The Road to Everywhere: Evolution, Complexity and Progress in Natural and Artificial Systems

Thomas Miconi

A thesis submitted
to the University of Birmingham
for the degree of Doctor of Philosophy

School of Computer Science
The University of Birmingham
August 2007

But that science is so fer us biforn,
We mowen nat, although we hadden it sworn,
It overtake, it slit awey so faste.
It wole us maken beggers atte laste!

Chaucer, *The Canon's Yeoman Tale*

Abstract

Evolution is notorious for its creative power, but also for giving rise to complex, unpredictable dynamics. As a result, practitioners of artificial evolution have encountered difficulties in predicting, analysing, or even understanding the outcome of their experiments. In particular, the concept of evolutionary "progress" (whether in the sense of performance increase or complexity growth) has given rise to much debate and confusion. After a careful description of the mechanisms of evolution and natural selection, we provide usable concepts of performance and progress in coevolution. In particular, we introduce a distinction between three types of progress: local, historical, and global, which we suggest underlies much of the confusion that surrounds coevolutionary dynamics. Similarly, we provide a comprehensive answer to the question of whether an "arrow of complexity" exists in evolution. We introduce several methods to detect and analyse performance and progress in coevolutionary experiments. We propose a statistical measure (Fitness Transmission) to detect the presence of adaptive Darwinian evolution in a reproducing population, based solely on genealogic records; we also point out the limitations of a popular method (the Bedau-Packard statistics of evolutionary activity) for this purpose. To test and illustrate our results, we implement a rich experimental system, inspired by the seminal work of Karl Sims, in which virtual creatures can evolve and interact under various conditions in a physically realistic three-dimensional (3D) environment. To our knowledge, this is the first complete reimplementation and extension of Sims' results. We later extend this system with the introduction of physical combat between creatures, also a first. Finally, we introduce Evosphere, an open, planet-like environment in which 3D artificial creatures interact, reproduce and evolve freely. We conclude our discussion by using Fitness Transmission to detect the onset of adaptive evolution in this system.

Acknowledgements

This research was mostly conducted at (and funded by) the School of Computer Science at the University of Birmingham. During the first six months, this research was conducted at (and funded by) the Intelligent Systems group, within the Electrical and Computer Engineering department at the University of Portsmouth, headed by Dr. David J. Brown.

This thesis was produced under the supervision of Alastair Channon, to whom I am immensely grateful for offering me the chance of taking up this research. Alastair created an ideal environment, allowing almost total freedom in the direction of research, while judiciously providing much-needed guidance at critical moments. Dr. David Brown initially co-supervised this thesis at Portsmouth, providing much valuable insight and leadership. This research also benefited greatly from challenging discussions with John Bullinaria, Richard Dearden and Aaron Sloman (who also provided a much-appreciated 17” monitor, though my initial hopes of reaping wisdom particles from the device did not materialise).

I am also much indebted to the many friendly people I have met over my years at Birmingham. Daily life was made much more enjoyable by the companionship of Rocio Gómez, Tebogo Seipone and Philip Rohlfshagen. Enormous thanks are due to Pascal Honoré for being, well, Pascal Honoré - and to Aubrey for being so nice and friendly. Much distraction, celebration, intoxication, and general good humour are owed to the above, as well as the Italian gang (Andrea Arcuri, Pietro Oliveto, Loretta Mancini, et al.), Arjun “Soothing Voice” Chandra, Isabel, Oksana Tymchyshyn, Michael Zillich, Andrea Soltoggio, Ramon, Per Christian, Kyriakos, the lovely Natalja, and many others. Thanks to Ronald Petiteau for keeping in touch and for much-enjoyed visits! Great thanks to Sidra Tul Muntaha for warm friendship and delicious meals, and to Mireille “Burning Touch” Yuan for (almost) the same. The general population of research students at the School of Computer Science, including the excellent Minbanks (Rob and Richard), the three Daves (G, B and P), fellow residents of “the best office on campus” Sayyed Shah and Peter Lewis, and all the others, proved every bit the delightful bunch that they are internationally renowned to be (yes, yes they are.) Best wishes to Steve and Marie, as well as to Thomas, Julie, and the irresistible Auguste. And of course, all the best to the infinitely lovable Gabriella.

What I owe to my parents is, of course, too ridiculously inexpressible to be put in words here.

Contents

List of Figures

Chapter 1

Introduction

1.1 The problem

1.1.1 The advent of artificial evolution

Evolution shapes the biological world. Ever since the beginning of life, what started as a peculiar chemical reaction between organic compounds has given rise to an astronomical diversity of entities, exhibiting features and behaviours of a complexity that baffle human engineers. Through evolutionary adaptation, life has invaded the world, colonising even the most hostile environments it had to offer. From the inner crust of the Earth to the top of mountains, life flourishes wherever it can, and even where we might think it can not.

The power of evolution is most visible in elaborate structures sported by living organisms, both at the microscopic and the macroscopic level. The bacterial flagellum, the acid guns of beetles, the vertebrate eye, the human brain, are testaments to the creativity of the evolutionary process. Explaining how these structures could have come into being without external design was precisely one of the great achievements of evolutionary theory.

Artificial evolution is an attempt to harness some of this power for the production of artificial entities. The idea of using computers to study concepts related to life and evolution is almost as old as computers themselves. John Von Neumann showed how Turing's universal computing automaton could be applied to the creation of self-replicating entities. His considerations on evolutionary processes led him to formalise the mechanisms of heredity and self-replication, and to investigate the question of evolutionary complexity (see McMullin's account [114]). Soon afterwards, in an apparently unrelated effort, engineers made use of evolutionary techniques to perform practical tasks such as optimisation and design. These two approaches to artificial evolution (theoretical and conceptual on one hand, and practical, engineering-oriented on the other hand) have grown into closely related lines of research in artificial life and evolutionary computation.

1.1.2 Discrepancies between natural and artificial evolution

Unfortunately, there seems to have been a lack of connection between students of evolution in Nature and in computers, which may be the cause for what has been described as a lack of theoretical grounding in artificial evolution. A notable result of this lack of connection is a discrepancy between what is expected from evolution, and what is actually obtained from it. The feeling that artificial evolution and artificial life systems often rely on implicit assumptions and ad hoc rules and machineries has been emphasised by Taylor [173]:

> The results of [our] work, and consideration of the existing literature on artificial evolutionary systems, leads to the conclusion that artificial life models (. . .) are lacking on a number of theoretical and methodological grounds. It is emphasised that explicit theoretical considerations should guide the design of such models, if they are to be of scientific value.

And further (citing Pattee [139]):

> I am certainly not the first person to criticise artificial life models on these grounds[. . .] For example, Howard Pattee warns that "simulations that are dependent on ad hoc and special-purpose rules and constraints for their mimicry cannot be used to support theories of life" [139][. . .] The ad hoc feel of [such] systems is a direct consequence of this lack of theoretical grounding. The unmanageable parameter space of many of them can also be attributed to this lack of direction. As a result of these weaknesses, even if interesting behaviours are observed in these systems, we are unlikely to be able to explain why.

This problem is particularly prevalent in *coevolution*, that is, when the reproductive success of individuals is determined at least in part by the outcome of their interactions, rather than by a fixed, external fitness function. Ficici [53, Chap. 1] writes:

> [The] gap between the hypothesized potential of coevolutionary algorithms and realized practice remains substantial - the successes of coevolution (of which there are now many) are balanced by frequently encountered and irksome pathologies (of which there are also many). These pathologies do not merely deprive us of a satisfying result (as might a local optimum), but more importantly they appear to violate our intuitions of how selection pressure in coevolution is supposed to operate [. . .] To grapple with this incongruity between expectation and reality, coevolution researchers have appropriated or invented a myriad of terms, such as *cyclic dynamic, mediocre stable-state, collusion, forgetting, disengagement,* and *focusing.*

1.1.3 The vagaries of progress

A notable example of this lack of grounding is related to the notion of *progress* in evolution. The idea that evolution by natural selection is fundamentally a progressive force, and that this progress is somehow linked to a general increase in complexity, pervades much of the literature. Let us cite a few examples:

> Intuitively, the distinctive mark or evolution is the spontaneous generation of innovative functional structures [. . .] The growth of adaptations causes the biosphere to increase in complexity, thus providing an arrow of time not implied by mere complex change, even if sustained through many generations (Bedau and Packard [10]).

> . . . the co-evolution of competing populations may produce increasingly complex evolving challenges. As discussed by Dawkins and Krebs [41], competing populations may reciprocally drive one another to increasing level of complexity by producing an evolutionary "arms race". . . (Nolfi and Floreano [133]).
>
> Since the parasites are also evolving with a fitness based on a competition's outcome, the success of a host implies failure for its parasites. When the parasites evolve to overcome this failure, they create new challenges for the hosts; the continuation of this may lead to an evolutionary "arms race" [39]. New genotypes arise to defeat old ones. New parasite types should serve as a drive towards further innovation, creating ever-greater levels of complexity and performance by forcing hosts to respond to a wider range of more challenging parasite test cases (Rosin and Belew [151]).

These passages express an implicit intuition that seems to be deeply entrenched within the general public, as emphasised by Gould [71]. In the (ironic) words of McShea [115], they state "what everybody knows".

However, as many of these same authors have pointed out, this expectation is not always borne out by experiment. Quite the contrary, as emphasised by Ficici above, a whole pandemonium of contrarian effects have been identified, studied and classified, leaving experimenters baffled with the apparently capricious dynamics of coevolution (these effects are described in more detail in Chapter 3).

This initial trust in the progressive nature of evolution can be put in contrast with the following quote by two leading biologists:

> The notion of progress has a bad name among evolutionary biologists (. . .) On the theoretical side, there is no reason why evolution by natural selection should lead to an increase in complexity, if that is what we mean by progress. At most, the theory suggests that organisms should get better, or at least no worse, at doing what they are doing right now. But an increase in immediate 'fitness' - that is, expected number of offspring - may be achieved by losing eyes or legs as well as by gaining them. Even if an increase in fitness cannot be equated with an increase in complexity, or with progress, it might seem at first sight that R. A. Fisher's 'fundamental theorem of natural selection' at least guarantees an increase in fitness. The theorem states that the rate of increase in the mean fitness of a population is equal to the genetic variance in fitness: since variances cannot be negative, the theorem states that fitness can only increase. (...) Unfortunately, the theorem holds only if the relative fitnesses of genotypes are constant, and independent of their frequencies in the population: for many traits, such constancy does not hold (Maynard-Smith and Szathmáry [165]).

Notice that there are really two concepts of progress in the above passages: one relates to an increase in quality and performance, regarding certain attributes. Another refers to the more open concept of increase in complexity, independently of any particular ability. Both concepts are a source of contention and controversy in biology [165, 71, 72, 40, 117, 153] and in artificial evolution [2, 12].

The concept of progress in evolution faces a basic difficulty: it is not immediately obvious how exactly the mechanism of evolution, based on differential reproductive success within a given *local* environment (local both in space and in time) might lead to *global* overall progress. Even less clear is the necessity for any such global improvement to be related to an overall drive towards complexification.

The problem is even more acute in coevolution. When an explicit fitness function is given, progress can at least be defined unambiguously - just compare the fitness of older and newer individuals. But in coevolution, the performance of an individual is necessarily *relative* to that of the individuals it interacts with. An individual is only "good" or "bad" in a certain context, defined by certain other individuals. But then how can we objectively compare any two individuals? How can we give an objective measure of performance and superiority that we might consistently apply across generations? And, more pressingly, if evolution is guided by current opponents that also change through evolution, what kind of long-term term trend (if any) can we expect to observe?

1.1.4 Lack of structural freedom in artificial evolution

Another restriction of most artificial evolution experiments is the severe limitations imposed on how much of the individuals' capacity to interact with their environment is under evolutionary control. Even when agents are *embodied*, i.e. when they are not represented as purely abstract entities, but as situated actors that can act upon (and react to) their environments, this body is often defined in such a way that the range of possible actions is strongly limited, usually within a pre-defined repertoire. Quite often, the only action an agent can perform is to move (e.g. predator-prey simulations). In artificial ecologies, an agent may perform a few pre-defined actions, such as exchange of resources (as in Echo [84]) , reproduction or combat (as in Polyworld [197]).

Ideally, a rich evolutionary system should not be limited to the adjustment of parameters for a limited repertoire of pre-defined behaviours (weights in a neural network, action rules, etc.) Rather, it should allow agents to evolve their own actions as well as the behaviours that regulate these actions: agents should be able to *construct* their own behaviours, rather than choosing between pre-defined behaviours. This is only possible if both control *and* morphology are under evolutionary control, in such a way that the morphology offers sufficient flexibility to permit the appearance of new behaviours.

Few experiments allow individual to construct their own actions through evolution, which requires that both morphology and control architecture be under evolutionary control. The most influential experiments in that domain are undoubtedly Karl Sims' virtual creatures [163], which despite ad hoc features in their machinery allowed for an unprecedented freedom in the evolutionary control over both morphology and behaviours of agents. It is notable that, at the time we started this work (a full decade after Sims), attempts at reproducing Sims' results had not been fully successful. The results presented by Sims certainly leave a lot of space for further investigation.

1.2 Objectives

Our main objective is to obtain a better understanding of evolution and (especially) coevolution, what we can expect from it, and how we can put it to use for our own interests. In particular, we seek to address the troublesome concept of evolutionary "progress", whether in the engineering sense of increase in "performance" (relative to a certain task of interest), or in the more subjective sense of growth in complexity. The questions we intend to tackle include the following:

- What is evolution?
- What is progress?
- What kind of progress occurs in (co-)evolution?
- What is the relationship between evolution and complexity?
- How can we detect, or measure, evolution and progress?

In other words, *we want to understand what exactly evolution is supposed to do, but also how to find out whether it's actually doing it.* This double objective of understanding and monitoring will be a constant concern throughout this work.

To support and illustrate our investigations, we will build an experimental system inspired by the seminal work of Karl Sims. In this system, physically realistic, three-dimensional (3D) agents will interact physically with their environment and with each other. Both the morphologies and controllers of these agents will be subject to evolution. We will use this system under various conditions, including evolution based on fitness functions, coevolution, and ultimately in the form of an open environment in which a population of free-living creatures interact (through physical combat) and evolve freely.

Thus, we can summarise our objectives in three broad points:

1. To produce a comprehensive description of evolution, coevolution, and progress, based on sound definitions, in order to minimise the potential for confusion and misunderstandings.
2. To devise tools and methods to detect and measure the activity of evolution.
3. To create a software platform in which artificial creatures evolve (and coevolve) in a physically realistic, three-dimensional environment, as an experimental system to apply and test our concepts and methods.

1.3 Contributions of the thesis

The following is a brief summary of the main contributions of this thesis (these are presented with more detail in the Conclusion).

1. We provide usable notions of superiority and progress in coevolution. This includes a previously overlooked distinction between *local*, *historical* and *global* progress, which we suggest lies at the root of much of the confusion that surrounds coevolution.

2. We address the question of whether or not an “arrow of complexity” [12] exists in evolution. After putting this question in a more usable form, we answer with a qualified positive, characterising this “arrow” as a *passive* trend. We propose a set of abstract conditions necessary for such a continuous trend to occur.

3. We describe several tools and methods to evaluate the performance of coevolutionary algorithms over time: Coarse-grained Master Tournaments (expanding on work by Cliff & Miller [32] and Nolfi & Floreano [133]), cross-validation of runs, and equal-effort comparisons between algorithms. We point out the distinction to be made between historical and global methods, which respectively detect historical and global progress.

4. Using our equal-effort comparison method, we measure the benefit obtained by increasing the number of competitions per evaluation. We show that the gain in performance peaks between 3 and 4, an intriguingly low value. We also show that using a sliding archive increases performance slightly, but significantly in the long term (though at the cost of a performance hit in the short run).

5. We introduce Fitness Transmission, a statistical signature of Darwinian evolution that can be computed from genealogical records. Perhaps equally importantly, we show that Bedau & Packard’s evolutionary activity statistics [11] cannot be used for this particular purpose.

6. We provide the first complete reimplementation and extension of Sims’ ‘virtual creatures’ system [163, 162]. We also introduce the first example of the evolution of actual physical (i.e. contact-based) combat between such virtual creatures.

7. We introduce Evosphere, an artificial world in which free-living creatures interact, fight, and evolve freely, without any explicit fitness function.

1.4 Outline of the thesis

The thesis is organised as follows:

- Chapter 2 provides a discussion of the Darwinian mechanism of evolution by natural selection. The terms “evolution” and “natural selection” are given definitions that will be used in the following sections.

- Chapter 3 investigates the notions of performance, superiority and progress in coevolution. We show that several notions of superiority and progress can be identified, depending on which opponents are used as a reference to evaluate this superiority. We suggest that much confusion about coevolution stems from a confusion between these various kinds of progress. We compare our concepts with those recently proposed by Ficici. [53, 54]

- Chapter 4 describes our experimental system, inspired by Sims’ virtual creatures. To our knowledge, this was the first complete replication (and extension) of Sims’ results.

- Chapter 5 describes an extension of our system, through which physical combat between creatures is implemented. To our knowledge, this is the first example of the evolution of actual physical (i.e. contact-based) combat between physically realistic 3D agents.

- Chapter 6 describes several tools and methods to track the performance of coevolutionary algorithms over time, building upon previous work, and making use of our improved understanding of coevolutionary progress.

- Chapter 7 illustrates some of the methods used in the previous chapter to tackle a simple question: in a coevolutionary algorithm, how many interactions per evaluation are optimal, and what is the advantage of using a sliding archive?

- Chapter 8 attempts to answer the question of whether or not an "arrow of complexity" [12] exists in evolution. Drawing from recent evolutionary research, the relation between evolution and complexity is examined in depth.

- Chapter 9 introduces Fitness Transmission, a statistical signature of adaptive evolution that can be computed from genealogical records. We explain how this measure addresses limitations of other statistics used for similar purpose, such as Bedau and Packard's evolutionary statistics. [11]

- Chapter 10 describes an open environment, called Evosphere, in which free-living creatures interact, fight, and evolve freely, without any explicit fitness function. Fitness transmission is used to demonstrate the onset of natural selection and adaptive evolution within this virtual world.

- Finally, the conclusion recapitulates the main contributions of the thesis, and summarises its main message in a short final paragraph.

Chapter 2

What is evolution?

> Another curious aspect of the theory of evolution is that everyone thinks he understands it.
>
> J. Monod, *Chance and necessity.*[129]

2.1 Introduction

Evolution is neither complicated not difficult to understand - at least, not in its basic principles. As we will see, these principles can be summarised in a couple of straightforward sentences. In fact, evolution is so simple that, as pointed out by Monod, we spontaneously tend to regard it a self-evident, and not really requiring any particular explanation. This, in fact, is perhaps the most important difficulty when discussing evolution. Although the basic principles are simple, they are also subtle, and very easy to "get wrong". One need only consider the persistence of misguided (but vocal) "critics" and "debunkers" of evolution to realise how much confusion the subject can generate.

Reasoning that such confusion is the necessary consequence of a lack of sound principles and clear concepts, we would like to come up with a satisfying answer to the question: "what is evolution?" In particular, we seek precise definitions that can be used as a basis for our inquiries. After all, when discussing a certain subject, it is generally useful to know just *what* one is talking about in the first place.

2.2 Evolution

2.2.1 Evolution is heritable change

Evolution denotes any change in the frequencies of heritable features within a population of reproducing entities, over generations. In practice, it is useful to regard it as a shorthand for "descent with modification", Darwin's own term.[1] This change may consist in loss of existing features, creation of new features, or alteration in the frequencies of existing features.

---

[1] It may be worth pointing out that the word "evolution" does not appear anywhere in the first editions of the *Origins of Species.* The word "evolved" appears once, as the final word of the entire text.

Importantly, the term "evolution" in itself does *not* carry any implication of adaptation, improvement or fitness. Evolution in the general sense is simply heritable change, nothing more, nothing less. The nature and direction of this change depend on the forces that drive it; these forces may be adaptive (e.g. natural selection), but also unguided (random genetic drift), or even potentially harmful (e.g. inbreeding through loss of diversity caused by genetic drift). It is important to distinguish between evolution on the one hand, and the mechanisms that cause or drive evolution on the other hand.

In a biological context, evolution is essentially the negation of the fixity of species: species change over time, rather than remaining fixed for all eternity as was commonly held before Lamarck and Darwin.

2.2.2 Conditions for evolution

In general, evolution occurs whenever three conditions are met, namely variation, multiplication and heredity.[2] To make these terms more precise, we may consider them in a different order, putting multiplication first. *Multiplication* is self-explaining in a biological context, but it may be useful to generalise the concept: it essentially means that new individuals are constantly being created, in such a way that they can somehow be associated with certain already existing individuals. New individuals are called "children" or "offspring", while the existing individuals with which they are associated are called "parents". Multiplication, or reproduction, is simply the creation of new individuals, each associated with a certain set of existing parents. An individual's parents, their own parents, their parents' parents, and so on, are this individual's ancestors. All individuals that have a certain individual as one of their ancestors are this individual's descendents.

Variation implies that the offspring will not be exactly identical to their parents, or to each other. Rather, they will occasionally differ from their parents and from each other in various ways, exhibiting certain characteristics which their parents and siblings do not possess. *Heredity*, the crucial factor which separates evolution from mere variation, means that some of this variation will in turn be transmitted to the children's own descendents, who will also exhibit these specific, different characteristics. Taken together, these three conditions ensure that the inherited characteristics of a population will change over time, and therefore, that evolution will occur.

Note that this account includes all possible pathways of heredity: genetic, epigenetic, cultural, etc. In practice, biological evolution is only concerned with biological features, transmitted through biological reproduction. Other fields, such as cultural evolution, may have different requirements.

Darwin did not invent the concept of evolution (as opposed to the fixity of species), though his work was instrumental in making it the generally accepted viewpoint rather than the fringe theory it had previously been. Darwin's fundamental contribution is to have identified a workable mechanism for evolution, that could explain the prodigious amount of adaptive complexity observable in Nature without invoking supernatural forces. This mechanism, of course, is natural selection.

---

[2] These three requirements have a long history: Wallace [182] already mentions them in this form.

2.3 Natural selection

2.3.1 What is natural selection?

Evolution is a change in the heritable features of a population of reproducing individuals. New features appear, some existing features are lost, and others become more or less common throughout the population. These changes in frequency may be caused by various factors, including random fluctuations. *Natural selection* is one of the forces that influence the direction of this change.

The general mechanism of natural selection is easily formulated: due to their interactions with the environment, certain heritable phenotypic characters have an impact upon the reproductive success (the "*fitness*") of their bearers, or of their bearer's children, or of other entities which also bear this character. But because these characters are heritable, they are themselves replicated through reproduction. As a result, heritable phenotypic characters which, through their interactions with the environment, favour the direct or indirect reproductive success of their bearers, will mechanically become more and more common within the population. Conversely, heritable phenotypic characters which impede the reproductive success of these entities tend to decrease in frequency. Ultimately, such characters may become fixated (that is, be present in every entity within the population), or on the contrary disappear entirely from the population.

Therefore, natural selection is the portion of the variance in fitness of entities, that is caused by the interaction between heritable phenotypic characters and the environment. This can be summarised (provided the accurate meaning for all terms is kept in mind, see next section) as *variance in fitness caused by heritable features.* Alternatively, the term "natural selection" may also refer to the set of forces in the environment which affect the replication of entities by interacting with their heritable phenotypic characters. These two viewpoints are, in a sense, mirror images of each other, and designate two perspectives on the same process.

Thus, to the three previous requirements of evolution, natural selection adds a further requirement: heritable variation in reproductive success, caused by interaction between the environment (in a very wide sense, see below) and heritable phenotypic characters. Whenever these four requirements hold, the Darwinian mechanism of evolution by natural selection will arise.[3]

2.3.2 On the importance of terms

Having described natural selection as the variation in "fitness" caused by the interaction between "heritable" "phenotypic" characteristics and "the environment" only takes us so far; several of these terms are highly ambiguous, and we must make clear what meaning we ascribe to each of them.

Importantly, the term "environment" here is understood in a very wide sense, covering absolutely everything that can interact with the considered characteristics. This includes other characteristics of the same individual or of others, or copies of the same characteristics in other individuals. It is quite clear that a certain characteristic will only have predictable results in

---

[3] Lewontin[106] famously offered a slightly different formulation of evolution by natural selection. Lewontin's formulation concentrates the four requirements into three, essentially by making "multiplication" implicit in the language. The formulation presented here seeks to emphasise the distinction between evolution, natural selection, and evolution by natural selection.

a certain background, both environmental and genetic: the exact same characteristic might have entirely different results if it occurs in a different environment, or together with different characteristics. The link between the characteristic itself, and the effect it produces on the replication of its bearers, can be highly indirect. [35]

It is also important to include the term “heritable”. The requirement of heritability is of course essential: natural selection is simply not concerned with characteristics that are not heritable. Overlooking the requirement for heritability is the root cause of a misguided, but persistent criticism of evolution by natural selection, namely the accusation that it amounts to a tautology: “better survivors survive better” (see Appendix C.)

The term “phenotypic” also requires some precision. The phenotype of an entity includes all its intrinsic observable characteristics, that is, the entire makeup of the entity - including “morphologies, physiologies, and behaviors” [106]. Importantly, the term “phenotypic” here is *not* in opposition to “genotypic”. Purely genic selection can certainly occur: genes, in their physically embedded form as nucleic acid, certainly do have their own “phenotype” that can interact with their environment (the cell) independently of their effects on the larger organism, replicating as selfish DNA (gene-level selection is discussed in Appendix D). Rather, the term “phenotypic” is used to exclude external, arbitrary labels that would somehow be heritable, but utterly independent of the entity’s makeup.

A simple scenario will immediately explain the necessity of this distinction. Consider a population of individuals, mating and reproducing in a completely random fashion. Now, within this population, we arbitrarily choose a given individual A, and then we decide that all descendents of A will be *favoured*, in that we will allow (through direct manipulation) these individuals to survive longer and reproduce more than others - independently of their features. We may call this process “aristocratic selection”, considering that individuals are favoured for no other reason than their ancestry. Now the fact of being a descendent of A is clearly a heritable characteristic. Moreover, this characteristic certainly does interact with a part of the environment - namely, the experimenter. Finally, considering the favouring rule, it clearly affects the reproductive success of those that possess it. However, we would not want to classify this scenario as a valid example of natural selection, because the actual makeup and features of the individuals would not be the source of their reproductive success. For example, in this scenario, no *adaptation* of any kind would be taking place. The Darwinian mechanism of evolution by natural selection bears on the intrinsic characteristics of entities. While the “aristocratic selection” process described above will certainly increase the proportion of descendents of A, and therefore, may also increase the frequency of some of A’s own heritable features, this increase would be entirely incidental and not caused by the features themselves.[4]

Perhaps unsurprisingly, the terms “reproductive success” and “fitness” require particular caution. In the above account, we have simply equated fitness with reproductive success, without any further precision. However, this will not suffice: in particular, we must emphasise that “fitness” here can *not* simply denote the number of children. Rather, fitness denotes the success of an individual in transmitting its inherited characteristics, not just to the next generation,

---

[4] Importantly, note that this scenario is very different from artificial selection as it is practised by breeders and cultivators. In the latter case, while successful individuals are chosen by external observers, this choice is indeed based on the features of the individual, rather than on its parentage alone.

but to *all* future generations as well - what Lewontin [106] calls "contribution [. . .] to future generations."[5]

To see why this is necessary, we need only consider the case of characteristics that do not affect the reproduction of their bearers, but do affect the reproduction of their children - independently of whether these children themselves possess these characteristics or not. An often cited example is the *grandchildless* mutant in Drosophila subobscura (mentioned in [106, p. 8]): homozygotes for this allele will have sterile children, independently of the children's own genotypes. Now it is clear that *grandchildless*-like genes are strongly disfavoured by natural selection. However, if we limit natural selection to variance in the number of children, then these effects would not fall within our definition, and thus we could not accept them as examples of natural selection - which would be absurd. This illustrates the fact that natural selection is not directly concerned about the reproduction of individuals, but more generally about the differential propagation of heritable characteristics, of which individual reproductive success is only a component. If natural selection optimises anything with regard to individuals, it is not their immediate reproductive success, but their capacity to transmit heritable characteristics to future generations. Contribution to future generations is thus the defining aspect of selective fitness.

Finally, we must stress that the heritable features we are talking about should *not* be equated with genes or alleles. This is for at least two reasons. First, evolution and natural selection can occur at many levels, and heritable features can have many different forms at various levels. Second, and most importantly, certain important biological traits do *not* have a genetic basis; endosymbiosis (the insertion of one organism into another, resulting in a single reproducing entity) is probably the most obvious example.[6]

2.3.3 Individual adaptation: a consequence of natural selection

One simple way in which a heritable feature can influence its own rate of propagation is by affecting the capacity of its bearer to survive and reproduce. If a certain heritable feature enhances the chances of its bearer to survive and reproduce, that is, its "individual fitness", it will mechanically increases its own chances of being propagated, precisely because it is heritable. It follows that novel characteristics which enhance the chances of survival and reproduction of their bearer (that is, their bearer's "individual fitness") will tend to become widespread, while those which decrease their bearer's fitness will tend to die off. As a result, evolution will be driven in the direction of *adaptation*: over time, as fitness-increasing features propagate, individuals will become increasingly "adapted" to their environment, in the sense that they will tend to become better at surviving and reproducing within it.[7]

---

[5] The term "fitness" is notoriously ambiguous [35, Chap. 10]. As a fitting illustration of this ambiguity, Lewontin [106] uses this same word to denote both "rates of survival and reproduction" and "contribution [. . .] to future generations" - in two successive sentences.

[6] See Appendix D for a closer look at the levels of selection and the gene's eye view of evolution.

[7] See also chapter 9, section 9.3.1.

2.3.4 Natural selection is not just individual adaptation

This aspect of evolution by natural selection (individual adaptation for survival and reproduction) is simple and intuitive, and corresponds to the everyday concept of Darwinian evolution: the process by which well-adapted individuals survive and reproduce more than others, and therefore, mechanically propagate the (heritable!) features that caused their success. However, it is important to note that this process of individual adaptation is *not* the whole story of natural selection - far from it. Heritable features can increase their overall rate of replication without directly enhancing the reproduction of their individual bearers. In fact, it is perfectly possible for natural selection to favour a certain characteristic which actually reduces, or even annihilates altogether, the probability that its bearer will survive and reproduce - if in doing so it can favour the production of other entities which also bear this feature, and the overall balance results in a net increase in replication.

This is readily seen in the widespread occurrence of *altruism* in Nature. Altruism occurs whenever certain entities perform a certain action which reduces their own individual fitness, but increases that of others. Social insects offer many striking examples of altruism, such as stinger bees that sacrifice themselves to defend the nest; but more generally, many eusocial animals curtail their own reproduction to benefit their colony (including mammalians such as naked mole rat). Sober and D.S. Wilson [167] use the lancet fluke (*Dicrocoelium dendriticum*) as a model case of altruism: when these parasitic worms infect ants, one of the parasites will migrate to the ant's brain and create a structure that will alter the ant's behaviour, causing it to move to exposed positions (such as the tips of grass blades) during nighttime. This increases the chance that the ant will be eaten by mammals, which is necessary for the next stage of the worms' life cycle. However, the migrating worm that created this behaviour by invading the ant's brain will *not* be able to infect the mammalian host: it effectively sacrifices itself for the good of fellow parasites within the same ant host.

Despite the (individual) fitness penalty that they incur, altruistic features can be maintained and propagated by natural selection. This occurs if, and only if, the sacrifice actually improves the fitness of other bearers of the same feature, and the overall balance turns out to be positive. Some of the ways in which this can occur are *kin selection* (the sacrifice is beneficial to close parents of the altruist, which are more likely to possess copies of the same feature), *group selection* (the sacrifice is beneficial to members of the group, so that groups which include altruists will thrive) and *reciprocal altruism* (the sacrifice is only directed towards other individuals that will also perform a similar sacrifice). All three mechanisms may result in a sacrificial feature being positively selected, due to the overall balance of positive and negative impacts on the fitness of the bearers.

At the other end of the spectrum, it must also be remembered that natural selection can act on individual genes, independently of (or in some cases, despite) their effects on individual organisms. Purely genic selection is readily seen in meiotic drive, whereby certain alleles manage to "cheat" the machinery of meiotic division to ensure their own presence in the gamete, defeating the fairness of meiosis. Such alleles can reach high frequency within a population, even if they decrease the fitness of homozygous individuals (the t-allele in mice, mentioned by Lewontin [106] is a well-known example).

Another clash between actual natural selection and intuitive notions of individual "adaptation" occurs in *sexual selection*, that is, the component of natural selection that is caused by mate preferences among the opposite sex. Through sexual selection, a detrimental feature (a feature that, all other things being equal, would decrease the chances of survival and reproduction of an individual) can still be selected for, because it increases the probability that its bearer will be chosen as a mate by members of the other sex. Such a process can result in a runaway process of self-reinforcing selection, as explained by Fisher [58, Chap. 6]: as soon as a given feature elicits even mildly higher interest from the other sex, then bearer of this feature will have higher reproductive success; this very fact, in turn, will create an incentive for other females to prefer this feature, since this will lead them to have more successful sons. This positive feedback loop proceeds until the feature grows to such extreme proportions that the handicap is too large for any further increase to occur. What's more, it has also been shown that in some cases, positive sexual selection for a detrimental feature can emerge precisely *because* of the handicap it incurs - the "handicap principle", suggested by Zahavi and put in formal terms by Grafen [73]. In the latter case, the detrimental feature is used as a necessarily honest marker of fitness. Miller [126] ascribes an extremely far-ranging role to sexual selection, describing the most striking features of human beings (large brains, complex language, art) as "courtship tools" (p.4) developed in response to (and in pursuit of) mating preferences.

In short, we must remember that natural selection is the differential reproductive success of entities caused by the effects of heritable features - no matter how indirect this effect (or the reproduction process itself) can be. Individual adaptation, in the usual sense of optimising the survival and reproduction of a particular individual, is just one of the ways in which this can occur.

2.3.5 Evolution, natural selection, and evolution by natural selection

It is important to point out the distinction between evolution, natural selection, and evolution by natural selection. The term "evolution" is commonly used as a shorthand for "evolution by natural selection",[8] to the effect that these three concepts are often blurred in public discourse. While shorthand expressions are of course useful, it is important to keep in mind that these three terms denote different things.

Evolution is change in the heritable features of a population. This change (which is most often caused by genetic mutations in nature) may occur in any direction and does not carry any connotation of improvement, adaptation or progress. Natural selection is one of the forces that operate on this change, and guide the trajectory of evolution. It is intrinsically an adaptive force, and does tend towards individual adaptation. Evolution by natural selection is, obviously, the component of evolution that occurs because of natural selection. This last concept is precisely the mechanism proposed by Darwin, often designated by the single word "evolution".

---

[8] For example, Taylor [173, Chap. 7.3.3] writes: "Evolution can explain how self-reproducers come to be adapted to their environment..."

2.4 The "Blind Watchmaker" effect: a pathway to the very improbable

In the preceding sections we saw how natural selection can guide evolution in the direction of higher adaptation, for individuals or other entities at many levels. However we have not emphasised what is perhaps the most momentous consequence of the Darwinian mechanism, namely the generation of features which are not only well-adapted, but also highly improbable. Evolution by natural selection does not just explain design: it explain enormously, extraordinarily, awe-inspiringly efficient design, relying on the the microscopic interactions of a multitude of complex chemicals. Dawkins [36, 35, 39, 38] has famously devoted much of his writing to this particular question, introducing the metaphor of a "Blind Watchmaker" to describe the activity of (evolution by) natural selection. Common examples of evolutionary excellence are animal organs, such as the vertebrate or molluscan eye, or the sonars of bats. Gould [68, Chap.3] discusses the remarkable adaptation of a certain anglerfish, which uses an appendage on its head as lure to attract prey: the lure imitates the form of a small fish with eerie precision.

How can blind, random mutations produce such a concentration of efficiency, such vastly improbable structures - precisely the very thing that would least be expected to occur out of random exploration? The answer, of course, is that evolution by natural selection is emphatically *not* random. Mutations and variation are random; natural selection is the very opposite of randomness. The (originally random) novelty created by mutations is non-randomly filtered by natural selection, which adds information to the process, and thereby turns this initially random novelty into genuine, adaptive *creativity*.

What makes this process so powerful is the fact that it is cumulative. Mutations create random heritable variations between entities. Those that are harmful are quickly discarded over generations, while those that are beneficial are retained. Once these small improvements are incorporated into the heritage of the population, *they become the new starting point*, and the process starts again: many random variations occur, most are discarded, occasionally one is retained and propagates, etc.

Thus, by slowly accumulating small, "believably lucky" jumps (and by non-randomly filtering them through natural selection), Darwinian evolution may in time generate arbitrarily improbable adaptations, which could never arise all at once. This principle will feature prominently in chapter 8, which discusses the relationship between evolution and complexity.

2.5 Conclusion

We have provided a usable answer to the question: "what is evolution?" In doing so, we sought to identify (and dispel) certain ambiguities and misunderstandings that could arise from basic evolutionary concepts. In particular, we have come up with the following definitions and principles:

1. Evolution is a change in the heritable features of a reproducing population, over generations.

2. Natural selection is the variance in reproductive success of entities that is caused by the interaction of heritable features with the environment. “Reproductive success” must be understood in the general sense of contribution to all future generations, rather than to the next one alone; “the environment” includes anything the feature might interact with, including other features of other individuals or of the same individual.

3. Alternatively, natural selection is the variance in replication (through biological reproduction) of features that is caused by the interaction between these features and the environment.

4. Evolution is change; natural selection is one of the mechanisms that guide this change. Evolution by natural selection, the portion of evolution that is caused by natural selection (as opposed to, say, random drift) is the Darwinian mechanism that explains the profusion of adaptive diversity in Nature.

These definitions provide a firmer basis for future inquiries. In the following chapters we will build upon these concepts to improve our understanding of evolution, and of how it can be put to work.

Chapter 3

Coevolution

3.1 Coevolution in nature and in computers

3.1.1 Coevolution: reciprocal evolutionary change and coupled fitness landscapes

Coevolution is a generic name covering all situations in which one or several species evolve in response to each other's evolution - that is, to situations of *reciprocal evolutionary change*. This occurs when the reproductive success of each individual is determined, at least in part, by evolvable characteristics of others, or in other words, when populations evolve on *coupled fitness landscapes*.

Consider two evolving species A and B, such that the reproductive success of each individual in one species is determined at least in part by the outcome of its interactions with individuals of the other species. If species A undergoes an evolutionary change, this may affect the fitness landscape of species B. This, in turn, may lead (through the agency of natural selection) to a change in the heritable characteristics of species B; this will in turn alter the fitness landscape of species A, possibly triggering another evolutionary change, etc. Thus coupled fitness landscapes may lead to reciprocal, or "echoing" evolutionary change.

Biologists define coevolution as reciprocal evolutionary change. In artificial evolution, however, the term seems to have been slightly extended, to encompass any situations in which lineages are evaluated after the outcome of their interactions.

