architectures to physical-manipulation tasks

This work adds further weight to the conclusion of chapter 5, that the hybrid neural architecture, when combined with a heterogeneous environmental complexification in the evolutionary system, has potential to continue to display increasingly complex, lifelike behaviours. Specifically addressing the question of whether motor control for manipulation can be evolved in the same platform, the conclusion that it can offers a further, necessary step to increasing the generality of virtual creature platforms, which in turn opens the possibility of producing increasingly complex behaviours.

7.4 Limitations and Further Work

Limitations of the work presented in chapter 4 are primarily around the generalisation of the result to wider classes of evolutionary algorithm and different problem domains. While the experiments conducted with the wavelength of oscillation in the heterogeneous strategies suggest a fundamental connection, further work is required to verify this result in other evolutionary problems. In addition, the effects of specific parameters of the evolutionary algorithm–for example population size and mutation rate–on the result are unknown. This further work would usefully tease out any interplay between these parameters and the optimal configuration of presentation strategies. More generally, the large number of assumptions built into the various models constitute a limitation that should be addressed with future work, a challenge that will continue to inform our understanding of the interplay of mechanisms that produce intelligent behaviour in artifacts. The assumption that EVCs could provide a mechanism to observe OEE in action should also be tested.

Ongoing work from this thesis concerns increasing the poverty of the initial conditions whilst maintaining the quality of the final results, following the grand a-life ambition outlined at various earlier points. The removal of pre-evolved oscillation, the shunting model representation and the omniscient world map are all candidates for this impoverishment. The inclusion of vision, more complex environments offering a greater variety of challenges, multiple species and very large scale environments could also offer promising research challenges.

Bibliography and References

Adami, C. and Brown, C. T. (1994). Evolutionary learning in the 2d artificial life systems 'avida'. In Brooks, R. and Maes, P., editors, *Proc. Artificial Life IV*, pages 377–381.

Agre, P. E. (1988). *The dynamic structure of everyday life*. PhD thesis, MIT Department of Electrical Engineering and Computer Science.

Allen, B. F. and Faloutsos, P. (2009). Evolved controllers for simulated locomotion. In *Motion in Games*, pages 219–230. Springer.

Annunziato, M. and Pierucci, P. (2002). Relazioni emergenti: experiments with the art of emergence. *Leonardo*, 35(2):147–152.

Arita, T., Joachimczak, M., Ito, T., Asakura, A., and Suzuki, R. (2016). Alife approach to eco-evo-devo using evolution of virtual creatures. *Artificial Life and Robotics*, pages 1–8.

Arnold, E. (2013). Experiments and simulations: do they fuse? In Durán, J. M. and Arnold, E., editors, *Computer simulations and the changing face of scientific experimentation*. Cambridge Scholars Publishing.

Ashby, W. R. (1948). Design for a brain. *Electronic Engineering*, 20:379–383.

Auerbach, J. and Bongard, J. C. (2009). How robot morphology and training order affect the learning of multiple behaviors. In *Evolutionary Computation, 2009. CEC'09. IEEE Congress on*, pages 39–46. IEEE.

Auerbach, J. E. and Bongard, J. C. (2010). Dynamic resolution in the co-evolution of morphology and control. In *Artificial Life XII: Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems*, pages 451–458. MIT Press.

Auerbach, J. E. and Bongard, J. C. (2012). On the relationship between environmental and morphological complexity in evolved robots. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, pages 521–528. ACM.

Auerbach, J. E. and Bongard, J. C. (2014). Environmental influence on the evolution of morphological complexity in machines. *PLoS computational biology*, 10(1):e1003399.

Ayala, L. and Forero-Romero, J. E. (2013). Computer simulations in a cosmological context. In Durán, J. M. and Arnold, E., editors, *Computer simulations and the changing face of scientific experimentation*. Cambridge Scholars Publishing.

Azarbadegan, A., Broz, F., and Nehaniv, C. L. (2011). Evolving sims's creatures for bipedal gait. In *Artificial Life (ALIFE), 2011 IEEE Symposium on*, pages 218–224. IEEE.

Back, T., Hammel, U., and Schwefel, H.-P. (1997). Evolutionary computation: Comments on the history and current state. *IEEE Transactions on Evolutionary Computation*, 1(1):3–17.

Bainbridge, C. J. (2010). *Digital control networks for virtual creatures*. PhD thesis, University of Edinburgh.

Bargh, J. A. and Chartrand, T. L. (1999). The unbearable automaticity of being. *American psychologist*, 54(7):462.

Barlow, G., Oh, C., and Grant, E. (2004). Incremental evolution of autonomous controllers for unmanned aerial vehicles using multi-objective genetic programming. In *Cybernetics and Intelligent Systems, 2004 IEEE Conference on*, volume 2, pages 689–694.

Bedau, M. A., McCaskill, J. S., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., Ikegami, T., Kaneko, K., and Ray, T. S. (2000). Open problems in artificial life. *Artificial life*, 6(4):363–376.

Bedau, M. A. and Packard, N. H. (1992). Measurement of evolutionary activity, teleology and life. In *Artificial Life II*, pages 431–461. Addison-Wesley.

Beer, R. D., Chiel, H. J., and Sterling, L. S. (1990). A biological perspective on autonomous agent design. *Robotics and Autonomous Systems*, 6(1-2):169 – 186.

Beer, R. D. and Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. In *Adaptive Behavior 1*, pages 91–122.

Bentley, P. (1999). Aspects of evolutionary design by computers. In *Advances in Soft Computing*, pages 99–118. Springer.

Boden, M. A. (2000). Autopoiesis and life. *Cognitive Science Quarterly*, 1(1):117–145.

Boeing, A. (2008). Morphology independent dynamic locomotion control for virtual characters. In *Computational Intelligence and Games, 2008. CIG'08. IEEE Symposium On*, pages 283–289. IEEE.

Bongard, J. (2008). Behavior chaining-incremental behavior integration for evolutionary robotics. In *Proceedings of ALIFE XI, the 11th International Conference on the Synthesis and Simulation of Living Systems*, pages 64–71.

Bongard, J. (2010). The utility of evolving simulated robot morphology increases with task complexity for object manipulation. *Artificial life*, 16(3):201–223.