3.1.2 Enter the Red Queen

In 1973 Leigh Van Valen tried to determine the distribution of survival times for different taxa ([177], described in [150]). He found this distribution to be roughly inverse-exponential: a *constant proportion* of taxa become extinct after any duration of existence. Thus, for Van Valen, the probability of extinction is independent of the taxon's age. Regardless of how long it managed to survive, a taxon essentially has the same probability of going extinct over time.

To Van Valen, a plausible cause for this fact was coevolution. In Nature, species do not merely evolve against a fixed or slowly changing environment: they coevolve with each other, because their fitness is based in part on their interactions with other evolving species. In particular, a consequence of coevolution is that, as a given species evolve and improves its fitness, other species evolve too, which may negate this improvement. Thus each species is pitted against a constantly deteriorating environment, i.e. an environment that constantly changes so that the current state is not as optimised as it was[1]. Thus, like the Red Queen in Lewis Caroll's *Through the Looking Glass*, species must constantly run (that is, adapt) just in order to survive; those who fail, become extinct.

3.1.3 Arms races in coevolution

However, if constant mutual adaptation occurs, should this mean that lineages are becoming in some way 'better' over time ? Does constant mutual adaptation between coevolving lineages lead to ever-increasing levels of performance and efficiency in the long run ?

This assumption is indeed as old as evolutionary theory itself, and seems to have been suggested by Darwin as a direct consequence of his theory of natural selection; for example, in passages such as Chapter IV of the *Origins of Species* [34], Darwin described how species indigenous to small, isolated environments such as islands, in which competition is limited, are often displaced by new species originating from a larger piece of land in which they had been exposed to more competition. This assumption of progress through mutual adaptation is the basis for what Dawkins & Krebs have called the *arms race* hypothesis [41]:

> As the arms race progresses and predators "improve", this does not necessarily mean that they catch more prey. The prey lineage, after all, is improving too. There seems to be no general reason to expect the average success of animals at out-running or out-witting contemporary enemies, victims, prey or competitors, to improve over evolutionary time. Van Valen has put this point more generally in his "Red Queen Hypothesis". But if modern predators are in general no better at catching modern prey than Eocene predators were at catching Eocene prey, it does at first sight seem to be an expectation of the arms race idea that modern predators might massacre Eocene prey. And Eocene predators chasing modern prey might be in the same position as a Spitfire chasing a jet.

In summary, while coevolution implies constant mutual adaptation and ever-changing environments to which lineages must constantly adapt, the 'arms race' concept posits that this constant mutual adaptation may at times result in some form of *cumulative* improvement over time.

This apparently simple proposition actually expresses a rather bold hypothesis. Natural selection is fundamentally a *local* process, both in space and time: it is based on differential success in the propagation of genes within a certain environment. When the environment changes (and the basis of coevolution is precisely that it changes constantly), so does the direction of the selection pressure: through natural selection, organisms become increasingly well-adapted to their *current* conditions, regardless of previous or future conditions. These adaptations, in turn, represent a change in the environment of co-evolving entities: the task for which they must be optimised has changed. It is not immediately obvious that this succession of local

---

[1] This description of the Red Queen effect owes much to the online lecture notes of David Rand at Brown University.

optimisations should mechanically lead to long-term, overall progress. At the very least, a more precise description of the concepts of "progress" and "performance" would help clarify the issue.

3.2 Coevolution in computers

3.2.1 Sorting networks and chasing robots

Coevolution has been introduced in artificial evolution as an alternative to traditional evolutionary methods based on fixed, explicitly defined fitness functions such as the genetic algorithm. The application of coevolutionary methods to optimisation problems was bolstered by initial successes such as those reported by Axelrod [8] (on the coevolution of strategies for the Iterated Prisoner's Dilemma), and Hillis [80] (on the coevolutionary optimisation of sorting networks). Hillis' experiments on sorting networks, in particular, drew much attention, because they took on a well-known problem in computer science (optimising sorting networks) and showed how coevolution brought a clear improvement over a straightforward genetic algorithm based on a hand-designed fitness function.

It is worth pointing out several aspects of Hillis' work. Firstly, we note that in this situation, coevolution is not imposed by the nature of the problem, but is used as an alternative to an existing 'naked' evolutionary method: it is possible to construct a simple hand-designed fitness function, but coevolution is used instead for efficiency reasons. This stands in contrast to much work in coevolution (e.g. predator-prey competition or games) in which there is no obvious way to build an external fitness function by hand, and the nature of the problem naturally suggests a competitive evolutionary mechanism. In particular, it implies that the fitness landscape of this particular problem does possess a simple underlying scale of performance (captured by the fixed fitness function), upon which individuals can be meaningfully and linearly ranked (according to their score under this fitness function). As we will see, the existence of such a unidimensional scale of performance to be optimised is not obvious a priori in coevolutionary situations - at least not until the concept of superiority has been cleared up.

Another aspect of this work is that the experimental settings reported by Hillis were not identical between evolutionary and coevolutionary experiments. In particular, the coevolutionary experiment made use of topological information in the selection of opponents, laying both populations on one single grid and matching individuals at the same location on the grid only.

3.2.2 Arms race in artificial coevolution

These considerations did not temper the enthusiasm of the artificial evolution community. Coevolution was adopted as an optimisation method offering specific advantages upon standard evolution based on externally defined fitness. This view of coevolution as an optimising mechanism was based mostly on a widespread application of the "arms race" hypothesis. Rosin & Belew [152] summarise the transposition of the arms race concept to artificial evolution:

> Since the parasites are also evolving with a fitness based on a competition's outcome, the success of a host implies failure for its parasites. When the parasites evolve to overcome this failure, they create new challenges for the hosts; the continuation of

> this may lead to an evolutionary “arms race” (. . .) New genotypes arise to defeat old ones. New parasite types should serve as a drive toward further innovation, creating ever-greater levels of complexity and performance by forcing hosts to respond to a wider range of more challenging parasite test cases.

Ficici & Pollack [55] posit that “the key to successful coevolutionary learning is a *competitive arms race* between opposed participants.” Nolfi & Floreano [133] use an even bolder interpretation, by mentioning that “competing populations may reciprocally drive one another to increasing levels of complexity by producing an evolutionary arms race.[2]” They also attempt to describe explicitly some of the advantages brought by coevolution:

> First, the co-evolution of competing populations may produce increasingly complex evolving challenges. As discussed by Dawkins and Krebs, competing populations may reciprocally drive one another to increasing level of complexity by producing an evolutionary “arms race”(. . .) As Rosin and Belew point out, it is like producing a *pedagogical* series of challenges that gradually increase the complexity of corresponding solutions (. . .) This nice property overcomes the problem that if we ask evolution to find a solution to a complex task we have a high probability of failure while if we ask evolution to find a solution first to a simple task and then for progressively more complex cases, we are more likely to succeed (. . .)
>
> Secondly, because the performance of the individual in a population depends also on the individual strategies of the other population which vary during the evolutionary process, the ability for which individuals are selected is more general (i.e. it has to cope with a variety of different cases) than in the case of an evolutionary process in which co-evolution is not involved (. . .)
>
> Finally, competing co-evolutionary systems are appealing because the ever-changing fitness landscape, due to changes in the co-evolving species, is potentially useful in preventing stagnation in local minima. From this point of view, co-evolution may have consequences similar to evolving a single population in an ever-changing environment.

Let us try to clarify each of these points. Firstly, coevolution is expected to create a gradual selective environment, in which evaluation is initially mild and increases in difficulty over time as coevolving populations improve. The rate of increase in difficulty is guided by evolution alone. This stands in contrast with fixed hand-designed fitness function, or adaptive fitness functions in which a specific parameter is modified over time according to a pre-defined mechanism. This argument, however, disregards the possibility that one of the populations could out-evolve the other in such a way that any of its member would be able to utterly defeat any opponent, which would lead to a *disengagement*, or *loss of gradient* (a term used by Watson & Pollack [183]) and negate the ’adaptive evaluation’ argument. In Nature, this would correspond to an extinction

---

[2] The concept of arms race does not imply any increase in complexity, merely in overall mutual efficiency between two or more lineages. For the relationship (or lack thereof) between evolution and increases in complexity, see chapter 8.

event; we note that the basis of Van Valen's argument is precisely that such extinctions occur routinely throughout evolutionary history, by the very effect of coevolution.

The second argument states that exposure to many different opponents will produce a selective pressure towards more general behaviours, that is, strategies that are able to perform well against a wide array of opponents. We note, however, that in experiments 'in which coevolution is not involved', such generality is often irrelevant for the simple reason that there is only one problem to solve. The very concept of 'generality' implies the existence of many possible adversaries, which naturally suggests a competitive evaluation mechanism.

The third argument is that constant evolutionary activity may help escape local optima: since the fitness landscape is constantly changing, it is expected that coevolving species will not be stuck in local optima because such optima will dissolve under the constant action of the Red Queen effect. While this is a direct consequence of the Red Queen hypothesis, it is not clear how this translate in terms of global, overall progress: coevolution certainly helps maintain movement, but movement to *where*? It is quite possible for a population of coevolving individuals, or a group of coevolving populations, to get stuck into a small region of the search space, corresponding to a local optimum of the coevolutionary process (a phenomenon which some authors have called "collusion" or "mediocre stable states" [142]).

3.2.3 Finding opponents

A recurrent question in coevolutionary algorithms is to determine which individuals should be matched for each run. It is clear that if an individual is to be evaluated after the result of its competition with another one, individuals competing against weak (resp. strong) opponents will be unduly overrated (resp. underrated).

The most reliable evaluation method would be to evaluate each agent against each other, but in a single population of size N, this would require $(N^2 - N)/2$ evaluations, which may quickly become intractable for larger populations. Inversely, pitting each individual against only one opponent would just require $N/2$ evaluations, but would only give a very unreliable estimation of each individual's capacities.

A solution to this problem is to match each individual with only a few opponents chosen at random (e.g. the players for the game of tag evolved by Reynolds, [149], or Gomez and Miikkulainen's ESP method for the cooperative coevolution of neurons in a neural network [66]). In cases where only one species is considered, i.e. when all individuals are functionally equivalent, a direct elimination tournament can also be held. [6]

A more popular method is to pit all individuals from a group against the best individual found at previous generation, which implies only N matches. This is the method chosen by Sims [162], which has since been called the Last Elite Opponent (LEO) algorithm by Cliff and Miller [32] (see also Appendix B). In this algorithm, each creature from each species is evaluated by competing against the current champion from the other species. The resulting score is used as a fitness value for selection and reproduction among individuals from this species, as well as to choose the new champions of this species. The same process is then applied to the other species, and starts again. This algorithm has proven very successful and has been used in several other studies, as we will see.

These methods offer various trade-offs between fairness (offering the same level of difficulty to each individual), generality (the competitive environment should be as diverse as possible) and computational cost. While the number of matches is important, the average number of evaluations undergone by each individual is also significant. For example, in a direct tournament, while $N-1$ matches are needed, each individual undergoes between 1 and $log_2(N)$ evaluations. In the "all-versus-previous-best" cycle, N matches are needed, and each individual is evaluated only once; however, the evaluation process is much fairer than a direct elimination tournament, in which competing against a very good individual in the first round will lead to an unfairly low score. See Reynolds [149] for an overview of these differences between matching choices.

In order to maintain as much diversity as possible in the competitive landscape, one might want to select opponents that defeated particularly efficient individuals, even though they may not have defeated as many individuals as others. This principle of "competitive fitness sharing" is described by Rosin & Belew [152]. Note that this process can be applied to all competing populations simultaneously.

3.2.4 Cooperative coevolution

Note that coevolution needs not be competitive. A cooperative coevolution method was introduced by Potter and De Jong [143] under the name of Cooperative Coevolutionary Genetic Algorithm (CCGA). The CCGA and its derivatives try to optimise a system by coevolving its components. They share many problems of usual coevolutionary algorithms, such as the choice of partners: Wiegand and colleagues [187] point out that the choice of partners for evaluation can affect the outcome of the process in unpredictable ways, illustrating this with simple experiments. However, in these methods, the centre of attention has changed: we do not simply expect to see "improvement" from each individual, we want to optimise a whole system, according to an existing global fitness function. This difference is important, because it may simply be more efficient to directly evolve entire teams (with suitable genetic operators), rather than coevolve individuals. We discuss this point, with an experimental illustration, in a previous paper [121]. An authoritative discussion of cooperative coevolution is provided in Wiegand's Ph.D. thesis. [186]

3.3 Pitfalls in coevolution

3.3.1 Roadblocks on the march to progress

While the arms race concept provided a justification for the use of coevolution in optimisation, it soon became clear that the expected march to progress was not guaranteed. The most obvious threat to the buildup of cumulative evolutionary adaptation is the problem of *intransitivity* in the global fitness landscape: : if an organism A can be said to be superior to B, and B is superior to C, it is not necessarily the case that A should always be superior to C. If B evolved against C, then A evolved against B, it does not mechanically follow that the adaptations that allowed A to defeat B would allow it to defeat C.

It has been pointed out by several authors that such intransitivities could lead to "cycles" [133] or "circularities" [170] in which three or more strategies exhibit a circular superiority relationship ($A1 > A2, A2 > A3, A3 > A1$; or in the case of two competing species, $A1 > B1, B2 > A1, A2 > B2, B1 > A2$), leading to cycles in the trajectory of evolution To quote Nolfi & Floreano, "co-evolving populations may cycle between alternative class of strategies that, although they do not produce advantages in the long run, may produce a temporary improvement over the co-evolving population." [133]

While such cycles may appear (and may even have been identified in Nature, e.g. in the case of Californian lizards with orange, blue and yellow throats [164]), the problem of intransitivity is more general. In coevolution, the reproductive success of individuals is determined by their interactions with their current opponents. The particular features of these opponents may elicit certain specific adaptations from the population. However, there is no guarantee that these adaptations, developed against current opponents, will prove helpful against future opponents. It is perfectly possible that coevolving populations keep mutually adapting to their current competitors alone, without any improvement in their general capacities. This phenomenon has been encountered several times and is known under many names, including "relativism" [183], "mediocre stable states" or "collusion" [142], and "Red Queen dynamics" [138] (see section 3.1.3).[3]

Other pitfalls have been identified in coevolutionary progress. We have already mentioned the possibility for a dramatic "loss of gradient" [183] or "disengagement" if one of the coevolving lineages out-evolves the other: if one of the populations becomes so strong that any of its member would be able to utterly defeat any opponent, no gradient remains to guide natural selection.[4]

Yet another threat is that of *parasitism* or *opportunism*: some individuals may exploit one specific weakness of others, thereby reaping an 'easy gain' (which Watson & Pollack [183] call "focusing on the wrong thing"); this may allow primitive, sub-optimal strategies to thrive at the expense of more complex opponents. Although this mechanism has often been mentioned as a temporary hindrance on the road to progress (e.g. by Stanley & Miikkulainen [170]), in the course of our experiments we have been able to show that this may have dramatic consequences and often leads to the displacement of promising sophisticated strategies by trivial opponents, especially if the algorithm considered is prone to such exploitation. [125]

3.3.2 When progress is not progress

Even when an arms race does actually occur, we are not at the end of our troubles. Nolfi & Floreano [133] have shown that sustained arms races do *not* necessarily result in optimum performance. They performed two coevolutionary experiments based on a predator-prey scenario, with one important difference: in one run, coevolution occurred in a straightforward manner, by pitting individuals of a given generation against the current champion of the opposing population

---

[3] The term "Red Queen" dynamics may be slightly inappropriate: Van Valen's original Red Queen hypothesis does not preclude any long-term progress, and indeed this point is precisely the basis of Dawkins and Krebs' arms race hypothesis. For Van Valen, "staying at the same place" really means "staying alive", rather than a stagnation in overall abilities.

[4] Notice that if competitions result in real-valued scores, rather than binary "victory/defeat" outcomes, then a gradient may still exist even though one species consistently outperforms the other: we simply need to compare the results of two individuals against a given opponent. For example, using Sims' algorithm on the tasks presented in this work, we have never encountered loss of gradient - simply because the real-valued scoring functions always offered *some* kind of gradient even to hopelessly maladaptive species. We explored this effect further in another paper. [125]

(essentially Sims' Last Elite Opponent algorithm [162], see above).

In the other run, however, individuals of a given generation were pitted not only against the current opposing champion, but also against the *previous* champions of the opponent population. This "Hall of Fame" technique, suggested by Rosin & Belew [152], forces new individuals to prove their worth not only against the current opponents but also against previous champions. Clearly we can expect that such a selection regime will favour the onset of arms race, that is, situations in which current individuals outperform their ancestors, since with this method, superiority over ancestors becomes part of the selection criteria.

Unfortunately, as Nolfi & Floreano have pointed out, the very fact of maintaining a gallery of permanent opponents in the selective process decreases diversity in the challenges that must be overcome by new individuals, possibly leading to what we might call an "ossification" of the fitness landscape. This makes the algorithm increasingly closer to a standard genetic algorithm based on a fixed fitness function, as time goes on and the addition of each new individual represents a smaller and smaller proportional increase of the archive (and thus a smaller and smaller injection of novelty).

In some circumstances, when the authors compared the results of coevolution with and without a Hall of Fame, they found out that individuals evolved *with* a Hall of Fame were defeated by individuals evolved *without* it. This occurred despite the fact that the Hall of Fame method had demonstrably enforced a *more* reliable evolutionary progress, in the sense that individuals evolved with a Hall of Fame were significantly better at defeating their own ancestors than those without. While progress had been more steady and continuous, it had also been more limited in scope. In short, this shows that a better arms race does *not* necessarily lead to better individuals.

3.3.3 Implicit assumptions

From this perspective it seems that coevolution is fraught with many disparate, unrelated problems, which prevent arms race from occurring or, if they do occur, reduce their supposed efficiency. Coevolution does not do what early practitioners expected it to do. But at this point it might be useful to ask a basic question: what is it *exactly* that we expect it to do?

The discrepancy between expectations and outcomes does not result exclusively from technical limitations, which could be addressed by some kind of algorithmic manipulation. In recent years various authors have pointed out fundamental conceptual ambiguities and assumptions in early coevolutionary research. In short, much of the early work in coevolution has overlooked the highly ambiguous nature of the concept of 'progress' in evolution, a notoriously controversial topic in mainstream evolutionary theory (as was emphasised in recent high-profile debates [71, 72, 40]). This ambiguity has silently percolated into the argumentation of several authors, leading to implicit assumptions that have only recently been emphasised. [53]

3.4 The meaning of "better": superiority relations

3.4.1 Comparing coevolving individuals

If we are to obtain progress through our algorithms, and if we wish to design algorithm which reliably produce "increasing levels of performance", it would be useful to define exactly what it means to have 'better performance'. What does it mean for a given individual to be 'better' than another one? This is clearly a prerequisite for evaluating the outcome of a given evolutionary process, or for comparing the outcomes of two different evolutionary processes. It is striking that until recently most authors did not explicitly describe the notion of performance that they had in mind, though as we will see the implicit underlying notion can be inferred easily. It is only recently that Ficici attacked the subject by introducing the notion of a "solution concept", which we will describe in section 3.7 (and in more detail in Appendix E). First, we introduce a somewhat different treatment, based on the notion of a *superiority relation.* Intuitively, to have a clear idea of what we want coevolution to do, we need to make clear exactly which individuals should be preferred to which others, that is, define a superiority relation according to which we can decide whether one individual is to be preferred over another.

3.4.2 Superiority is always relative to a certain set of opponents

In the domain of symmetric competitive coevolution, all individuals are evaluated by direct opposition against each other. This seems to provide a clear basis for comparing any two individuals: if A defeats B, then A is in some sense "better" than B. However, this apparent simplicity is tainted by the problem of intransitivity: let us imagine a situation in which a given individual A is able to defeat many (or even all) opponents, except for one single individual Z, which cannot defeat any opponent except A itself.[5] If we accept direct competition as the fundamental criterion of superiority, we must say that Z is better than A, and therefore that a coevolutionary algorithm that would converge to Z should be preferred to one that would converge to A. This does not seem to fit the intuitive notion that underlies the word "performance" as used in the sentence above.

This hints at the notion that superiority between any two individuals cannot be estimated by isolating these two individuals, independently of their results against other opponents. Rather, the superiority of an individual A over an individual B can only be determined (or even meaningfully discussed) with regard to a common set of opponents. Only by comparing their performances against certain other opponents can we usefully call A or B "better" than the other. The particular set of opponents *over which* this superiority is decided is as important as the particular criterion which is used to calculate this superiority. As we will see, using different categories of opponents produces different notions of "progress" (which are often implicitly confused).

3.4.3 The need for a superiority criterion

While a common "reference" set of opponents is necessary to establish superiority between two individuals A and B, it is not sufficient: we also need a criterion according to which this

---

[5] As Stanley & Miikkulainen point out [170], parasite-host relationship offer frequent examples of large, successful organisms exploited by highly specialised organisms.

superiority must be calculated. The criterion determines how the results of the competitions between A, B and the set of opponents are combined to perform a comparison.

It seems that the intuitive notion of superiority is best expressed by the following, simple idea: A is superior to B if A can obtain significantly better scores against significantly more opponents than B. This loose definition can be made more formal: De Jong [94] (who calls this concept 'maximisation of expected utility') defines as a solution "the individuals that maximise the expected score against a random opponent". Ficici calls this the "Best Scoring Strategy" (BSS), calling it the "conventional" criterion. When taken against all possible opponents (that is, when the reference set of opponents is the entire search space), this criterion seems to cover the intuitive notion of global performance that underlies many papers in coevolutionary research, either explicitly or implicitly.

Other simple superiority criteria exist. One prominent example is Pareto dominance [56, 22]: an individual A dominates an individual B if A performs at least as well as B against all possible opponents (within a certain set), and better against at least one. It seems reasonable enough to say that A is better than B in such a situation. Unfortunately a problem with Pareto dominance is that it is often undecidable: typically, in many cases, an individual A will be able to defeat some individuals that defeat B, but will also be defeated by some individuals that B can defeat. This means that the set of non-dominated individuals will tend to be quite large, providing little information. Ficici & Pollack [56] report over 75% of the total explored space for some experiments and resort to diversity-maintenance techniques to introduce an additional selective gradient.[6]

Let us mention one last criterion for comparing individuals, namely the Nash equilibrium criterion [54]: In this context, a strategy S is a Nash equilibrium if it is its own best answer, that is, no strategy has a better expected score against S than S itself. Unfortunately, the existence of a such a Nash equilibrium is only guaranteed if we allow *mixed strategies*: a mixed strategy is a set of individual ("pure") strategies endowed with a probability distribution over these pure strategies, such that for each competitive event one of these pure strategies is picked from the set according to the probability distribution.[7] Mixed strategies may appear somewhat awkward, but Ficici [54] shows that Nash equilibria possess a highly interesting property with regard to progress, that of *monotonicity*. We discuss this point, together with other results obtained by Ficici, in section 3.7 and Appendix E.

3.4.4 Summary

To summarise this section, we note that in order to define a superiority relation between any two individuals in a coevolutionary experiment, one needs to define two components:

1. The criterion which is to be calculated to determine superiority (best average score, Pareto

---

[6] It must be noted that Pareto coevolution is especially vulnerable in the face of parasitism. For any non-trivial problem, successful individuals are likely to be defeated by simplistic, opportunistic opponents that exploit some of their weaknesses (but would easily be defeated by many other individuals). Thus, for non-trivial domains, the "non-dominated set" might not be much smaller that the search space itself.

[7] It is easy to see that if mixed strategies are not allowed, then the criterion becomes exceedingly harsh; for example any strategy that is consistently defeated by even one single opponent in the entire search space cannot be a Nash equilibrium, so only strategies which can consistently defeat all possible opponents may qualify. For even moderately complex problems, such an 'all-powerful' strategy is unlikely to exist.

dominance, Nash equilibrium, etc.)

2. The set of "reference" opponents against which this criterion is to be calculated and compared for both individuals.

3.5 Local, historical, and global progress

Once we have chosen a concept of superiority, it might seem that progress is easily defined: progress occurs if new individuals are superior to previous ones. However, as was pointed out in the previous section, useful concepts of 'superiority' can only be measured by comparing performance against a given set of opponents; for example maximum expected opportunity or Pareto dominance can only be decided between any two competitors if we make it clear over which set of opponents they should be measured. This point (also stressed by Ficici [54]) has given rise to ambiguities and confusions in the literature, to the point that some authors have used the same word ("progress", "monotonicity") to express different concepts.

Depending on the set of "reference" opponents being used, the results of the comparison will not only differ, but actually reflect different concepts of superiority. It turns out that these different concepts of superiority, in turn, underlie different concepts of progress. Part of the confusion arising in the literature can be reduced to a confusion between these various types of progress.

Let us describe the problem more specifically: we have a population of interest, composed of individuals which may be grouped into one or more species. This population changes over time according to a particular algorithm, by finding new individuals and incorporating them (or not) into the population. It turns out that in the course of an evolutionary process, at any time, the term "progress" can have three distinct meanings:

1. Local Progress: The new individuals are superior to their ancestors, when compared against their *current* competitors. This means that when newer individuals and their ancestors are evaluated (according to a defined superiority criterion) against the current set of opponents, the newer individuals are found to perform better.

   This local form of progress is the one that guides natural selection; because natural selection is a *local* process, both in space and time, local progress is the one type of progress that is directly and mechanically brought about by natural selection. Any further effect must rely on additional causes, which should be made explicit if these effects are expected to occur.

2. Historical Progress: The new individuals are superior to their ancestors, when compared against *all the competitors that have been encountered so far* by the evolutionary process. This is essentially the concept that is expressed in the term "arms race"[8].

   The local progress brought about by natural selection can easily be harnessed to drive historical progress by the use of an *archive*. An archive is a collection of previously

---

[8] Actually the canonical concept of an "arms race" as expressed by Dawkins and Krebs [41] may be understood in a more restrictive way, in that it seems to imply historical progress based on the solution concept of Pareto-dominance: new individuals must be able to defeat any opponent that their ancestors could, and then some.

encountered individuals which is maintained and enriched throughout the evolutionary process. The purpose of an archive is simply to keep *past* opponents within the *current* selective environment, thereby making local and historical progress identical. Archive algorithms include Rosin & Belew's Hall of Fame method [152], or De Jong's Incremental Pareto Coevolution Archive [42] and MAXSOLVE algorithm [94] (intended respectively for the superiority criteria of Pareto dominance and maximum expected score).[9] Archive mechanisms have also been used for tracking and analysis purposes, for example Cliff & Miller's Current Individual vs. Ancestral Opponents (CIAO) method [32]. Master Tournament matrices [133] and Dominance Tournaments [170] have been suggested both for evolutionary selection and post-evolutionary analysis.

As mentioned in section 3.2.1, several authors have claimed that maintaining a successful arms race is the goal of coevolutionary algorithms, and naturally leads to ever-increasing levels of "performance": because the use of coevolution is based on the arms race concept, a successful arms race and a successful coevolutionary run are one single concept, and the result of a consistent arms race is mechanically the expected result of a successful coevolutionary run. In this view, historical progress *is* progress, period. This explains the popularity of archive-based mechanisms, not only to guide evolution, but to assess its performance a posteriori, as in Master Tournament matrices and Dominance Tournaments.

However we have already pointed out that arms races do not seem to be absolutely identical to progress in the sense that we are really expecting. Indeed, as we will see, Ficici demonstrated that for many solution concepts (including maximum expected utility / best-scoring strategy) a perfectly successful arms race may lead the evolutionary process *away* from the overall optimum, at least temporarily (but then in this case 'temporarily' may mean 'until the whole search space is exhausted'). Furthermore, as mentioned in section 3.3.2, Nolfi & Floreano's experiments showed that using a Hall of Fame could improve historical progress, while decreasing overall performance. This highlights a discrepancy between the concept of historical progress and a third notion of progress, which seems to correspond more to the implicit objective of artificial coevolution:

3. Global progress: The new individuals are superior to their ancestors, when compared against *all possible competitors across the entire search space.* We believe that this global notion of "progress" is the one that captures the intuitive goal of artificial coevolution: it describes quite simply what we would like coevolution to do. Unfortunately it is not directly related to the other two. It is not clear that an algorithm could be devised that would directly impose this kind of progress as mechanically as archives can enforce historical progress: this would involve knowledge of unknown opponents, which is absurd. Therefore, a practical option is to examine the conditions under which local and historical

---

[9] IPCA maintains an archive of learners and an archive of tests. At every cycle, new candidate learners and tests are generated. Candidate learners are only included in the archive if no member of the archive can defeat all the tests (including newly generated candidate tests) that this candidate learner can defeat. Candidate tests are only included in the test archive if they can be defeated by some of the newly admitted candidate learners, but not by any previous member of the learner archive [42]. MAXSOLVE is similarly organised, but candidates are only included in the archive if they can defeat more tests than some of the current members; candidate tests are only kept if they are defeated by one or more learners.[94]

progress can be harnessed to obtain global progress with some degree of reliability.

It is clear that these three notions of progress are essentially based on three notions of superiority: local, historical and global superiority, that is, superiority when comparing performance against the *current* opponents, *current and past* opponents, and *all possible* opponents respectively[10]. This last notion of global superiority is especially interesting, since it naturally establishes a global scale of comparison, a single measure of "goodness" through which any two individuals could be (theoretically!) compared, and which possesses one or more optima. The difficulty, of course, lies in the fact that unless the entire search space can be sampled, this global superiority cannot be calculated directly: at any time, all that we can safely estimate is local and historical superiority. But even though this global fitness landscape is hidden, it still exists, in contrast to more pessimistic opinions which express the impossibility of such a global unambiguous scale in coevolutionary situations (see section 3.7).

The distinction is particularly important in that historical and global progress are not only different, but may even be *contradictory.* We already mentioned Nolfi & Floreano's experiment, in which an archive brought about a steadier arms race, but poorer final performance against unseen opponents. In other words, the use of an archive improved historical success, but damaged global progress. On the theory side, Ficici demonstrated that for many "solution concepts" (including maximum expected utility / best-scoring strategy), and even without discarding any information, a perfectly successful arms race may lead the evolutionary process *away* from the overall optimum, at least temporarily (but then in this case 'temporarily' may mean 'until the whole search space is exhausted'). We will discuss this in detail in section 3.7.

Let us take an example: De Jong [94] asserts that the MAXSOLVE algorithm ensures "monotonic progress" in terms of how many opponents the population can defeat. However, MAXSOLVE involves discarding candidates that might potentially be superior (in a global sense) to those currently known. De Jong justifies this as follows:

> We note that a candidate discarded in this manner could potentially solve more unseen tests than the candidates that are maintained. This does not violate the monotonicity of the archive, however, since progress is measured with respect to the set of tests seen so far. If the rejected candidate solution does indeed solve more tests, this will eventually be discovered given sufficient further search, at which point it will be included.

Thus MAXSOLVE does enforce historical progress, but does not enforce global progress. De Jong points out that any superior solution will eventually be recognised after enough information has been gathered. This is true, but not necessarily usable: it expresses the obvious fact that over time, the set of opponents "seen so far" becomes closer and closer to the entire space of possible opponents, and thus that superiority against previously seen opponent will necessarily become more and more equivalent to superiority against all possible opponents. In this trivial

---

[10] To avoid any misunderstandings, we reiterate that we are not talking about being superior *to* past, current or unseen opponents; we are talking about three different relationships of superiority, determined between any two individuals, and which differ in the fact that they are calculated *against* different groups of opponents; say, for the maximum expected score concept, any two individuals can be compared by evaluating the result of their competitions against current opponents, against past and current opponents, or against the whole opponent space.

sense, historical progress does converge to global progress - but only after a significant portion (or in fact the totality) of the search space has been explored. In the meantime, as shown by Nolfi & Floreano's experiments and Ficici's theoretical results (described below), perfect historical progress may well lead the algorithm *away* from global progress, by discarding individuals that are historically inferior, but globally superior. This highlights the fact that the term "progress", when used on its own (as it usually is), is thoroughly ambiguous.

3.6 Two fragile assumptions in early coevolution research

We can apply these concepts to some of the earlier work in coevolution. It is clear that, in addition to the absence of a clearly defined concept of superiority (as pointed out by Ficici), some of the early works on artificial coevolution are based on two implicit assumptions which may be difficult to justify in the general case:

1. Local progress implies historical progress: Coevolution is expected to lead, mechanically, to an arms race in which newer individuals are superior to their ancestors (which may mean that they are being able to defeat these ancestors, or to defeat more opponents than them, or to defeat all the opponents that they could - and then some). This assumption was qualified early on with appropriate restrictions: Cliff & Miller's efforts to "track the Red Queen" [32] or Nolfi & Floreano's discussion of intransitive cycles [133] come to mind.
2. Historical progress implies global progress: If an arms race occurs, that is, if newer individuals are consistently superior to their ancestors, then mechanically some form of global optimum is being approached. In fact this assumption often takes the form of a lack of distinction between the two forms of progress: in this view, historical progress *is* progress. While, as we mentioned, Nolfi & Floreano [133] already invalidated this assumption, Ficici was apparently first to demonstrate that it was not justified for many "solution concepts" - that is, for many superiority criteria.

3.7 Solution concepts, monotonicity and progress

Recently, Ficici has made important contributions to the understanding of coevolution. Ficici's results are summarised in a recent paper [54] and in his Ph.D. thesis [53]. We provide an extensive description and discussion in Appendix E. Here we will only summarise the main points of interest in the context of the present chapter.

Ficici introduces the notion of a *solution concept*. A solution concept is essentially the partition of the search space (or any subset of it) into two parts, one of which contains those individuals that we call 'solutions', while the other contains 'non-solutions'. Thus defining a solution concept is equivalent to stating what exactly we are looking for, before starting to look for it. To quote Ficici, "while this point may seem obvious, years of coevolutionary practice indicate otherwise." [53, Conclusion] For our purposes, "solution concepts" can be thought of very roughly as what we call superiority criteria.[11] Solution concepts include maximum expected score, Pareto dominance, Nash equilibria, etc.

---

[11] See Appendix E for a more precise explanation.

An important point of Ficici's formalism is the notion of *monotonicity*: certain solution concepts are monotonic, while others are not. Using our terminology, we can easily summarise Ficici's notion of monotonicity. A monotonic solution concept is one in which *historical progress implies global progress.* That is, for every new individual, if we know with certainty that the new individual is historically superior to the current champion, then monotonicity guarantees that the new individual will also be globally superior to the current champion. Thus, when using a monotonic superiority concept, a new individual that performs better than the current champion (under this solution concept) against known opponents, will also perform better against all possible opponents - including unseen ones.

It is important to stress that this guarantee only applies to the sequence of solutions: at any time, even if we have enforced perfect historical progress, the current champion is not necessarily the best individual found so far. It is possible that a previously encountered individual, that was not previously seen as a champion, will eventually replace the current champion after more information is gained [Ficici, personal communication]. However, we can be certain that the replaced champion will *not* become a champion again, no matter how much information is gained: demoted champions will never be reinstated. Thus, with a monotonic solution concept, evolution will never need to backtrack.

In addition, Ficici demonstrates that the "conventional" solution concept of maximum expected score, or "Best Scoring Strategy" (BSS) is *not* monotonic. Thus, with this solution concept, perfect historical progress does not imply global progress, even if we never discard any information. However, Ficici also demonstrates that the solution concept of Nash equilibrium (as described in section 3.4.3 above) *is* monotonic.

Ficici's monotonicity is clearly a desirable property. However, monotonicity comes at a cost. The most obvious is that monotonicity only "works" if we never discard any information. In addition, we must constantly compare older and newer individuals against all previously encountered opponents. Using our terminology, we can easily see the reason: these are two obvious requirements of perfect historical progress. This may quickly become intractable in practice.

Second, only two monotonic solution concepts are known: Nash equilibria and (under certain conditions) Pareto dominance. While Nash equilibria appear conceptually simple, they imply the use of mixed strategies, that is, sets of strategies from which the individual chooses randomly (according to a certain distribution) at every interaction. This requirement is rather unwieldy from an implementation viewpoint, especially considering that neither the number of strategies in the set nor the probability distribution, are fixed. As for Pareto dominance, we have already highlighted its specific problems (especially the prevalence of non-dominated individuals) in section 3.4.3. By contrast, the Best Scoring Strategy concept is simple, and intuitively corresponds to the common notion of superiority in coevolutionary optimisation. This is precisely why Ficici calls it the "conventional" concept. Thus, while Ficici's results are of considerable theoretical importance, in practice, the intuitive appeal of conventional superiority criteria may overcome the theoretical cost of non-monotonicity.

3.8 Summary and conclusion

In coevolution, the concepts of superiority and progress are not as self-evident as in fitness-based evolution and must be carefully defined. To evaluate the superiority of an individual A over another individual B, we must define two things:

1. A superiority criterion (maximum expected utility, Pareto dominance, etc.)
2. A common set of opponents over which this criterion will be assessed and compared between A and B.

In the course of coevolution, different sets of "reference" opponents characterise different notions of progress, namely local progress (superiority of newer individuals over their predecessors, evaluated against current competitors), historical progress (superiority of newer individuals, evaluated against all previously encountered competitors) and global progress (superiority of newer individuals, evaluated against all possible competitors). The first is all that natural selection is concerned with. The second can be brought about by means of an archive or similar algorithmic devices. The third is the most desirable form of progress in coevolution, but is difficult to achieve reliably. In particular, it does *not* mechanically follow from any of the previous two.

Ficici has shown that, for some solution concepts (a notion closely related to that of superiority criteria), historical progress is demonstrably equivalent to global progress: newer champions are better than previous ones, not just against previously encountered opponents, but against the entire search space. However, monotonic solution concepts (such as Nash equilibria and Pareto non-dominated fronts) carry additional requirements and limitations which may make them rather unwieldy in practice. The conventional superiority criterion (Best Scoring Strategy, or maximum expected score against a random opponent), while non-monotonic, may be seen as preferable due to its intuitive appeal and its conceptual and technical simplicity.

Chapter 4

A software platform for evolving virtual creatures

4.1 Introduction

4.1.1 Evolving (better) virtual creatures

In this chapter we introduce the experimental system that we will use to test and illustrate the concepts and methods devised in the remainder of this thesis. This system is based on autonomous creatures in a physically realistic, three-dimensional (3D) environment, similar to those introduced by Karl Sims [163, 162] (described in Appendix B). Creatures are articulated structures composed of rigid blocks and controlled by embedded neural networks. The main differences with Sims are, first, the use of standard McCulloch-Pitts neurons (instead of a set of ad hoc, complex functional neurons) and, second, an improved genetic encoding and developmental system (allowing fine-grained control of neural connections in duplicated morphological features, and replication-exaptation processes). We also describe some results of early experiments with this system. To the best of our knowledge, this work is the first replication of Sims' efforts to achieve results comparable to Sims' in efficiency and complexity, with standard neurons and realistic Newtonian physics.

4.1.2 Sims' creatures and their descent

It has now been more than a decade since Karl Sims presented the results of his experiments on the evolution of virtual creatures in a three-dimensional (3D), physically realistic environment [163, 162] (see Appendix B for a detailed description). These experiments have had a rich legacy. The evocative morphology of these creatures, and the immediate realism of their behaviour, bears an uncanny resemblance to biological organisms that can hardly be found in any other experiment in the field.

While there has been a significant amount of work in projects related to the simulation of 3D creatures, much of it has concentrated on specific areas of research such as developmental

---

Most of the material in this section has been published in previous papers. [123, 124] systems or modular robotics. For example, Bongard & Pfeifer [18] used such a system to build a model of genetic regulation in development, based on genetic regulatory networks. Because the study of regulation networks was the central subject of this work, resulting creatures were rather limited in their behaviours and did not exhibit the efficiency and nature-like aspect of Sims' - neither of which were objectives of this study. More recently, Bongard and colleagues have explored new directions in the joint evolution of morphology and behaviour: real-world robots engage in continuous self-modelling and self-simulation, in effect evolving models of themselves. This allows the robot to recover from random damage, e.g.: "when a leg part is removed, [the robot] adapts the self-models, leading to the generation of alternative gaits." [17]

Applications to modular robotics were also reported. Marbach and Ijspeert [111] and Mesot [120] simulated modular robotics systems, in which identical elementary modules were assembled into larger robots. In these two systems, either the controllers or the morphology (or both) were strictly constrained, with controllers being typically reduced to coupled oscillators.

Hornby & Pollack [83] offered a more computational model, in which development was directed by the application of a genetic grammar based on L-systems. The system involved creatures made of sticks arranged in complex 3D shapes. Their physics engine was strongly restricted, in that it assumed that the simulation is stable at each timestep, thereby preventing any dynamic behaviour such as running or jumping. The same can be said of the GOLEM project [107], in which creatures made of rigid cylinders were first evolved in simulation, then built in the real world using 3D printing. Taking a very different route, Ray [145] chose to develop Sims' model toward aesthetic evolution. Creatures were evaluated on subjective aesthetic performance, in order to encourage the generation of "interesting" appearances and behaviours.

Full replications of Sims' results were a long time coming. Taylor & Massey have attempted to replicate this work [174], with limited results. In fact, the present work (as originally introduced in [123] and [122]) apparently constitutes the first example of such a complete replication and extension. More recently, Shim and Kim [160] have evolved *flying* creatures. However, they used parametrised sinusoidal functions rather than neural networks, and morphologies were more constrained than Sims' (which is understandable, considering their objective of designing functional wings.) Chaumont and colleagues [29] have reimplemented Sims' model and successfully applied it to the evolution of catapults. Lassabe and colleagues [104] also implemented a Sims-like system, using classifier systems selecting among pre-set activation patterns rather than neural networks, and used it to evolve various locomotive behaviours in rugged environments (including relief, trenches, etc.) and simple tasks such as block-pushing.

4.1.3 Ad hoc machinery in virtual creatures

Besides physics simplification, another common feature in the simulation of 3D creatures is the widespread use of high-level, ad hoc elements in the controllers. Sims' creatures, in particular, were controlled by functional networks, including arithmetic functions, tunable oscillators and logic operators (among others) as elementary building blocks. This represents a significant amount of *a priori* knowledge given to the system. Sims [163] acknowledged the ad hoc aspect of this choice in which, in his own words, led "a creature's brain [to] resemble a dataflow computer program more than a typical neural network". Hornby & Pollack [83] and Framsticks [97], among others, also resort to high-level functions such as oscillators.[1] High-level, ad hoc elements have the obvious advantage of facilitating the emergence of certain types of behaviours, especially with regard to locomotion; however this may come at the expense of generality. We believe that using simple McCulloch & Pitts neuron may provide more freedom to the system in the type of behaviours it generates. At any rate, it would reduce the amount of a priori knowledge in the system, and thus offer a higher evolutionary challenge.

4.2 The system

Our system is broadly similar to Sims', but with important differences. Our work brings two contributions with respect to Sims':

1. Our creatures are controlled by standard neural networks, based on classical McCulloch & Pitts neurons with sigmoid or radial activation functions. This is in contrast with the ad hoc functional neurons used by Sims. While Sims' approach was entirely justified given the seminal aspect of his work, we believe that using standard neurons provides a higher level of generality to our model: creatures cannot rely on complex neurons to generate behaviours, they must build these behaviours 'from scratch' (including simple, vital behaviours such as oscillations).

2. We introduce extensions to the genetic-developmental system described by Sims. First, we address a problem not mentioned by Sims: structures which are replicated by the developmental system, either through symmetry (reflection) or through recursion (segmentation), initially possess identical neural information and thus cannot be independently controlled or provide distinct information to ancestor limbs. We solve this problem by adding genetic flags which control the actual wiring.[2] Second, we make it possible for developmental duplications of genetic nodes to be transcribed back into the genome, creating several (initially similar) genetic nodes which may then evolve independently, in analogy with the duplication-exaptation process found in Nature.

Other minor differences with Sims exist. For example, our system did not originally include any contact sensor (these were introduced later, see chapter 5.) Our sensors were limited to proprioceptors measuring the angle of a given joint, as well as sensors detecting the distance of an object of opponent along the x and y axis of the limb's frame of reference (both of which were present in Sims.) Also the actuators of our creatures specified a desired angular speed, rather than torque. As in Sims, inter-penetration between limb was forbidden, with a limited exception for adjacent limb. However, while in Sims adjacent limbs can penetrate each other up to a limit set by one half of each limb, we simply introduce a *limit* on the amplitude of each joint. The fact that joints can only rotate within a given range of angles has consequences for the emergence of sustained motion, as we will see.

---

[1] Oscillators in general are ubiquitous in physics and biology, and might be argued to be fundamental elements in their own right; however the type of oscillators referred to here are high-level oscillating 'neurons', with a tunable frequency calculated by temporal integration of inputs, arguably a high-level mechanism.

[2] Although Sims does not discuss this problem in his papers, one anonymous reviewer informed us that Sims *did* in fact address it with undocumented, unpublished features. This information is apparently based on a personal communication from Sims himself.

We also provide a complete description of our system, as well as the original source code. The lack of information on crucial aspects of Sims' system has been an obstacle to replication. At the time of writing, the software described in this chapter is available (together with video samples) at the following location:
`http://www.cs.bham.ac.uk/~txm/creatures/`

We note that several authors (including Sims himself, but also Taylor & Massey [174] and Chaumont and colleagues [29]) have pointed out the problem of *stability* in physics simulation. Physics simulation may become unstable, generating non-physical behaviours that are quickly exploited by evolution. However, we have managed to find (through perspiration more than inspiration) adequate values of certain parameters that made the system *reliably stable*. Since then, we did not observe any troublesome case of non-physical behaviour exploited by evolution to its advantage: the system is consistently realistic. These parameters are described in Appendix A, section A.5, and may be of interest to users of the ODE package.

A complete technical description of our system can be found in Appendix A, as well as in two previously published papers [123, 124]. Here we only provide a brief description of the system's features and of simple experiments.

4.3 Experiments

4.3.1 Locomotion: a (not-so-)simple evolutionary task

Our early experiments focused on a simple evolutionary task (locomotion), based on a straightforward fitness function (distance covered within a fixed period of time). In these early experiments, we used a simpler version of the platform (lacking certain niceties such as "segmentation" - see Appendix A).

While locomotion is a conceptually simple task, performing it with any efficiency is not trivial in our system - in fact it is significantly more difficult than in Sims' model. The reason is that, in our system, any efficient locomotion must of necessity include oscillatory behaviours, and therefore oscillatory neural activation: in the absence of oscillation, any creature will end up stuck in place after all limbs have reached the maximum angle allowed at their respective joints. Now in Sims' model, oscillatory neural output is provided "for free" by using complex neurons with a sinusoidal output. In our model, however, exclusive use of standard McCulloch-Pitts neurons implies that outputs are strictly growing functions of their (weighted) inputs. Any oscillation must occur through the suitable organisation of connections and feedbacks within the network: oscillators must be *discovered* by the system, using neural connections and simple sigmoid activation functions. This implies an additional level of complexity in comparison to Sims's model concerning the evolution of robust locomotion (or any sustainable behaviour at all).

Our experiments involved populations of 500 individuals, which were evaluated for 10000 timesteps, each timestep corresponding to 0.01 second of simulated time. We chose to use a simple steady-state genetic algorithm (SSGA, as described in e.g. Mitchell [128]) with "triple tournament" selection: after a whole population of individuals is randomly generated, and each individual is evaluated, the algorithm randomly picks three individuals and replaces the one

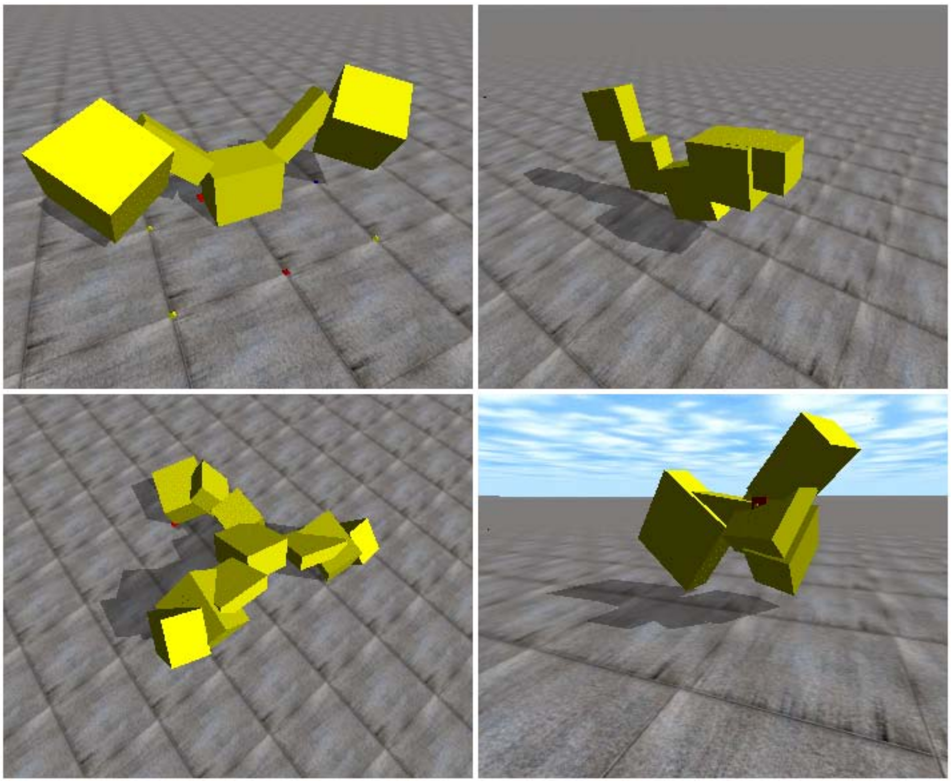

Figure 4.1: Evolution of locomotion: four creatures from four different runs. from left to right: a simple crawler, a "snake", a tailed crawler and a bouncer.

with the lowest score with a recombination of the other two, or a heavily mutated version of one of the other two. If the resulting offspring is non-viable (i.e. two non-adjacent limbs intersect, or the developed phenotype contains too many limbs), this step is repeated as often as necessary until a viable creature is produced. The resulting offspring is evaluated, and the cycle starts again.

Our system was successful in consistently evolving successful locomotion. We note that this was achieved through a wide range of different behaviours: "crawlers". "Crawlers" use two symmetric limbs to propel themselves forward by direct contact with the ground, "snakes" undulate in a coordinated way to obtain efficient locomotion, "bouncers" use very quick oscillations at one or several joints and exploit the resulting momentum to jump forward, etc. An important feature is proper exploitation of dynamics: the most efficient creatures are able to remain in a constantly dynamic, unstable state, in other words, they are constantly "jumping".[3] This diversity is reflected in Figure 4.1, which represents four creatures obtained in four different runs.

4.3.2 Coevolution: the "box-grabbing" task

After these encouraging early results, we performed more complex experiments, this time pertaining to coevolution. The task being considered is the "box-grabbing" contest described by Sims [162]. Two creatures compete to gain control of a cubic box. At the beginning of each

---

[3] Note that such behaviour could not have occurred in a simplified simulation system that only allows stable states, such as the one used in the GOLEM project. [107]

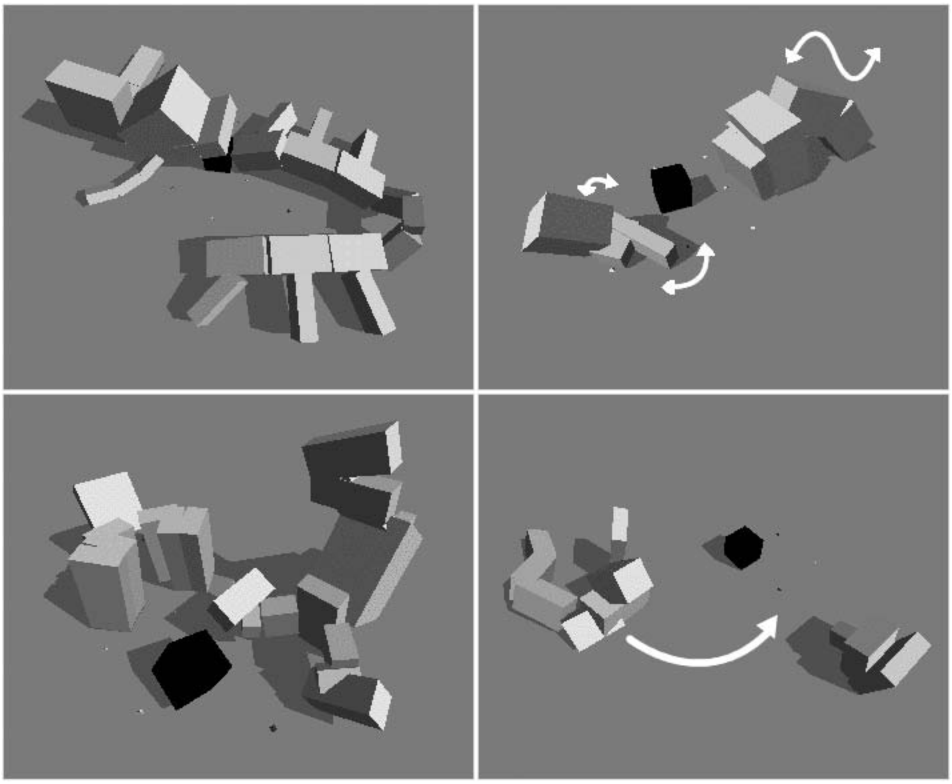

Figure 4.2: Creatures from four different runs for the “box-grabbing” coevolutionary task. In the first run, a two-armed ‘centipede’ creature tries to wrestle the box from a massive creature which uses its mass to pin it down (the centipede uses both segmental and symmetric duplication of limbs). In the second picture the left creature uses horizontal wiggling to crawl forward, while the right creature uses ample vertical sinusoidal movement from its ’tail’ to propel itself against the ground. In the third picture one creature has just knocked the ball away, using the constant rotatory movement from its right ‘arm’ (the left arm is fixed, a clear example of differential neural control over two duplicated limbs); but the other creature, using its accordion-like segmented appendages, will eventually displace it and interpose itself between its opponent and the cube. In the fourth picture, the creature uses a sinusoidal crawling movement from its tail to move past the cube, and relies on its sensor to turn around it, thereby cutting off its opponent.

contest, the box is placed at the centre of the environment. Both competing creatures are placed on opposite sides of the box, at a certain distance from it. The creatures are then left to act for a given period of time. At the end of the evaluation period, the score for each creature is determined by their relative distances from the box: if one creature is at distance $d1$ from the box (as defined by the distance between the centre of the creatures’ closest limb to the box, and the centre of the box) and the other at distance $d2$, then the former creature’s score is $d1 - d2$, while the latter’s is $d2 - d1$. Creatures have four kinds of sensors, measuring the x and y distances of either the opponent or the box, within the frame of reference of the limb in which the sensor exists.

The evolutionary algorithm that we used for these particular experiments is a “3-strikes-out” algorithm, a simple steady-state genetic algorithm for coevolution of our devising. In general, the 3-strikes out algorithm simply consists in performing competitions between randomly picked creatures, and removing individuals that have been defeated 3 times over their entire history. Our

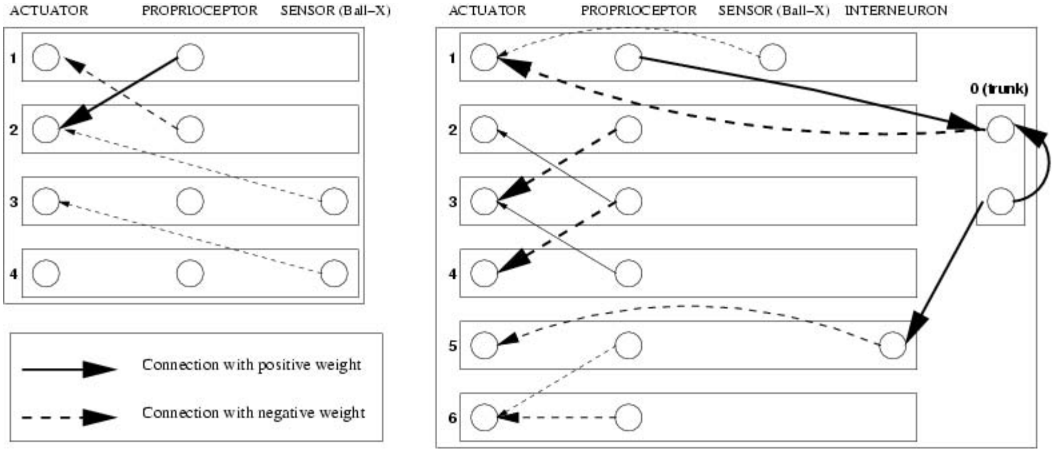

Figure 4.3: Functional subnetworks extracted from the neural networks of two creatures. The network on the left is taken from one of the accordion-shaped appendages of the creature seen in the third picture of Fig. 4.2. The network on the right is taken from the "encircling" creature seen in the fourth picture of Fig. 4.2. Thickness of lines is proportional to connection weight. See text for details.

preferred reproduction strategy, which we use in the present experiments, is simply to replace the removed individual with either a heavily mutated copy of itself, or a heavily mutated copy of the individual by which it has just been defeated, or some form of recombination between both. Importantly, when the problem involves two species, then victory and defeat are assessed by comparing the performance of two individuals of one species against a common opponent in the other species. This involves two simulated competitions; however, the result of one of these competitions is used in the next round (when two individuals of the other population are compared), so each new comparison requires the simulation of just one new competition. A complete description and justification of N-strikes-out algorithms can be found in a previous publication.[125]

Some examples of evolved creatures can be seen in Fig. 4.2. These runs involved two populations of 100 individuals each. Interactions between two individuals were simulated over 15000 timesteps, each timestep corresponding to 0.01 second of simulated time.

From visual inspection it is clear that our system managed to come up with diverse solutions to the problem. Locomotive behaviours emerged rather easily. A locomotive behaviour requires some sort of cyclic, oscillating movement, which in turn implies some coordination between proprioceptors and actuators. Additionally, sensorimotor coordination was observed, for example the creature in the third picture of Fig. 4.2 can adjust its trajectory depending on the position of the box. A complete description of the system and results (with analysis of evolved neural controllers) can be found in the related paper. [124]

Comparing these results to those obtained by Sims is not easy for several reasons. First, Sims apparently used far fewer interactions (100 generations with two populations of 100 individuals, which is much less than the $\tilde{1}$00K evaluations used in the present experiment). Furthermore Sims used complex neurons which automatically provided behaviours such as oscillations 'for free'. Restrictions imposed by Sims on the morphology of creatures (maximum numbers of genes, of blocks, etc.) are unknown.

The diversity of morphological plans is reflected in the neural controllers. Figure 4.3 describes the functional subnetworks of two creatures. The network on the left corresponds to one appendage of a crawling creature (the other appendage, being a symmetric replication, contains a similar network). It is easy to notice the feedback loop between the proprioceptors and actuators of limbs 1 and 2, leading to an oscillating movement of the corresponding joint which propels the creature forward. Other, small-weight connections originating from sensors lead to a folding of the appendage in the direction of the box.

The network on the right, however, is more complicated. It describes the functional subnetwork of the "encircling" creature shown in Fig. 4.2. The basic core of the neural controller is easily isolated: the "feedback cascade" in limbs 2, 3 and 4 provides the coordinated sinusoidal movement which propels the creature forward (limbs 2, 3 and 4 are a recursive replication sequence which corresponds to the segmented motile appendage of the creature). The sensor in limb 1 influences the shape of the creature in order to change the direction of motion depending on the relative position of the box.

However the behaviour of the creature is also dependent on all other connections and neurons, even though their exact function is difficult to isolate individually. Collectively, they ensure that the creature assumes the right shape and posture (and maintains them) to follow the desired trajectory. This occurs through coordination of activities which are tuned to the morphology of the creature (mass and orientation of limbs, etc) in order to skew the trajectory correctly. Experimental lesions in this network reduce the efficiency of the creature in various ways, which range from a slowdown to a loss of trajectory. Several other connections were present in the network (including between the neurons depicted) but were found to have no impact on the behaviour of the creature. All in all, it appears that an intricate collection of characters (neurons, connections, shapes of limbs) with no obvious individual effect, collectively provide a useful and reliable function - a good example of evolutionary "bricolage" (tinkering).

4.4 Conclusion

We have described a software platform that constitutes the first complete replication of Sims' result. In contrast to Sims', this platform uses standard McCulloch-Pitts neurons, rather than complex functional neurons; it also extends the developmental system to increase flexibility. In the next sections we will extensively build upon this platform to perform various experiments and improve our understanding of evolution.

Chapter 5

Evolving physical combat among artificial creatures

5.1 Limitations of the box-grabbing task

The box-grabbing task introduced by Sims has many advantages, not least simplicity: it is easy to understand, easy to evaluate numerically, and easy to implement. It also has the less obvious advantage of offering a fitness function that can "work" at all stages of the evolutionary process, in that it can offer an informative evaluation both to very poor and very advanced competitors. This is due to the fact that it is based on relative distances, and that even the most primitive creatures will possess some heritable variance in this characteristic (if only by falling down).

However, this simplicity can also be seen as a limitation. While there are several ways to grab a box, the variety of efficient behaviours is necessarily limited. Another problem is that it is not easy to see how this task could be extended to large numbers of competing individuals. We might imagine box-grabbing competitions involving a few creatures; we might even fancy the evolution of "rugby-playing" creatures, in which teams of individuals would compete against each other. But there does not seem to be any obvious way in which box-grabbing could meaningfully be used as a competition involving many independent individuals, constantly competing against each other, with varying lifespans and asynchronous births and eliminations.

As a consequence, it appears that box-grabbing is not easily applicable to systems based on unconstrained natural selection among a population of free-living creatures. For this type of system, we need another type of task - something more direct, less constrained than the box-grabbing setup.

5.2 Evolving fighting creatures

5.2.1 Physical combat: the appeal of (virtual) violence

Physical combat between creatures appears intuitively appealing as a basis for evolution. This comes in no small part from the fact that physical combat is ubiquitous in nature. Predation, sexual competition among males and other forms of fighting have been fruitful sources of evolutionary creativity in many lineages, producing remarkable examples of arms races and mutual adaptations. Another attractive feature of physical combat is that it is a very direct form of interaction, requiring no mediating device or instrument (as opposed to box-grabbing, and therefore box-requiring, experiments). This means that it can be used in many different settings with relatively few constraints. Crucially, while it may be used in head-to-head competitions very much like the box-grabbing task, it should be transposable to free-living creatures in an open environment. We can begin to envision a virtual world, in which creatures would roam freely and fight each other in an unconstrained fashion.

5.2.2 Related work

Many evolutionary experiments use some idealised form of "fighting" or "killing" behaviour as part of a range of pre-defined behaviours. These include Geb [27], Echo [84], Polyworld [197], Framsticks [97] and others (see Appendix B). However, in these systems, the actual process of fighting is essentially *abstract*. It corresponds to a pre-defined rule, hard-coded into the program, such as "eliminate the individual with lowest energy level", or even simply "eliminate the targeted individual, no matter what" (as in Geb). Evolution bears on when and how to use the abstract fighting behaviour, not on how to fight.

In fact, despite the possibilities offered by physical combat, we have only been able to find one attempt at evolving physical combat in a 3D environment: O'Kelly and Hsiao [135] have implemented a modified version of Sims' model, based on a very simple form of combat. In this system, "the first creature to touch its enemy's root node is deemed the winner." This simplified form of combat is easy to implement and assess, and avoids the difficulties described in the following sections. However, it is also less flexible in many ways, not least in being an "all-or-nothing" measure of success. To provide a gradient for evolution, O'Kelly and Hsiao add another component to their fitness function: at the end of each round, both creatures are rewarded with a value inversely proportional to the final distance between the two. This is expected to favour the emergence of simple approach behaviours in the early stages of evolution. Of course this has the drawback that the corresponding reward is equally given to both creatures, independently of how much each creature contributed to reducing this distance.[1] Another problem with this method of combat, especially for our own block-based creatures, is that it has an obvious weak point: simply protecting the root limb makes a creature effectively invincible.

We would like to create a more realistic system, relying on a much less abstract form of combat. Essentially, we want to obtain behaviours that are as lifelike as possible, and ideally where combat itself is the sole basis of evolution. This immediately leads to several difficulties.

5.2.3 Difficulties of physical combat: Newton vs. Darwin

The central question in physical combat is to determine how damage should be evaluated: when do we say that an individual has somehow hurt, or otherwise dominated, its opponent? This apparently simple question turns out to pose significant problems.

---

[1] A simple way to reward creatures more fairly would be to calculate, at each timestep, the modification in the distance between the position of each creature and the *previous* position of the other. In this manner, creatures that actually move towards their opponent could be rewarded, while those which stay put or move away from their opponents would not.

The most obvious answer is simply to use impacts (and some measure of kinetic energy at the time of impact) as the basis of combat: essentially, to let individuals trade blows with each other. However, this introduces a difficulty caused by Newton's third law (often summarised as "action equals reaction"). If two rigid blocks come into collision, and suffer some damage as a result, then both blocks will suffer *equivalent* damage. This is because physical damage is mostly related to kinetic energy. Clearly the relative velocities of each limb with regard to the other are equal in magnitude (and of opposite signs), and the resulting kinetic energy (and associated impact damage) will therefore be equal for both. The consequence is that when a creature hits another, the creature dealing the blow will suffer the *same* damage as the one receiving it. Clearly this is not conducive to the evolution of fighting behaviours.

In nature, the main reason why physical combat can occur is simply the heterogeneity of materials. Flesh, bones, teeth, skin, horn, etc., have different properties that make it possible to inflict damage on an opponent without suffering too much as a result. The cheetah's claws are harder than the gazelle's skin and flesh, and can therefore damage it more than they are damaged by it. Martial arts fighters attempt to throw their fists and heels at their opponent's face and stomach - rather than the other way round - because the bone structure of those parts favour (closed) hands and feet in collisions against the nose and the belly. Additionally, the geometry of object plays a roles: sharp, pointy objects will behave differently than flat or dull objects in collisions - hence the variety of mammalian tooth shapes.

Implementing such variety of materials in our simulation would clearly be cumbersome and difficult to "get right." In addition, we would need to impose some cost on the toughness of materials, to prevent evolution from turning into a simple maximisation of toughness. In nature, such runaway escalation in armour is simply prevented by the trade-offs imposed by available resources and other tasks. This would not be readily transposable in our simple model.

5.2.4 Solution: favouring the aggressor

To overcome this difficulty, we chose to evaluate the damage inflicted by a creature upon another by measuring "how much" this creature contributed to the occurrence and intensity of the collision. The result is that the creature that initiates contact more than the other (that is, the creature that is "dealing the blow") is favoured in the interaction.

Collision intensity is estimated by penetration depth. How can we measure how much each of the colliding limbs contributed to this collision? This is estimated by suspending the simulation, and then letting each of the colliding block in turn move for one timestep at its current velocity, while the other one is kept fixed; the resulting *increase* in penetration depth, if any, is used as a measurement of how much this creature contributed to the collision - that is, how much it actually moved towards the other (see Figure 5.1). After this, all blocks return to their original positions, and the simulation proceeds normally.

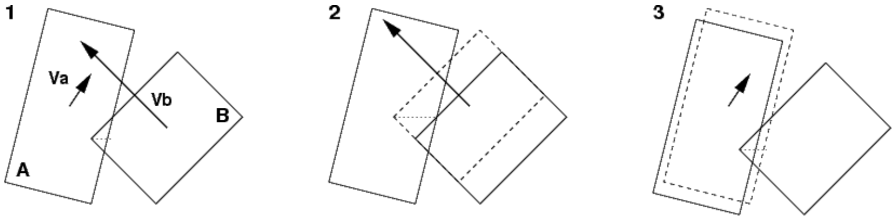

Figure 5.1: Damage calculation. 1: A collision occurs between limbs A and B, moving with velocities `Va` and `Vb` respectively. 2: Letting B move at its current velocity for one timestep (while keeping A fixed) results in a large increase in penetration depth. 3: By contrast, letting A move at its current velocity for one timestep (while keeping B fixed) results in a smaller increase in penetration depth. Thus, in this collision, B inflicts more damage upon A than A upon B. Note that if `Va` was pointing away from B, then letting A move for one timestep would actually *reduce* penetration depth, and thus A would not be inflicting any damage upon B at all.

5.3 Summary of platform modifications

5.3.1 Collisions and damage

In line with the preceding discussion, the main modification we brought to the platform was to implement damage in collisions. Anytime a collision occurs between two blocks, the following procedure is applied:

1. Keeping one block fixed at its current position, allow the other one to move on for one simulation step. Record the resulting penetration depth.
2. Restore block positions and do the same thing with the other block. Record the resulting penetration depth.
3. Whichever block caused the greater penetration depth in these “virtual futures” is deemed to have dealt the blow; the other limb (the victim of the blow) suffers an amount of damage proportional to the relative velocity between the two limbs.

5.3.2 New sensor types

In the previous sections we used sensors that perceived the distance of a certain object (another individual, or a box) in a given direction (the x or y direction). Here we replace these sensors with a more informative type of sensor, inspired in no small part by those introduced in Geb[27].

The new sensors produce an output that is proportional in magnitude to the inverse of the Euclidean distance of a certain object. The sign of this output is determined by whether the object in question lies on the “left” or “right” of the limb in which the sensor is located, relative to the axis of rotation. More precisely, the sign of the output produced is such that applying an input of the same sign to the motor joint of the limb would make the limb rotate towards the object. If a sensor is directly linked to the motor joint of the limb that contains it, and if a certain object (to which this sensor is receptive) is placed in the vicinity of the limb, then the limb will tend to reach towards the object.

Two types of sensors are used. One of these sensors detects the closest *individual* (more precisely, the trunk of the individual whose trunk is closest to the limb in which the sensor exists). The other sensor type detects the closest *limb*, that is, the direction and distance of the closest limb of any other individual in the simulation. If no other individual exists then both sensor return zero output.

5.3.3 Contact proprioceptors

In previous experiments, we did not use any contact sensors. However, considering the importance of physical contact in the present experiments, we introduce a contact sensor, the output of which is one if the limb in which the sensor exists is in contact with a limb of another creature, and zero otherwise.

5.3.4 Rules of engagement

In the box-grabbing task, creatures were pushed away from the box by a certain distance, then further pushed behind a slanted plane inclined by 45 degrees so that they could not exploit their length to get closer to the box (by simply falling down.) Here we chose to dispense with this latter precaution.

Competitions between two creatures are organised as follows: first, creatures are put on each side of a vertical plane, and then pushed away from each other by a very small distance to avoid any contact. Then creatures are allowed to move according to their controllers' output. Over the first 10% of evaluation time, creatures benefit from an "immunity period", during which they can neither hurt nor be hurt by each other; this allows creatures a safe period to reach their stable configuration, and discourages trivial behaviours such as just falling over one's opponent (in this sense it is an equivalent to Sims' "inclined plane" in the coevolution of box-grabbing [162]). After this immunity period has elapsed, damage is evaluated according to the previously described method, and accumulated over the entire evaluation period.

The fact that creatures are initially close favours the probability of contact occurring, even in the very early stages. This provides an immediately exploitable gradient for natural selection to act upon.

At the end of the evaluation period, each creature is given a final score equal to 1 + (Damage inflicted - Damage suffered) / (Damage inflicted + Damage suffered). This calculation is inspired by Sims [162]. Note that this score always falls within the $[0, 2]$ range.

5.4 Experiments and qualitative results

We used various algorithms (mostly variants of Sims' Last Elite Opponent algorithm) with this platform. The following chapters describe these algorithms in full, and provide quantitative results. In this section, we only provide a short, qualitative description of our results.

Useful creatures consistently evolved within a couple of generations. The system generated a wide range of morphologies, as shown in Figure 5.2. In general, the creatures evolved made very little use of external sensors. Their initial proximity gives them little incentive to track the movements of any opponent. Rather, the strategies focused on how to inflict maximum damage without moving away too much, implying a balance between fast motion (necessary to inflict damage) and stability.

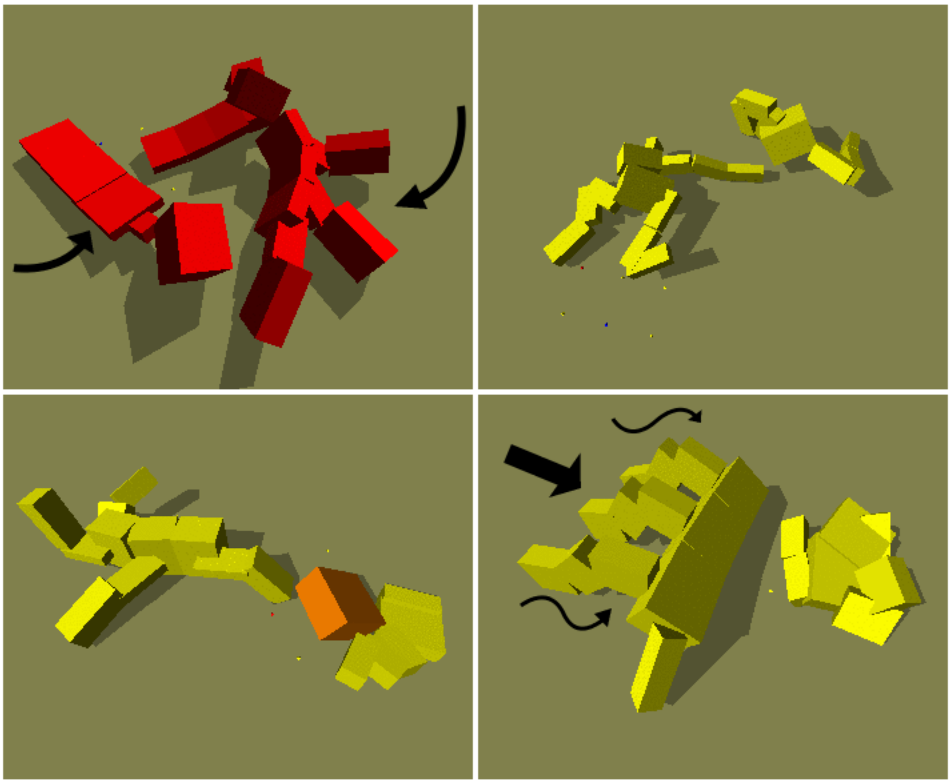

Figure 5.2: Four pairs of fighters obtained in the course of the experiments described in the next chapter. In the top-left corner, one simple creature uses its rotating cubic head to perform a "compass" motion, while the other creature uses three rotating appendages both as flails and legs. The dark colour indicates that the creatures are still within their immune period. In the top-right corner, a three-legged creature confronts a two-legged creature. In the bottom-left corner, a creature with three rotating appendages has just dealt a blow to its opponent (as indicated by the darker colour of the opponent's head). In the bottom-right corner, a large creature uses three undulating appendages as powerful legs to "steamroll" its opponent.

One commonly observed strategy was the "compass" method, in which one extremity of the creature remains fixed on the ground (mostly through sheer mass) while the other extremity features a constantly rotating structure, behaving like a rolling arm, with a plane of rotation perpendicular to the axis of the creature. As the rolling arm pushed the head sideways, while the tail remains fixed, the creature undergoes a compass-like motion: the head describes a circle around the end of the tail. In addition, the rotating arm serves as a striking implement to inflict damage upon opponents. This simple strategy proves very effective, as the creature literally laminates anything that passes within its radius. Notice that the rotating arm requires the co-operation of at least two oscillating limbs (one limb oscillating vertically, the other horizontally) to achieve overall rotation.

A variant on this strategy is the "flail" method, in which the head and single arm are replaced with a group of heads and arms, which may vary widely in size and complexity. More generally, "whipping appendages" were widespread. A different, less common approach is the "steamroll" method, which simply consists in repeatedly bumping into the opponent at full speed, constantly pushing it away in the process.

5.5 Conclusion

We have implemented a system for evolving physical combat among 3D creatures. The consistent success in evolving efficient fighters indicates that the system can be used as a platform for further experiments. This system is important for two reasons. First, it provides us with a rich experimental device with which we can study coevolution in a complex environment. Second, having implemented physical combat among creatures, we can use it as a basis of interaction for *free-living* creatures, dwelling upon the surface of an open environment. As mentioned in section 5.2.1, this is in contrast with previously proposed coevolutionary tasks, such as box-grabbing.

5.6 Looking back, looking ahead: summary and future directions

So far in this thesis we have mainly been concerned with a "buildup" phase, in which we introduced a set of concepts, definitions and tools to serve as a sound basis for further investigations. This basis includes usable definitions of evolution, natural selection, and superiority and progress in coevolution. On the practical side, it also includes the software platform discussed in this chapter and chapter 4.

In the remainder of this thesis, we will move to an "exploitation" phase, in which the concepts and tools previously developed will be used extensively. First, building upon our concepts of superiority and progress in coevolution, we will discuss methods to analyse the dynamics of coevolution and assess its performance. We will use our software platform as an experimental device to demonstrate and validate the use of these methods.

Then, we will discuss the notoriously vexing relationship between evolution and complexity, using the definitions laid out in the first chapters of these thesis. We will introduce a method, called Fitness Transmission, for detecting the activity of evolutionary adaptation within a system (that is, a signature of Darwinian evolution); in particular, we will also point out that previously suggested methods (such as the Bedau-Packard measure of evolutionary activity [11]) should not be used for this purpose.

Finally, building upon our implementation of physical combat among virtual creatures, we will introduce an artificial environment in which a population of free-living creatures interact, reproduce and evolve. This will lead us to a discussion of what it means for evolution to occur "without a fitness function" (also known as "intrinsic adaptation" [137] or "natural" selection, as opposed to "artificial" selection [28]), in the face of the necessity for any artificial system to possess rules for determining survival and reproduction of any given individual (that is, precisely, a "fitness function"). We will conclude our discussion by applying fitness transmission to demonstrate the onset of adaptive evolution within our system.