Bongard, J. C. (2011). Morphological and environmental scaffolding synergize when evolving robot controllers: artificial life/robotics/evolvable hardware. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 179–186. ACM.

Bongard, J. C. (2013). Evolutionary robotics. *Communications of the ACM*, 56(8):74–83.

Booker, L. B. (1982). *Intelligent behavior as an adaptation to the task environment*. PhD thesis, University of Michigan.

Borg, J. and Channon, A. (2011). Discovering and maintaining behaviours inaccessible to incremental genetic evolution through transcription errors and cultural transmission. In *Advances in Artificial Life, ECAL 2011: Proceedings of the Eleventh European Conference on the Synthesis and Simulation of Living Systems*, pages 101–108.

Box, G. E. (1957). Evolutionary operation: A method for increasing industrial productivity. *Applied Statistics*, pages 81–101.

Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. MIT Press/Bradford Books.

Bremermann, H. J. (1962). Optimization through evolution and recombination. *Self-Organizing Systems*.

Brooks, R. (1986). A robust layered control system for a mobile robot. *IEEE journal on robotics and automation*, 2(1):14–23.

Brooks, R. (1989). A robot that walks; emergent behaviors from a carefully evolved network? MIT AI Lab Memo No 1126.

Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence*, 47:139–159.

Brooks, R. (2000). Artificial life: From robot dreams to reality. *Nature*, 406(6799):945–947.

Brooks, R. A. (1990). Elephants don't play chess. *Robotics and autonomous systems*, 6(1):3–15.

Buendia, A. and Heudin, J.-C. (2000). Towards digital creatures in real-time 3d games. In *Virtual Worlds*, pages 44–53. Springer.

Cantú-Paz, E. and Kamath, C. (2005). An empirical comparison of combinations of evolutionary algorithms and neural networks for classification problems. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 35(5):915–927.

Channon, A. (2001). *Evolutionary Emergence: The Struggle for Existence in Artificial Biota*. PhD thesis, Department of Electronics and Computer Science, University of Southampton.

Channon, A. and Damper, R. I. (1998a). Evolving novel behaviors via natural selection. In Adami, C., Belew, R., Kitano, H., and Taylor, C., editors, *Proceedings of Artificial Life VI*, pages 384–388.

Channon, A. D. and Damper, R. I. (1998b). Perpetuating evolutionary emergence. In *Proceedings of SAB98*, pages 9–16. MIT Press.

Charlesworth, W. R. (1976). Human intelligence as adaptation: an ethological approach. In Resnick, L. B., editor, *The nature of intelligence*. Lawrence Erlbaum.

Chaumont, N., Egli, R., and Adami, C. (2007). Evolving virtual creatures and catapults. *Artificial Life*, 13(2):139–157.

Cheney, N., MacCurdy, R., Clune, J., and Lipson, H. (2013). Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, pages 167–174. ACM.

Chiel, H. J. and Beer, R. D. (1997). The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in neurosciences*, 20(12):553–557.

Christensen, A. L. and Dorigo, M. (2006). Incremental evolution of robot controllers for a highly integrated task. In Nolfi, S., Baldassarre, G., and Calabretta, R., editors, *Proceedings of the Ninth International Conference on the Simulation of Adaptive Behavior (SAB06)*, pages 473–484.

Clark, A. (1998). *Being there: Putting brain, body, and world together again*. MIT press.

Clark, A. J., Moore, J. M., Wang, J., Tan, X., and McKinley, P. K. (2012). Evolutionary design and experimental validation of a flexible caudal fin for robotic fish. In *Artificial Life*, volume 13, pages 325–332.

Cliff, D. (1991). Neural networks and visual behaviour: Flies; panned eyes; and statistics. In Steels, L. and Smith, B., editors, *Proceedings of the Eighth Conference of the Society for Artificial Intelligence and the Simulation of Behaviour (AISB91)*.

Corucci, F., Cheney, N., Lipson, H., Laschi, C., and Bongard, J. (2016). Material properties affect evolution's ability to exploit morphological computation in growing soft-bodied creatures. In *Proceedings of ALIFE XV, the 15th International Conference on the Simulation and Synthesis of Living Systems*.

Damer, B., Marcelo, K., and Revi, F. (1998). Nerve garden: A public terrarium in cyberspace. In *Virtual Worlds*, pages 177–185. Springer.

Damer, B., Marcelo, K., and Revi, F. (2003). Nerve garden: germinating biological metaphors in net-based virtual worlds. *Kybernetes*, 32(1/2):174–183.

de Castro, E. C. and Gudwin, R. R. (2010). *An Episodic Memory Implementation for a Virtual Creature*, pages 393–406. Springer, Berlin, Heidelberg.

Dennett, D. C. (1981). *Brainstorms: Philosophical essays on mind and psychology*. MIT press.

Dobzhansky, T. (1973). Nothing in biology makes sense except in the light of evolution. *The American Biology Teacher*, 35(3):125–129.

Druhan, K. (2004). Genetic programming for evolution of virtual creatures. *DCSI'04*, page 21.

Dubbin, G. A. and Stanley, K. O. (2010). Learning to dance through interactive evolution. In *Applications of Evolutionary Computation*, pages 331–340. Springer.

Duthen, Y., Luga, H., Lassabe, N., Cussat-Blanc, S., Breton, T., and Pascalie, J. (2010). An introduction to the bio-logic of artificial creatures. In *Intelligent Computer Graphics 2010*, pages 1–23. Springer.

Eccles, J. S., Dickerson, J. A., and Shao, J. (2000). Evolving a virtual ecosystem with genetic algorithms. In *Evolutionary Computation, 2000. Proceedings of the 2000 Congress on*, volume 1, pages 753–760. IEEE.

Eiben, A. E. (2014). Grand challenges for evolutionary robotics. *Frontiers in Robotics and AI*, 1:4.

Eisenmann, J., Schroeder, B., Lewis, M., and Parent, R. (2011). Creating choreography with interactive evolutionary algorithms. In *Applications of Evolutionary Computation*, pages 293–302. Springer.