Chapter 6

Methods for coevolutionary analysis and evaluation

6.1 Introduction

In the previous sections we saw that progress and performance need to be defined carefully in the context of coevolution. This leaves us with another question, namely the practical problem of detecting, and estimating, performance and progress in coevolutionary experiments. We need practical methods to analyse the outcomes of coevolution. In particular, we seek to address the following problems:

- How can we detect the presence or absence of (historical or global) progress in a coevolutionary process?
- Given two coevolutionary algorithms, how can we determine which one produces the best individuals (for equivalent computational effort)?

Obviously, the latter question is especially important from a practical viewpoint: if we want to devise efficient methods for coevolution, we need to compare the efficiency of different algorithms in a reasonably unambiguous manner.

6.2 Methods for analysing and monitoring coevolution

The problem of tracking progress and performance in coevolution has been recognised early on. The Red Queen effect, by making it difficult to assess the intrinsic quality of an individual (see chapter 3), also complicates the task of assessing the behaviour of a particular algorithm. Several types of statistics have been proposed for analysing the results of coevolutionary processes, with a stress on the identification of progress.

Much of this chapter has been published in a separate paper [122]. Also some extracts from another paper [125] have been included.

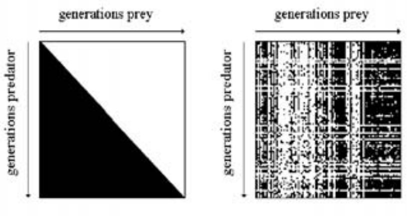

Figure 6.1: Master Tournament matrices, taken from Nolfi & Floreano [133] in a predator-prey experiment. Black (resp. white) dots indicate a victory for the champion of the predator (resp. prey) population. The picture of the left represents an "ideal" situation of perfect progress, in which each champion is able to defeat all previous opponents. The picture on the right represents the results of a real experiment.

6.2.1 CIAO and Master Tournament

Cliff & Miller's "Current Individual vs. Ancestral Opponents" method (CIAO) [32] and Nolfi & Floreano's "Master Tournament" method [133] both use a similar principle. At every generation, a champion is selected (often a champion is naturally provided by the algorithm at hand.) Then, we pit the champion of every generation against the champion of all other generations (or, for CIAO, of all previous generations.) The result of each confrontation is represented by a coloured dot within a grid, such that the dot at coordinates (n, m) in the grid is darkened if the champion of generation n in the first population defeats the champion of generation m in the other population, and left blank otherwise (in tasks where draws are possible, a three-colours scheme can be used.) The CIAO method pits the champion of a population at generation n against the champions of each *previous* generation in the opposing population, thus resulting in a triangle of dots. The Master Tournament method collates the results of confrontation between all champions of all generations, producing a square matrix of dots. The Master Tournament square corresponds to the collation of two CIAO triangles (one for each of the competing populations) joined along their common hypotenuse, the diagonal of the square.

Both methods can be applied to two-species or one-species coevolution. If we draw a Master Tournament grid for single-species coevolution in a symmetric problem, then the grid will be symmetric along the diagonal, since (m, n) represents the same confrontation as (n, m) (namely, between the current champion of the species at generations m and n.) However, if the task is not symmetric (that is, if pitting A against B may produce different scores for each individual than B against A), then the dot matrix will in general not be symmetric either.

The data from a Master Tournament matrix can also be summarised in a simple graph, in which, for each species, the percentage of opponents defeated by the current champion is shown as a function of generation number. We call such a graph a Master Tournament summary graph, by contrast to Master Tournament matrices. Importantly, both are visual representations of the data generated by the Master Tournament method: we must not reduce this method to either the matrices or the summary graphs.

These methods have the advantage of providing reasonably complete information about an entire evolutionary run. However this completeness comes at a price. An obvious problem with these methods is their computational cost. Let N be the total number of generations. Since N^2 evaluations are needed to obtain a complete Master Tournament grid ($(N^2 - N)/2$ in the single-population, symmetric case), as soon as N becomes even moderately large, calculating the figure is a time-consuming process. Of more concern to the analyst is the fact that the resulting figures are often somewhat obscure: although "ideal" conditions of progress lead to a very simple figure, these ideal conditions are rarely met in practice. Real experiments often produce disorderly arrangements of dots from which it may be difficult to extract any meaning at all (see Figure 6.1.)

Cartlidge [26, Chap. 7] has proposed a graphical method to facilitate the analysis of CIAO plots. He applies image processing techniques (essentially blurring and binarisation) to the plot, which enhances the visibility of patches and strips within the graph. This makes it easier to identify the presence of *regular cycling* between strategies, and to measure the period and regularity of this cycling: regular cycling produces regular diagonal bands on the graph. However, while the method seems ideally suited for detecting regular cycling, it is not clear whether other patterns have a direct signature on the resulting graphs. The method has only been applied to domains with an inherently cycling structure (Rock-Paper-Scissors) or to populations undergoing random drift; distinguishing between the two occasionally proved difficult [26, 145]. The method may also produce spurious structures [26, 139,144]. In addition, note that the method requires a full CIAO plot, so there is no reduction in computational cost.

6.2.2 Dominance Tournament

A more recent technique for observing progress in coevolution has been proposed by Stanley & Miikkulainen under the name of Dominance Tournament [170]. Dominance Tournament was developed for single-population coevolution, but can be readily extended to multiple populations. In a dominance tournament analysis, one must keep track of every new individual that defeats all previously dominant individuals. Dominance is defined recursively: The first dominant strategy d_1 is the champion of the first generation; then, at every generation, the current champion becomes the new dominant strategy d_i if it can defeat all previous dominant strategies $d_{j<i}$.

Thus the dominance tournament method concentrates on a sequence of individuals which are seen as particularly important, due to their recursive superiority relationship. Dominance, in this context, is not synonymous with superiority: dominant strategies need not be historically or globally superior to their predecessors. Clearly, there is an implicit assumption behind that in practice, dominance should be more or less correlated with (at least) historical superiority. This assumption is reasonable if the search space is not too complex or chaotic, though estimating this would in itself be an interesting problem.

Dominance Tournament has the advantage of being much easier to compute than Master Tournament, since at any time the total number of dominant strategies against which candidates are to be tested is significantly lower than the total number of generation champions. It is also much easier to analyse, since it can be represented as a one-dimensional series of ticks along a time-coordinate axis, each tick corresponding to the appearance of a new dominant strategies. However, the massive simplification of the statistics eliminates a lot of information, and it is not clear exactly how precisely the Dominance Tournament captures the global trajectory of a given run.

6.2.3 Limits of historical methods

Note that these methods (CIAO, Master Tournament, Dominance Tournament) must be applied to the history of a particular run: they essentially rely on "single record" analysis. They are useful in studying the trajectory of evolution and the presence (or absence) of *historical* progress. However, they do not bring direct information about general progress in the sense of overall superiority when evaluated against the whole search space - that is, *global* progress. We will expand on this in the next section.

6.2.4 Comparing the performance of algorithms

A simple method to compare the efficiency of two algorithms or techniques is simply to perform an "all-against-all" comparison: we pit all champions generated by a given algorithm, against all champions generated by another, and we simply tally the results. Clearly, if champions generated by algorithm A_1 consistently defeat many more champions generated by algorithm A_2 than the other way round, we may deduce that A_1 is superior to A_2. Crucially, because the individuals being compared have not coevolved together, they may serve as impartial reference for each other. This simple method is used by Nolfi and Floreano [133]. However, it has the drawback of being a single-number method, providing aggregate data rather than detailed information.

More importantly, since it includes all champions generated by either method over time, it may produce undesirable results due to the inclusion of early, non-representative champions. Imagine an algorithm that would consistently produce individuals that defeat 50% of all possible opponents; now imagine another algorithm, that would initially produce very poor individuals with 0% success rate, but would consistently end up producing optimal individuals with 100% success rate against unseen opponents after a certain period of time. If we are ready to give enough time to the algorithm, we would certainly prefer the latter "late-optimal" over the former "stable-mediocre" algorithm. However, it is very possible that the "all-against-all" comparison method would produce the opposite result, favouring the stable-mediocre algorithm over the late optimal one. This is simply because the early, poorly-performing champions would overwhelm the results of the latter, optimal champions, and their overall average may still prove lower than the average of constantly mediocre champions generated by the other algorithm.

6.3 Historical methods cannot assess global progress

As we saw in Chapter 3, distinction has not always been made between historical and global progress. Methods based on the data of a single run (such as Master Tournament, CIAO, Dominance Tournament, etc.) essentially measure historical progress: they compare the performance of individuals, when opposed to the set of opponents they have evolved against. Saying that such methods can assess "progress", without any qualification, only emphasises the troublesome ambiguity of the word "progress."

In fact, inferring global progress (or its absence) from historical methods is akin to a well-known error in statistics, namely that of *using the training set as a test set*. In historical methods, all individuals involved are supposed to represent populations which have been used as a training set for each other. In fact, for certain algorithms (including Sims' algorithm and its extensions based on a Hall of Fame), the champions of each populations *are* each other's training sets. Observing continuous, "perfect" historical progress with these methods, is equivalent to obtaining classifiers that successfully classify all of their training data. But it is well-known that such a feat cannot be used as a measure of the general performance of the classifier: rather, we need a test set, distinct from the training set, against which the classifier will be evaluated. This allows us to assess the performance of the classifier against unknown examples, and therefore its actual usefulness in practical use.

In the case of the Master Tournament method, this must be qualified: each individuals is not just evaluated against past opponents, but also *future* opponents, which neither it nor its ancestors ever encountered. This is outside the scope of strictly historical progress. However, if B is a future opponent of A, then A is a previous opponent of B - that is, an opponent of one of B's ancestors (again, a Master Tournament matrix is a collation of two purely historical CIAO matrices, each describing the history of one of the coevolving species.) Furthermore, future opponents are descended from previous opponents, which means that they are clearly not independent of them. Thus, in a Master Tournament matrix, for each species, half the test set is composed of the training set, and the rest is derived from it.

Of course, in practice, strong historical progress may well carry some evidence of global progress. For example, in Master Tournament matrices, if the champion of generation 1 defeats no opposing champion, while the champion of generation N defeats *all* opposing champions, then we may suspect that the champion of generation N is indeed not just historically, but globally superior to the first champion. This is not a logical necessity, but a reasonable inference that we can make if the difference is large enough. Similarly, in classification, if one system correctly classifies no example of the training set while another system correctly classifies all training examples, then we may suspect that the latter will indeed prove more powerful than the former, not just against the training set, but also against the test set or indeed any set of examples. However, we must keep in mind that such a generalisation is not always justified, and that lesser differences in success over the training set may not carry any information regarding performance against random examples. Similarly, historical superiority may not imply global superiority, and therefore historical progress may not imply global progress.

6.4 Coarse-Grained Master Tournament

6.4.1 Coarse-Grained Master Tournament matrices

Master Tournament matrices are an informative method to evaluate historical progress in co-evolution. However, as previously mentioned, they are computationally expensive and may be difficult to read due to their fine texture.

To monitor the historical progress of evolution, we introduce a modified version of the Master Tournament method. Our method simply consists in performing a regular sampling operation - a "coarse graining" of the matrix. Instead of performing a full tournament between the champions of all N generations, we simply pick a fixed number k of champions, evenly sampled through time, and use this sample to produce a coarse-grained Master Tournament matrix. This means that we only perform tournaments between champions of generations which are integer multiples of the N/k fraction. In our example, we chose to sample the champions of every 10th generation, condensing a run of 500 generations into a 50x50 grid. Note that because of the N^2 cost of Master Tournaments (graphically illustrated by the square aspect of the resulting matrices), by sampling 10% of the generations, computing costs are divided by 100.

Like any sampling process, coarse-graining incurs a loss of information. However, the information which is lost by coarse-graining is expected to be short-term, small-scale information. When analysing the results of a coevolutionary experiment, we are usually interested in higher-level trends, especially regarding evolutionary progress. Coarse-graining should preserve much of this type of larger-scale information. In particular, the question of whether or not a given individual can generally outperform older ones, which is the crucial aspect of the "arms race" concept, is not affected by coarse-graining. Moreover, coarse-graining can actually make a Master Tournament matrix more descriptive by suppressing spurious, irrelevant information: as we make clear in the following paragraphs, coarse-grained Master Tournament matrices may exhibit discernible, informative features which are often difficult to observe in full matrices.

6.4.2 Experimental setup

We apply our method to coevolution among artificial creatures, as described in the previous chapters. We use the same task as Sims [162]: two creatures compete for control of a single cube. The setup is similar to that described in the previous chapter for this task, except that we do use Sims' original LEO algorithm. For every run, creatures are divided into two populations. At every generation, creatures of each population are evaluated against the current champion of the opposing population. The creature which obtains the best score becomes the new champion of this population. Survival rate is 50%, which means that half the population is replaced at every generation.[1] Selection of parents occurs by direct tournament selection based on score. New individuals are created with equal probability by one of three operations: grafting between the two individuals, crossover between the two parents, or three successive applications of the mutation operator to one of the parents. Then the mutation operator is applied to the resulting creature and produces the final offspring. Each run covers 500 generations - clearly a prohibitive number for an original Master Tournament matrix, but easily manageable after coarse-graining by 10% sampling.

6.4.3 Reading a coarse-grained Master Tournament matrix

Figure 6.2-left shows a coarse-grained Master Tournament matrix for a particular run. Each (m, n) location is marked with a dark square point if the champion of population 1 generation $10 * m$ defeats the champion of population 2 at generation $10 * n$, or with a light cross mark otherwise. The $y = x$ line, drawn in a lighter shade, provides a time axis for the actual run. Points on this line indicate how the actual run went along, indicating the victorious population at each generation. On a coarse-grained Master Tournament matrix, vertical patterns are related to individuals from population 1, while horizontal patterns are related to individuals from population 2.

---

[1] This is a much more conservative value than Sims' 20%. In later experiments, we actually found that Sims' more aggressive parameter proved much more successful.

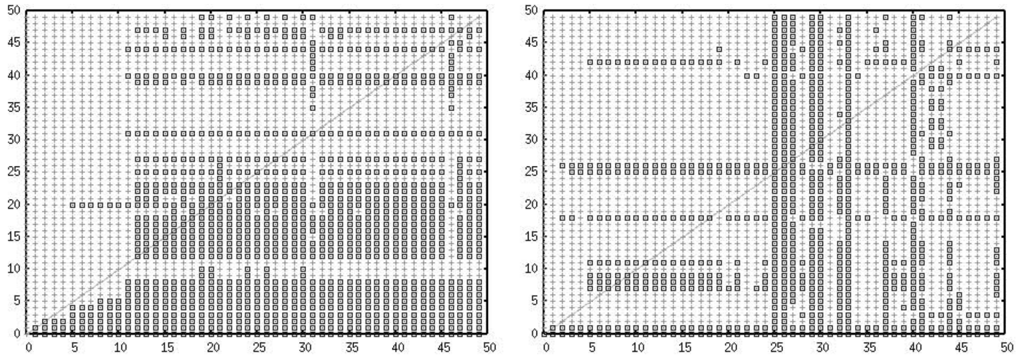

Figure 6.2: Coarse-grained Master Tournament matrix for two different runs. Squares indicate victory for the x species, while dots indicate victories for the y species. See text for details.

A first observation for this run is that the $y = x$ line goes through several regions of different colour. This means that the champions of the two populations successively outperform each other, an indicator of healthy competition. However, the particular patterns of this alternation provide a better insight about the course of evolution in this run.

Identifying similar phenotypes from their competitive profiles: The graph in Figure 6.2-left contain many similar lines and columns. In particular, it may be seen that many columns offer strikingly similar patterns of dark and light marks, although with appreciable variation. Each column, however, corresponds to the competitive profile of a champion in population 1: it accounts for its successes and failures against every champion of population 2. Two identical columns denote two individuals that defeat the same opponents, and are defeated by the same opponents. It is not too far-fetched to assume that similarity in competitive profile is linked to similarity in phenotypes.

Similarity is not identity, and much variation can be seen. However there are at least two columns which offer a significantly different profile to the neighbouring columns, namely columns 31 and 46. These two columns can be said to represent different types from their neighbours, due to the difference in competitive profiles. In particular they are unique in being able to defeat the opponents in rows 41-43.

The significance of the high similarity in columns after 12 can be seen as an indication that, at least in population 1, evolution seems to have settled on a particular type of creature, which is marginally "fine-tuned" in the later course of coevolution. This capacity to indicate phenotypical convergence is an interesting property of coarse-grained Master Tournament matrices.

Evidence of breakdowns in arms race: The arms race concept implies that newer individuals are consistently able to outperform their ancestors. Breakdowns in the arms race are associated with a loss of adaptive function, since an ability (to defeat some individuals that could be defeated by ancestors) has been lost by the newer individual [57].

Breakdowns or interruptions in arms races are easy to locate on a Master Tournament matrix. Any dark mark occurring immediately above a light mark, or any light mark occurring immediately on the right of a dark mark, indicates such an interruption: it means that a given individual (from population 2 in the first case, from population 1 in the second case) was unable to defeat an opponent that could be defeated by its ancestor. Such breakdowns may be very short events, indicating a prompt recovery. Alternatively they may result in a long-term loss, or even permanent loss of the capacity to defeat some particular individuals.

Let us take the example of the first horizontal stripe of lighter marks, at rows 9-11. These rows correspond to particularly fit individuals which are able to defeat a large number of opponents (all of them for row 11). In particular, they have no difficulty defeating the champions of generations 12 to 18 in population 1, as can be seen from the fact that their rows are void of dark marks in the section between columns 12 and 18.

Yet the same graph shows that, from generation 12 onwards, the first bisectant encounters a series of dark marks, indicating superiority of the champions of population 1. This indicates that by generation 12, the current champion of population 2 had become unable to defeat individuals that earlier ancestors could defeat. How did this come to be? If the champion of population 2 at generation 11 was good enough to defeat all opponents that population 1 would ever come up with, why was it displaced with one that would prove to be inferior?

Causes for breakdowns in arms races: This alternation between a lighter stripe and a sudden block of dark marks indicates a dramatic example of a breakdown in the arms race. In this particular case, the cause can be identified as *over-specialisation.* While it is true that the champion of generation 12 in population 2 was potentially able to defeat a large number of opponents, population 2 managed to come up with a new individual which was even better (that is obtained a lower differential distance to the cube) against the current champion of population 1. Unfortunately this change, while beneficial in the short term, proved fatal when population 1 managed to evolve a counter-strategy which defeated this specialised opponent. This allowed the newly evolved type of individual in population 1 to take the lead, even though previous champions from population 2 would have been able to defeat it.

Figure 6.2-right shows the results of a different run. This figure exemplifies several other informative patterns. In particular, let us look at the centre of the matrix, at row and column 25. At that point, we see that the first bisectant encounters a kind of wedge, composed of two stripes of dark marks - one vertical, one horizontal. Can we infer some meaning from this pattern? The wedge shape indicates that a successful change in population 1 (indicated by the appearance of a different competitive profile, leading to a distinct, darker series of columns) has led to a dramatic breakdown in the arms race on the side of population 2. The appearance of this new champion in population 1 has upset the hierarchy in population 2: the previous champion was no longer the best possible candidate against this new opponent. Confronted with the new, successful champion of population 1, population 2 has settled on a new champion, which happened to perform better, or at least less badly, than others against this particular new opponent (though not well enough to actually defeat it). This new “champion of fortune”, however, was not particularly well-rounded and performed badly against a large range of opponents. Innovation in population 1 has caused a *confusion* in population 2.

This idea of new individuals breaking down the arms race by upsetting the hierarchy and voiding previous adaptations in their opponents is not necessarily linked to wedge-like patterns, but simply to the appearance of a new type of opposing champions. For example, Figure 6.2-left contains several dark horizontal lines, apparently isolated. In particular, the individual in row 31 indicates that this champion suddenly lost much of its ancestors’ aptitudes against opposing champions. What is the cause of this loss? If we track the point at which this new, poorly performing champion occurs (by locating its intersection with the first bisectant) and observe the corresponding column, we notice that the individual from population 1 at column 31 has a subtly different pattern from its predecessors. The poor performance of population 2 at generation 31 is thus caused by the emergence of a new opponent which upsets the hierarchy in population 2 and propels an apparently poor individual to the rank of “champion”.

These interruptions in the arms race (temporary or long-term) that can be observed on the coarse-grained Master Tournament matrix are an indication of the local nature of co-evolution. Because co-evolution is only concerned about the immediate present, it may directly induce a loss of ability against past or future opponents. This loss may occur spontaneously (as in over-specialisation) or may be provoked by a change in the opposing population (as in “confusion”).

6.5 From historical to global progress: Cross-validation of co-evolutionary runs

Master Tournament matrices and summary graphs are a powerful method for analysing the historical behaviour of coevolutionary processes. However, as we saw, we cannot use them to reliably assess any global progress, even though the latter usually constitutes the true objective of coevolutionary optimisation. It would be desirable to possess a method that allows us to detect global (rather than just historical) progress, while still being able to perform historical comparison between successive generations.

A simple generalisation of the Master Tournament method is to include other opponents taken from other evolutionary runs than the one to which the evaluated population belongs. Because opponents from different runs do not coevolve together, they do not influence each other’s evolution, and therefore offer test elements distinct from the training set. In other words, we use individuals from different runs to *cross-validate* each other, expecting that this would provide more reliable information about their global efficiency. Thus, by evaluating the successive champions of a given run against a set of champions from different runs, we can actually obtain a valuable estimation of *global* progress. The reliability of this estimation is of course directly related to the number of other individuals against which this cross-validation is performed.

The basic principle of cross-validation is to pick N runs, and test each successive champion of each population against all champions of all other populations. Clearly this imposes a much higher evaluation cost. In fact, for N runs, and sampling k successive champions per run, the cost grows with $(kN)^2/2$. This quadratic cost can quickly become unacceptable. Fortunately, we can still gain valuable information from very low values of N. We will now demonstrate this by applying 2-runs cross-validation to the same evolutionary runs described in previous sections.

Figure 6.3 shows, for each champion of all 50 generations in each population, the number of champions of all other opposing populations that it is able to defeat. 13-A and 13-B are population 1 and 2 from the left-hand side matrix in Fig. 6.2, while 3-A and 3-B are population 1 and 2 from the right-hand side matrix. This graph is interesting both for its similarities and its difference with the individual Master Tournament matrices in Figure 6.2.

Within this larger context, the best performing individuals are the champions of generations

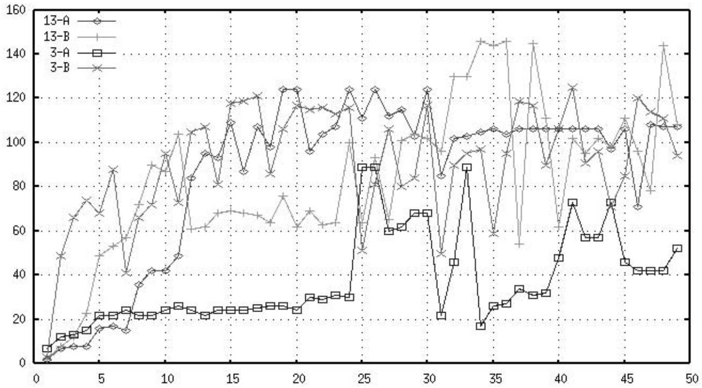

Figure 6.3: Cross-validation of each individual in each of the four populations shown in Fig. 6.2. 13-A and 13-B are population 1 and 2 from the left-hand side matrix, while 3-A and 3-B are population 1 and 2 from the right-hand side matrix in Fig. 6.2.

34 and 36 from population 13-B, with a score coming close to the maximum 150, meaning that they can defeat almost all other champions. Looking at the corresponding rows in Fig. 6.2, we observe that they indeed obtain 'clean sheets' against all their opponents. However, this is also the case with rows 11 and 28-30, yet these ones obtain a much lower score on the cross-validation graph. This indicates a difference in performance that could not have been deduced from Master Tournament (coarse-grained or not) or Dominance Tournament analysis, nor indeed from any single-record analysis method alone.

Similarly, we see that population 3-A seems to perform rather poorly when compared to others. Specifically, after generation 10, all champions of population 3-A obtain much lower performance that champions in population 13-A. This is in contrast with the corresponding Master Tournament matrices, in which it can be seen that some champions of population 3-A are able to defeat all opposing champions from population 3-B (columns 25-26 and 33), while no champion in population 13-A shows such a perfect record. Again, a single-record analysis could not have detected this apparent superior performance of individuals from population 13-A.

6.6 Equal-effort performance comparisons between algorithms and methods

6.6.1 How to compare algorithms?

In the previous sections we described various techniques to evaluate progress (whether historical or global) in the trajectory of a coevolutionary process. Here we seek to address a different issue, namely that of comparing the efficiency of different algorithms and methods for coevolutionary optimisation. The question here is not: "do later individuals perform better than early ones", but rather: "does algorithm/method A create better individuals, for a given computational effort, than algorithm/method B?"

In section 6.2.4 we mentioned a simple way to compare the efficiency of two methods, that is, simply pitting all champions generated by one method against those generated by the other, and count the number of victories for each. However, we noted that this method provides rather crude information, and may even be misleading under certain conditions.

Here we introduce a different method, that we will use extensively in the next section. This method simply consists in opposing individuals generated by different algorithms (in many different runs), with the restriction that all competitions should only oppose individuals generated after *equivalent computational effort* - that is, after the same total number of competitions have been performed.

More precisely, the procedure for equal-effort comparison is as follows:

1. Take several runs generated by two algorithms or methods A and B.
2. For each run, pick the current champions of this run after successive multiples of a certain number of evaluations have been performed; that is, pick the current champion of the run after $i * K$ evaluations have been performed, for a certain pre-defined value of K and successive values of $i \in [1, N]$. This provides a sample of N "evenly spaced" champions for each run.
3. For each rank $i \in [1, N]$, oppose the i^{th} champion of each population generated by algorithm A against the i^{th} champions of all populations generated by algorithm B. Count the total number of victories for champions generated by either algorithm. The resulting values indicates the relative efficiency of each algorithm after iK evaluations have been performed.

Here we associate computational effort with number of competitions, because we assume that both algorithms being compared use the same task and the same parameters for each competition; therefore, competition time is independent of the algorithm being used. In more complex situations (for example if we want to compare vastly different methods which use incomparable evaluation processes), we may use elapsed time directly, but this requires a guarantee that both methods are run on similar machines with similar loads, etc.

Note that equal-effort comparisons provide a "real-time" picture of the relative performance between the two algorithms: we can track the relative performance of each algorithm over time, and identify periods over which either of them is superior. As we will see, some methods may prove superior in the early stages, and then be overcome as the process goes on.

In the next chapter we will use equal-effort comparisons to assess the effects of increasing the number of competitions per evaluation, as well as of "sliding archives". This application will provide an illustration of how equal-effort comparisons can be used in practice.

6.7 Conclusion

It has long been recognised that the subtle dynamics of coevolution require specific methods of analysis and evaluation. Existing methods tend to focus on historical progress, providing little information about the more desirable property of global progress. Equipped with a consistent notion of "progress" and with a clear distinction between historical and global progress, we can refine our understanding of existing tools and design new ones, with a clearer picture of their respective properties. In this section we have mainly introduced:

1. Coarse-grained Master Tournament matrices, which provide not only an obvious gain in computational costs, but also improved readability.
2. Cross-validation (multi-run Master Tournaments), which extend Master Tournaments to include information about global progress. As a result, these methods provide information that cannot be obtained from standard Master Tournaments alone.
3. Equal-effort comparison, a practical method to evaluate the relative performance of two different algorithms or methods, and track variations in relative performance over time.

In the next chapters we will make use of some of these methods to gain information about the properties of various algorithms.

Chapter 7

Application: Evaluation of extended LEO and archive methods

7.1 How many competitions is enough?

It is impractical to use many competitions to evaluate artificial creatures. This is mostly due to the obvious computational costs of 3D physical simulations. Typically, using recent machines, a 200000-competitions run takes more than ten days. Faced with such constraints, frugality in the evaluation process is clearly paramount.

However, many methods in coevolution require high numbers of competitions. For example, Rosin and Belew [151] use hundreds of competitions to evaluate any single individual. Stanley and Miikkulainen [170] use twelve competitions (each consisting of two contests) to evaluate an individual. While this is clearly less than an "all-against-all" evaluation method (where every individual competes against every other individual within the population), the computational cost for any single evaluation is still high. The problem is even more acute in theoretical algorithms which focus on specific properties rather than sheer efficiency, such as MAXSOLVE [94] and the IPCA [43] (see Section 3.5, note 9).

The intuitive justification for using many competitions is that any single competition tells us little about the (global) performance of an individual. Pitting an individual A against an opponent B will only tell us whether A defeats B or vice-versa; it brings no direct information as to how either A or B would compete against any other individual. However, if we pit both A and B against many opponents, and A turns out to defeat significantly more of them than B, then (provided the fitness landscape is not utterly chaotic) we may begin to suspect that A is indeed superior to B in a global sense. The more competitions we use, the more we can justify this extrapolation. But how *much* more?

In this chapter we take the approach of enriching Sims' original algorithm, which Cliff and Miller [32] have since called the Last Elite Opponent (LEO) algorithm. As previously mentioned, this algorithm only requires one competition (against the champion opponent of the previous generation) to evaluate an individual. We will extend it by increasing the number of champions used at every generation, as well as by using a small "sliding" archive, in the manner of Nolfi & Floreano [133]. We will then use the method of equal-effort comparison (introduced in Chapter 6) to evaluate the effect of these additions. Thus this chapter provides both an informative comparison between various algorithms, and an illustration of the usefulness of equal-effort comparisons.

7.2 Experiments and results

7.2.1 N-LEO: Extending Sims' algorithm

Let us recapitulate the mechanism of Sims' LEO algorithm[162]: at every generation, each individual from population A is pitted against the current champion of population B. The resulting score is used as a fitness value for selection and reproduction among individuals from population A; it is also used to choose a new champion for population A (the individual from A which obtained the best score against the current champion of B). Then the same process is applied to population B, using the new champion of population A for evaluation, etc.

The obvious extension of using several champions at each generation immediately suggests itself.[1] In this N-LEO algorithm, at every generation, each individual from population A is pitted against the N current champions of population B. individuals from A are then sorted according to their average score against all N opposing champions. The N individuals with the highest average scores are selected as the new champions of population A. The same cycle is applied to population B, then again to population A, etc.

7.2.2 Details of equal-effort comparisons

Using equal-effort comparison, we can measure the effect of increasing the number of champions N. To compare two versions of the algorithm (based on two different values of N), we first generate a certain number R of runs with each version. Then, for each run, we pick the current absolute champions in this run after multiples of a certain number of competitions have been performed - that is, the current absolute champions of this run after k, $2k$, $3k$, etc. competitions have been performed (for a pre-defined value of k). This provides us with a sample of n "regularly spaced" champions for population of each run.

We then pit each sampled champion against all champions of equal rank (i.e. obtained after equal number of competitions) generated by the *other* algorithm. Because each run contains two populations, at each rank, this results in $2R^2$ competitions[2] We then simply count the number of victories for each algorithm at every rank. This provides a picture of which algorithm is more efficient after k, $2k$, $3k$... evaluations.

In this chapter, for all experiments, 14 runs are generated for each algorithm. The value of 14 is an arbitrary choice based on practical considerations and on hardware resources. Each run is allowed to reach 200000 competitions, and 50 champions are sampled from each population of each run (thus selecting the current champions after 4000, 8000, 16000, etc. competitions have been performed.) At each rank, we therefore perform $2 * 14 * 14 = 392$ competitions.

---

[1] Such a process is indirectly hinted at by Rosin and Belew [151, sec. 3.2], but is not used by them (or indeed by anybody else that we know of.)

[2] Each species of each run is arbitrarily labelled A or B, and competitions are only performed between members of A and B species. A full, all-against-all tournament would result in $4R^2$ competitions for each rank.

7.2.3 Results

Figures 7.1 and 7.2 describe the result of equal-effort comparisons between various instances of the N-LEO, using successive values of N. The 1-LEO is simply Sims' original LEO algorithm.

The first result of these graphs is that the N-LEO does indeed provide a gain in performance against Sims' original LEO algorithm. This is confirmed by the difference in performance between the 1-LEO and both the 2-LEO and 3-LEO algorithms. Similarly, the 3-LEO improves upon the 2-LEO, as indicated both by their direct comparison and by comparing their results against the 1-LEO. In particular, the long-term difference between 3-LEO and 1-LEO is found significant at $p < 0.013$. The difference between 2-LEO and 4-LEO falls just short of significance at $p < 0.064$ (see the Appendix of this chapter for details about the calculation of significance).

However, if we increase N beyond 3, things become less clear-cut. At N=4, it appears that the 4-LEO is initially *less* efficient than the 3-LEO, but catches up quickly; the long-term difference is not found to be significant ($p < 0.463$). Increasing N further does not seem to bring any significant advantage. In fact, if anything, the 6-LEO appears slightly *inferior* to the 4-LEO, though the long-term difference was not found to be significant ($p < 0.136$).

In short, it appears that using more than one champion significantly increase performance, but that this increases quickly reaches a plateau between 3 and 4 champions. Further increases in the number of champions do not bring any improvement, and might actually decrease performance.

7.3 Sliding archives

In this section we seek to quantify the advantage (if any) of using an *archive.* As has been mentioned, archives are a popular device used to enforce progress in coevolution, though as Nolfi & Floreano (and others) reported, they may have the opposite result. We explained this fact by pointing that what archives enforce is historical progress, which does not imply global progress.

Here we will consider what we call a "sliding" archive, a mechanism introduced by Nolfi & Floreano [133]. For each population, we maintain a sliding archive that retains the absolute champions of this population for the 15 previous generations. This archive is updated at every generation. To evaluate individuals of a given population, we pit them against the current champion of the opposing population, *and* against a certain number K of opponents, randomly picked within this archive of the 15 previous champions. In other words, the algorithm corresponds to a 1-LEO, augmented with K opponents randomly picked from the sliding archive. Thus, every individual undergoes $1+K$ competitions for each evaluation. Of course, the process could be extended to including more current champions, resulting in $N+K$ evaluations, but this will not be discussed here. Rather, we will only consider 1+2 and 1+3 processes (1 current champion, plus 2 or 3 from the archive), due to the fact that the 3-LEO and the 4-LEO seemed to bring optimal results in the previous section.

Figure 7.3 provides the results of equal-effort comparisons between the 3-LEO, the 4-LEO, and (1+2) and (1+3) sliding archives. These graphs indicate that the sliding archive provides an advantage in the long run, though not in the short term. First, the (1+3) is significantly

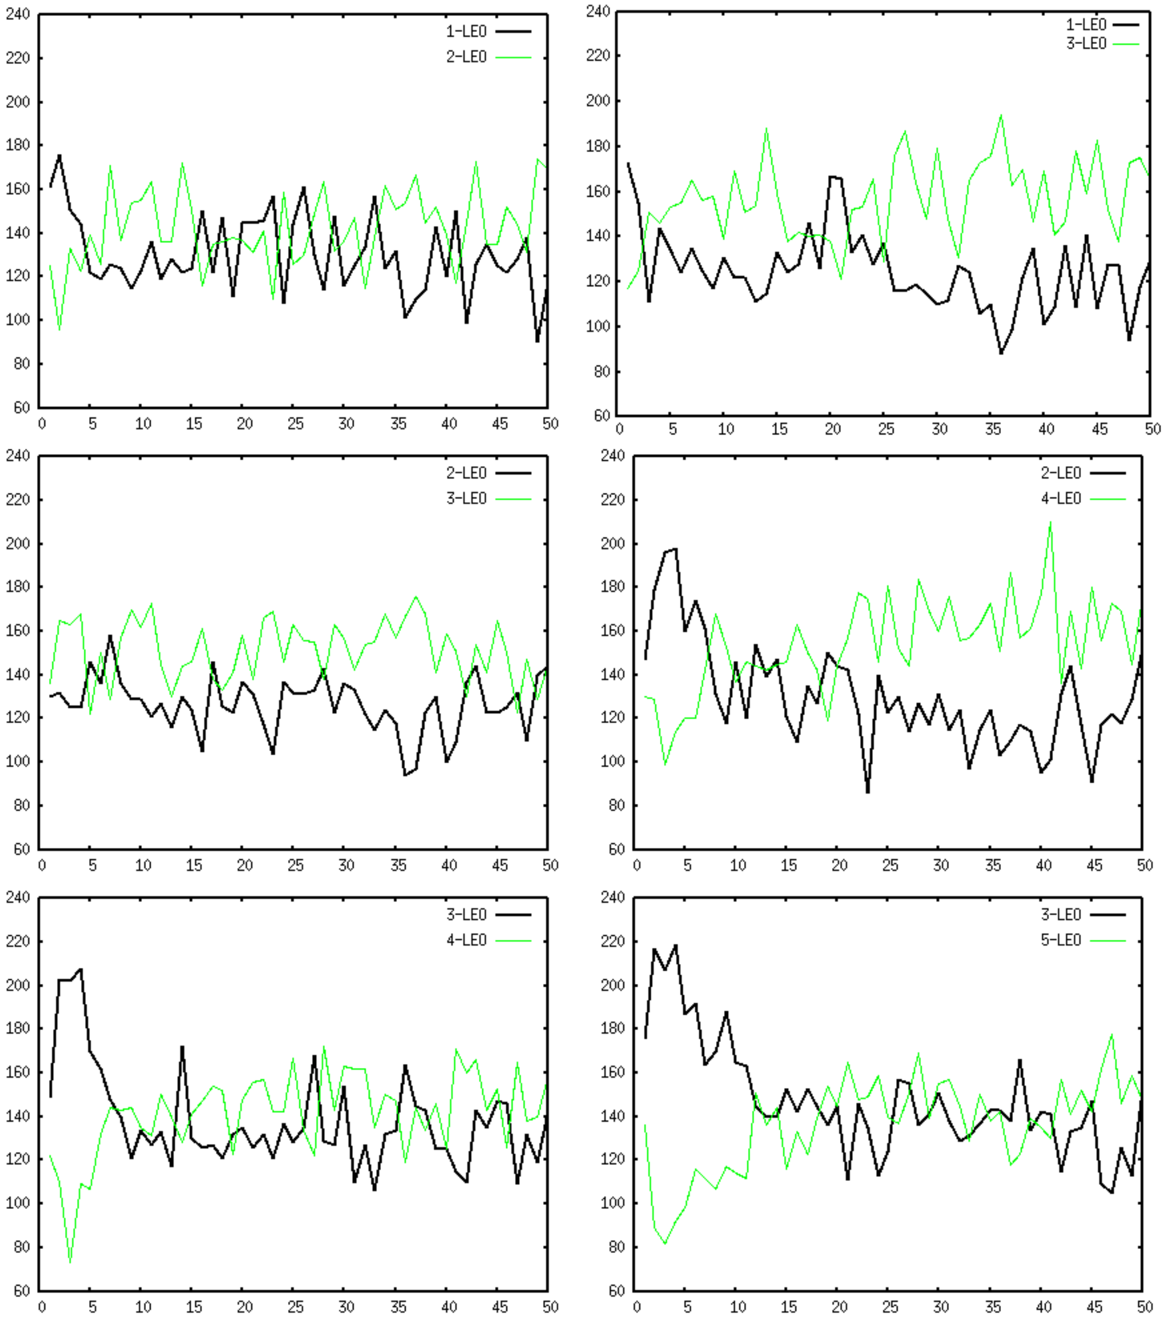

Figure 7.1: Results of equal-effort comparisons between various instances of the N-LEO. 50 champions are sampled every 4000 evaluations, for a total of 200000 evaluations per run. While 392 competitions are performed at each rank, the number of victories for each algorithm at each rank do not sum up to 392 due to the widespread occurrence of draws (usually when fighters avoid each other and no contact occurs.)

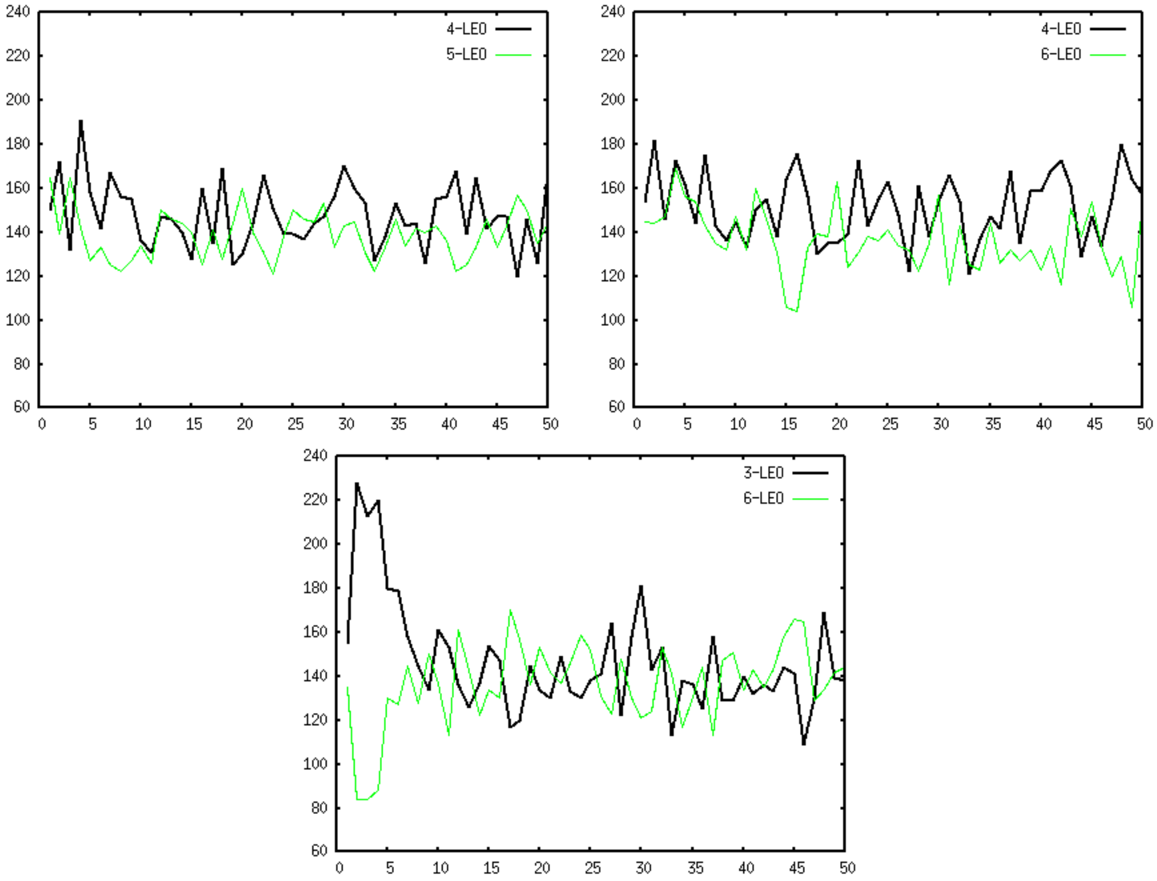

Figure 7.2: Continuation of Figure 7.1.

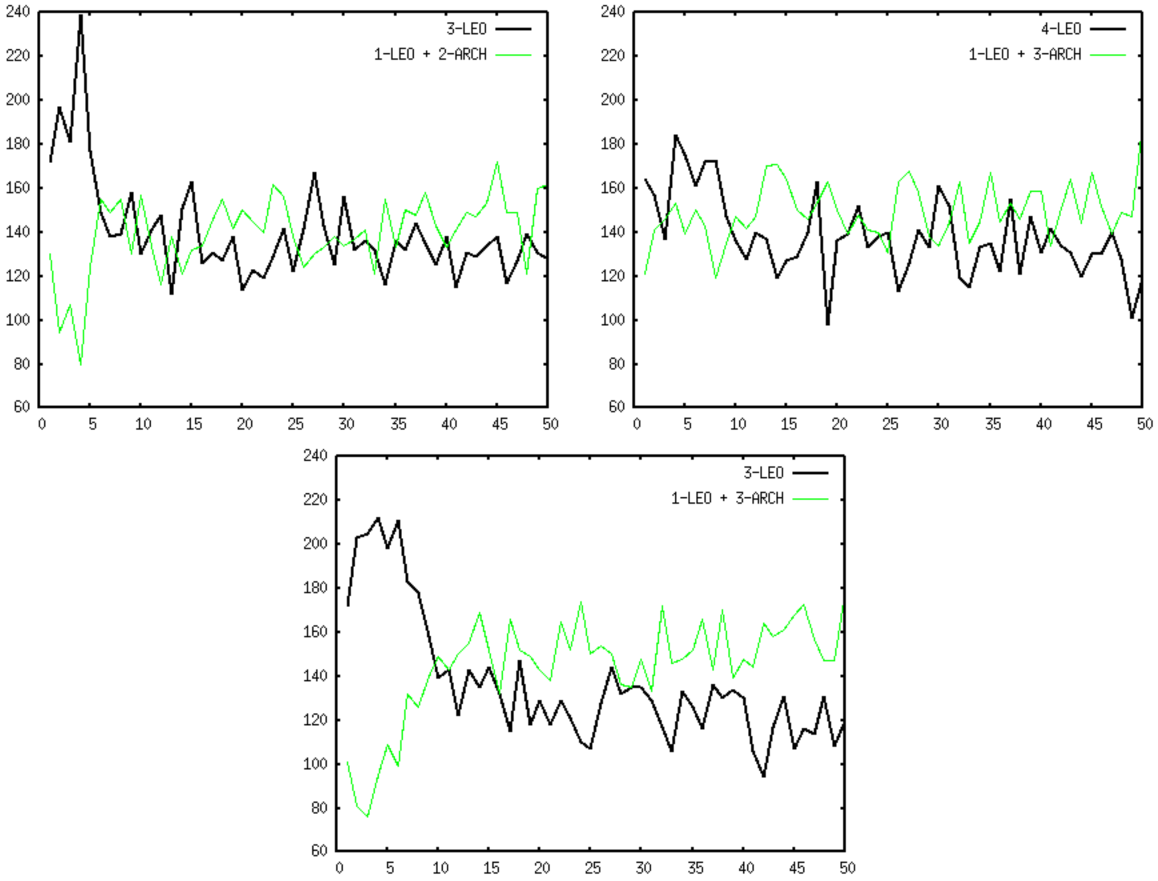

Figure 7.3: Results of equal-effort comparisons between the 3-LEO, the 4-LEO, and (1+2) and (1+3) sliding archives. Note the clear superiority, at least in the long run, of the 1+3 archive process (1 current champion, 3 from the sliding archive.)

superior to the 3-LEO ($p < 0.032$). There is also indirect evidence that the (1+3) process may be superior to the 4-LEO in the long run; this evidence is indirect because the small, apparent long-term difference between (1+3) and 4-LEO was not found to be significant ($p < 0.194$). However, the possible superiority of the (1+3) archive process over the 4-LEO is supported by the fact that the (1+3) process does obtain a significant advantage over the 3-LEO, whereas the 4-LEO did not (as observed in the previous section).

In the very short term, sliding archives actually seem to fare worse than simple N-LEO methods; this is compounded by the similar effect caused by increasing the raw number of competitions (observed in the previous section), as demonstrated by the comparison between the 3-LEO and the (1+2) sliding archive.

7.4 Conclusion

These results support three conclusions:

- Using several champions rather than just one significantly improves the performance of the algorithm.
- The benefits of increasing the number of champions peaks early between N=3 and N=4, with the latter being more efficient in the long run at the cost of a performance hit in the short term.
- Using a sliding archive improves the performance in the long run, especially when using 4 (1 current + 3 archive) champions. In the long run, this latter method is the best of all considered methods. However, in the short term, the archive does bring a performance hit.

The first result is clearly not surprising. By increasing the number of opponents used for each evaluation, we expect that the reliability of these evaluations will increase. However, the fact that this improvement seems to peak at N=3-4 deserves attention. This result suggests that large numbers of competitions are not necessary to obtain reliable evaluation. This is in contrast with many coevolutionary methods that require large numbers of competitions. The positive results (in the long run) of using a sliding archive validate the use of this method, though the short-term performance hit also deserves attention.

Whether these conclusions are specific to our particular domain (physical combat between virtual creatures) or hold over a wide class of problems can only be ascertained through further experiments. However, due to the rich possibilities offered by physical combat, we suspect that similar results would be found for many other non-trivial domains.

Appendix: tests of statistical significance

Information from graphical methods is only useful if it can be confirmed with formal significance testing. To test the significance of our data, we applied the following method. Recall that for each experiment, our data consisted of 14 runs of each algorithm (let us call these algorithms I and II), each run involving two populations arbitrarily labelled A and B. First, we arbitrarily match runs of both algorithms in pairs: to each run of one algorithm, we randomly associate one unique run of the other algorithm. This gives us 14 pairs of runs, each pair containing one unique run from algorithm I and one unique run from algorithm II. The purpose of this pairing is to obtain independent samples.

Then, for each such pair, we pit the current A-champion of each algorithm against the current B-champion of the other, after every multiple of 4000 evaluations. Thus, after every multiple of 4000 evaluations, we have two contests per pair of runs (A_I vs. B_{II}, and A_{II} vs. B_I). Note that this data was already collected during equal-effort comparison.

Importantly, we only consider the information obtained in the second half of all runs (that is, after 100000 evaluations have been performed), since we are interested in the long-run behaviour of algorithms.

Then, for each pair, we simply add up the number of victories obtained by either algorithm, from 100000 evaluations onwards. This gives us 14 pairs of victory counts. We perform a two-tailed Wilcoxon Matched-Pair Signed Rank test using these 14 pairs of values. The resulting value of p is used to assess the significance of the difference, as reported in the text.

Chapter 8

The relationship between evolution and complexity

8.1 Introduction: the "Arrow of Complexity" hypothesis

It appears that the maximum level of organismal complexity in the biosphere has increased, in some sense and at a certain scale, over time. Bacteria are simpler than humans, if only because humans contain millions of (interacting) cells, each much more complex than bacteria. Sponges are simpler than flatworms, which are simpler than trilobites, which are simpler than birds. This crude, large-scale growth of the maximum of complexity seems to hold for any notion of organismal complexity which we know of (and as we will see there are quite a few of them). To our knowledge, no serious author has explicitly denied an overall increase in the maximum of complexity at a global scale, though some have urged caution [117]. At a slightly lower level, there is some evidence that the maximum number of cell types (a common measure of organismal complexity) has increased over time within metazoans [19, 178].

Faced with such a remarkable increase in complexity, some authors have suggested a fundamental tendency of evolution to create more and more complexity. This position actually predates Darwin, being found in Lamarck. One variant of this notion is the "Arrow of Complexity" hypothesis [12]:

> . . . the hypothesis of the arrow of complexity asserts that the complex functional organization of the most complex products of open-ended evolutionary systems has a general tendency to increase with time.

It turns out that this apparently straightforward statement has generated significant amounts of confusion and misunderstandings. We need not go very far to find one example: in the same article, the author goes on to express that this hypothesis is not accepted by Gould[71], Maynard Smith and Szathmáry[165] and McShea [117]. Yet these authors, and those works, do *not* deny that the complexity of the most complex organism increases in evolution. At most, McShea

---

This chapter is based on a paper that has been accepted for publication in a special issue of the *Artificial Life* journal on the evolution of complexity. It benefited greatly from the comments of several anonymous reviewers.

professes “emphatic agnosticism”. Gould, however, writes (p. 169): “I do not challenge the statement that the most complex creature has tended to increase in elaboration through time” - an almost verbatim statement of the arrow of complexity hypothesis[1]. As for Maynard Smith and Szathmáry, their book is a description of several successive “major transitions” in the history of evolution, many of which can reasonably be interpreted as increases in complexity (more precisely hierarchical complexity[117]).

How can such confusion arise? The crux of the matter, of course, is that while nobody seems to utterly deny the *fact* that maximum complexity has increased over the entire course of evolution, there is much debate regarding the *interpretation* of this fact: as we will see, what these authors explicitly question is not that increasingly complex individuals are found in Nature, but rather the notion of an built-in evolutionary *drive*, or *bias*, towards increased complexity in evolution.

It is one of the achievements of recent evolutionary theory to have shown that this controversy results in no small part from a conjunction of misunderstandings. The crux of the matter is that the question of whether ‘there is a tendency for complexity to increase in evolution’ is an extremely ambiguous question. This is because almost every word in it can be understood in many different ways, leading to much unnecessary confusion. When fuller attention is given to details (and to the implicit assumptions which may lurk behind apparently innocuous terms), it is realised that the sustained growth of maximum or average complexity does not necessarily imply a pervasive bias towards complexity in evolution.

8.2 What is complexity?

8.2.1 Descriptive and functional complexity

If we are to study the relationship between evolution and complexity, we need to explain what we call “complexity” in the first place, if only to avoid unnecessary confusion. It turns out that definitions and measures of complexity abound throughout the literature (for a broad panorama see [50, 117, 2]). We will concentrate specifically on two notions of complexity, which we believe are of particular interest when discussing the effects of Darwinian evolution: descriptive (or morphological) complexity, and functional complexity.[2]

Descriptive or morphological complexity is easily defined: it is the amount of information that is necessary to describe an object. In the words of Dawkins [40]: “A crab is morphologically more complex than a millipede because, if you wrote a pair of books describing each animal down to the same level of detail, the crab book would have a higher word-count than the millipede book. The millipede book would describe a typical segment then simply add that, with listed exceptions, the other segments are the same. The crab book would require a separate chapter

---

[1] In a different, related piece [69], Gould writes: “No one can doubt that more complex creatures arose sequentially after this prokaryotic beginning - first eukaryotic cells, perhaps about two billion years ago, then multicellular animals about 600 million years ago, with a relay of highest complexity among animals passing from invertebrates, to marine vertebrates and, finally (if we wish, albeit parochially, to honor neural architecture as a primary criterion), to reptiles, mammals and humans.”

[2] Adami and colleagues [3, 2] have introduced a specific concept called “physical complexity”. They claim that this notion of complexity “is forced to increase” in evolution. However, physical complexity is quite distinct from the everyday concept of complexity. See Appendix F for details.

for each segment and would therefore have a higher information content.” However, very much in the same way as with Shannon entropy (another measure of information content), it turns out to be maximal for completely disordered objects. A pile of rubble consists of many more parts, each much more different from any other, than a standing wall of bricks, so its descriptive complexity would be higher.

What about functional complexity? We choose to simply define functional complexity as the descriptive complexity of a *functional* system. A ‘system’ is any set of interacting parts. We call it ‘functional’ if the interaction of the parts produces a certain external result, which would not be obtained if some (or any) of the parts failed to behave just as it does. Note that the result must come from the actual interaction of active parts, not just from their mere existence. Functional complexity is simply descriptive complexity applied to any such functional system, that is, a growing function of the number of parts and of their mutual differences.

While we admit that this is a rather crude definition, we believe that it does capture much of the intuitive notion of functional complexity within a system. In particular, it does retain the idea of ‘difficulty’ (in a design sense), or ‘improbability’, which we associate with highly complex biological features: clearly a certain system which depends on the precise behaviour of many parts is somehow more improbable or surprising than one which only depends on the collaboration of a few. Similarly, it stands to reason that systems relying on the interactions of many similar parts are somehow easier to arrive at than systems relying on the precise features of many differentiated parts.

This notion of functional complexity may be contrasted with that of Heylighen [79], which emphasises the joint occurrence of separation and interconnection; or to that of Dawkins [39], which stresses the improbability of the result (rather than the ‘emerging improbability’ resulting from the highly constrained interactions of many parts). For example, in our view, ‘Rube Goldberg’ contraptions (phantasmagoric constructions in which highly intricate machinery is used to obtain trivial results) would be regarded as functionally complex, whereas Dawkins’ view would regard them as comparatively non-complex, due to their underwhelming output. ‘Functional complexity’, in our sense, is more about the functioning than the function being performed.

8.2.2 Implications of the definitions

With this notion of functional complexity in mind, we can already make a few inferences. The most important one is that adaptive functional complex systems are *rare* in the space of all possible systems. Any given set of parts is unlikely to be actually performing any meaningful function, let alone an adaptive one. This unlikeliness increases with the descriptive complexity of the system: the more parts there are, the more improbable it is that they will harmoniously cooperate (a requirement for functional complexity). This is important, because it implies that generating random systems would be a very inefficient way to come up with adaptive complex systems, while other search methods (such as Darwinian evolution) might, under the right conditions, prove more successful.

Another observation is that although adaptive functional complexity is rare in the space of all systems, it comes in ‘loose clusters’: coming up with a new, more complex, yet still successful design is typically much easier if we use an already complex, successful design as a starting point. More precisely, slightly altering a successful complex system has a higher chance of yielding another system of equivalent or higher complexity than random search. If we have access to the mechanism of a working clock, random alterations to this existing mechanism are much more likely to result into other working clocks than random assemblages of gears. Importantly, this is not a logical necessity; rather, it is a reasonable inference which we may assume holds for many system spaces.

8.3 What is a tendency?

While having a working notion of complexity is important, it appears that much of the confusion actually derives from the notion of *tendency.* What do we mean when we say that the maximum or average complexity "tends" to increase in evolution?

8.3.1 Trends, passive trends and driven trends

McShea [116, 117, 119] has devoted much work to the conceptual clarification of trends and mechanisms in evolution. The most widely known result of this work is probably the operative distinction between passive and driven trends. A *trend* is any sustained directional change in a summary statistic of a certain quantity (such as maximum or average size, complexity, etc.) within a particular clade. Such a trend is called *driven* if it can be shown to result from a pervasive, consistent *bias*, applying to all sub-clades; for example, if speciation or extinction events are consistently biased towards higher values of the quantity. *Passive* trends, on the other hand, are those which occur mechanically in globally unbiased processes, possibly constrained by a local inhomogeneity - for example, a 'hard' boundary which the value cannot cross. Figure 8.1 illustrates this distinction graphically.

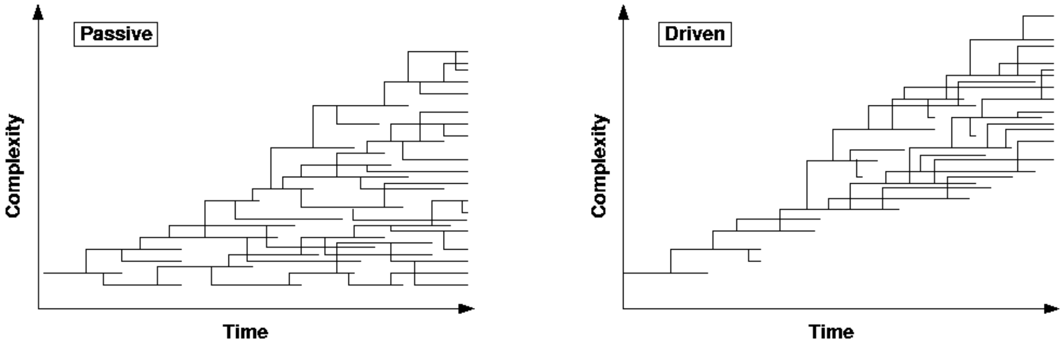

Figure 8.1: Passive and driven trends. In both of these graphs, the maximum and the average increase sustainedly through times, creating a trend. In the left-hand graph, these trends are *passive*: the branching process is not biased towards higher or lower values. A trend in the maximum emerges mechanically by envelope expansion. A trend in the average also emerges, because the process is constrained by a 'hard' lower bound which cannot be crossed. Both trends, resulting from a globally unbiased process, are passive. In the right-hand graph, however, the trends of increase in the average and the maximum are *driven*: the underlying branching process is clearly biased towards higher values.

It is important to note that passive trends may well be enforced by active, external forces, such as natural selection. For example, we might have a situation where no member of a certain clade is selectively viable under a certain size, or complexity level. In this case, provided there is no constraint on the maximum, the average would mechanically grow over time.

Furthermore, even in a globally passive trend, the small-scale behaviour of each individual lineage may well be everywhere governed by a local, driven trend - as long as these many local trends do not exhibit any consistent global direction. In other words, *passive trends may well be everywhere driven*[3] [116, p.1751]. Conversely, non-selective forces may enforce driven trends. Thus the important distinction between passive and driven trends does not lie in the source or mechanism of the bias which creates this trend, but solely on whether this bias is *local* or *pervasive*, respectively, throughout the space explored by the clade [119].

8.3.2 Benefits of the terminology

This sound terminology goes a long way to identify (and prevent) possible sources of misunderstandings. Restated in McShea's terminology, the "arrow of complexity" hypothesis simply states that evolution spontaneously generates a trend of increase in the maximum of organismal complexity over time. It is silent on the question of whether this trend is driven (the result of a pervasive bias in Darwinian evolution towards higher complexity) or passive (the result of a direction-blind branching process, mechanically pushing the external envelope).

We can now see where much of the controversy arises: what some authors (especially Gould [71]) reject is not the existence of the trend, but the idea that it should necessarily be driven, rather than passive. This, in turn, seems to be misconstrued by others as denying the existence of the trend itself, or the role of evolution in its emergence: the implicit assumption is that any trend must of necessity be driven, and therefore, that denying bias is tantamount to rejecting the trend itself. This seems to explain why Gould, for example, is portrayed as a "skeptic" regarding the existence of a trend [12], or even as denying that "complexity increases in evolution" [2]. Gould, as we saw, certainly agreed that complexity increases in evolution: what he rejected was the idea of a general, pervasive bias or drive towards complexity. This illustrates the high potential for confusion that lies behind such deceptively straightforward terms as 'trend' and 'tendency'.

8.4 Reformulating the question

Using McShea's terminology, we can now reframe our previously ambiguous question in a more satisfying manner. We admit that, at a sufficiently large scale, there has been a trend of increasing complexity in evolution: there has (so far) been an "arrow of complexity" from exclusively simple forms to a variety of simple and complex forms. This point is uncontroversial and widely accepted, though appropriate caution is necessary [117]. We are now left with two

---

[3] This point deserves particular attention: in everyday language, the words "passive" and "driven" seem to be mutually exclusive. But in McShea's terminology, a trend is passive if it is *globally* unbiased, which does not imply that its local behaviour at any point must be: even in a globally unbiased (passive) branching process, each particular branch may well be strongly biased (driven) in its own local direction. Any overall trend will still be called passive if the many local biases are not correlated, resulting in a globally unbiased process.

questions:

1. Has this trend been driven (whether by selection or by some other mechanism), or passive? Is there a built-in bias in evolution to push average or maximum complexity preferentially upward, in general?
2. What is the role of the evolutionary mechanism in the emergence of this trend? What are the necessary conditions for the arrow of complexity to arise in Darwinian evolution?

The following sections will address both questions in turn.

8.5 Theoretical justifications for a driven complexity growth

It is interesting that, independently of the evidence for or against a general, driven trend of increasing complexity in evolution, several theoretical justifications have been suggested for such a driven trend. We will discuss a few typical examples.

8.5.1 The "complexity advantage" assumption

An apparently common (though often implicit) viewpoint is that functional complexity should be favoured by selection, in general, because living organisms can generally be made more efficient by increasing their functional complexity. A statement of this position is apparently offered by Darwin himself, in the crucial fourth chapter of the *Origins of Species* [34] (quoted with this intent by Rosslenbroich [153]):

> If we take as the standard of high organisation, the amount of differentiation and specialisation of the several organs in each being when adult . . . , natural selection clearly leads towards this standard: for all physiologists admit that the specialisation of organs, inasmuch as in this state they perform their functions better, is an advantage to each being; and hence the accumulation of variations tending towards specialisation is within the scope of natural selection.

However, the continuation of this passage (*not* quoted by Rosslenbroich) paints an altogether different picture:

> On the other hand, we can see . . . that it is quite possible for natural selection gradually to fit a being to a situation in which several organs would be superfluous or useless: in such cases there would be retrogression in the scale of organisation. . .
>
> But it may be objected that if all organic beings thus tend to rise in the scale, how is it that throughout the world a multitude of the lowest forms still exist; and how is it that in each great class some forms are far more highly developed than others? Why have not the more highly developed forms everywhere supplanted and exterminated the lower? . . . On our theory the continued existence of lowly organisms offers no difficulty; for natural selection, or the survival of the fittest, does not necessarily include progressive development - it only takes advantage of such variations as arise

> and are beneficial to each creature under its complex relations of life. And it may be asked what advantage, as far as we can see, would it be to an infusorian animalcule - to an intestinal worm - or even to an earthworm, to be highly organised. If it were no advantage, these forms would be left, by natural selection, unimproved or but little improved, and might remain for indefinite ages in their present lowly condition. And geology tells us that some of the lowest forms, as the infusoria and rhizopods, have remained for an enormous period in nearly their present state.

To the naïve view that natural selection spontaneously leads to complexification, Darwin opposes a more neutral notion of evolutionary progress, according to which species simply adapt to their own particular conditions. This may, or may not, lead to increases (or decreases) in complexity over time. In other words, the flaw of the 'complexity advantage' assumption is that, while it rightly recognises the possible advantages of complexity, *it fails to consider the possible advantages of simplicity.* The world is sufficiently complex that there is an almost infinite number of ways to 'earn a living', and on the whole, evolving lineages may just as readily gain advantage through simplicity as through complexity.

8.5.2 Modern endorsements of the "complexity advantage" assumption

Despite Darwin's own objections, several authors have used the complexity advantage assumption as an implicit justification for the idea that Darwinian evolution by natural selection entails a drive towards complexity.

Coevolution and "arms races"

One way in which evolution is sometimes said to favour the increase of complexity is through coevolution between lineages which must adapt to each other's evolving features. As seen in chapter 3, this is expected to lead to an "arms race"[41], in which coevolving lineages should come up with more and more adaptive features to defeat each other. Let us recall Rosin and Belew's summary of the argument [151]:

> When the parasites evolve to overcome this failure [against the hosts], they create new challenges for the hosts; the continuation of this may lead to an evolutionary "arms race" [39]. New genotypes arise to defeat old ones. New parasite types should serve as a drive toward further innovation, creating ever-greater levels of complexity and performance by forcing hosts to respond to a wider range of more challenging parasite test cases.

It is interesting that these authors, like several others, cite Dawkins as a source. Dawkins, co-author of the original "arms race" hypothesis, has explicitly and emphatically denied that the progress brought about by arms race was intrinsically linked with complexification - as we shall see in section 8.5.4.

But at any rate, the problem with this position is easily identified: it implicitly assumes that the most straightforward way to cope with an opponent's increase in complexity is a corresponding, 'countering' increase in complexity. In other words, complexity calls for complexity, because only complexity can cope with complexity. Unfortunately, as many practitioners of artificial coevolution (including Rosin and Belew themselves) have remarked [183, 55, 125], the opposite is often true: simplicity is often quite efficient at dealing with complexity, with devastating effects. The most obvious example is the phenomenon of *opportunism*, or parasitism, whereby promising complex organisms are systematically defeated by simplistic opponents which successfully exploit some of their 'weak points' (what Watson and Pollack call "focusing on the wrong thing" [183]). We may also mention "mediocre stable states" [55] (also called "collusion" [142] or "immature inductions" [6]), where competitors remain stuck in a small, uninteresting section of the search space from which none has any incentive to budge - the coevolutionary equivalent of a local optimum. As a special case of the latter, we may mention the possibility of cycles [133], in which competitors cycle through a short, predictable sequence of simple behaviours without any long-term gain in general performance or complexity[4].

In general, these problems can be summarised by this simple observation: evolution is interested in immediate advantage, rather than in complexity or generality, and there seems to be no general relationship between the two: to reliably obtain the latter, one must forcefully connects it to the former by means of algorithmic contrivance [133, 151, 170, 162, 125].

To be sure, there are many examples of specific situations where certain coevolutionary pressures have clearly driven increases in complexity; some will be discussed in section 8.6. What must be doubted is not the existence of coevolutionary-driven increases in complexity, but the fact that these should be a general rule rather than interesting specific cases.

A cybernetic argument for the complexity advantage

Heylighen [79] uses a cybernetic argument to justify the complexity advantage assumption. In short, natural selection should favour organisms which are able to cope with a larger variety of perturbations; but, by application of Ashby's law (Heylighen argues), this in turn requires higher 'countering' variety in the behaviour of these organisms, and thus more functional complexity:

> [Ashby's Law of Requisite Variety] states that in order to achieve control, the variety of actions a control system is able to execute must be at least as great as the variety of environmental perturbations that need to be compensated. The larger the variety of available counteractions, the larger the set of disturbances that can be corrected, and the larger the domain of potential environmental situations in which the control system can survive. All other things being equal, greater control variety implies greater fitness... Therefore, evolution through natural selection will tend to increase control, and thus, because of Ashby's law, internal variety. This can be interpreted as a functional differentiation, i.e. the appearance of more diverse activities or functions...
>
> All other things being equal, a system that can survive situations A, B and C, is absolutely fitter than a system that can only survive A and B. Such an increase in absolute fitness is necessarily accompanied by an increase in functional complexity. Thus, evolution will tend to irreversibly produce increases of functional complexity.

---

[4] Interestingly, examples of such coevolutionary cycles have actually been reported in Nature [164].

This leads to interesting questions. For example, by any measure of organismal complexity, beetles are arguably less complex than chimpanzees. Yet there does not seem to be many "situations" that could be survived by chimpanzees, but not by beetles - as opposed to the other way round. More generally, it is not immediately clear that more complex individuals can cope with more environments than simpler individuals.

Quite possibly, the crucial difficulty in this argument is that Ashby's law applies when one needs to actively counter certain variations in order to homeostatically control certain variables; but in Nature, *adapting* to a change in an external variable does not necessarily entail *countering* this change. Evolving lineages may well cut their dependency on certain variables (as in previously carnivorous lineages turning into herbivores), or simply ignore the variation in question. To use Heylighen's own example of "warm-bloodedness," we note that birds and mammals are the only surviving lineages of warm-blooded animals: insects, whose range, diversity and total biomass vastly exceed that of either mammals or birds, have hardly suffered from their lack of internal temperature control mechanisms. If Nature is any guide, surviving more "situations" can be obtained by simplicity just as readily as by complexity, which contradicts the premise of the argument[5].

To summarise, it seems that the cybernetic argument, like the arms race argument, attributes a general advantage to complexity by disregarding the potential advantages of simplicity: it essentially assumes that only complexity (in the evolving lineage) can deal with complexity (in the environment).

8.5.3 Cope's rule and the "size-complexity insight"

Bonner [19] has proposed a way in which natural selection might indirectly drive an increase in the maximum of complexity. The basis of this argument is Cope's rule, that is, the generalisation that there is a tendency for lineages to grow in size over time. Bonner provides a justification for this rule by pointing out that "there is always room at the top": "by becoming larger the organisms enter new size worlds where, among other things they avoid predation and competition." Bonner also notices that size is (very roughly) correlated with complexity: on the whole, larger organisms tend to be more complex. Bonner takes the practical approach to define complexity as the number of cell types in an organism.

Bonner argues that the relationship between size and complexity is explained by considerations of engineering and efficiency: creating and maintaining a larger body is likely to require higher division of labour. Thus, a selective drive for size creates an indirect selective drive towards complexity, though Bonner is very clear that this is only one component among many of selective action on complexity (p. 221). He repeatedly points out that this selection-driven increase in size and complexity only applies at maximum values: intermediate forms have certainly been maintained, and "every biological community on the globe today contains a complete spectrum from the smallest to the largest organism" (p. 221). In addition, Bonner emphasises that selection for increase and decrease in size and complexity have occurred concomitantly throughout evolution.

---

[5] The author extends the cybernetic argument to coevolutionary situations and arms races, which can be answered in a similar manner: adapting to an "increase in variety" in a competing species does not entail matching this variety. Adaptation is simply not a problem of control in Ashby's sense.

This argument is not without difficulties. First, it relies on Cope's rule, which is itself the subject of recurrent controversy within evolutionary research [89, 5, 67]. But more generally, the argument does not seem to provide any compelling evidence to show that the observed trend is indeed driven (biased), rather than passive (diffusive). In fact, the very possibility of a passive trend is simply not considered at all. It should be pointed out that the book was written before McShea and Gould began to draw attention to the question of trend mechanisms. Because the possibility of a passive trend is not considered, the increase in maximum value is seen as requiring an explanation, a causative force, which happens to be the "room at the top" principle[6].

In fact, with hindsight, the wealth of data and evidence provided by Bonner appears to be not only compatible, but strikingly suggestive of a passive trend. This applies particularly to ancestry data taken from the fossil record (such as Fig. 8, 9 in the book), which do not seem to show any pervasive upward bias, but rather an apparently undirected branching process constrained by a lower bound. Also Gingerich's evidence [64, 63], which Bonner discusses several times (including in the passage quoted above), is suggestive of a passive trend: when no pervasive bias exists, but rather species evolve both up and down indifferently along the scale and the sole constraint is a local boundary, we may expect that the overall rate of change will decrease as we average it over longer periods[7]

8.5.4 Expert opinion against the complexity advantage assumption

More generally, theoretical justifications of a complexity advantage face a difficulty: the variety of environments (both biotic and abiotic) faced by evolving lineages, as well as their remarkable plasticity, implies that at any time there is a very wide range of ways to 'earn a living' in Nature. Optimising the efficiency of a population in its chosen way of life (or, if the opportunity arise, switching trades altogether) is what Darwinian evolution does. There does not seem to be any compelling reason why the majority of natural lifestyles should benefit from continuous increases in complexity.

This point has been admirably expressed by Dawkins, in an unduly overlooked piece [40] from which we will now quote at length:

> Why should any thoughtful Darwinian have expected a majority of lineages to increase in anatomical complexity? Certainly it is not clear that anybody inspired by adaptationist philosophy would. . . Our human line happens to have specialised in complexity, especially of the nervous system, so it is only human that we should define progress as an increase in complexity or in braininess. Other species will see it differently, as Julian Huxley. . . pointed out in a piece of doggerel entitled "Progress": -
>
> The Crab to Cancer junior gave advice:

---

[6] Incidentally, this may show that Gould's and McShea's argument is not quite as "obvious" or "routine" as it may have seemed to some [40].

[7] This interpretation of Gingerich's data now seems consensual: see Weiner[184], quoted in pp. 106-109 of [20]. Valentine and colleagues [178] provide refined data on the increase of cell types among metazoans, and model this increase with a passive system.

> Know what you want, my son, and then proceed
> Directly sideways. God has thus decreed
> Progress is lateral; let that suffice.
>
> Darwinian Tapeworms on the other hand
> Agree that Progress is a loss of brain,
> And all that makes it hard for worms to attain
> The true Nirvana peptic, pure and grand.
>
> Man too enjoys to omphaloscopize.
> Himself as Navel of the Universe...

> ...Huxley's tapeworms, using a parasite-centred definition of progress, see the point with opposite sign [from human authors, who focus on increasing complexity]. A statistically minded swift would search in vain for evidence that a majority of evolutionary lineages show trends towards improved flying performance. Learned elephants, to borrow a pleasantry from Steven Pinker..., would ruefully fail to uphold the comforting notion that progress, defined as a driven elongation of the nose, is manifested by a statistical majority of animal lineages.

In this quote, the opinions of three authors (all of them well-known evolutionists, all of them staunch supporters of the power of natural selection) are cited. It is interesting that this passage is taken from a review of Gould's own book on the subject of progress [71], which on the question of complexity adopts pretty much the same position. In the beginning of this section we showed that Darwin had already expressed very similar feelings. Let us also cite Maynard Smith and Szathmáry, who unambiguously reject the notion that evolution should be biased towards complexity in the first chapters of their book [165] (this chapter carries the rather explicit title "The fallacy of progress"). Similarly we may quote Mayr [113] (see also [9, sec. 3.2]):

> Among all the forms of life, neither the prokaryotes nor protists, fungi or plants has evolved smartness, as it should have if it were "better." In the 28 plus phyla of animals, intelligence evolved in only one (chordates) and doubtfully also in the cephalopods. And in the thousands of subdivisions of the chordates, high intelligence developed in only one, the primates, and even there only in one small subdivision. So much for the putative inevitability of the development of high intelligence because "it is better to be smart."

All in all, we note that there is much expert opinion against the idea that natural selection should drive a pervasive trend towards increasing functional or behavioural complexity in evolution. Rather, the consensus seems to point towards Darwin's original conception of evolution by natural selection: a recklessly opportunistic process, which will constantly exploit any possible way to make a living, independently of whether it leads to increases or decreases in complexity.

8.6 The evidence: complexity in evolution

8.6.1 The enduring dominance of simple lifeforms

A significant argument against the idea of an active, pervasive, selection-driven bias towards complexity growth over the course of evolution is that simpler life forms still dominate the planet. In a typically engaging section, Gould [71] famously argued that "on any possible, reasonable or fair criterion, bacteria are - and always have been - the dominant forms of life on Earth." Prokaryotes (Bacteria and Archaea) are more ancient, more numerous, more diverse, have conquered more environments and apparently compose vastly more of the cellular biomass than all other living entities *combined* (a position only reinforced by later estimates [185]). Simpler prokaryotes have clearly not been displaced by their more complex eukaryotic cousins.

We note that it is possible to extend this pattern to other taxonomic levels. The most functionally complex animals, besides humans, are arguably great apes - of which, it turns out, many happen to be on the verge of extinction. We may compare this with the success (whether in numbers, range, or conservation prospects) of 'lesser'[8] simians, especially the genus of macaques, the most widespread of all primate genera. Non-human primates in general are significantly *less* successful, by any measure of 'success', than the placid artiodactyls (ruminants, pigs, etc.). Chordates (the phylum which includes vertebrates) are vastly exceeded in number, diversity and total biomass by arthropods, the truly dominant form of metazoan life [14]. Land animals in general are simply overwhelmed by land plants in terms of total biomass. The dominance of prokaryotes, Gould's "modal bacter", merely crowns what appears to be a general pattern, or rather the absence of a pattern: at any given taxonomic level, more functionally complex entities are simply not more successful than their simpler counterparts. In fact, a lot could be said for the opposite position, though of course counter-examples abound.