Faíña, A., Souto, D., Orjales, F., Bellas, F., and Duro, R. (2014). Towards feasible virtual creatures by using modular robots. In *Modular and Swarm Systems. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

Featherstone, R. (1987). *Robot Dynamics Algorithms*. Kluwer.

Floreano, D. and Keller, L. (2010). Evolution of adaptive behaviour in robots by means of darwinian selection. *PLoS biology*, 8(1):e1000292.

Floreano, D. and Mondada, F. (1994). Automatic creation of an autonomous agent: Genetic evolution of a neural network driven robot. In *Proceedings of the third international conference on Simulation of adaptive behavior: From Animals to Animats 3*, pages 421–430.

Floreano, D. and Nolfi, S. (1997). God save the red queen! competition in co-evolutionary robotics. In *Genetic Programming 1997: Proceedings of the Second Annual Conference*, pages 398–406. Morgan Kaufmann.

Fogel, L. (1962). Autonomous automata. *Industrial Research*, 4:14–19.

Fontijne, D. (2000). Rigid body simulation and evolution of virtual creatures. Master's thesis, University of Amsterdam.

French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135.

Friedberg, R. M. (1958). A learning machine: Part i. *IBM Journal of Research and Development*, 2(1):2–13.

Froese, T. et al. (2010). Life after ashby: ultrastability and the autopoietic foundations of biological autonomy. *Cybernetics & Human Knowing*, 17(4):7–49.

Funes, P. and Pollack, J. (1998). Evolutionary body building: Adaptive physical designs for robots. *Artificial Life*, 4(4):337–357.

Funes, P. and Pollack, J. (1999). Computer evolution of buildable objects. *Evolutionary design by computers*, 1:387–403.

Furukawa, M., Morinaga, M., Watanbe, M., Suzuki, I., and Yamamoto, M. (2010). Artificial flying creature by flapping. In *Systems Man and Cybernetics (SMC), 2010 IEEE International Conference on*, pages 624–629. IEEE.

Gardner, M. (1970). The fantastic combinations of John Conway's new solitaire game "life". *Scientific American*, 223:120–123.

Geijtenbeek, T., van de Panne, M., and van der Stappen, A. F. (2013). Flexible muscle-based locomotion for bipedal creatures. *ACM Transactions on Graphics (TOG)*, 32(6):206.

Giraud-Carrier, C. (2000). A note on the utility of incremental learning. *AI Communications*, 13(4):215–223.

Glasius, R., Komoda, A., and Gielen, S. (1995). Neural network dynamics for path planning and obstacle avoidance. *Neural Networks*, 8(1):125–133.

Glette, K. and Hovin, M. (2010). Evolution of artificial muscle-based robotic locomotion in physx. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 1114–1119. IEEE.

Gomez, F. and Miikkulainen, R. (1997). Incremental evolution of complex general behaviour. *Adaptive Behaviour*, 5:317–342.

Griffith, V. and Yaeger, L. (2006). Ideal free distribution in agents with evolved neural architectures. In Rocha, L., editor, *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*.

Groß, R., Magnenat, S., Küchler, L., Massaras, V., Bonani, M., and Mondada, F. (2011). Towards an autonomous evolution of non-biological physical organisms. In *Advances in Artificial Life. Darwin Meets von Neumann*, pages 173–180. Springer.

Grossberg, S. (1988). Nonlinear neural networks: Principles, mechanisms, and architectures. *Neural networks*, 1(1):17–61.

Grzeszczuk, R. and Terzopoulos, D. (1995). Automated learning of muscle-actuated locomotion through control abstraction. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 63–70. ACM.

Guajardo, J. A., Weber, R., and Miranda, J. (2010). A model updating strategy for predicting time series with seasonal patterns. *Applied Soft Computing*, 10(1):276–283.

Haflidason, S. (2010). *On the significance of the permutation problem in neuroevolution*. PhD thesis, University of Manchester.

Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1-3):335–346.

Harvey, I. (1992a). Species adaption genetic algorithms: A basis for a continuing saga. Technical Report CSRP-221, University of Sussex School of Cognitive and Computing Sciences.

Harvey, I. (1992b). Untimed and misrepresented: Connectionism and the computer metaphor. Technical Report CSRP-245, University of Sussex School of Cognitive and Computing Sciences.

Harvey, I. (1997). Artificial evolution for real problems. In Gomi, T., editor, *Proceedings of the 5th International Symposium on Evolutionary Robotics*.

Harvey, I. (2001). Artificial evolution: A continuing saga. In *Proceedings of the International Symposium on Evolutionary Robotics From Intelligent Robotics to Artificial Life*, ER '01, pages 94–109, London, UK, UK. Springer-Verlag.

Haykin, S. (1994). *Neural Networks: A Comprehensive Foundation*. Prentice Hall.

Hayles, N. K. (1999). Simulating narratives: What virtual creatures can teach us. *Critical Inquiry*, pages 1–26.

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological approach*. John Wiley & Sons.

Heinen, M. R. and Osório, F. S. (2009). Evolving morphologies and gaits of physically realistic simulated robots. In *Proceedings of the 2009 ACM symposium on Applied Computing*, pages 1161–1165. ACM.

Heleno, P. and dos Santos, M. P. (1998). Artificial animals in virtual ecosystems. *Computer networks and ISDN systems*, 30(20):1923–1932.

Helmreich, S., Roosth, S., and Friedner, M. (2015). *Sounding the Limits of Life: Essays in the Anthropology of Biology and Beyond*. Princeton Studies in Culture and Technology. Princeton University Press.

Hiller, J. D. and Lipson, H. (2010). Morphological evolution of freeform robots. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pages 151–152. ACM.

Hillis, D. (1998). *The Pattern on the Stone: The Simple Ideas that Make Computers Work*. New York: Basic Books.

Hodgkin, A. L. and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Physiology*, 117(4):500–544.

Hoile, C. (2014). Black shoals: Evolving organisms in a world of financial data. Black Shoals Stock Market Planetarium.

Hokkanen, J. (1999). Visual simulations, artificial animals and virtual ecosystems. *Journal of Experimental Biology*, 202(23):3477–3484.

Holland, J. H. (1975). *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. U Michigan Press.

Horgan, J. (1995). From complexity to perplexity. *Scientific American*, 272:104–109.