The important fact here is not the enduring existence of some primitive lifeforms, per se. What is significant is that at almost any given scale, and for any sub-tree of the Tree of Life, we observe that simpler organisms not only endure, but often seem to dominate the more complex organisms. This leads to the following question: if evolution is indeed generally biased towards complexity, then why do only a minority of lifeforms become more complex? Why did all those simple forms fail to feel the supposed "pull" of the complexity drive? This pattern would suggest that those taxa which did reach high complexity did so because of certain particular, specific selective pressures, to which they somehow adapted by increasing their complexity.

8.6.2 Examples of particular selection-driven trends towards complexity

Considering the amount of confusion which has been pointed out by McShea, Gould, Dawkins and others, it may not be entirely futile to stress the obvious: nobody denies that in many cases, active, selection-driven increases in complexity do in fact occur. Complex adaptations often do confer a definite selective advantage over relatives which do not possess them. The result is that they will mechanically propagate throughout the population. Such trends have happened at all levels, ranging from individual sub-species to the entire biosphere, both within and across

---

[8] Quantitative measures of behavioural complexity and neocortex development among primates are given by Dunbar. [48]

lineages. We will only give a few examples which we believe are of particular interest.

For a specific example, we need look no further than the evolution of our own species, *Homo Sapiens* [20, 105]. Modern humans are the latest stage in a sequence of various groups, starting with the early australopithecines (essentially bipedal apes), going through the early Homo branch (the 'Habilis group': *H. Rudolfensis, H. Habilis. . .*), the 'Erectus group' (*H. Ergaster, H. Erectus. . .*), the 'archaic' *Homo Sapiens* (*H. Heidelbergensis, H. Neanderthalensis*), and eventually modern humans, presumably descended from the same branch as *H. Heidelbergensis* (in parallel to the Neanderthal derivation). Now this sequence shows two notable features. First, each new group showed demonstrably more complex behaviour than its predecessors, as can be deduced from artifacts. A trend of increase in brain size is also apparent (especially in certain stages, such as the Erectus and Sapiens groups [105]). Second, at a sufficiently large scale, these groups tended to thrive while their predecessors were displaced. It is therefore reasonable to argue that this sequence constitutes a trend, driven by natural selection, whereby increasingly complex groups successively displace their predecessors. Is it possible that the specific case of human evolution, in which natural selection did drive a sequential increase in complexity, influenced our view of evolution in general?

However, even in this most indisputable of trends, it is important to exert caution. First of all, the appearance of early Homo certainly did not displace the australopithecines. On the contrary, australopithecines continued to expand and radiate, with a trend towards 'robust' forms (*A. Basie, A. Robustus*), for at least another million years [61]. The divergence somewhere along the line of the early Homo branch (the Habilis group) was apparently a "much more isolated and insignificant trend" [105, p.328] which locally reversed the general tendency. Only with the 'Erectus group' did Homo finally achieve prominence. Furthermore, it is not absolutely clear that the displacement of each group by its successor was invariably caused by the latter's higher complexity: climate changes have clearly had an influence on the evolution [101] (and apparently the extinction [60]) of various human lineages. Thus, even in the clearly driven trend of human complexification, the pluralism of evolution shows its signs.

More generally, the evolution of complex organs constitutes a remarkable, ubiquitous example of selection-driven increase in complexity. Historically, the paradigm of complex organ evolution is the eye, a high feat of engineering that has been discovered many times in the course of evolution. Dawkins [39] has famously described the highly complex sonar of certain bats, using it as a replacement for the often-used vertebrate eye as an example of complex organ evolution. The evolution of such complex organs, each composed of several, highly heterogeneous, interdependent parts delicately assembled to obtain a certain result (and thus functionally complex according to our definition) has necessarily involved a long sequence of gradual increases in complexity, each of them preserved by selection. In the case of the eye, many of these intermediate stages have actually been found in living species [150, Fig. 13.2]. Of course, we should keep in mind that many creatures (perhaps most impressively ants) do very well with 'low fidelity' eyes.

The massively enlarged primate brain must of necessity be the result of a selection-driven trend; this is because primate brains are costly structures, especially in terms of energy. As Dunbar [48] explains, "because the cost of maintaining a large brain is so great, it is intrinsically unlikely that large brains will evolve just because they can." Dunbar provides a wealth of evidence which indicates that the relative size of the neocortex in primates (including humans) correlates remarkably with measures of social behaviour (such as the size of the social group), but not with ecological aspects such as diet, foraging behaviour or range size. This would suggest that social pressures, rather than ecological factors, were the proximal drive for larger neocortices in primates - a clear example of a coevolutionary-driven complexity increase. Because courtship and mating are important aspects of social interactions, there is clearly much overlap between this position and that of Miller [126], which strongly emphasises sexual selection as a driving factor in the enormous growth of the human brain.

The evolution of complex chemical cycles within cells may be seen as a counterpart (and predecessor) to the evolution of complex organs in multicellular organisms. Perhaps the most remarkable case is aerobic respiration, an intricate oxygen-based cycle which extracts energy from complex organic molecules. This highly efficient process, yielding up to 36 ATP molecules for each glucose molecule (against 2 for anaerobic fermentation), is now extremely widespread among living creatures: exclusively anaerobic creatures, once dominant, have been relegated to a secondary role (though by no means eliminated altogether) by creatures endowed with the more complex, but more powerful aerobic respiration apparatus[9]. Of course, anaerobic fermentation is still present as a component of the aerobic respiratory cycle, to the effect that aerobic cells may revert to anaerobic fermentation in the absence of oxygen (much to the satisfaction of certain higher mammals not entirely adverse to the effects of fermented beverages).

More generally, it is widely acknowledged that the general architecture of modern living systems (based on nucleic acids and proteins) is too complex to have arisen spontaneously, in its present form, from abiotic reactions. Therefore, intermediate forms (now extinct) must have preceded the appearance of modern life. The utter elimination of these early systems, which we may surmise to have been quite varied, can hardly be explained by pure contingency. In fact, it is reasonable to posit that these 'ur-organisms' were actually displaced by modern life, which simply proved more efficient in assimilating precious carbon atoms and free energy. From this viewpoint, the total domination of modern life can be seen as a selection-driven trend, through which a sequence of increasingly complex systems successively displaced their simpler, less efficient ancestors.[10] This process, however, has not endured into the evolution of modern life, as can be shown by the massive dominance of prokaryotic forms in the present biosphere.

8.6.3 Examples of particular selection-driven decreases in complexity

Having shown particular examples of driven trends towards complexity, we feel obliged to emphasise that trends often run in the other direction as well.

One obvious form of selection towards lower complexity is the emergence of parasitism in previously free-living species: an organism which was originally able to survive on its own, finds itself dependent on its host. This often leads to a reduction in complexity, as Gould [71, p.220]

---

[9] Even then, things are not quite clear-cut: the fermenters' higher rate of consumption may actually allow them to overcome respirators in certain situations of direct topical competition [141].

[10] This account is agnostic with regard to which processes led to the emergence of early life, as long as it is accepted that some now-extinct intermediate form(s) preceded modern organisms. However, if we follow Fernando's interesting hypothesis[51, 52] of the appearance of life as the culmination of a natural selection process among early pre-biotic chemical systems divided by external forces, then the very origin of life itself would be the result of such an active selection-driven trend towards increasing complexity.

predictably observes[11]. This phenomenon is well-known and does not require further emphasis.

Another common source of selectively driven decrease in complexity is miniaturisation. Bonner [19], pursuing his argument that size and complexity often go hand in hand, mentions the case of 'dwarf' rotifers which, because of new biomechanical constraints imposed by their nanism, replace their gut system with a continuous cytoplasm nested with small food-carrying vacuoles. The result is a decrease both in structure differentiation and in number of cell types. A well-known example of miniaturisation (and associated loss of complexity) is that of acari, particularly among the most minute mites. For example, dust mites have shed off much of their digestive structure (including their stomach), and as a result must ingest the same piece of food several times and use external digestion.

More generally, in parallel with the constant elaboration of complex organs, an equally constant process of complexity reduction occurs in Nature. We need only cite a few well-known examples. The most obvious tell-tales of evolutionary simplification are vestigial structures. The vermiform appendix in humans is the remnant of a structure which allows its bearer to digest cellulose in plants, something that humans cannot do: this is clearly a diminution in functional complexity, by any measure. The small pink pouch in the inner corner of our eyes is the non-functional remnant of the nictitating membrane (the 'third eyelid') found in many other animals, including cats and dogs. Modern whales retain small, undifferentiated bone structures as a memory of what used to be perfectly functional pelvis and hind legs[175]. This also affects the paradigmatic eye: in many creatures which have reverted to dark environments, the functional complexity of their visual systems has decayed into non-functionality, sometimes culminating into the complete elimination of the eye itself (the cavefish *Astyanax* is interesting in that various stages of eye loss can be observed in distinct modern populations [93]). Even the pinnacle of animal complexity, the vertebrate brain, has undergone significant morphological and functional simplification in certain lineages: let us mention salamanders[154] or certain bats[158][12].

Finally, we will mention a rather spectacular form of successful complexity decrease. This example, tentatively suggested for this purpose by McShea [118] (and recently confirmed by further studies [47]), concerns the appearance of infectious tumour cells. In at least two separate examples (one affecting dogs in a rather benign way, the other affecting Tasmanian devils with a usually fatal outcome), it has been found that a certain lineage of tumour cells could be transmitted from host to host through contact, with the result that the tumours effectively behaved as populations of single-cell parasites of the host species. In both cases, the tumour cells were shown to share a common origin (being genetically similar to each other across hosts, and different from their hosts), and thus can be considered a distinct species. For our purpose, the key point is that these single-cell species originated from mutated cells within multicellular, 'higher' organisms. In other words, in at least two separate occurrences, a *major transition* of evolution[165] (namely the invention of multicellularity) was successfully reversed. This dramatic

---

[11] However Gould also reminds us that in many case, parasites may also exhibit high complexity at some stages of their life: *Sacullina* is mentioned as an example both of extreme simplicity (in its almost amorphous host-bound form) and complexity (in its remarkably intricate development cycle).

[12] Brain evolution in vertebrates is a complicated business, despite an overall highly conserved structure. Many fishes have similar brain size to mammals and birds, relative to their body size. Many independent increases in brain size have been identified, but stability and decreases have also occurred[134].

drop in complexity from highly complex mammals to single-cell organisms created new, simpler, yet successful species, especially in the case of the dog tumour lineage which is thought to be at least several centuries old.

8.7 Conclusion: the passive trend of complexity growth in evolution

8.7.1 Passive growth of complexity

In summary, both theory and evidence suggest the same interpretation: there has been a trend of increase in the maximum of complexity in evolution, and this trend seems to have been resolutely passive - that is, globally unbiased. At all scales, individual lineages will be driven to whatever form of adaptation evolution may encounter, which may or may not lead to increases, or decreases, in complexity. There does not seem to be any overarching bias, neither in evidence nor in theory, which would drive a majority of lineages towards higher complexity. The Tree of Life does not 'seek' complexity as actual plants seek sunlight.

Again, it is important to note that this does *not* contradict the hypothesis of an "arrow of complexity"; it does, however, contradict certain interpretations, namely those which surmise a built-in tendency for Darwinian evolution to drive lineages towards increasing complexity. This does not imply a rejection of the existence of a trend: a passive trend is still a trend.

8.7.2 Non-triviality of passive trends

In addition, we must stress that a passive trend does not in any way imply that Darwinian evolution is neutral or irrelevant to complexity growth. There is a tendency to see passive trends as somehow trivial, or not requiring any particular explanation. That this is not the case is of fundamental importance. Complexity has indeed risen to extraordinary levels, and the ease with which Nature comes up with complex forms demands explanation.

The difficulty of obtaining a sustained arrow of complexity is amply demonstrated by the field of artificial evolution, in which thorough use of Darwinian evolution has indeed created appreciable levels of complexity [162, 107], but nothing that could even be compared to the enormous growth of complexity in Nature. In fact, the difficulty of initiating "open-ended" evolution in artificial experiments is widely noted [110]. Consequently, the very fact that evolution seems unbiased towards higher or lower complexity is in itself worthy of note, because in artificial settings the process seems very much biased *against* continuing complexity growth. The apparent absence of a bias for or against complexity in Nature only seems trivial if we forget how strong the bias against complexity seems to be in computers.

8.7.3 New questions

This casts a new light on our problem. If a passive trend has occurred in the evolution of complexity, the question then becomes: what makes such a trend *possible*? What conditions are necessary for the appearance and sustained increase of complexity (that is, for the emergence of an "arrow of complexity") in a system? What is the role of the Darwinian mechanism in this emergence? This is the second part of our reformulated question (as put forth in section 8.4), with which the remaining sections of this chapter will be concerned.

8.8 The role of Darwinian evolution in complexity growth

8.8.1 Systems, designs and complexity

In this section we want to evaluate the effects of Darwinian evolution on complexity, in the most abstract, general sense. Our perspective is unabashedly conceptual: we want to gain insight into the relationship between evolution and complexity, rather than devise technical methods to be used in actual experiments. To do this, we ask which constraints on complexity exist in abstract non-evolutionary generative systems, then we ask how the introduction of an evolutionary mechanism alters these constraints.

Let us consider an abstract generative system, of which we only know two things:

- A certain 'design generator' is being used to generate new designs; the particulars of the design generator define a certain space of possible designs which can be explored.
- New designs are being generated according to a certain heuristic; in other words, the space of available designs is being sampled in a certain sequence, as specified by the heuristic.

For now, we assume that the heuristic used to generate new designs is unknown: it could be random, evolutionary, hill-climbing, or anything else. We now ask the question: what is necessary for such a system to produce unbounded growth in functional complexity?

A moment's thought suggests that in principle, at least two (rather obvious) conditions are absolutely necessary. First, the system generator must in principle be able to generate arbitrarily complex designs. Second, the search heuristic must not impose a hard upper bound on adaptivity or functional complexity of the designs being sampled. It need not favour complexity, it may even discriminate against it, but not to the point of forbidding it entirely above a certain level. Provided these two conditions are met, over time, the system should mechanically produce arbitrarily complex designs, regardless of which heuristic is being used.

8.8.2 Limitations of non-evolutionary complexity growth

Inefficiency

However another moment's thought will make it clear that, without any further device, such a system might well take impossibly long times to come up with any such complexity. This is because, as we have pointed out in section 8.2.2, functionally complex designs are *rare* in the space of all possible designs. A random design generator may come up with arbitrarily complex designs, given enough time, but what good would that be if "enough time" exceeds the thermodynamic limit of universal heat death?

Finite populations and finite lifetimes

Another difficulty occurs when designs are eventually removed from the system. For example, designs might have a specified maximum lifetime (which would mean that all designs are eventually eliminated), or the population could have a maximum size (which implies that some designs must be eliminated).

In this case, while the system will still come up with designs of unbounded functional complexity over time, we have no assurance that the maximum of complexity at any time will *sustainedly* increase. This is because, without any further assumption, we have no way to know that the rate at which complex designs are being generated balances the rate at which they will be removed: in fact, the opposite is much more likely. While arbitrarily complex designs will eventually occur, these designs may very well be mere ‘blips’ in an otherwise consistently low background: if the current maximum of complexity is removed, we have no assurance that it will be replaced by designs of similar complexity. Again, this is because functionally complex designs are rare within the space of all possible designs. The consequence is that, barring further assumptions, the population may well spend most of its time at a very low level of functional complexity.

Consequence: no arrow of complexity

We must therefore conclude that, while the two “obvious” requirements expressed above are sufficient to generate arbitrary complexity over (much) time, they are not sufficient to create what we would call an “arrow of complexity.” This would imply a *sustained*, *noticeable* increase in the maximum complexity, rather than the highly occasional appearance of fleeting, isolated complexity spikes.

8.8.3 Impact of evolution on complexity

The Darwinian heuristic

Now let us introduce a new requirement: we posit that our system is evolutionary, that is, it follows the “Darwinian heuristic” as elegantly formulated by Wagner and Altenberg [181]: “the Darwinian heuristic is to choose sample points by perturbing the more fit ones among those sampled thus far.” In other words, new designs are not generated arbitrarily: rather, existing designs are evaluated after a certain metric (their ‘fitness’), and those designs which exhibit a relatively high fitness are selected as starting points for small ‘jumps’ towards newer designs. Our question is: how does this new specification affect the emergence of functional complexity?

The Darwinian mechanism restricts exploration

The first thing to note is that the Darwinian heuristic actually imposes a *restriction* on the search process: the search process will now follow certain directions at the expense of others. This restriction constrains both the scope and the direction of the search process. The restriction in scope comes from the fact that, because evolution proceeds by small jumps from existing designs, newer designs can only be chosen in the neighbourhood of the current population. The restriction in direction comes from the fact that only those designs which exhibit a high fitness will be used as starting points for further exploration.

In unfavourable conditions, this restriction may severely limit the exploration of design space. For example, if the population happens to be in a certain domain of the search space, which is entirely surrounded by relatively unsuccessful designs, then the process may well end up being 'trapped' into that domain. Conversely, potentially highly fit designs may never be reach if they are surrounded by a region of unfit designs. This suggests at least two conditions for complexity to appear in evolutionary systems, namely, that it should be *successful* and *reachable*[13].

Benefits of the restrictions

On the other hand, this restriction has important potential benefits. First of all, when successful designs are found, they will tend to be *preserved*, even beyond the lifetime of individuals - or at the very least, they will stand an immensely higher chance of being preserved than under random search, as long as they remain successful within the context of the current population. This permits the persistence of functional complexity, despite its rarity in design space, even in the face of finite lifetimes: complexity is no longer condemned to occur only in short 'blips' over an otherwise simple background.

Furthermore, as we have already noted in section 8.2.2, adaptive functional complexity tends to come in 'loose clusters' in design space: if we want to obtain a successful complex design, we stand a much higher chance to obtain it by slightly altering another successful design of lower complexity, rather than by randomly poking at the entire design space. The consequence is that, by preserving existing (successful) complexity, the Darwinian heuristic also facilitates further increases in complexity, provided such increases are possible within the neighbourhood of the population: a successful complex design may act as a *'stepping stone'* towards even higher complexity. This is, in essence, the basis of the "Blind Watchmaker" mechanism, so vividly expounded by Dawkins [39]: by slowly accumulating small beneficial improvements, each working as a stepping stone towards the other, the Darwinian heuristic may eventually produce enormous amounts of adaptive functional complexity, with incomparably higher probability than under random search.

Thus we see that certain requirements must hold if evolution is to generate designs of unbounded complexity: it is necessary that, at all times, successful designs of higher complexity exist within the current neighbourhood of the population. At first sight, these conditions might seem rather straightforward. However, are they really sufficient?

Enter the Red Queen: variable fitness and complexity

So far we have not considered a certain aspect of fitness, namely the fact that it is not necessarily stable over time. A design which is successful at a certain time may later prove unsuccessful if conditions have changed. This is particularly the case in situations of coevolution, that is, when the fitness of a given design depends at least in part on the features of other existing designs within the population. In this case, species evolve constantly in order to adapt to each other's evolution - and those which fail to do so, disappear. This is the well-known Red Queen effect [177, 32].

For our purposes, the basic consequence is that even if all previously mentioned conditions are

---

[13] In the context of this discussion, a design is 'successful' if its fitness is high enough, relatively to the current population, to be used as a starting point for further samples. A point in design space is 'reachable' if certain successful designs within the current population might be altered into the corresponding design.

met, an evolutionary process may still fail to produce a sustained increase in complexity. Even if, *at any time*, there exists potentially successful complex designs in the reachable neighbourhood of the current population, complexity will *not* necessarily grow over time. This may seem counter-intuitive: at first sight, if at any time successful complex designs exist within the reachable neighbourhood, then we would expect that eventually evolution would find them, and thus complexity would grow.

The reason why fitness instability may prevent such continuous growth is that it may disrupt the 'stepping stones' effect on which this growth relies. If fitness is variable, then designs of intermediate complexity may become unsuccessful before they have a chance to be altered into more complex designs, putting the latter beyond reach of the process. Thus, despite the constant presence of reachable, successful designs of higher complexity, no continuous process of complexification may occur. If moderately complex designs are not sufficiently robust to serve as stepping stones towards higher complexity, the Blind Watchmaker fumbles, and the arrow of complexity is broken. The effects of this phenomenon in artificial coevolution have been mentioned in section 8.5.2.

Thus we see that for complexity to appear in evolutionary system, it needs not only be successful and reachable: it also needs to be *robust*, even in the face of possible instabilities (including new competitors, parasites, etc.), at least sufficiently so to provide a stepping stone for further complexity growth. This is an additional requirement which, if not observed, may actually prevent the rise of complexity in evolutionary systems.

Putting it all together

How can we express these requirements as concisely as possible? We suggest the following formulation:

- At any time, there are more possible jumps towards successful designs of higher or equal complexity, than towards designs which would make the starting points of these jumps unsuccessful.

Here "higher or equal complexity" is to be understood in comparison to the current maximum of the population. A jump is the creation of a new design, using an already existing design (the 'starting point' of the jump) as a parent. Clearly, by definition, the starting point of a jump must be successful.

This condition is clearly statistical: even if such jumps exist and constitute the majority of possible jumps, we have no guarantee that they will be taken in the next step, so the corresponding complexity increase may not occur immediately. Conversely, even if jumps towards higher complexity are rare, we may still occasionally encounter one of them. What this condition determines, however, is the presence or absence of a *trend* - that is, a long-term tendency. If this condition does not hold over time, then statistically we do not expect any long-term trend towards higher complexity to emerge; similarly, if it does hold consistently, then a long-term trend towards higher complexity is statistically expected.

More importantly, this condition is very much local (both in time and in design space). If the condition holds at a certain time, it does not mean in any way that it will keep doing so in the future: any new jump runs the risk of breaking the condition. For example one particularly unlucky jump might discover a niche of devastatingly efficient simplistic individuals (in the extreme, we might conjure images of ‘grey goo’ bacteria). Should this happen, then the probability of jumps might become highly biased towards these simplistic individuals, and thus the condition would not hold any more. Conversely, when the condition does not hold, one particularly lucky jump may stumble upon a new area of robust complex adaptations, making the condition valid. Therefore, maximum complexity will tend to increase as long as the condition is consistently met - but only that long. Complexity will not tend to increase (or may even tend to decrease) if the condition consistently fails to hold.

8.8.4 Evolution and complexity: a double-edged sword

In summary, we suggest that the conditions under which Darwinian evolution will generate a sustained trend of increase in the maximum of complexity (that is, an “arrow of complexity”) can be summarised as concisely as possible in the following manner:

1. Arbitrarily complex adaptive designs must exist in the design space explorable by the design generator.
2. The search heuristic must not impose a hard upper bound on the functional complexity of designs.
3. At any time, there are more possible jumps towards successful designs of higher or equal complexity, than towards designs which would make the starting points of these jumps unsuccessful.

It is in this way, and under these conditions, that Darwinian evolution can really be said to create an “arrow of complexity”. It does not mean that evolution will consistently ‘seek’ complexity, or that evolution is intrinsically a complexifying force - rather, it means that, *if* some advantageous complexity exists, and as long as it is successful, reachable and sufficiently robust (in the sense made explicit by the third requirement), then Darwinian evolution will indeed produce an arrow of complexity which could not emerge without it. But as soon as the requirements cease to hold, then evolution may actually *prevent* the growth of complexity altogether.

Thus we see that there is a deep ambiguity in the relationship between evolution and complexity. This ambiguity is captured in the third requirement above, which as we saw is very much specific to Darwinian evolution. Random search, for example, may create arbitrary amounts of complexity over (much) time, even if this requirement does not hold, while Darwinian evolution cannot. But if (and as long as) this requirement holds, then Darwinian evolution can actually create an “arrow of complexity” (a sustained, noticeable trend of increase in maximum complexity), which is utterly beyond reach of random search. In this sense, with regard to complexity, Darwinian evolution is very much a double-edged sword.

8.9 Conclusion

Clearly, at this point, we have gone a long way to clarify the statement that "open-ended evolutionary processes have an inherent, law-like tendency to create creatures with increasingly complicated functional organization". We have seen that the hypothesis of an "arrow of complexity" can be interpreted in several ways, with varying support for each interpretation. We have seen that the undeniable increase of maximum complexity in evolution is best explained as a passive (unbiased) trend, that is, a side consequence of an otherwise complexity-neutral process of reckless local adaptation.

We have seen that Darwinian evolution is capable of creating an arrow of complexity (that is, a sustained, noticeable trend of increase in maximum complexity) which could not exist without it, provided certain requirements hold. These requirements are that at any time, higher complexity should be reachable, successful and reasonably robust. They result from the restrictions imposed by the Darwinian heuristic on the exploration of search space. As long as these evolution-specific requirements are consistently met, then evolution will keep producing designs of increasing adaptive complexity in a sustained manner. But if they cease to be, then evolution will not only fail to produce an arrow of complexity - it may actually prevent complexity increases altogether. In other words, with regard to complexity growth, evolution is a double-edged sword, which may slash both ways.

Thus the appearance and growth of complexity in evolution results neither from a built-in, relentless march to "progress", nor from the miraculous conjunction of a series of vastly improbable coincidences. Rather, it is simply ascribable to the presence of certain objective conditions, applying to the capacities of the design generator and the structure of design space. Evidence from Nature and the fossil record indicates that the design space faced by eukaryotic life has met the appropriate conditions in a roughly constant manner - give or take an asteroid or two. The fundamentally explorative nature of evolution has been able to find successful, reachable and robust forms of complexity, repeatedly and sustainedly, over the entire history of modern life. No grand scheme of cosmological destiny, no poignant tale of triumph against all odds; just the pedestrian, plainly logical, yet amazingly powerful consequence of the properties of modern life in the face of the challenges and opportunities posed by the laws of physics, the geological environment - and itself.

Chapter 9

Fitness Transmission: detecting Darwinian evolution with genealogic records

9.1 Introduction

9.1.1 The problem: Are we evolving yet?

Consider the following problem: imagine that you are observing a simulation, in which a population of agents move, interact and reproduce. The simulation is complex, or its output is obscure (or both), and it is not easy to grasp what, if anything, is going on. Knowing that these agents reproduce, we may ask ourself the question: are they also *evolving*? Are they undergoing genuine natural selection and adaptive evolution? Or are they just perpetuating random genetic traits, following a chaotic trajectory through genotype space without ever undergoing any meaningful evolution?

This question arises from the fact that when a population of reproducing agents is observed, it is not always immediately clear whether the dynamics of the population result from Darwinian evolution, or merely from random variations and stochastic effects such as genetic drift. The particular system at hand may also introduce its own effects, which may bias or alter the dynamics of the population in unpredictable ways. When this system is sufficiently complex, determining whether a population is evolving in a Darwinian sense may not be a trivial task.

Besides its conceptual implications, the question is of practical interest. It is often desirable to determine whether natural selection and evolutionary adaptation are occurring within a given system, especially in the fields of evolutionary computation and artificial life. Indeed in some situations, the onset of significant adaptive evolutionary activity is by itself a major objective of the system: for example, artificial environments such as Echo [84] and Geb [28] (both described in Appendix B) were explicitly designed with the aim of exhibiting meaningful evolutionary

---

This chapter benefited greatly from the comments of several anonymous reviewers for the *Artificial Life* journal.

activity. Being able to detect the presence of genuinely adaptive evolution is a fundamental pre-requisite for the validation of such systems.

9.1.2 What is evolution?

As we saw in the first chapter, evolution is simply a change in the inherited characteristics of a population, from generation to generation. The word “evolution” only refers to change, in the general sense, and does not carry any implication of value or improvement. In particular, evolution can be random, as seen in genetic drift. However, the mechanisms which cause or drive evolution can certainly steer evolution towards increased adaptation. Of course, this is particularly true for natural selection, the mechanism originally purported by Darwin as a cause for the remarkable adaptation of living beings.

Note that this Darwinian mechanism, which explains adaptive evolutionary change, may also lead to local, temporary stability, for example if the population is currently located around a convenient local optimum: hence the concepts of ‘stabilising selection’ and ‘evolutionary stasis’ (see chapter 4 of [150]).

Here we are interested not just in evolution, but specifically in evolution by natural selection. Our aim is to provide a measure to detect the active presence of the Darwinian mechanism of evolution by natural selection. We want to detect the presence (or absence) of *Darwinian forces* which actively shape a population, as opposed to the mere accumulation of incidental, non-adaptive fluctuations - to catch the ‘invisible hand’ of natural selection in the act, as it were. We are specifically *not* interested in measuring ‘progress’ or complexification, two notions sometimes associated with the idea of “measuring evolution” [24]. Neither are we trying to determine how innovative or successful evolution is in terms of optimisation. Our objective is simply to find out whether the trajectory of a given population of reproducing agents through genotype space is the result of Darwinian evolution by natural selection - no more, no less.

One consequence is that our measure will give high marks for systems that are being actively molded by natural selection, even if this molding consists of actively maintaining the population around a given optimum. In other words it will give high marks to systems undergoing stabilising selection (except in extreme circumstance where stabilising selection is so strong that no mutation can survive at all, as discussed in section 9.6.) Importantly, our measure will differentiate such a system from ‘frozen’ systems in which no genetic variation occurs at all, which will be given a zero score. This is precisely what is required.

9.1.3 Fitness Transmission: A signature of Darwinian evolution

We posit that the defining characteristic of Darwinian evolution is the differential propagation of heritable, fitness-affecting characteristics, that steer the population in locally adaptive directions. From this rather uncontroversial basis, we deduce that it should be possible to detect the presence or absence of Darwinian evolution from the observed genealogical dynamics of evolutionary processes. Roughly, our reasoning is that if any transmission of heritable fitness-affecting traits is occurring, then on the whole, individuals sharing a common lineage, which tend to inherit similar traits, should also tend to have some degree of similarity in their fitnesses - that is, their reproductive success. For example, we might expect that fitter (resp. less fit) parents should tend to have fitter (resp. less fit) children. Thus if heritable, fitness-affecting characteristics are being transmitted, a tenuous, but persistent *correlation* should exist between the fitness of individuals sharing a common lineage: the differential transmission of fitness-affecting characteristics should result in a *differential transmission of fitness.* This differential transmission of fitness is the genealogical signature which we believe characterises the presence of active, Darwinian adaptation through natural selection based on heritable characteristics.

9.2 Related Work

9.2.1 Traditional methods for detecting natural selection

The problem of detecting natural selection has a long history in biology. Endler's authoritative treatment [49] describes the traditional (that is, non-molecular) methods for detecting natural selection. However, all these methods are based on *phenotypic observation* of chosen traits: they require collecting statistics on the frequencies of certain, pre-defined traits, and then performing some calculations to determine whether or not natural selection has acted on these traits. This is precisely what we seek to avoid here: we do not ask whether natural selection has acted on this or that trait, but simply whether it is active in the population. Also we want to dispense with detailed phenotypic observation.

The molecular revolution in biology has made it possible to collect vast amounts of genetic data. This creates new possibilities for the detection of natural selection, based on direct assessment of nucleotide variation (Sabeti and colleagues [157] provide a recent review.) But these approaches require access to a full genetic record. Furthermore, biological genomes are simple sequences of items in a four-letters alphabet; but artificial life models need not be so simple, and this may affect the applicability of these methods. By focusing on the genealogic record alone, we avoid any such difficulty.

9.2.2 The Bedau-Packard measure of evolutionary activity

Bedau and Packard [10, 11] have developed a groundbreaking set of concepts and methods to "discern whether or not evolution is taking place in an observed system". It should be noted that Bedau and Packard's aim in that work is somewhat different from ours here. Here we are interested in detecting the *presence* of natural selection, that is, in finding out whether the trajectory of a population in configuration space is being actively shaped by Darwinian evolution and natural selection. Bedau and Packard, on the other hand, are specifically interested in the *innovations* produced by evolution, and in the capacities of various systems to keep on producing adaptive innovations over time - or not. This requires a method to determine whether an apparent innovation is indeed adaptive or merely the result of random fluctuations, which clearly relates to our own concerns. To this end, Bedau and Packard introduce a set of methods to compute the "*evolutionary activity*" of components and, by extension, of systems. These methods form the basis of a classification of systems into various classes, each class being associated with a certain type of evolutionary dynamics.

The Bedau-Packard measures of evolutionary activity are based on *persistence* of adaptive innovations: they identify components that persist over time at a level that exceeds what would be expected under purely random conditions. In a short introduction to the framework [23], Bullock and Bedau provide the following rationale:

> A paradigm of the idea of evolutionary activity is the notion of measuring the continued persistence of elements in an adaptive system as a proxy for their adaptive significance – if a particular element persists in the system for a long time, this is likely to be because it is being maintained by selection.

In this short passage, we already perceive the basic reason why Bedau and Packard's statistics are not readily applicable to our purposes of detecting the presence of natural selection: by asserting that persistence "is being maintained by selection", we are already assuming the presence of Darwinian selection in the first place. It is not unthinkable that persistence could occur through other processes than natural selection and individual adaptation - and as we will see, it is easy to come up with (non-evolutionary) systems in which this occurs.

If we are to use persistence "for a long time" as a criterion for detecting evolution, we need a method to determine what "a long time" is. When do we decide that a given element has persisted long enough to be regarded as 'adapted'? It is clearly impractical to provide a general answer to this question *a priori*, which could apply to all systems and in all situations. To tackle this problem, Bedau and Packard introduced the idea of using a neutral "shadow" of the system under study: a replication of the original system, in which birth, reproduction and death of individuals occur in synchronisation with the real system, but are applied to randomly chosen individuals. More precisely, every time a new individual is being created in the real system under study, a new individual is also created in the shadow; but with the difference that, in the shadow, the parents of the new individual are chosen randomly. Thus the neutral shadow is expected to show the behaviour that would be seen in the system, in the absence of any selective pressure. By comparing the persistence data obtained in this "shadow" to that obtained in the real system, Bedau and Packard argue, it should be possible to detect whether selection and adaptive evolution are present.

Building upon the concept of enduring persistence as a measure of evolutionary activity, Packard and Bedau have developed a series of evolutionary statistics based on persistence information. These statistics include:

- Diversity D: The number of different components present at a given time in the population.
- Activity $a_i(t)$: The age of component i at time t, indicating how long it has persisted so far.
- Cumulated activity $A_{\text{cum}}(t)$: The sum of the ages of all individuals present at time t.
- New activity $A_{new}(t)$: The sum of the ages[1] of all 'new' (Bedau and Packard use the term '*adolescent*') components present in the system at time t, divided by diversity at

---

[1] In the original presentation [11], there is a discrepancy between the discrete formula (sum of all ages) and the continuous formula (integration of the distribution function, which would correspond to number of counts within that range rather than the sum of their values.) Perhaps interestingly, this dissonance is repeated in the literature, where some authors choose to sum activities [173] while others choose to count the number of individuals that reach the relevant activity [169].

time t. Entities are 'new' if their age a_i falls within a certain interval $[a_0, a_1]$. The idea is that a_0 is the minimum age at which a new component can be regarded as adaptively significant: if a component reaches age a_0, then we assume that it has persisted more than could be accounted for by mere random fluctuations, and that this persistence is a sign of its adaptive value. A suitable value for a_0 is to be determined empirically through comparison with a neutral 'shadow' system. $A_{new}(t)$ is used as a measure of significant *innovation* taking place in the system at a given time.

Bedau and Packard have used these statistics as a basis for classification of evolutionary dynamics into three classes: Class I ("no adaptive evolutionary activity"), Class II ("bounded adaptive evolutionary activity) and Class III ("unbounded adaptive evolutionary activity"). Class III, in particular, is expected to encompass truly 'lifelike' evolution, similar in creativity and endurance to that found in nature. It requires that the total evolutionary activity A_{cum} (the sum of the ages of components present at time t in the population) be unbounded, that is, grow at least linearly with t, and that new activity $A_{new}(t)$ be positive.

Bedau and Packard's measures are arguably the most widely known of their kind. They applied these measures to several systems, including artificial ecologies such as Echo, and (notably) the fossil record of the biosphere [11]. Other researchers have successfully applied them to various systems. [169, 27, 173, 136, 100]

9.2.3 The Bedau-Packard statistics and non-Darwinian systems

The Bedau-Packard method is a powerful tool for observing the activity and creativity of evolutionary systems. However, it is not suitable as a method to assess whether a system is specifically undergoing Darwinian evolution by natural selection, as opposed to other, non-evolutionary processes.

The basic reason why the Bedau-Packard statistics cannot be used as a detector of evolution by natural selection is that they may attribute a positive score to processes which are clearly not "evolutionary" in a Darwinian sense, because no heritable variance in fitness (and thus no evolutionary adaptation) occurs. Importantly, this is the case even if a shadow is used to normalise activity scores. The crux of the matter is that these statistics essentially track "excess" variance in the persistence of components, which is used as a proxy for selection and therefore (it is argued) for adaptive value. The shadow is used to define the level of persistence which can be termed "excess." But what if some high variance in persistence occurs that is not heritable? If we apply the Bedau-Packard statistics to such a system, we may find that the Bedau-Packard measure classifies such a system as adaptive, even though it is not - even if we use a shadow.

It is easy to devise examples of systems that illustrate this distinction. For instance, consider a population of individuals in which reproduction and removal occur completely at random, except for the following caveat: certain randomly chosen individuals (the set of which changes over time by losing and gaining random members) are allowed to survive longer and reproduce more than the others. Clearly this process is not "evolutionary" in a Darwinian sense, let alone adaptive. Crucially, we should not expect the children of these 'lucky' individual to show any superior reproductive success: no heritable variance in fitness would be observed. Furthermore, whichever notion of 'adaptation' we use, no novel adaptive ability would be introduced into the system. However, in this system, the activity $a(t)$ of the genetic components associated with the favoured individuals would actually be positive, even when normalised against a shadow. This is because the (randomly chosen) favoured individuals would survive much longer than others, giving rise to strong persistence scores; in addition, because they reproduce more, their genetic material would tend to dominate the population: even though their offspring would not exhibit remarkable fitness, their sheer number would alter the distribution of genetic material.

Another example is a system in which evolution occurs normally, except for the fact that fitness is randomly attributed to each individual at birth, independently of its genome. That is, while genetic material is transmitted as expected from parent to offspring, this genetic material has no influence over fitness, which is chosen randomly and arbitrarily for each new individual. Again, note that no heritable variance in fitness occurs, nor does any adaptive innovation take place. In such a situation, those individuals that happen to be highly fit (out of sheer luck) will tend to persist for a long time, and may flood the population with their (short-lived, but nevertheless genetically similar) offspring. No such thing will be observed in the shadow, where reproduction and survival will be random, leading to random diffusion of the genetic material throughout genotype space. Therefore, a difference will occur between the activity counts (and diversity counts) of the shadow and of the real system, creating a positive signal on the Bedau-Packard measure and associated tests.

In figure 9.1 we describe the results of Bedau-Packard statistics applied precisely to such systems.[2] In both cases, the system is a simple steady-state genetic algorithm in which, at every "generation", 10 out of the 100 individuals are eliminated and replaced by new individuals, created by copying and mutating a surviving parent.

In the first system, selection of survivors and of parents occur randomly, except that at any time one half of the population is "favoured": individuals in this set are 5 times "more likely" than others to be selected for survival and for parenthood. More precisely, an individual is randomly picked, and immediately selected if it belongs to the favoured set; if not, a random number between 0 and 4 is chosen, and if this number is 0 the individual is selected, otherwise the process is repeated until an individual is selected. Any given individual can only be remain "favoured" for a finite number of generation, randomly and uniformly chosen between 1 and 1000 (note that the graph of cumulative frequency distributions, or "waves", demonstrates that some genomes routinely persist much longer than this limit.)

In the second system, selection of survivors occurs by fitness ranking, and selection of parents occur through tournament selection, very much as in a normal genetic algorithm. However, the fitness of individuals is randomly chosen at birth, independently of their genome. The actual method to "calculate" fitness is to increase a certain counter repeatedly until a random number picked between 0 and 9 is equal to 0 (thus the distribution of fitnesses is exponential.) In the shadow systems for both experiments, selection of survivors and parents are random (thus

---

[2] In these experiments we have applied the Bedau-Packard statistics to entire genotypes, in order to follow the authors' method [11]. However we are not at all certain that whole genotype persistence is a reliable indicator of evolution. In particular, Bedau and Packard [11, pp. 7-8] seem to state that, when parents consistently differ from their children (due to mutation), no adaptive evolution occurs (as shown by the absence of persistence signal.) We note that in nature, it is very unlikely that any genotype ever persists for more than one generation. However this should not affect our discussion: this problem may prevent the detection of adaptive evolution where it exists, whereas we want to show that the Bedau-Packard statistics can actually return a positive signal even when natural selection is clearly absent.

the shadow systems for both experiments are essentially identical, which predictably results in similar graphs).

Note, in particular, the onset of high average activity, the flattening of the cumulative distribution of persistence counts (with an order of magnitude difference between the longest-living genotypes of real and shadow systems), and perhaps most significantly the appearance of large "telic waves" [10] (tall, lengthy lines) in genome frequency plots, despite the decidedly non-teleological nature of these environments.

Surely many other examples could be found. More generally, these simple systems illustrate the fact that high variance in persistence can be caused by many other processes than natural selection. "Random" systems, in which no meaningful evolution or adaptation occurs, can still obtain high marks on the Bedau-Packard measure if they produce high variance in genetic persistence.

Of course, in these systems, it is easy to see (just by looking at their rules) that variance in persistence is due to random fluctuations, and that no true natural selection exists. But this is precisely the heart of the matter. First, when we study a real system, we may not have access to its internal rules, so clearly in this case we cannot use the Bedau-Packard statistics as a test of Darwinian evolution. But even if we *do* have full access to the rules of the system, the complexity of even mildly elaborate systems may prevent us from asserting with absolute certainty whether or not a "random force" generates strong variance in persistence. For example, considering a system similar to Echo, can we really exclude, a priori, that such a factor could come into play? Can we offer absolute guarantee, simply by looking at the rules of the system, that no weird effect will arbitrarily and significantly favour certain individuals rather than others (without being based on these individuals' heritable features)? The answer, of course, is that we cannot. It follows that, if we apply the Bedau-Packard statistics on such a system and obtain a positive result, we cannot (in the absence of further information) use this fact alone to conclude that adaptive evolution is active in the system.

It is important to be clear about the meaning of these results: this should not be interpreted as a minimisation of the importance of Bedau-Packard statistics. Rather, this is a reminder that these statistics should not be used to detect adaptive, Darwinian evolution, even by normalising against a shadow. *If* we know, a priori and through other means, that the entity under study is indeed affected by natural selection that can causally drive genuine adaptation, and if we can rest assured that "weird" effects will be nil or negligible, *then* we can fruitfully apply the Bedau-Packard measure to assess the dynamics of long-term evolutionary innovation within this system. The valuable contribution of these statistics in this regard has often been pointed out. However we cannot use these statistics to determine the presence of evolution by natural selection within a system, as opposed to any system-induced dynamics which create high variance in persistence: the Bedau-Packard statistics are not designed to distinguish the former from the latter, even by using a shadow system.

In a sense, the Bedau-Packard statistics illustrate the distinction between "evolution" in general, and the specific Darwinian process of evolution by natural selection. By focusing on varying persistence, the Bedau-Packard statistics track changes in the genetic makeup of the population, that is, "evolution" in the most general sense. The shadow filters out evolution

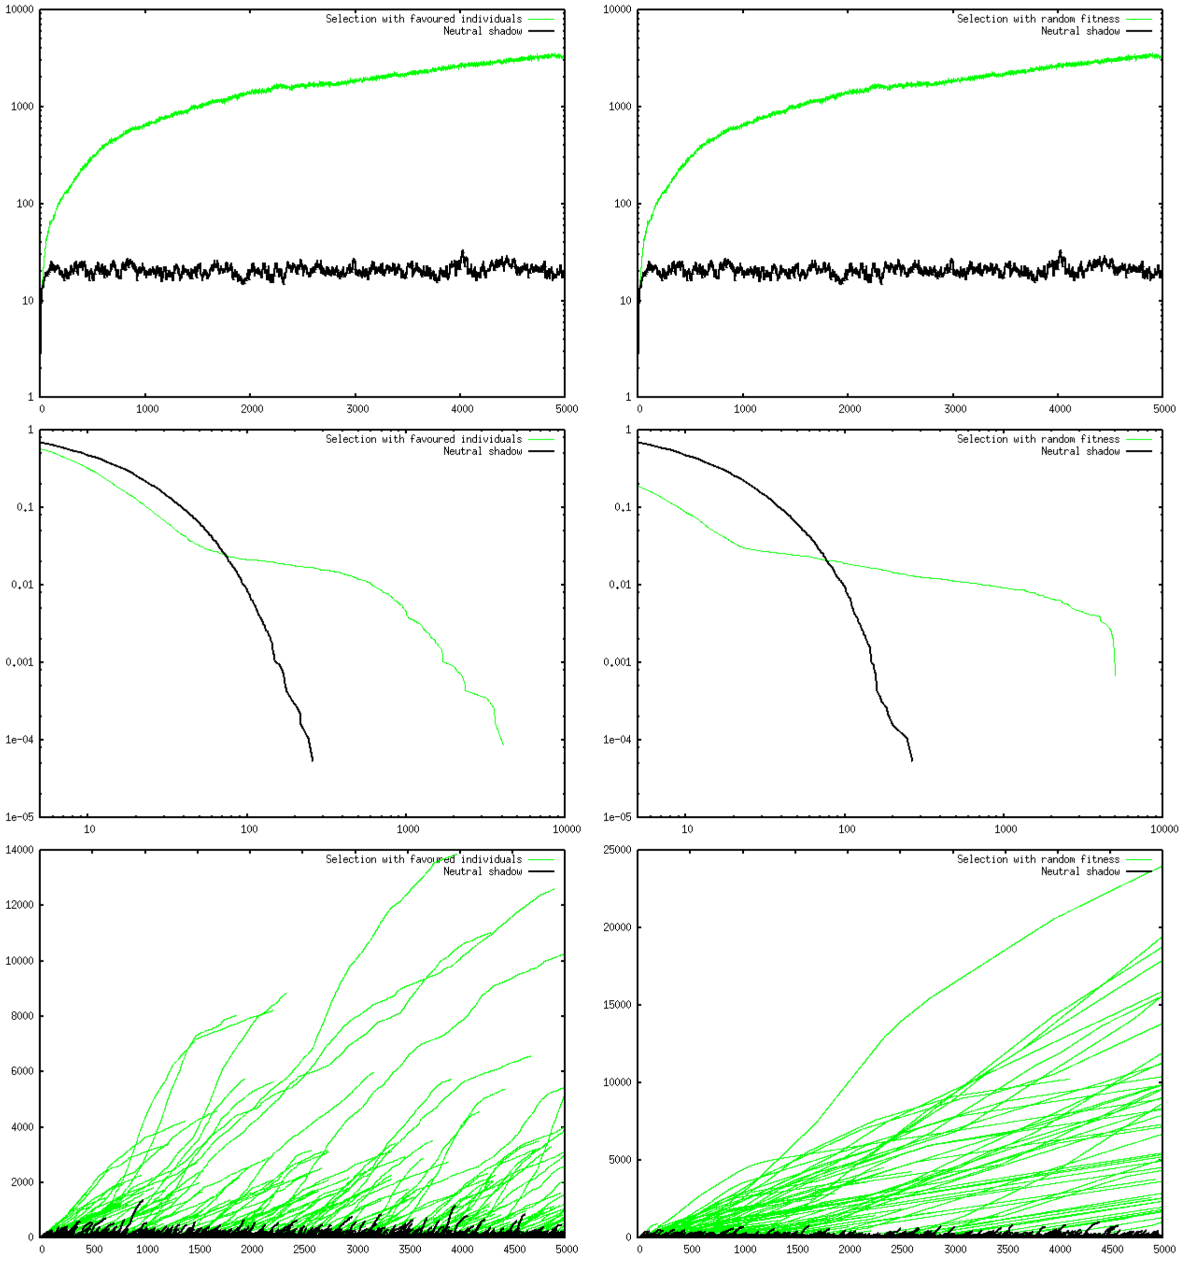

Figure 9.1: Graphs showing the results of Bedau statistics for two non-Darwinian systems, as well as for corresponding shadow systems. For each experiment, the topmost graph indicates average cumulative activity $\bar{A}_{\mathrm{cum}}(t) = \frac{A_{\mathrm{cum}}(t)}{D(t)}$ - the sum of all persistence counts of genomes present at a time, divided by the number of genomes present at that time. The middle graph indicates the cumulative distribution of persistence counts for all genomes over the history of a run (that is, for each value, the number of individuals that survived longer than this value.) The lowest graph shows the cumulative frequency counts for each genome over time (that is, the running sum of the frequency of each genome within the population at each generation.) These graphs are consistent with what is expected from real evolutionary systems [10, 11].

caused by random drift, but cannot determine whether the remaining change is actually the result of natural selection, rather than some other, unknown effect. The measure proposed here, by contrast, takes pretty much the opposite approach: here we focus on the question of whether or not a system is undergoing active Darwinian evolution, rather than the question of actual genetic change or innovation.

The following section provides a description and justification of our measure, based on an analysis of the Darwinian mechanism of evolution by natural selection.

9.3 Fitness Transmission: A test statistic for natural selection

9.3.1 Darwinian evolution: randomness, selection and heredity

As we saw in chapter 2, evolution in its broadest sense is simply a change in the heritable characteristics of a population over generations. Natural selection is variance in reproductive success caused by heritable characteristics. The Darwinian mechanism of evolution by natural selection occurs when natural selection is sufficiently strong to produce noticeable evolution.

Darwin saw that adaptive evolution automatically results from the existence of fitness-impacting, heritable variations. Variations that improve fitness will be propagated quickly, initiating thriving lineages; while those that reduce fitness will hinder their own propagation, creating feeble (or even quickly extinct) lineages. Thus lineages constantly branch out into variants, and the uneven distribution of these branches, being dramatically skewed towards those which result from fitness-enhancing variations, will result in the overall effect that the newer descendants of the original lineage will tend to be those better adapted to their current, local environment: heritable fitness-affecting variation will have "steered" the original lineage towards adaptive directions among all those encountered by mutational variations.[3]

Note that although this process will usually result in a modification of the species over sufficiently long periods of time, it will also often result in temporary stasis. If a species happens to be located at a convenient local optimum in the fitness landscape, then many variations which depart from the optimum will tend to reduce the fitness of the individual. In this case the differential transmission of characteristics enforced by natural selection will actively maintain the population around the optimum: the population will be constantly steered back towards its current position. This also occurs when being close to the population average is by itself advantageous. This phenomenon, known as 'stabilising selection', is actually regarded as more common than directional selection (see [150], Chap. 4.4). Lungfish and horseshoe crabs, which have remained very similar (though not exactly identical) over the last 200 million years, are well-known examples of evolutionary stasis. Indeed, distinguishing the continuing action of Darwinian forces when the population remains globally stable (as opposed to mere absence of variation or reproduction, which is a passive, non-evolutionary phenomenon) is one of the challenges faced by any detector of natural selection.

---

[3] Or, in short: as creatures replicate, traits propagate, adaptations proliferate, and species originate.

9.3.2 Fitness Transmission: A genealogic signature of Darwinian forces in a reproductive system

From this discussion we can deduce a method to detect the active presence of natural selection. If fitness-impacting, heritable traits are actually being transmitted and propagated, then this should have an impact on the genealogical record: individuals sharing a common lineage, being more likely to inherit common fitness-impacting characteristics, should therefore tend to exhibit slightly similar fitnesses in comparison to the rest of the population. In other words, if some fitness-affecting traits are being transmitted, then there should be some degree of *correlation* between the fitnesses (that is, the reproductive success) of individuals from a common lineage: the transmission of heritable, fitness-affecting traits should result in some degree of *differential transmission of fitness*.

Fitness transmission is our proposed signature for natural selection. It is, quite simply, the statistical correlation between the fitness of children and parents. The basic idea of fitness transmission is that, when natural selection is active in a population, parents and children should exhibit a tenuous, but persistent correlation in fitness.

This correlation would certainly be quite weak. Partly this is because the impact of any single trait on a particular individual's fitness is usually small. Reproductive success is determined by very many factors (chance being often an important one). In addition to this, children will usually not inherit all of their parent's traits, so the fitness-impacting traits may not be transmitted to all children. The result is that in general, the fitness of a particular child cannot be precisely estimated from the fitness of its parents alone. Heritable fitness-impacting traits do not predictably alter evolution at the level of a particular individual; but when many individuals are considered, then over a large enough sample a tenuous correlation should emerge between the reproductive successes of individuals from a common lineage, very much in the same way that thermodynamical properties such as pressure and temperature, while not discernible (or indeed definable) at the level of single molecules (which only exhibit widely varying kinetic energies), emerge as consistent and important properties when a sufficient number of molecules are considered. More than its actual value, which should be quite low, the *enduring presence* of this correlation over time should be the real indicator of natural selection.

9.4 Calculation of Fitness Transmission

9.4.1 Number of grandchildren as a measure of fitness

The term "fitness" is notoriously ambiguous and can be a significant source of confusion [35, Chap. 10] (see also chapter 2, section 2.3.2 in the present thesis). In general, it is supposed to denote the ability of an individual to pass on its genes. This is usually equated with "reproductive success" in a wide sense. However, this must not be taken to the letter: using direct reproductive success (that is, number of children) as an estimator of evolutionary fitness is dangerous. For example, it is possible to imagine genetic traits which would allow an individual to have many children - but would make those children severely unfit or unviable. Such an adaptation would not be favoured by natural selection, and would not correspond to the usual notion of fitness.

A common practical measure of an individual's fitness is its number of grandchildren, rather than number of children (see e.g. [85]). To have many grandchildren, an individual must not only have many children, but these children themselves must also be successful in reproducing; this corresponds to the intuitive notion of fitness as ability to pass on one's genes. We will use the number of grandchildren as a measure of individual fitness. Therefore, to measure fitness transmission, we measure the statistical correlation between the number of grandchildren (NOGC) of an individual, and that of its children.

9.4.2 Fitness correlation as a local measure

Fitness correlation should be used as a local measure in time. That is, rather than calculating fitness transmission in bulk over the entire genealogical record, it is preferable to divide the record in time periods, or "slices", and calculate fitness transmission independently for each period. This is done by only considering individuals born within this time period for the "child" data set of each period (the parents of these individuals are then collected in the "parents" data set, independently of their time of birth).

Clearly this method has the advantage of allowing for a more precise analysis: we can track periods of selective activity over time, identify periods of stronger and weaker selection, and potentially locate important selective events over the history of the population. But, perhaps more importantly, this also allows us to track the persistence and endurance of the correlation. As we said earlier, the fitness transmission signal created by natural selection is likely to be low: the enduring persistence of this signal, rather than its absolute value, should be the important defining criterion for the presence of natural selection.

Importantly, this does not mean that the data itself should be strictly local to the time period. For example, reproductive success for a given individual may be collected over its entire history. The time-slicing only applies to the individuals included in each local data set for calculation (in the "child" section of the data set), rather than the underlying data itself.

9.4.3 Comparing what is comparable

As usual when calculating statistical correlations, care should be taken in only comparing what is comparable: conflating data from widely different distributions may result in artificial, spurious correlations. Selective conditions may change widely over the course of an evolutionary run, even with a fixed fitness function. This may wreak havoc on undiscerning evaluations of statistical correlation. For example, if strict ranking is used, surviving and reproducing entails dislodging a previous survivor; but as evolution proceeds towards an optimum, and new champions are increasingly well-adapted, it becomes increasingly difficult (and thus rare) for new individuals to dislodge previous champions. This means that the children's fitness will tend to go down (because more of them disappear without a descent) and the parent's fitness will tend to go up (because they remain in the population longer) over time. This alone is sufficient to create a strong, *negative* correlation between the fitnesses of parents and children over the whole process: earlier parents would have a moderate number of grandchildren, each with a good chance to reproduce; while later parents would accumulate enormous numbers of grandchildren, which would have comparatively low reproductive success.

To avoid this, we must ensure that we only consider quantities (that is, fitnesses) obtained under similar conditions. Collating data from widely different times would violate the requirement for a 'fair comparison'. To this end, the periods over which reproductive successes are measured should start at the same point in time, so that we can ensure that they are obtained over equivalent conditions. In practice, this means that when we compare the NOGC of an individual X and its parent, we should only consider the grandchildren of the parent that were born *after* X (or, if applicable, after X's maturity). This ensures a "fair game" between the parent and the child: both scores will be obtained under similar circumstances, and results obtained by the parents in earlier (possibly harsher or milder) circumstances will not spoil the data.

9.4.4 Necessary normalisations

Unfortunately, the choice of using NOGC as a measure of fitness introduces an obvious problem: the NOGC of an individual and that of its children are clearly not independent quantities. Saying that A has many grandchildren is saying that A's children have many children, and therefore, *out of this fact alone*, are likely to have many grandchildren themselves, even with random reproduction. This problem can be easily addressed by normalisation to make the considered values independent. To do this, we do not use the raw NOGC for the parents; rather, for every parent-child pair, we consider the parent's NOGC *minus the number of children of this particular child.* This modified NOGC is an estimation of the parent's fitness that is not biased by this particular child's own success, and thus any correlation represents a true correlation in fitness.

Another, less significant problem is that, in general, the population of interest will be finite. The consequence is that the reproductive successes of individuals living during the same period of time are not independent: any child for a given individual is one less opportunity for another individual to have a child. Even with random mating and reproduction, if one individual happens to have more children than average, then any other randomly picked individual is mechanically more likely to have fewer children than average. In other words, limited population introduces a slight negative correlation between the modified NOGC of parents and children. This effect is much less important than the previous one, but may be noticeable, especially with small populations. A simple solution to this problem is to normalise the modified NOGC of the parent: for every parent-child pair (P_i, C_i) from the slice, we divide the modified NOGC of P_i by the total sum of all grandchildren of all other parents within the slice - minus C_i's children. The resulting proportion is independent of this child's own success.

Those normalisations are made necessary by the fact that the quantities under scrutiny are not independent. They would become unnecessary if, instead of evaluating fitness transmission from parents to children, we attempted to calculate it between grandparents and grandchildren. The problem, of course, is that any signal would be much weaker due to the increased indirection - often to the point of being drowned in noise.

9.4.5 Calculation method for fitness transmission

Where does this leave us? From all these considerations, we can deduce the following calculation method for fitness transmission:

- Divide the entire genealogic record into discrete periods of time. If the system is generational, generations may be used as time periods.
- For every time period within the genealogic record, perform the following operations:
  1. For every individual C_i born during this time period, find its parent Pi (which may be born at any time before C_i, not necessarily during this time period) and store the resulting parent-child pair (P_i, C_i). Note that any given individual may occur in several pairs.
  2. For every stored parent-child pair (P_i, C_i), retrieve their respective total number of grandchildren (NOGC) $N(P_i)$ and $N(C_i)$, born during or after (not before) this time period.
  3. Elimination of dependency: for every pair (P_i, C_i), subtract the number of children of C_i from the $N(P_i)$, resulting in the new value $N'(P_i)$.
  4. Normalisation: for every parent P_i in the set of parent-children pairs for this time period, divide $N'(P_i)$ by the sum of all grandchildren of all other parents $P_{j \neq i}$ - carefully excluding C_i and its descendants from the count. This results in a final value $N''(P_i)$
  5. Calculate the statistical correlation between the $N''(P_i)$ and the $N(C_i)$ variables over all parent-child pair for this time period, using the standard Pearson formula:

$$\text{Corr}(X, Y) = \frac{\sum_{i=1}^{N}(x_i - \bar{X})(y_i - \bar{Y})}{(N-1)\sigma_X \sigma_Y}$$

The resulting value $\text{Corr}(N''(P_i), N(C_i))$, for every time period, is our estimator for the intensity of fitness transmission during that time period.

9.5 Experiments

9.5.1 Experimental settings

Our purpose in this section is to set up a couple of experiments in order to determine whether fitness transmission is indeed a reliable indicator of Darwinian evolution. To do this, we will use simple evolutionary systems with predictable dynamics, in which the presence or absence of evolution can be easily controlled. We will apply our calculation method to these systems and determine whether the presence or absence of Darwinian evolution was successfully detected.

To perform our experiments, we used genetic algorithms involving a population of 1000 individuals, over 100 generations. We considered two optimisation problems: the Rosenbrock function $100(x^2 - y^2)^2 + (1 - x)^2$ (using genomes of 2×12 bits) and a very simple OneMax problem over 20 bits. The Rosenbrock function is a commonly used test function in the field of optimisation. The purpose of the simple OneMax problem is to examine the behaviour of different algorithms on very easy problems, when the the global optimum is discovered quickly.

In our algorithms, at each generation, a new population is created either by applying bitwise mutation to a parent selected from the previous generation, or (with 66% probability) by applying one-point crossover between two parents, and then applying bitwise mutation to the resulting offspring. The probability of mutating (flipping) each bit is the reciprocal of the total number of bits in the genome, rounded to the closest higher percent; thus, on average, each genome should undergo about one mutation. As explained below, we tested different methods of selection and replacement.

In a normal genetic algorithm, Darwinian evolution can be expected to occur. As a point of comparison, we need to compare a “neutral” version of the genetic algorithm, which preserves as many features of the algorithm as possible, while effectively removing Darwinian evolution. We could simply use a normal genetic algorithm with random selection, similar to the shadow systems used for the Bedau-Packard measures. However, in this case, the absence of fitness transmission is a rather obvious result which hardly requires experimental verification. Can we find a more challenging test case?

We chose to use genetic algorithms with uncorrelated landscapes, in which every new individual was attributed a random genotype (and therefore a random phenotype) at birth, regardless of the genetic make-up of its parents. This is different from purely random selection in that selection still occurs, and is still based on fitness; however the randomness of the reproductive process turns the problem into a random search on a random landscape, thus preventing any meaningful evolution: no innovation can propagate throughout the population. Fitness-affecting traits are still present, but not heritable. A satisfactory measure of evolutionary activity should be able to detect the absence of real evolution and return a zero value for this situation. Note that, as discussed in section 9.2.3, Bedau-Packard statistics produce a positive signal when faced with such a system.

9.5.2 A simple genetic algorithm

We first describe the calculation of fitness transmission in a standard genetic algorithm, quite similar to the original model proposed by Holland [82], except that we use tournament selection rather than roulette wheel selection. In this algorithm, each new individual is created by selecting parents from the previous generation (using tournament selection), and generating offspring as previously described. The process is iterated until the new population is filled.

Figure 9.2 shows the results of these calculations, applied to the “fossil record” generated by our simple genetic algorithm. This figures shows the results for the Rosenbrock function optimisation problem with 20 bits, both with normal reproduction and with reproduction based on random phenotypes. The top graph shows the results of 5 different run for each of these reproduction methods, while the bottom graph shows average curves over 50 runs. Figure 9.3 shows the same data for the OneMax problem. In the normal selection case, the correlation between the number of children of parents and children is distinctly positive (especially at the very beginning at the run) and stabilises to a positive plateau. The enduring positive value indicates that the population is constantly and actively maintained in the vicinity of the global optimum (which is reached quite early in the OneMax problem) through active evolutionary forces. Even though the optimum has been reached, mutation constantly disperses the population, and Darwinian evolution constantly drives it back. Stabilising selection results in a positive value for differential fitness transmission.

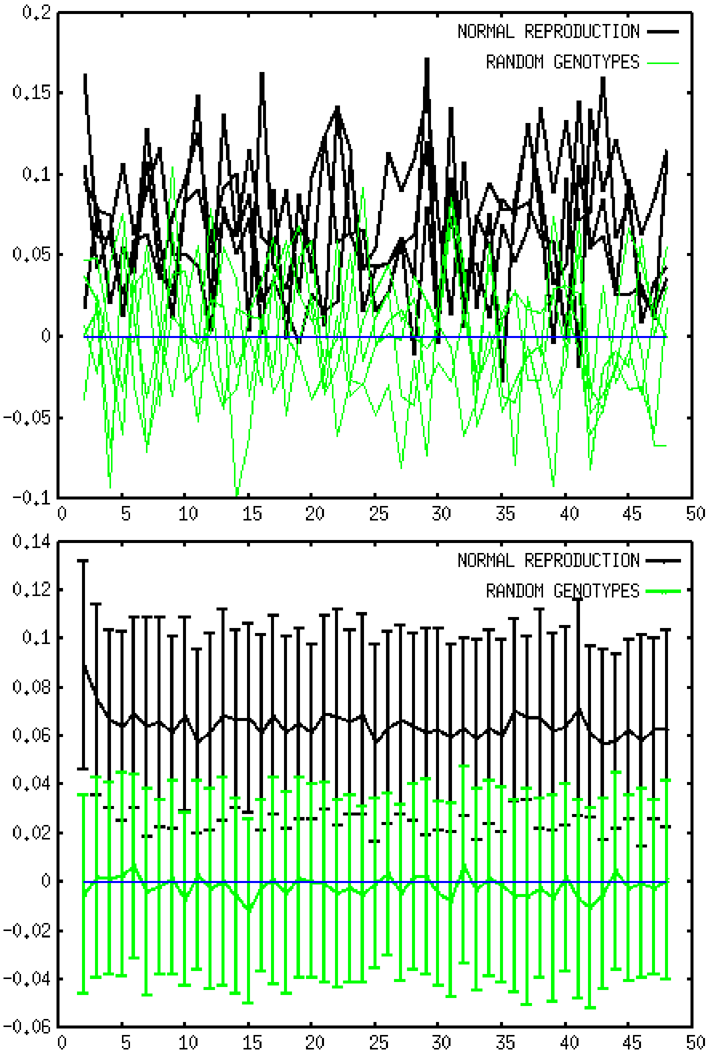

Figure 9.2: Rosenbrock function, non-overlapping generations, 5 different runs (top) and average of 50 different runs with standard deviation at each point (bottom).

In the case of random genotypes, as expected, no meaningful fitness transmission occurs.

9.5.3 Removing mutation

It is easy to show that the enduring presence of fitness transmission in this case is caused by mutation. If we set the mutation rate to zero (see figure 9.4), then the population converges totally: all individuals end up sharing the exact same genome, and diversity disappears. From this point on, all individuals having exactly the same genotype, reproduction patterns become effectively random. The result is that evolutionary activity, as indicated by fitness transmission, quickly goes to zero (with noise oscillations) after an initial phase of high activity: Natural selection exerts no active force upon the system, because the population has basically frozen and no genetic variation exists to drive evolution. This illustrates the capacity of fitness transmission to distinguish between active stabilising selection on the one hand, and passive stillness caused by absence of genetic variation on the other (though this ability breaks down in extreme situations, as discussed in section 9.6.)

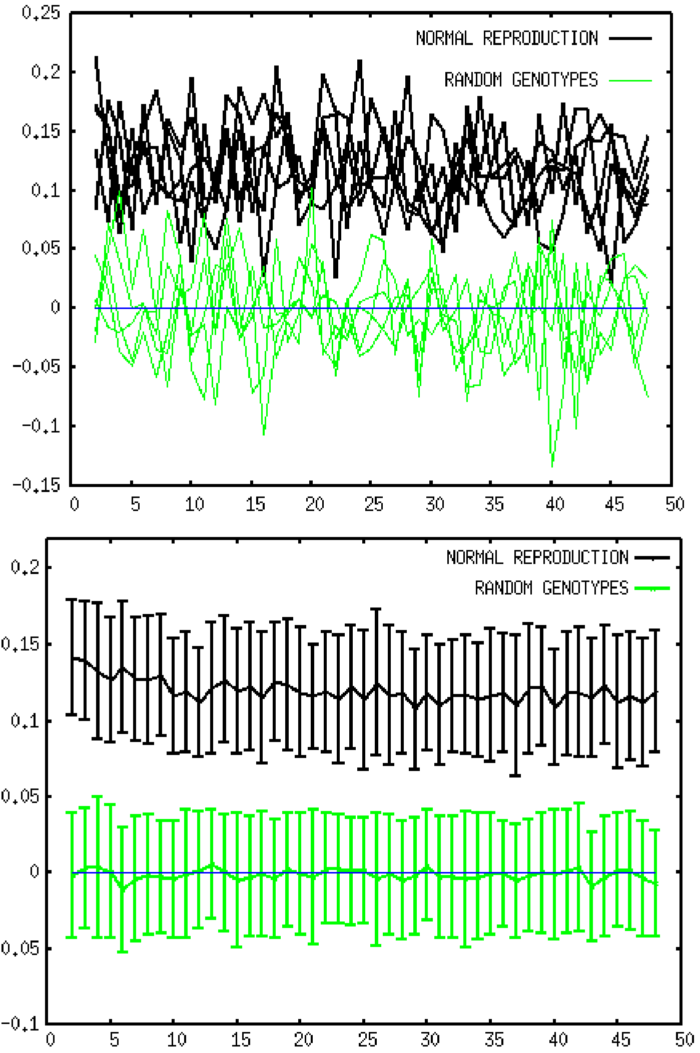

Figure 9.3: OneMax function, non-overlapping generations, 5 runs (top) and average of 50 runs (bottom).

9.5.4 Removing selective gradient among parents

Here we try to make the problem more challenging problem by reducing the scope of selection. To do this, we modify our algorithm as follows: at every generation, a small set of survivors is selected from the population through strict ranking selection, and the parents for the next generation are then *randomly* selected from among this set of survivors. Offspring are created using previously mentioned methods (60% crossover, mutation, etc.)

The effect of this modification is to effectively remove any selective gradient among parents. This is because the only effect of selection is to decide which individuals become parents in the first place. Once individuals have been selected as parents, their number of children is random, and as a result is not affected by natural selection. In particular, note that if we had tried to evaluate fitness by the number of children alone, then no fitness transmission could be detected: no correlation can exist between the number of children of parents and children, simply because all parents have a random number of children. However, as shown in figures 9.6 and 9.5, our measure for fitness transmission is able to detect the signal created by this more indirect form

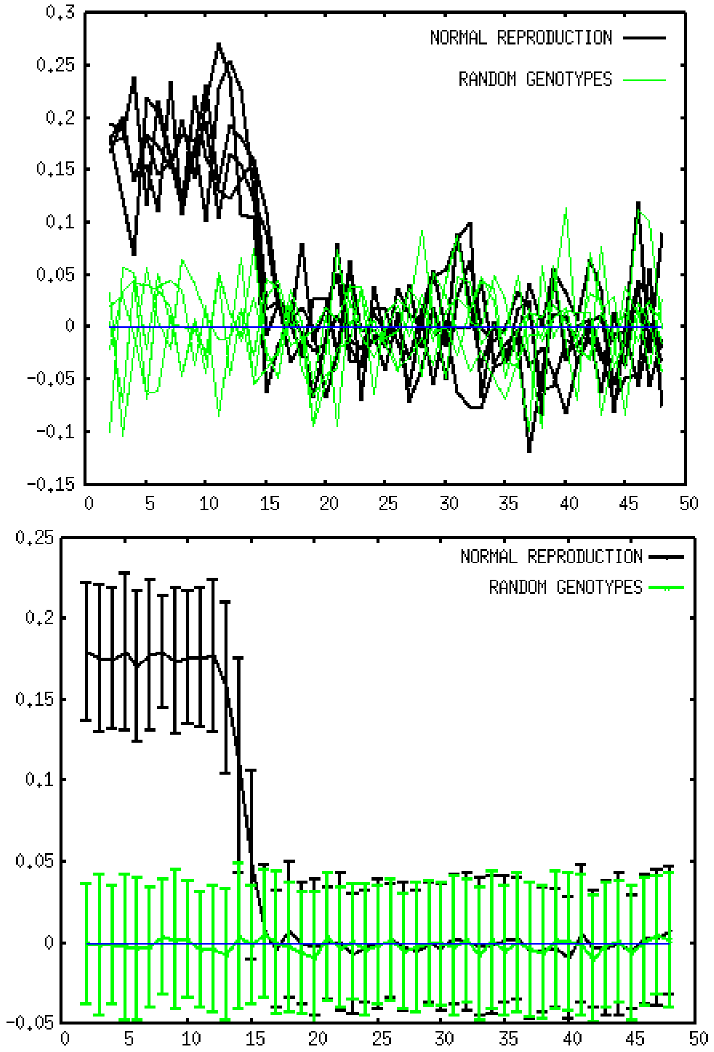

Figure 9.4: OneMax function, non-overlapping generations, without mutation, 5 runs (top) and averages of 50 runs (bottom).

of natural selection.

Note that in the case of the Rosenbrock problem, the resulting values are very low. We use two significance test to confirm our interpretation of the graphical result. First, we want to confirm that the expected fitness transmission is indeed non-zero for the selective case in the long term (we arbitrarily define "long term" as "after cycle 30"). To do this, for each of the 50 runs, we calculate the average value of fitness transmission from cycle 30 onwards. Out of the resulting 50 values, we find that 46 are above zero, while only 4 are under zero. A binomial distribution calculation indicates that under the null hypothesis of equal probability of being under or over zero, the probability p of obtaining at most 4 successes out of 46 trials is enormously low (on the order of $10^{-}10$).

Second, we want to confirm that the values of fitness transmission for the selective process are indeed superior to those for the non-selective process. To do this, we calculate the average value of fitness transmission from cycle 30 onwards, for all 50 runs of both the selective and non-selective process. We then perform a two-tailed Mann-Whitney test between the two samples.

We obtain a significance value of $p < 2 * 10^-6$.

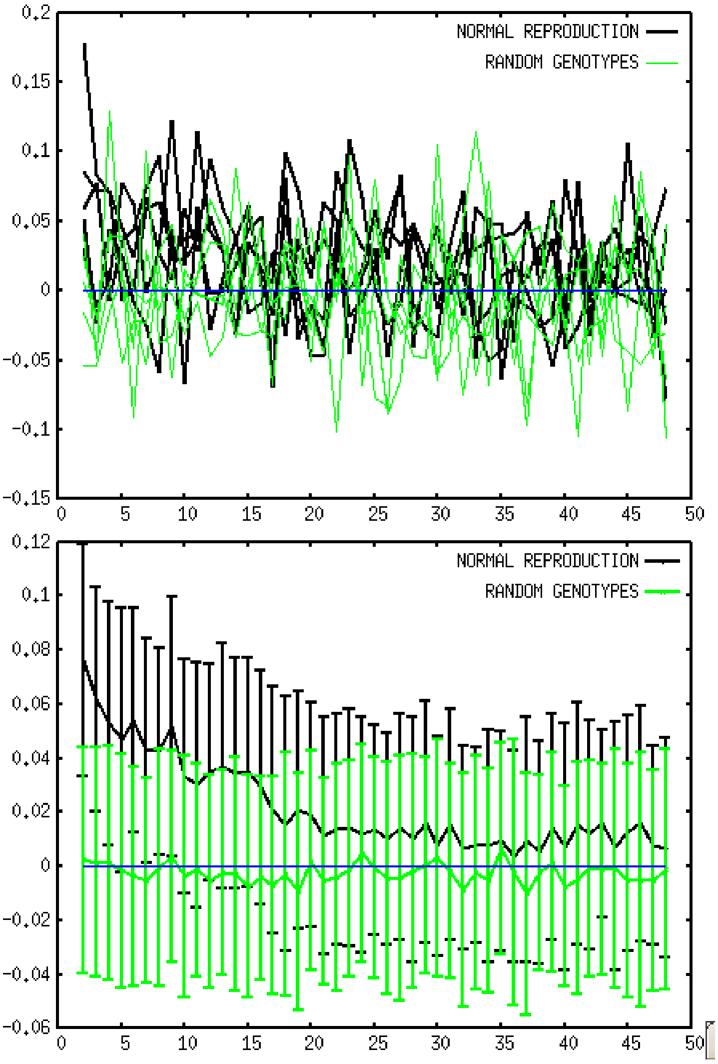

Figure 9.5: Rosenbrock function, non-overlapping generations with ranking-based survival and random parent selection, 5 runs (top) and averages over 50 runs (bottom). The initially high signal goes to a very low, but still noticeably non-zero value.

9.6 Limitations of fitness transmission

9.6.1 Extreme stabilising selection

We have seen how fitness transmission can be used to detect the presence and activity of natural selection, even when this activity consists in maintaining the population around a local optimum (stabilising selection). However, it must be pointed out that this last property does not hold in the extreme case when stabilising selection is absolute - that is, when only those individuals exactly at the optimum are able to reproduce, with no significant difference in reproductive success among them, and individuals which deviate in any way from the optimum cannot reproduce at all.

In this case, it is clear that no fitness transmission occurs: there is no heritable characteristic

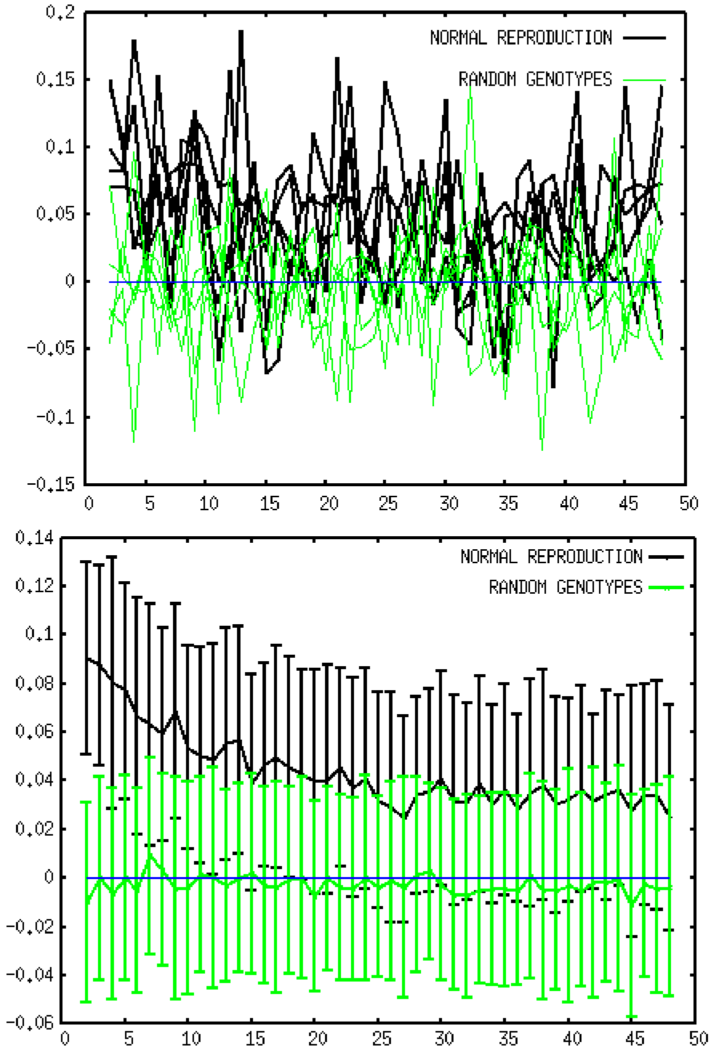

Figure 9.6: OneMax function, non-overlapping generations with ranking-based survival and random parent selection, 5 runs (top) and averages over 50 runs (bottom).

that can alter an individual's reproductive success, except by destroying it entirely. Fitter parents do not have fitter children. The only feature that can alter reproductive success (mutation away from the optimum) actually eliminates it altogether, and therefore is not transmissible[4]. In this situation, stabilising selection has the effect of effectively freezing the reproducing population, and therefore becomes invisible to fitness transmission. In fact, it is impossible to distinguish such a situation from purely random selection without accessing the genetic record (or an indirect proxy, such as the record of fitness function evaluations for artificial systems.)