Hornby, G. S. and Pollack, J. B. (2001). Evolving l-systems to generate virtual creatures. *Computers & Graphics*, 25(6):1041–1048.

Hornby, G. S. and Pollack, J. B. (2002). Creating high-level components with a generative representation for body-brain evolution. *Artificial life*, 8(3):223–246.

Hornby, G. S., Pollack, J. B., et al. (2001). Body-brain co-evolution using l-systems as a generative encoding. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 868–875.

Ito, T., Pilat, M. L., Suzuki, R., and Arita, T. (2013). Alife approach for body–behavior predator–prey coevolution: body first or behavior first? *Artificial Life and Robotics*, 18(1-2):36–40.

Ito, T., Pilat, M. L., Suzuki, R., and Arita, T. (2014). Population and evolutionary dynamics based on predator-prey relationship in 3d physical simulation. In *ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems*, volume 14, pages 105–112.

Jakobi, N. (1998). *Minimal simulations for evolutionary robotics*. PhD thesis, University of Sussex.

Jensen, J. F. (1999). 3d inhabited virtual worlds: Interactivity and interaction between avatars, autonomous agents, and users. In *Virtual interaction: Interaction in virtual inhabited 3D worlds*, pages 23–47. Springer.

Jones, B., Jin, Y., Sendhoff, B., and Yao, X. (2008). Evolving functional symmetry in a three dimensional model of an elongated organism. In *ALIFE*, pages 305–312.

Jones, B. H. D. (2010). *The evolutionary emergence of neural organisation in computational models of primitive organisms*. PhD thesis, University of Birmingham.

Kadota, T., Yasuda, T., Matsumura, Y., and Ohkura, K. (2012). An incremental approach to an evolutionary robotic swarm. In *System Integration (SII), 2012 IEEE/SICE International Symposium on*, pages 458–463.

Kim, K.-J. and Cho, S.-B. (2006). A comprehensive overview of the applications of artificial life. *Artificial Life*, 12(1):153–182.

Kitano, H. (1990). Designing neural networks using genetic algorithms with graph generation system. *Complex Systems Journal*, 4:461–476.

Klein, J. (2003). Breve: a 3d environment for the simulation of decentralized systems and artificial life. In *Proceedings of the eighth international conference on Artificial life*, pages 329–334.

Komosinski, M. (2000). The world of framsticks: simulation, evolution, interaction. In *Virtual worlds*, pages 214–224. Springer.

Komosinski, M. (2003). The framsticks system: versatile simulator of 3d agents and their evolution. *Kybernetes*, 32(1/2):156–173.

Komosinski, M. and Polak, J. (2009). Evolving free-form stick ski jumpers and their neural control systems. In *Proceedings of the National Conference on Evolutionary Computation and Global Optimization, Poland*, pages 103–110.

Komosiński, M. and Rotaru-Varga, A. (2001). Comparison of different genotype encodings for simulated three-dimensional agents. *Artificial Life*, 7(4):395–418.

Komosinski, M. and Ulatowski, S. (1998). Framsticks – artificial life. In Nédellec, C. and Rouveirol, C., editors, *ECML 98 Demonstration and Poster Papers*, page 7–9, Chemnitz. Chemnitzer Informatik-Berichte, Chemnitzer Informatik-Berichte.

Komosinski, M. and Ulatowski, S. (1999). Framsticks: Towards a simulation of a nature-lke world, creatures and evolution. In Floreano, D., Nicoud, J.-D., and Mondada, F., editors, *Proceedings of the Fifth European Conference on Artificial Life*, pages 261–265. Springer-Verlag.

Komosinski, M. and Ulatowski, S. (2004). Genetic mappings in artificial genomes. *Theory in Biosciences*, 123(2):125–137.

Kou, K. and Kawaguchi, Y. (2012). Emergent gait evolution of quadruped artificial life. In *Advances in Computer Entertainment*, pages 287–296. Springer.

Krčah, P. (2007). Evolving virtual creatures revisited. In *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 7, page 341.

Krčah, P. (2008). Towards efficient evolutionary design of autonomous robots. In *Evolvable Systems: From Biology to Hardware*, pages 153–164. Springer.

Krčah, P. (2012). Solving deceptive tasks in robot body-brain co-evolution by searching for behavioral novelty. In *Advances in Robotics and Virtual Reality*, pages 167–186. Springer.

Krčah, P. (2014). Virtual creatures evolved using speciation and historical markings. In *Genetic and Evolutionary Computation Conference*.

Langton, C. G. (1986). Studying artificial life with cellular automata. *Physica D: Nonlinear Phenomena*, 22(1-3):120 – 149.

Langton, C. G. (1989). *Artificial Life: Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.

Lassabe, N., Luga, H., and Duthen, Y. (2007). A new step for artificial creatures. In *Artificial Life, 2007. ALIFE'07. IEEE Symposium on*, pages 243–250. IEEE.

Lehman, J. and Stanley, K. O. (2011). Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 211–218. ACM.

Lenski, R. E. (2001). Twice as natural. *Nature*, 414(6861):255–255.

Lenski, R. E. (2004). The future of evolutionary biology. *Ludus vitalis*, 12(21):67–89.

Lessin, D., Fussell, D., and Miikkulainen, R. (2013). Open-ended behavioral complexity for evolved virtual creatures. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, pages 335–342. ACM.

Lessin, D., Fussell, D., and Miikkulainen, R. (2014a). Adapting morphology to multiple tasks in evolved virtual creatures. In *ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems*, volume 14, pages 247–254.

Lessin, D., Fussell, D., and Miikkulainen, R. (2014b). Trading control intelligence for physical intelligence: Muscle drives in evolved virtual creatures. In *Proceedings of the 2014 conference on Genetic and evolutionary computation*, pages 705–712. ACM.

Lewis, M. A., Fagg, A. H., Solidum, A., and Bekey, G. A. (1992). Genetic programming approach to the construction of a neural network for control of a walking robot. In *Proceedings of the IEEE Conference on Robotics and Automation*.

Lipson, H. and Pollack, J. (2006). Evolving physical creatures. In *Artificial Life VII: Proceedings of the Seventh International Conference on Artificial life*, pages 282–287.

Liu, J., Qin, H., Tang, Y. Y., and Wu, Y. (1997). Adaptation and learning in animated creatures. In *Proceedings of the first international conference on Autonomous agents*, pages 371–377. ACM.

Lowel, S. and Singer, W. (1992). Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity. *Science*, 255(5041):209.

Lund, H. H. (2003). Co-evolving control and morphology with lego robots. In *Morpho-functional Machines: The New Species*, pages 59–79. Springer.

Luo, C., Gao, J., Li, X., Mo, H., and Jiang, Q. (2014). Sensor-based autonomous robot navigation under unknown environments with grid map representation. In *Swarm Intelligence (SIS), 2014 IEEE Symposium on*, pages 1–7.

Macinnes, I. and Di Paolo, E. (2004). Crawling out of the simulation: Evolving real robot morphologies using cheap reusable modules. In *Artificial Life IX: Proceedings of the Ninth Interational Conference on the Simulation and Synthesis of Life*, pages 94–99.

Mahdavi, S. H. and Bentley, P. J. (2003). Evolving motion of robots with muscles. In *Applications of Evolutionary Computing*, pages 651–660. Springer.

Marks, S., Conen, W., Kollakowsky, T., and Lux, G. (2006). Evolving autonomous locomotion of virtual characters in a simulated physical environment via neural networks and evolutionary strategies. *Preceedings of 3IA*, 2006:18.

Matsushita, K., Yokoi, H., and Arai, T. (2006). Pseudo-passive dynamic walkers designed by coupled evolution of the controller and morphology. *Robotics and Autonomous Systems*, 54(8):674–685.

Maturana, H. R. and Varela, F. J. (1980). Autopoiesis and cognition. *Dordrecht, Holland: D. Reidel*.

Maynard Smith, J. (1986). *The problems of biology*. Oxford: Oxford University Press.

Mazzapioda, M., Cangelosi, A., and Nolfi, S. (2009). Evolving morphology and control: a distributed approach. In *Evolutionary Computation, 2009. CEC'09. IEEE Congress on*, pages 2217–2224. IEEE.

McCarthy, J. and Hayes, P. J. (1969). Some philosophical problems from the standpoint of artificial intelligence. *Readings in artificial intelligence*, pages 431–450.

McCloskey, M. and Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. *The psychology of learning and motivation*, 24(109-164).

McCulloch, W. S. and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4):115–133.

McKenna, M. and Zeltzer, D. (1990). Dynamic simulation of autonomous legged locomotion. *SIGGRAPH Comput. Graph.*, 24(4):29–38.

McOwan, P. W. and Burton, E. J. (2005). Sodarace: Adventures in artificial life. In *Artificial Life Models in Software*, pages 97–111. Springer.

Meng, M. and Yang, X. (1998). A neural network approach to real-time trajectory generation [mobile robots]. In *Robotics and Automation, 1998. Proceedings. 1998 IEEE International Conference on*, volume 2, pages 1725–1730. IEEE.

Miconi, T. (2008a). Evosphere: evolutionary dynamics in a population of fighting virtual creatures. In *Evolutionary Computation, 2008. CEC 2008.(IEEE World Congress on Computational Intelligence). IEEE Congress on*, pages 3066–3073. IEEE.

Miconi, T. (2008b). In silicon no one can hear you scream: Evolving fighting creatures. In *Genetic Programming*, pages 25–36. Springer.

Miconi, T. and Channon, A. (2005). A virtual creatures model for studies in artificial evolution. In *Evolutionary Computation, 2005. The 2005 IEEE Congress on*, volume 1, pages 565–572. IEEE.

Miconi, T. and Channon, A. (2006a). Analysing co-evolution among artificial 3d creatures. In *Artificial Evolution*, pages 167–178. Springer.

Miconi, T. and Channon, A. (2006b). An improved system for artificial creatures evolution. In Rocha, L. M., editor, *Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems (ALife X)*, pages 255–261.

Mignonneau, L. and Sommerer, C. (2001). Creating artificial life for interactive art and entertainment. *Leonardo*, 34(4):303–307.

Miller, G. S. P. (1988). The motion dynamics of snakes and worms. *SIGGRAPH Comput. Graph.*, 22(4):169–173.

Minsky, M. (1985). *Society of mind*. Simon and Schuster.

Minsky, M. and Papert, S. (1969). *Perceptrons*. MIT press.

Miranda, F. R., Kogler Jr, J., Netto, M. L., and Del Moral Hernandez, E. (2001). Arena and woxbot: first steps towards virtual world simulations. In *Computer Graphics and Image Processing, 2001 Proceedings of XIV Brazilian Symposium on*, pages 208–215. IEEE.

Mitchell, M. (1998). *An introduction to genetic algorithms*. MIT Press.

Montana, D. J. and Davis, L. (1989). Training feedforward neural networks using genetic algorithms. In *IJCAI*, volume 89, pages 762–767.

Moore, J., Gutmann, A., McGowan, C., and McKinley, P. (2013). Exploring the role of the tail in bipedal hopping through computational evolution. In *Advances in Artificial Life, ECAL*, volume 12, pages 11–18.

Moore, J. M. and Clark, A. J. (2014). Hold the spot: Evolution of generalized station keeping for an aquatic robot. In *ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems*, volume 14, pages 200–201.

Moore, J. M. and McKinley, P. K. (2014a). Evolving joint-level control with digital muscles. In *Proceedings of the 2014 conference on Genetic and evolutionary computation*, pages 209–216. ACM.

Moore, J. M. and McKinley, P. K. (2014b). Investigating modular coupling of morphology and control with digital muscles. In *ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems*, volume 14, pages 148–155.

Moreno, R., Faiña, A., and Støy, K. (2015). Evolving robot controllers for structured environments through environment decomposition. In *Applications of Evolutionary Computation*, pages 795–806. Springer.

Mouret, J.-B. and Doncieux, S. (2009). Overcoming the bootstrap problem in evolutionary robotics using behavioral diversity. In *IEEE Congress on Evolutionary Computation*.

Muybridge, E. (1887). *Animal Locomotion: An Electrophotographic Investigation of Consecutive Phases of Animal Movements. 1872-1875*. University of Pennsylvania.

Nakamura, K. (2015). Emergence of adaptive behaviors for artificial creature using a combined artificial neural network. *Artificial Life and Robotics*, pages 1–7.

Nakamura, K., Suzuki, I., Yamamoto, M., and Furukawa, M. (2011). Virtual fluid environment on behavior ability for artificial creature. In *ECAL 2011 Proceedings of the eleventh European Conference on the Synthesis and Simulation of Living Systems, T. Lenarets, M. Giacobini, H. Bersini, P. Bourgine, M. Dorigo, and R. Doursat*, pages 581–588.

Nelson, A. L. (2014). Embodied artificial life at an impasse can evolutionary robotics methods be scaled? In *Evolving and Autonomous Learning Systems (EALS), 2014 IEEE Symposium on*, pages 25–34. IEEE.

Newell, A. and Simon, H. A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19(3):113–126.

Nikovski, D. (1995). Evolving legged locomotion of virtual creatures. In *Proceedings of the Midwest Artificial Intelligence and Cognitive Science Society Conference (MAICS)*. Citeseer.

Nishino, H., Aoki, K., Takagi, H., Kagawa, T., and Utsumiya, K. (2004). A synthesized 3dcg contents generator using iec framework. In *Systems, Man and Cybernetics, 2004 IEEE International Conference on*, volume 6, pages 5719–5724. IEEE.

Nogueira, B., Lenon, Y., Eduardo, C., Vidal, C. A., and Cavalcante Neto, J. B. (2013). Evolving plastic neuromodulated networks for behavior emergence of autonomous virtual characters. In *Advances in Artificial Life, ECAL*, volume 12, pages 577–584.

Nolfi, S. (1998). Evolutionary robotics: Exploiting the full power of self-organization. *Connection Science*, 10:167–183.

Nolfi, S. and Floreano, D. (1998). Coevolving predator and prey robots: do "arms races" arise in artificial evolution? *Artificial life*, 4(4):311–335.

Nolfi, S. and Parisi, D. (1997). Learning to adapt to changing environments in evolving neural networks. In *Adaptive Behavior*, pages 75–98.

Oh, S.-H. and Suk, J. (2013). Evolutionary controller design for area search using multiple uavs with minimum altitude maneuver. *Journal of Mechanical Science and Technology*, 27(2):541–548.

Ohono, S., Yagi, S., and Kawaguchi, Y. (2007). Simulation techniques for generating behavior for unitable artificial life. In *Computer Graphics, Imaging and Visualisation, 2007. CGIV'07*, pages 133–137. IEEE.

O'Kelly, M. and Hsiao, K. (2004). Evolving simulated mutually perceptive creatures for combat. In *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems (ALIFE9)*, pages 113–118.

Ooe, R., Suzuki, I., Yamamoto, M., and Furukawa, M. (2013). Study on evolution of the artificial flying creature controlled by neuro-evolution. *Artificial Life and Robotics*, 17(3-4):470–475.

Ouannes, N., Djedi, N., Duthen, Y., and Luga, H. (2012). Toward the construction of a virtual ecosystem by evolving virtual creature's behaviours. In *Multimedia Computing and Systems (ICMCS), 2012 International Conference on*, pages 350–355. IEEE.

Packard, N. H. (1989). Intrinsic adaptation in a simple model for evolution. *Artificial Life*, VI:141–155.

Pennock, R. T. (2012). Negotiating boundaries in the definition of life: Wittgensteinian and darwinian insights on resolving conceptual border conflicts. *Synthese*, 185(1):5–20.

Penny, S. (2009). Art and artificial life–a primer. *Digital Arts and Culture 2009*.

Pfeifer, R., Iida, F., and Bongard, J. (2005). New robotics: Design principles for intelligent systems. *Artificial life*, 11(1-2):99–120.

Pfeifer, R., Iida, F., and Lungarella, M. (2014). Cognition from the bottom up: on biological inspiration, body morphology, and soft materials. *Trends in Cognitive Sciences*, 18(8):404 – 413.

Pichler, P.-P. (2009). *Natural Selection, Adaptive Evolution and Diversity in Computational Ecosystems*. PhD thesis, University of Hertfordshire.

Pilat, M. L., Ito, T., Suzuki, R., and Arita, T. (2012). Evolution of virtual creature foraging in a physical environment. In *Artificial Life*, volume 13, pages 423–430.

Pilat, M. L. and Jacob, C. (2008a). Creature academy: A system for virtual creature evolution. In *Evolutionary Computation, 2008. CEC 2008.(IEEE World Congress on Computational Intelligence). IEEE Congress on*, pages 3289–3297. IEEE.

Pilat, M. L. and Jacob, C. (2008b). Infusing virtual creatures with vision. Technical report, University of Calgary.

Pilat, M. L. and Jacob, C. (2010). Evolution of vision capabilities in embodied virtual creatures. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pages 95–102. ACM.

Pollack, J. B. and Lipson, H. (2000). The golem project: Evolving hardware bodies and brains. In *Evolvable Hardware, 2000. Proceedings. The Second NASA/DoD Workshop on*, pages 37–42. IEEE.

Prophet, J. (2001). Technosphere:“real” time,“artificial” life. *Leonardo*, 34(4):309–312.

Rada, J. and Aguilar, J. (2012). Morphology and behaviour evolution of virtual creatures: A model based on swarm inteligence. *Latin America Transactions, IEEE (Revista IEEE America Latina)*, 10(2):1653–1660.

Raibert, M. H. and Hodgins, J. K. (1991). Animation of dynamic legged locomotion. *SIGGRAPH Comput. Graph.*, 25(4):349–358.

Ray, T. S. (1992). An approach to the synthesis of life. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, pages 371–408. Addison Wesley Publishing Company.

Ray, T. S. (1993). An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artificial Life*, 1(1_2):179–209.

Ray, T. S. (1997). Evolving complexity. *Artificial Life and Robotics*, 1(1):21–26.

Ray, T. S. (2001). Aesthetically evolved virtual pets. *Leonardo*, 34(4):313–316.

Rechenberg, I. (1973). *Evolutionsstrategie Optimierung technischer systeme nach prinzipien der biologischen evolution*. Friedrich Frommann Verlag, Stuttgart-Bad Cannstatt.

Reil, T. and Husbands, P. (2002). Evolution of central pattern generators for bipedal walking in a real-time physics environment. *IEEE Transactions on Evolutionary Computation*, 6(2):159–168.

Rieffel, J., Knox, D., Smith, S., and Trimmer, B. (2014). Growing and evolving soft robots. *Artificial life*, 20(1):143–162.

Robinson, E., Ellis, T., and Channon, A. (2007). Neuroevolution of agents capable of reactive and deliberative behaviours in novel and dynamic environments. In *Advances in Artificial Life*, pages 345–354. Springer.

Rossi, C. and Eiben, A. (2014). Simultaneous versus incremental learning of multiple skills by modular robots. *Evolutionary Intelligence*, 7(2):119–131.

Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Learning internal representations by error propagation. In Rumelhart, D. E. and McClelland, J. L., editors, *Parallel Distributed Processing*, volume 1. MIT Press.

Russell, S. and Norvig, P. (2003). *Artificial intelligence: a modern approach*. Prentice Hall, 2 edition.

Samuelsen, E. and Glette, K. (2014). Some distance measures for morphological diversification in generative evolutionary robotics. In *Proceedings of the 2014 conference on Genetic and evolutionary computation*, pages 721–728. ACM.

Samuelsen, E., Glette, K., and Torresen, J. (2013). A hox gene inspired generative approach to evolving robot morphology. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, pages 751–758. ACM.

Saul, G., Xu, C., and Gross, M. D. (2010). Interactive paper devices: end-user design & fabrication. In *Proceedings of the fourth international conference on Tangible, embedded, and embodied interaction*, pages 205–212. ACM.

Schut, M. C., Haasdijk, E., and Eiben, A. (2009). What is situated evolution? In *Evolutionary Computation, 2009. CEC'09. IEEE Congress on*, pages 3277–3284. IEEE.

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and brain sciences*, 3(03):417–424.

Shanken, E. A. (1998). Life as we know it and/or life as it could be: Epistemology and the ontology/ontogeny of artificial life. *Leonardo*, pages 383–388.

Shim, Y.-S. and Kim, C.-H. (2003). Generating flying creatures using body-brain co-evolution. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 276–285. Eurographics Association.

Shim, Y.-S. and Kim, C.-H. (2006). Evolving physically simulated flying creatures for efficient cruising. *Artificial Life*, 12(4):561–591.

Shim, Y. S., Kim, S. J., and Kim, C. H. (2004a). Evolving flying creatures with path-following behavior. In *ALife IX: Proceedings of the 9th International Conference on the Simulation and Synthesis of Living Systems*, pages 125–132. Citeseer.

Shim, Y.-S., Shin, S.-Y., and Kim, C.-H. (2004b). Two-step evolution process for path-following virtual creatures. *Computer Animations and Social Agents 2004*, pages 85–93.

Sims, K. (1994a). Evolving 3d morphology and behavior by competition. *Artificial life*, 1(4):353–372.

Sims, K. (1994b). Evolving virtual creatures. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, SIGGRAPH '94, pages 15–22, New York, NY, USA. ACM.

Skipper, J. (1992). The computer zoo – evolution in a box. In Varela, F. J. and Bourgine, P., editors, *Toward a practice of autonomous systems, proceedings of the first European conference on Artificial Life.*, pages 355–364.

Solé, R. V. and Valverde, S. (2013). Macroevolution in silico: scales, constraints and universals. *Palaeontology*, 56(6):1327–1340.

Sommerer, C. and Mignonneau, L. (1999). Art as a living system: interactive computer artworks. *Leonardo*, 32(3):165–173.

Spector, L., Klein, J., and Feinstein, M. (2007). Division blocks and the open-ended evolution of development, form, and behavior. In *Proceedings of the 9th annual conference on Genetic and evolutionary computation*, pages 316–323. ACM.

Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8(2):131–162.

Stanley, K. O. and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127.

Stanton, A. and Channon, A. (2013). Heterogeneous complexification strategies robustly outperform homogeneous strategies for incremental evolution. In *Advances in Artificial Life, ECAL*, volume 12, pages 973–980.

Stanton, A. and Channon, A. (2015). Incremental neuroevolution of reactive and deliberative 3d agents. In *Advances in Artificial Life, ECAL 2015: Proceedings of the Thirteenth European Conference on the Synthesis and Simulation of Living Systems*, pages 341–348. MIT Press.

Stanton, A. and Channon, A. (2016). Neuroevolution of feedback control for object manipulation by 3d agents. In *Proceedings of the Fifteenth International Conference on the Simulation and Synthesis of Living Systems (ALIFEXV)*, pages 144–151. MIT Press.

Steels, L. (1993). The artificial life roots of artificial intelligence. *Artificial life*, 1(1_2):75–110.

Steels, L. (2008). The symbol grounding problem has been solved. so what's next. *Symbols and embodiment: Debates on meaning and cognition*, pages 223–244.

Stredwick, J. M. (2005). Evolutionary dynamics of 3d digital constructs. Master's thesis, Michigan State University.

Tan, J., Gu, Y., Turk, G., and Liu, C. K. (2011). Articulated swimming creatures. In *ACM Transactions on Graphics (TOG)*, number 4 in 30, page 58. ACM.

Tange, O. (2011). Gnu parallel - the command-line power tool. *;login: The USENIX Magazine*, 36(1):42–47.

Taylor, T. (2000). Artificial life techniques for generating controllers for physically modelled characters. In *GAME-ON*, page 89.

Taylor, T., Bedau, M., Channon, A., Ackley, D., Banzhaf, W., Beslon, G., Dolson, E., Froese, T., Hickinbotham, S., Ikegami, T., McMullin, B., Packard, N., Rasmussen, S., Virgo, N., Agmon, E., Clark, E., McGregor, S., Ofria, C., Ropella, G., Spector, L., Stanley, K. O., Stanton, A., Timperley, C., Vostinar, A., and Wiser, M. (2016). Open-ended evolution: Perspectives from the oee workshop in york. *Artificial Life*, 22(3):408–423.

Taylor, T. and Massey, C. (2001). Recent developments in the evolution of morphologies and controllers for physically simulated creatures. *Artificial Life*, 7(1):77–87.

Taylor, T. J. and Hallam, J. (1997). Studying evolution with self-replicating computer programs. In Husbands, P. and Harvey, I., editors, *Proceedings of the Fourth European Conference on Artificial Life (ECAL97)*. MIT Press/Bradford Books.

Teo, J. and Abbass, H. A. (2002). Trading-off mind complexity and locomotion in a physically simulated quadruped. In *Proceedings of the 4th Asia-Pacific*

Conference on Simulated Evolution And Learning (SEAL'02), volume 2, pages 776–780. Citeseer.

Teo, J. and Abbass, H. A. (2004). Automatic generation of controllers for embodied legged organisms: A pareto evolutionary multi-objective approach. *Evolutionary Computation*, 12(3):355–394.

Terzopoulos, D. (1999). Artificial life for computer graphics. *Communications of the ACM*, 42(8):32–42.

Terzopoulos, D., Rabie, T., and Grzeszczuk, R. (1997). Perception and learning in artificial animals. In *Proc. of the 5th Int. Workshop on Artificial Life: Synthesis and Simulation of Living Systems (ALIFE-96)*, pages 346–353.

Terzopoulos, D. and Rabie, T. F. (1995). Animat vision: Active vision in artificial animals. In *Computer Vision, 1995. Proceedings., Fifth International Conference on*, pages 801–808. IEEE.

Terzopoulos, D., Tu, X., and Grzeszczuk, R. (1994). Artificial fishes: Autonomous locomotion, perception, behavior, and learning in a simulated physical world. *Artificial Life*, 1(4):327–351.

Thalmann, D., Noser, H., and Huang, Z. (1996). How to create a virtual life. Technical report, Swiss Federal Institute of Technology.

Thompson, A. (1996). An evolved circuit, intrinsic in silicon, entwined with physics. In *International Conference on Evolvable Systems*, pages 390–405. Springer.

Tibermacine, A. and Djedi, N. (2014). Neat neural networks to control and simulate virtual creature's locomotion. In *Multimedia Computing and Systems (ICMCS), 2014 International Conference on*, pages 9–14.

Todd, P. M. (1996). The causes and effects of evolutionary simulation in the behavioral sciences. In *In Belew and Mitchell*. Citeseer.

Tschacher, W. and Bergomi, C. (2011). *The implications of embodiment: Cognition and communication*. Imprint Academic.

Tuci, E. (2004). *An Exploration on the Evolution of Learning Behaviour Using Robot-based Models*. PhD thesis, University of Sussex.

Turing, A. M. (1946). Letter from Turing to Ashby concerning use of ACE. Private communication.

Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236):433–460.

Valsalam, V. K. (2010). *Utilizing symmetry in evolutionary design*. PhD thesis, University of Texas.

Van de Panne, M. and Fiume, E. (1993). Sensor-actuator networks. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 335–342. ACM.

Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and brain sciences*, 21(05):615–628.

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

Vassie, K. and Morlino, G. (2012). Natural and artificial systems: Compare, model or engineer? *From Animals to Animats 12*, pages 1–11.

Von Neumann, J. (1966). *Theory of Self-Reproducing Automata*. University of Illinois Press, Champaign, IL, USA.

von Uexküll, J., von Uexküll, M., and O'Neil, J. D. (2010). *A Foray into the Worlds of Animals and Humans: with A Theory of Meaning*. University of Minnesota Press.

Vonasek, V., Saska, M., Winkler, L., and Preucil, L. (2015). High-level motion planning for cpg-driven modular robots. *Robotics and Autonomous Systems*, 68:116–128.

Wallace, A. R. (1858). On the tendency of varieties to depart indefinitely from the original type. Private communication.

Walter, W. G. (1950). Imitation of life. *Scientific American*, 182(5):42–45.

Wampler, K. and Popović, Z. (2009). Optimal gait and form for animal locomotion. *ACM Transactions on Graphics (TOG)*, 28(3):60.

Wang, F. and Mckenzie, E. (1998). *Virtual life in virtual environments*. Citeseer.

Watson, R. A., Ficiei, S., and Pollack, J. B. (1999). Embodied evolution: Embodying an evolutionary algorithm in a population of robots. In *Evolutionary Computation, 1999. CEC 99. Proceedings of the 1999 Congress on*, volume 1. IEEE.

Webb, B. (2001). Can robots make good models of biological behaviour? *Behavioral and brain sciences*, 24(6):1033–1050.

Werbos, P. (1974). *Beyond regression: New tools for prediction and analysis in the behavioral sciences*. PhD thesis, Harvard University.

Wilson, R. A. and Keil, F. C. (2001). *The MIT encyclopedia of the cognitive sciences*. MIT press.

Wilson, S. (1985). Knowledge growth in an artificial animal. In Grefenstette, J. J., editor, *Proceedings of an International Conference on Genetic Algorithms and Their Applications*.

Yaeger, L. (1993). Computational genetics, physiology, metabolism, neural systems, learning, vision, and behavior or polyworld: Life in a new context. In *Artificial Life III, Vol. XVII of SFI Studies in the Sciences of Complexity*, pages 263–298. Addison-Wesley.

Yaeger, L. (2009). How evolution guides complexity. *Human Frontier Science Program*.

Yaeger, L., Griffith, V., and Sporns, O. (2008). Passive and driven trends in the evolution of complexity. In Bullock, S., editor, *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*.

Yaeger, L. and Sporns, O. (2006). Evolution of neural structure and complexity in a computational ecology. In Rocha, L., editor, *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*.

Yang, S. X. and Meng, M. (2000). An efficient neural network approach to dynamic robot motion planning. *Neural Networks*, 13(2):143 – 148.

Yonekura, S. and Kawaguchi, Y. (2008). Adaptive body use of multipede driven by coupled neural oscillators. In *18th International Conference on Artificial Reality and Telexistence (ICAT 2008)*.

Zaera, N., Cliff, D., and Bruten, J. (1996). (Not) evolving collective behaviours in synthetic fish. In *Proceedings of International Conference on the Simulation of Adaptive Behavior*.

Ziemke, T. (1998). Adaptive behavior in autonomous agents. *Presence: Teleoperators and Virtual Environments*, 7(6):564–587.