9.6.2 General limitations of genealogic analysis

More generally, there are pathological situations in which genealogic methods cannot be used to detect natural selection. Such a situation occurs when all individuals present at any given

---

[4] Notice the difference with the situation evoked in section 9.5.4, in which, despite the fact that all parents for a generation had equal chances at reproduction, they could pass on genes that could influence the probability that their own descendants would also be selected as parents.

time share the exact same genealogic tree - in other words, when there is never more than one lineage in the population.

Genealogic analysis basically consists in comparing the relative growth of various lineages. This implies that there should actually be several lineages within the population, otherwise no comparison is possible. However, we may imagine a system in which all individuals present at any time within the population share the same entire genealogic tree. In this case, there is only ever one single lineage present at any time within the population, and no comparison is possible. For example, consider a non-overlapping generational system, such that all individuals from every new generations have exactly the same parents from the previous generation: at every generation, two individuals are chosen to serve as parents for the next generation, and *all* the individuals from the new generation are children of *both* of those selected parents. In this case, Darwinian evolution may certainly occur, but genealogic analysis will not be able to detect whether the selection of parents is adaptive or random. This is because all individuals will always share the exact same set of parents, grandparents, and so on.

We believe that this situation is sufficiently exotic to preserve the usefulness of genealogical analysis. At any rate such situations can be easily detected in any system for which a genealogical record exists.

9.7 What Fitness Transmission measures, and what it doesn't

Having proposed differential fitness transmission as a signature for Darwinian evolution, it is important to make it clear what exactly this signature reveals and what it doesn't, in order to avoid confusion.

Fitness transmission is not a measure of performance or optimisation. This is readily seen from the fact that even random search can produce some amount of optimisation, even though no evolutionary activity or fitness transmission is present. Conversely, we saw that for some algorithms, fitness transmission can remain positive even after the global optimum has been reached, as long as Darwinian forces keep acting on the population.

Fitness transmission is not a measure of inventiveness, creativity or novelty. Whether or not the newly created individuals are novel or not is not considered: only their fitness (and the correlation between their fitness and their parents' and children's fitness) is taken into account. Again, this is seen from the fact that positive fitness transmission is compatible with the fact that the population is being maintained around a single optimum, where no adaptively significant innovation is introduced.

A more subtle aspect of fitness transmission is that it detects natural selection in general, including natural selection bearing on non-genetic traits. Any kind of heritable trait that affect reproductive success (genetic, epigenetic, cultural, etc.) will be detected by fitness transmission. If the objective is to detect biological natural selection only, then fitness transmission should not be used alone.

Fitness transmission is simply the answer to this question: "all other things being equal, is it true that individuals from a common lineage will tend (ever so slightly) to have similar fitness?" Since nothing is assumed about the nature of fitness, except for its basic observational property (reproductive success), we can attempt the following, general statement: fitness transmission simply measures *the presence or absence of some force which actively favours some lineages at the expense of others* (over at least two generations). We believe that in general, this constitutes a good signature of Darwinian evolution by the operation of natural selection on heritable characteristics.

9.8 Conclusion

We have shown that differential fitness transmission is a useful signature of Darwinian evolution, which can be detected in genealogical records by using simple statistics. We believe that this signature may be more suitable for this purpose than previously suggested methods for detecting evolution. We have applied these statistics to the genealogical records generated by real evolutionary algorithms, demonstrating their capacity to detect the presence or absence of adaptive evolution. In chapter 10 we will apply this measure to a much more complex system, namely a "virtual world" in which virtual creatures evolve freely.

Chapter 10

Evosphere: Emergence of natural selection and Darwinian evolution among free-living artificial creatures

Wherefore he made the world in the form of a globe, round as from a lathe, having its extremes in every direction equidistant from the centre, the most perfect and the most like itself of all figures; for he considered that the like is infinitely fairer than the unlike.

Plato, *Timaeus.*

10.1 Introduction: evolution unbound

In previous chapters we have used our platform to implement and study various algorithms. Our discussion was very much oriented towards performance, optimisation and efficiency: we sought to obtain "good" individuals, where "good" is understood to relate to a given task specified *a priori* (locomotion, box-grabbing, combat, etc.) In fact, we devoted much time to the definition of useful measurements of performance, both for individuals or for algorithms.

However, it has long been argued that evolution as it occurs in nature is fundamentally different from the most common type of evolutionary process encountered in artificial evolution - namely, selection based on human-defined fitness functions. There is much speculation that the use of a pre-defined fitness function constrains the evolutionary process into a predictable pathway, thereby curtailing the creativity of evolution as it occurs in nature. In this line of reasoning, evolution based on human-defined fitness functions is sometimes equated to the process of *artificial selection*, as applied by human breeders and farmers to crops and animals in order to favour certain desirable characteristics; this is contrasted with "real" natural selection, which is bound neither by our selfish interests nor by our lack of insight or imagination. The conclusion is that only "free", unguided natural selection can produce the dazzling amount of complexity and diversity evident in nature. For example, Ray [147] writes:

> I suggest that the most likely way to achieve complexity increase in digital evolution is through evolution by natural selection in an ecological community. No attempt

> should be made to provide fitness functions, or artificial selection, to guide evolution towards useful products. Rather, evolution should be free to explore the possibilities without the burden of human "guidance".
>
> Traditionally we have managed evolution through manipulating selective forces. In this new approach, our role is to create the conditions for complexity increase, rather than trying to guide it through artificial selection.

Ray prudently points out that "this is an interesting scientific challenge, as the conditions that generate complexity increase are unknown."

In a similar vein, Channon [28] writes (emphasis original):

> *In the context of evolutionary emergence, any artificial selection system used constitutes just one of the parts of a system.* Artificial selection can only select for that which is specified to. Therefore anything that emerges during evolution *must* be due to another aspect of selection, which must in turn be due to the innate dynamics of the system - natural selection.

Packard [137] has attempted to make this distinction more explicit, by introducing the concepts of *intrinsic* and *extrinsic* adaptation (emphasis original):

> A common approach in modelling evolutionary processes is to regard a member of a population as fit or unfit according to how it is evaluated by a fitness function. The fitness function is a map that assigns a real number (the fitness) to every possible member of the population, and it is generally specified as an a priori feature of the model. . . I will call any adaptive dynamics that uses such an a priori fitness function *extrinsic adaptation*. . .
>
> The biosphere does not appear to have any *a priori* fitness function defined on a space of possible organisms; in fact, one of the most amazing features of biological evolution is that the biosphere evolves automatically, with each organism (or population of an organism) adapting to an environment made up of both external natural features and all other organisms it interacts with. . . I will call adaptation that occurs as a result of a population of subsystems changing in response to interactions between them (without an a priori fitness function) *intrinsic adaptation.*

10.1.1 Explicit and implicit fitness functions

There may be an objection to this view. In any evolutionary simulation, it is necessary that the human programmer define how survival and reproduction occur: there must be rules to determine who lives and who reproduce, otherwise no evolution could occur at all. But a set of rules that determine which individuals survive and reproduce (or not) is precisely what we call a fitness function. This being so, how could we build a system "without any fitness function"?

The issue is resolved by making a distinction between *explicit* and *implicit* fitness function. It is true that the programmer must define the rules that determine when and how an individual dies, or reproduce, or is born. Together, these processes define the *explicit* fitness function -

the narrow specification of which conditions will mechanically cause death, or reproduction, or birth. However, this under-determines the question of *which* individuals will turn out to actually fulfil these narrow conditions. This is decided by the dynamics of the process, which occur as a (potentially unpredictable) result of the rules of the system. The important point is that, if the fulfilment of these conditions is somehow dependent on heritable characteristics, then these characteristics (which are not directly involved in the explicit fitness function) will become the target of adaptive evolution.

For example, imagine that the rule for reproduction is simply that any individual that reaches a certain energy level is allowed to produce one offspring. This is clearly a fitness function. However, it does not determine *how* such a level of energy may be reached. This depends on the mechanics of the system, which may (or may not) be highly complex and unpredictable.

Thus, out of the interaction between the narrow, explicit fitness function, the genetic representation (what is under genetic control) and the dynamics of the system, an *implicit* fitness function may spontaneously emerge, that may not be immediately predictable from the explicit rules of the system.

10.1.2 Different types of implicit fitness functions

However, even if an implicit fitness function does emerge that is not directly specified in the explicit fitness function, it is perfectly possible that this function will turn out to be utterly boring and static. If such a thing happens, then all we have done is indirectly recreating a fixed (possibly multimodal) fitness function, that the system will duly optimise. To carry on with the previous example, suppose that the only way for an individual to gain energy is simply to move around and cover as much ground as possible, or to be located in a certain portion of the environment. If this is so (and assuming that motion is affected by hereditary characteristics), then the implicit fitness function reduces to “move very fast” or “go to that place and stay there”. In this case, there is not much difference between the implicit and explicit fitness function; both will be defined *a priori*, creating a fixed fitness landscape that will be determined even before the experiment starts.

However, things are different when *coevolution* is introduced in the fitness landscape; that is, when the survival and reproduction of individuals is linked to the interactions between their own heritable characteristics, and those of other individuals. If it is so, then the fitness landscape becomes *self-generated*: it arises spontaneously from the interactions between individuals, and changes over time as individuals themselves change through evolution and thereby offer new adaptive opportunities for other individuals. In other words, when the system is coevolutionary, the population constantly generates its own implicit fitness functions.

However, even in this case, it is perfectly possible that this emerging, self-generated fitness landscape turn out to be entirely unimodal and predictable: the coevolutionary process may well turn out to possess one global optimum, towards which evolution will converge. In fact, in the context of coevolutionary optimisation, this is precisely what we want: for example, if we want to design a chess player through coevolution, we use coevolution as a substitute for a (non-existent) fitness function that would capture the essence of “good play”. We are not interested in mutual adaptation or ecological effects - in fact, these are seen as a nuisance that harms progress in terms of general abilities. Indeed, much of this thesis has concentrated precisely on defining unambiguous, *a priori* scales of "goodness"- and on ensuring that algorithms actually follow those scales, instead of getting trapped in endless mutual adaptations.

But in the context of artificial ecologies, it might be preferable to observe dynamic fitness landscapes, in which various implicit fitness function would continuously emerge *in response to each other*. That is, the evolving lineages would enter into a constant cascade of innovations, adaptations and counter-adaptations, perpetually generating (and filling) new niches and thereby offering new opportunities for others. Such processes, which are readily apparent in nature, seem to capture the intent of the authors quoted above. Through such a process, artificial evolution might attempt to imitate nature's ability to generate "endless forms most beautiful" from "simple beginnings".

Note that we are not simply talking about dynamism and multimodality in themselves. It is of course easy to generate multi-modality and dynamic fitness function through entirely external means, simply by adjusting and altering the parameters of the simulation over time, so that the fitness optima vary over time. However, the really interesting case is when the dynamism and multi-modality of the implicit fitness function is a direct result of internal evolutionary effects - that is, when new fitness landscape emerge in response to each other through coevolution itself.

10.1.3 Summary and classification

To sum things up, the distinction between "extrinsic" and "intrinsic" adaptation cannot be summarised as the (impossible) absence of any fitness function. Rather, this distinction hinges on the following parameters:

1. Whether the implicit function is fixed a priori, or self-generated; the latter occurs if there is coevolution within the system.
2. Whether the implicit fitness function is fixed and unimodal, or dynamic and massively multi-modal.
3. Whether the dynamism and multi-modality of the implicit fitness function is determined by entirely external factors, or whether they are also internally generated by the system, in that implicit fitness function constantly emerge *in response to each other.*

This discussion suggests a natural classification of evolutionary environments, according to their degree of evolutionary "creativity".

- Level 0: The implicit fitness function is entirely defined a priori. Coevolutionary effects play no significant role in the fitness landscape.
- Level 1: The fitness landscape is self-generated, being affected by coevolutionary effects, but results in a single, stable implicit fitness function. If several peaks exist, or if the implicit fitness function(s) change(s) over time, this is due to entirely external, "climatic" variations in the parameters of the environment.[1]

---

[1] To distinguish between levels 0 and 1, a simple test suggests itself: let us alter the system so that, at any time, only one individual is subject to evolution (that is, the evolving genome is expressed by only one individual

  - Level 1b: Similar to Level 1, but the implicit fitness function that emerges is *unpredictable* and can vary from run to run.

- Level 2: The fitness landscape is affected by coevolutionary effects, and results in several implicit fitness functions that do emerge in response to each other. This process goes on for a finite period of time, after which the system settles in a fixed equilibrium in which no novel implicit fitness function appears. This may include constant cycling through previously found fitness function. In short, the total number of implicit fitness function generated (and optimised) by the system is *bounded* over time.

  - Level 2b: Similar to Level 2, but the implicit fitness functions that emerge are unpredictable and can vary from run to run.

- Level 3: The fitness landscape is coevolutionary and the system constantly generate new implicit fitness functions which emerge in response to each other. No equilibrium is found, and novel fitness functions are constantly being generated. This corresponds to a situation of truly "open-ended evolution".

10.2 Objectives

Several open evolutionary systems have been described. Tierra [146, 148], Polyworld [197], and Geb [28, 27] are well-known examples (see Appendix B). However, these systems generally use immaterial entities, devoid of a tangible "body" through which they could physically interact with their environment and neighbours. By contrast, in the present chapter we seek to produce an intrinsic adaptation environment that is based on embodied individuals, endowed with realistic physical morphologies. In fact, we intend to make physical interactions the very basis of evolution. Previous implementations of evolution among physically realistic individuals have been centred on explicit fitness functions or head-to-head coevolution, perhaps most famously in Sims [163, 162]. In the case of Framsticks (described in Appendix B), despite the physical realism of the system, interactions between individuals are entirely *non*-physical; for example, "fighting" is implemented simply as a comparison between energy levels, such that the individual with the highest energy level "wins".[2]

The objectives of the present chapter can therefore be summarised as follows:

- Creating an environment in which a population of free-living creatures physically interact with each other.

- Initiating a process of spontaneous natural selection and Darwinian evolution in this environment, without any explicit fitness function.

---

within the system) while the rest of the population is composed of *fixed* individuals. If the results of evolution are the same in this altered system as in the original system, then it is clear that the trajectory of the system is not dependent on coevolutionary effects.

[2] See the "Deathmatch tutorial" by W. de Back and M. Komosinski, apparently not formally published but available on the website of the University of Utrecht Virtual Life Lab.

- Demonstrating the occurrence of Darwinian evolution, not just by visual inspection, but more formally by using the method previously described in chapter 9 - namely, Fitness Transmission.

By applying Fitness Transmission to our system, we expect to detect the onset of natural selection and Darwinian evolution (should they occur), but we also seek to illustrate the application of Fitness Transmission to a "real" system (as opposed to the toy models used in chapter 9).

10.3 Evosphere: An open environment for free-living creatures

10.3.1 The "micro-planet" environment

Because our environment is expected to contain many free-living creatures, we would like it to be *open*, that is, without boundaries. The imposition of barriers could introduce unpredictable effects into the dynamics of the system. While certain systems (such as Polyworld) choose to embrace the existence and effects of such barriers, others (such as Geb) solve the problem by using a "toroidal" 2D environment where the top connects with the bottom and the left with the right.

In a 3D environment, the simplest structure that dispenses with boundaries is obviously the sphere. We therefore choose to implement our environment in the form of a sphere, or "micro-planet". Creatures dwell on the surface of this planet, which is sufficiently large to allow for free movement of each creature.

10.3.2 Interaction: physical damage

The creatures themselves are similar to those presented in chapter 5. The main mode of interaction between creatures (besides sensors) is through physical contact, and the damage that it may create. When two creatures come into contact, they may suffer damage, in a way similar to that described in chapter 5; however, damage is now proportional to the absolute speed of each contacting limb towards the other. In this way, at any point of contact between two creatures, one or both contacting limbs may actually suffer damage.

Let us call both contacting limbs L_1 and L_2. Since there is contact, L_1 and L_2 inter-penetrate by a very small amount P (the sign of which is unimportant). We first run an additional step of simulation by allowing L_1 to pursue its own motion for one step, while L_2 is fixed. This results in a new penetration depth P_{12}. We then restore their initial positions and do the opposite, allowing L_2 to move for one step while L_1 is kept fixed. This results in another penetration depth P_{21}. The damage suffered by L_1 is proportional to $P_{21} - P$, while the damage suffered by L_2 is proportional to $P_{12} - P$. If any of these values is negative, the corresponding limb suffers no damage at all, because it implies that the other limb was moving "away" at the time of contact.

10.3.3 Energy and death

Every new creature is born with a certain amount of "energy" E_0, which is a constant of the system. At every timestep, this energy is being decreased by a (small) amount δE_{time}. Additionally, energy is also lost through physical damage: if a creature suffers damage D, its energy is decreased by an amount δE_{damage} proportional to D. The parameters are chosen in such a way that physical contact is much more damaging than the passage of time.

When a creature's energy falls to zero or under zero, it is removed from the simulation and immediately replaced with a new creature: the population size is therefore constant. How the new creature is generated depends on the manner in which the removed creature met its end.

10.3.4 Reproduction

At any timestep, if a creature's energy becomes nil or negative, there are two possibilities: either the creature has suffered damage in this very timestep, which has contributed to its demise; or it has not suffered any damage in this timestep, and therefore the annihilation of its energy is solely the result of δE_{time}. In the first case, it has been "killed". In the latter case, it has simply "died". Of course, a dying creature may well have suffered extensive damage in previous timesteps, but this is not relevant: only the very last timestep of the creature is considered.

If a creature dies "peacefully", that is, without having suffered any damage in this timestep, then the new creature that replaces it is generated by selecting two parents at random from the population, and applying one of the previously mentioned genetic operators (mutation, genetic crossover, grafting; see chapter 4 and Appendix A) to one or both of the parents as suitable.

On the other hand, if a creature has been "killed", it is replaced with a mutated clone of its killer, that is, of the creature that delivered the fatal blow. In the unlikely case that two or more creatures have inflicted damage upon the dying creature during the same timestep, the one that inflicted the highest damage is deemed the "killer".

There is the question of how to insert the new creature into the ongoing simulation, while minimising disturbances. We simply "drop" the new creature, from a fixed height, over a random point of the sphere. In addition, each new creature is granted an initial "grace delay" during which it cannot suffer or inflict damage. This delay is long enough to cover the fall to the ground and a short period beyond.

10.3.5 Consequences of the rules

From these rules we can infer certain properties of the system:

1. The initial amount of energy E_0 sets a maximum lifetime (equal to $E_0/\delta E_{time}$) that no creature can exceed. This maximum is that of a creature that does not undergo any contact. Any action taken by the creature can only *reduce* this lifetime (through physical damage), or leave it unaffected. In particular, in the early stages when damage will be infrequent, it is expected that many poor creatures will reach the maximum lifetime - simply by standing still and not having any contacting neighbour.

2. A creature can favour its own reproduction in two ways: by killing others, or (to a lesser extent) by avoiding damage better than others. In the first case, it can generate clones. In

the second case, it will have more chances than others to be selected as a random parent for dying creatures.

3. Due to the manner in which damage is apportioned, any contact is risky. When making contact with another creature, if the other creature's limb is coming towards you, you will suffer damage.
4. A creature that inflicts much damage upon another will only benefit if this damage is deadly. If the creature goes on to live for just another timestep, and then dies, the aggressor will not benefit.[3]
5. Even in the earliest stages of evolution, a small degree of damage will be inflicted due to semi-random contact. As a consequence, a small portion of deaths will be caused by killings.

10.4 Results

10.4.1 Emergence of natural selection and evolution

Natural selection and evolution consistently arise in the system, as observed both by visual inspection of creatures and by numerical evaluation of fitness transmission.

The system generated a wide range of morphologies.[4] Though some broad categories may be found ("worms", two-armed "crawlers", etc.), some creatures were altogether quite bizarre in their construction. Figure 10.2 shows four pairs of creatures evolved from four different worlds.

The majority of behaviours were purely motile behaviours and did not make use of sensor input: these creatures evolved efficient movements for "hitting without being hit." However, some creatures did include sensory input, including contact sensors and detection of other individuals. For example, the "crested" creature in the top-right corner of Figure 10.2 uses a y-sensor in its head limb (see chapter 4) to approach the closest other creature in a "swirling" motion (the creature moves by undulation).

In general, the dominant strategies were variants of a "roaming" behaviour: move fast in an undulating manner to hit whatever comes close.

10.4.2 Fitness transmission analysis

Figure 10.3 shows a graph of fitness transmission as a function of time. A genealogic record of the system is divided into periods of 20000 timesteps (a rather arbitrary value corresponding to

---

[3] Actually, the aggressor does benefit marginally, because the weakened creature dying sooner provides another opportunity for reproduction of random parents, in which the aggressive creature has a small probability of being chosen as a parent. However, this "benefit" is enjoyed by all creatures alive at the time, since parents for the new creature are randomly chosen among all the living population. On the other hand, the risk induced by damaging the dying creature were taken solely by the aggressor itself. Consequently, theory indicates (and experiments confirm) that this benefit does not create a positive selection for aggression; another, more focused benefit (in this case, cloning-by-killing) is necessary.

[4] All creatures described in this chapter had either the highest or second-highest number of children of all creatures alive at the time of observation. In open evolution, maximum number of children is the closest approximation to the notion of "champion" that exists in experiments based on fitness function or head-to-head coevolution.

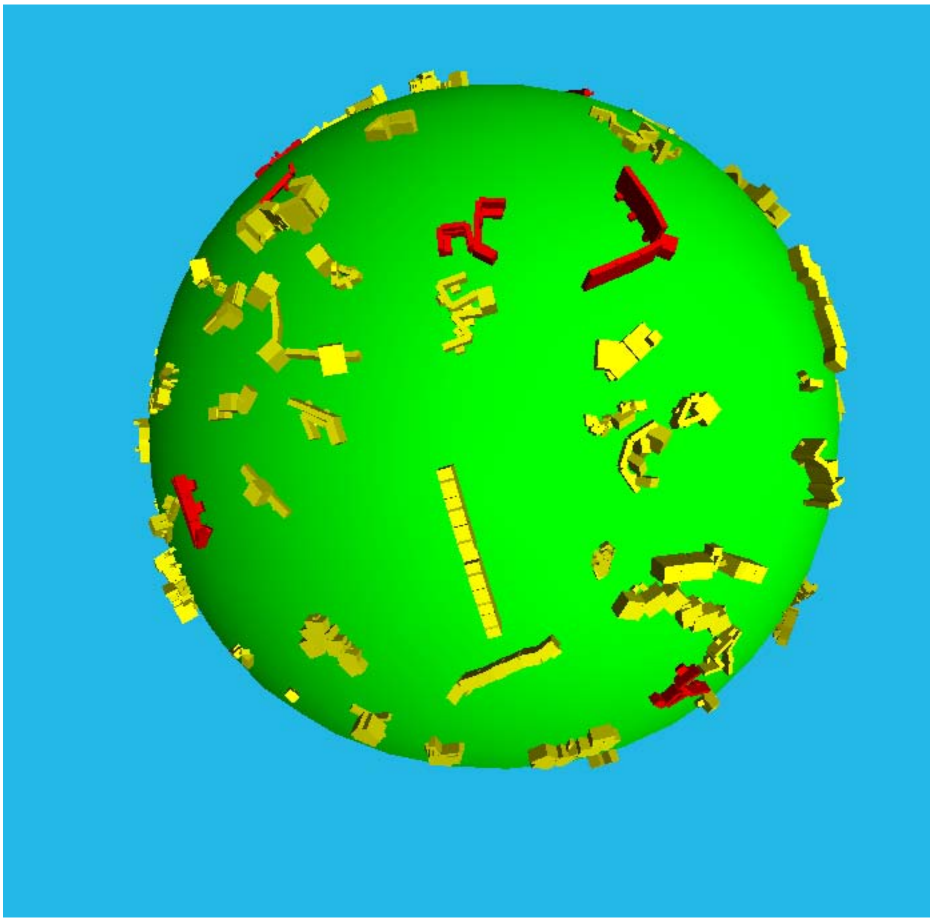

Figure 10.1: Evosphere, as seen from “space”.

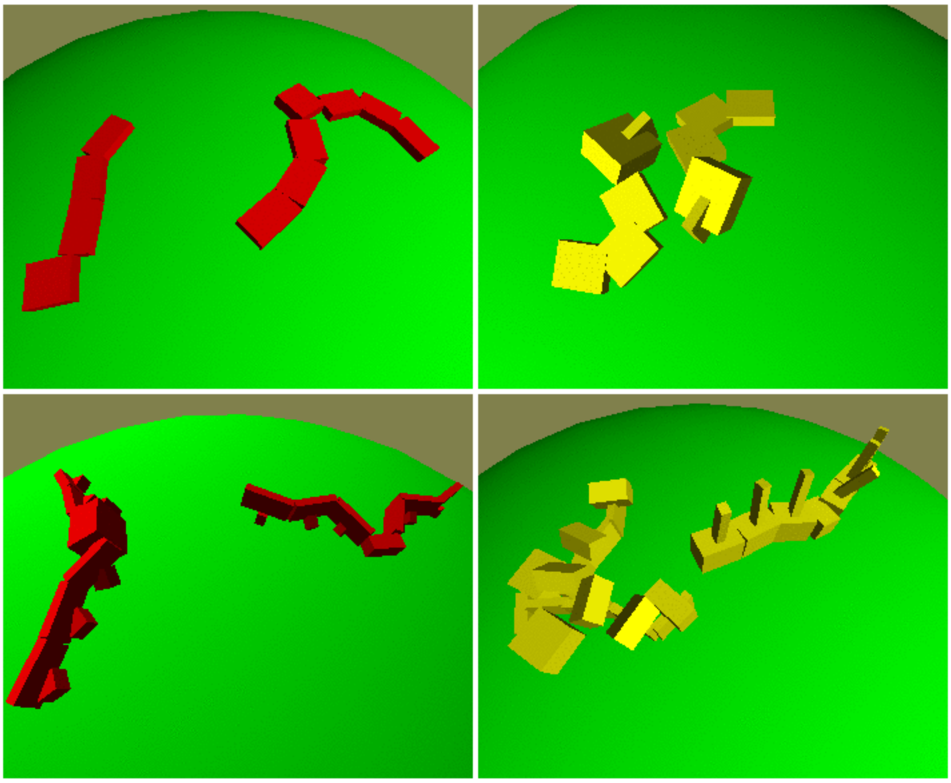

Figure 10.2: Four pairs of creatures from various runs.

twice the maximum possible lifetime of a creature). The graph in this figure provides the fitness transmission calculated for each of these periods, as well as a rolling average over 5 successive values.

As indicated by this graph, the system usually goes through an early phase of very low or absent fitness transmission, denoting low or non-existent natural selection. In this phase, creatures are so ill-adapted that any damage they inflict or suffer arises mostly through chance rather than because of their characteristics. However, certain characteristics can be faintly advantageous, such as temporary motion as opposed to absolute immobility (depending on the presence of close neighbours).

This phases ends with the appearance of persistent, oscillating motion. As discussed in chapter 4, persistent motion requires oscillating neural activation and therefore specific connection patterns. Oscillating limbs are highly advantageous for two reasons. First, they allow the creature to displace itself on the surface of the world, rather than remaining in its initial location. Second, constant motion of limbs implies constant possibility of inflicting damage upon other creatures. Once this innovation is introduced in the system, natural selection has something substantial to operate on (that is, fitness-impacting heritable characteristics), and Darwinian evolution begins in earnest. This is reflected in a sudden increase in the value of fitness transmission. Roaming behaviours quickly emerge, and thrive.

As an additional confirmation of the validity of fitness transmission analysis, we provide the number of killings per period of time (where the division in periods is identical to that used for fitness transmission). The onset of high fitness transmission occurs precisely at the time

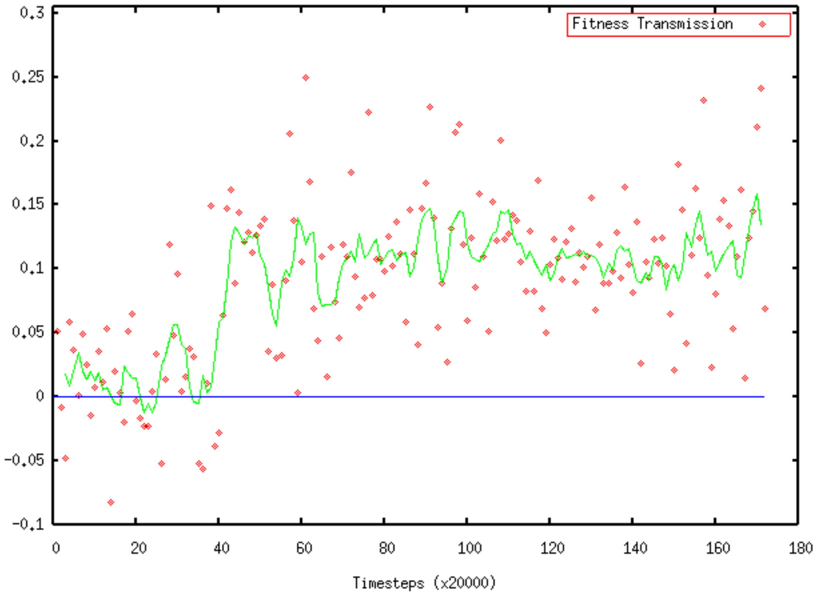

Figure 10.3: Fitness transmission as a function of time, together with a rolling average over 5 values. An initial inactive phase suddenly gives way to a phase of consistently positive activity.

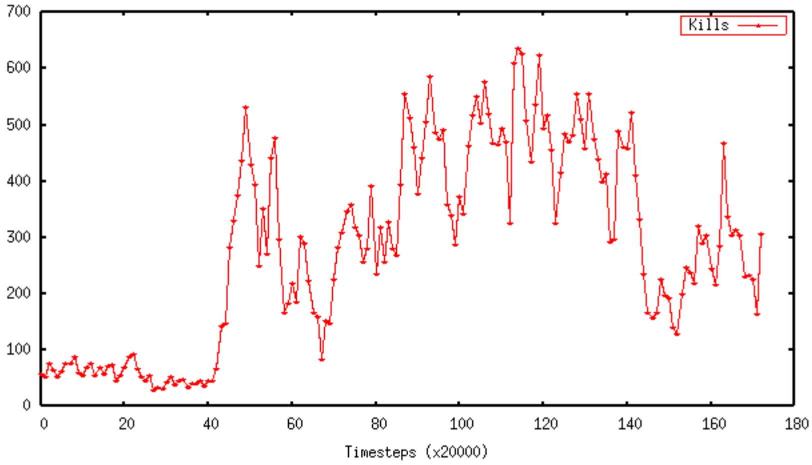

Figure 10.4: Number of kills within the population of the same run as in Figure 10.3. A sudden rise in the number of kills coincides with the onset of positive fitness transmission.

when the number of killings starts to increase significantly, indicating that new capacities are spreading throughout the population.

10.4.3 Panmictic models and diversity

While the previous experiments demonstrated the onset of adaptive evolution, the end result of this evolution tends to be predictable and uniform: large individuals equipped with oscillating limbs, roaming over the sphere and hitting anything they can. Also, while there is variability between runs, each run tends to converge towards one superior morphology, leaving little room for diversity. Thus our system (as it stands) seems to fall in class 1, according to the classification discussed above.

We note, however, that the rules for reproduction and replacement are not really favourable to diversity: we used a panmictic model, in which everyone can mate with everyone. Once a dominant strategy emerges, it can easily take over the population. Thus we ask the following question: is predictability of outcome a fundamental aspect of Evosphere, or is it simply an effect of panmictic reproduction? To answer this question, we eliminate the panmictic aspect of the system by introducing *species*, that is, reproductively isolated populations.

10.5 Coevolution between species

10.5.1 Species implementation

We want to implement species in such a way as to enforce genetic separation, but also leaving evolution as free as possible. In particular, we want species numbers to fluctuate over time, according to their respective success - to the point of going extinct if such is their fate.

To do this, we set a fixed number of species N_s. Each creature belongs to a certain species. If a creature dies (without being killed), it is replaced by a new creature from a randomly picked species. Two parents from this species are randomly selected, and the genetic operators mentioned above are applied. The resulting new creature is included in the same species as its parents. On the other hand, if a creature is killed, it is replaced by a mutated clone of its killer. The new creature (the mutated clone) is included in the same species as the killer. Thus species can increase and decrease through deaths and replacements.

If a species' population falls to 3 individuals, then this species is deemed extinct. Its remaining creature are removed, and the freed slots in the population are attributed to the largest species. The largest species is then *split* into two separate species, of equal size. One of these to daughter groups continues the split species, while the other replaces the extinct species. Thus, as each extinction creates a split, the number of species is kept constant.

10.5.2 Results

As in the panmictic model, the first behaviour to emerge is consistently a “roaming” behaviour: creatures that use oscillating limbs to move fast and hit at close-by creatures. Discovery of these behaviours is the key event that generates active natural selection, as indicated by the fitness transmission graph (compare Figure 10.7 and Figure 10.6).

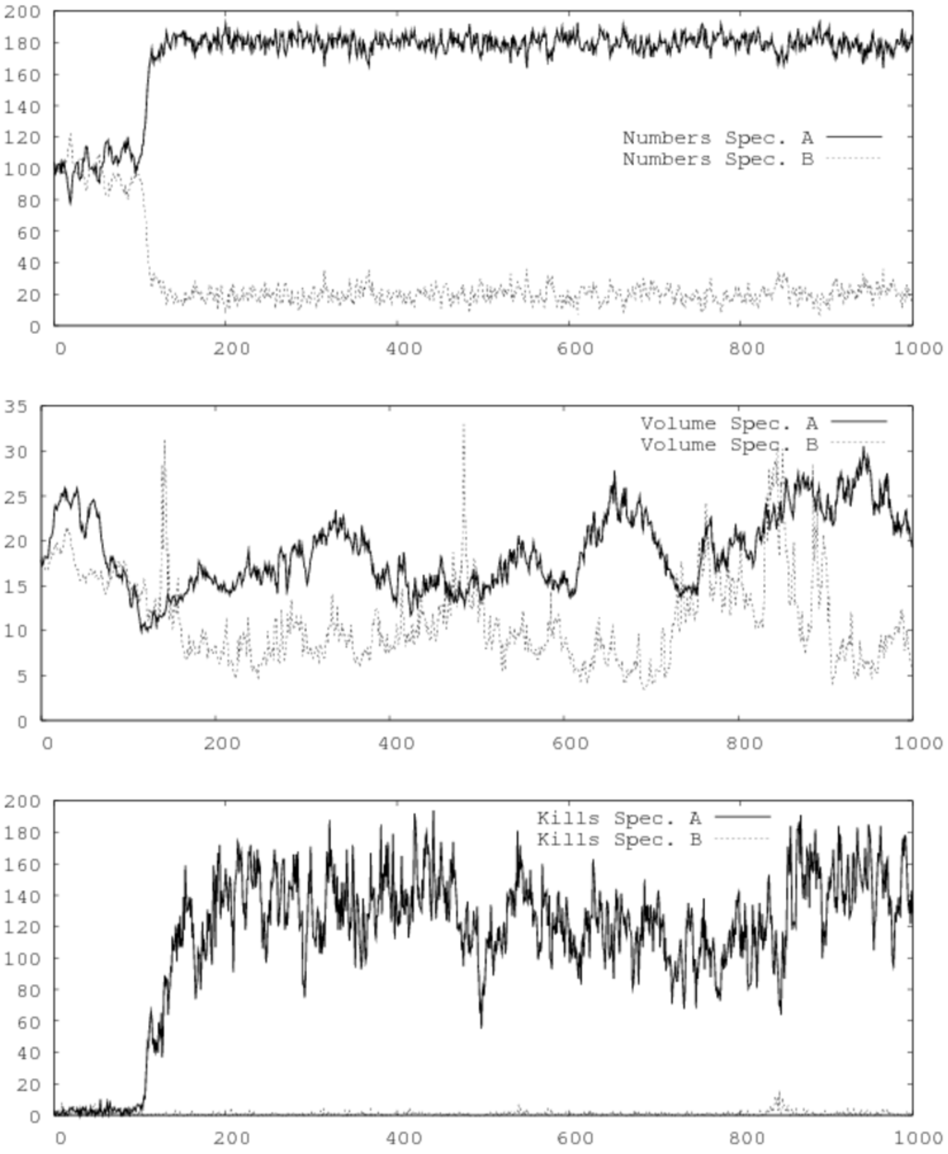

Figure 10.5: From top to bottom: population sizes, average volume, and number of kills inflicted for two species A and B. The x-axis (time dimension) is graded in "cycles", where each cycle corresponds to 5000 timesteps. Notice the initial indistinct period, followed by a division between a larger, active species and a smaller, passive species. Notice the diminution in volume of (active) species A around cycle 500: this is due to the invention of a highly successful "worm" form, which is smaller in volume than previous killer morphologies.

This behaviour then consistently triggers an evolutionary response from other species, namely *miniaturisation*: competing species tend to get smaller in volume. The advantage of miniaturisation is obvious: the smaller you are, the lower the chance of getting hit by roamers. When the number of roamers is high enough, small creatures can obtain a strong survival advantage, which leads to increased reproductive success as long-lived individuals have more opportunities to be chosen as parents.

Figures 10.5-10.6 depict the population size, the number of kills inflicted, and average creature size for runs involving two or three species, as a function of time. In these pictures, the x-axis corresponds to "cycles", where each cycles covers 5000 timesteps. While the dynamics of these runs vary, certain common features a readily seen: after a short initial period, one species reaches an efficient (roaming) behaviour, resulting in a high number of kills and population increase. The other species shrink rapidly, not only in numbers, but also in average size. In the end a steady state emerge, in which one species has specialised as "active" roamers, reproducing mostly through kills, while the others specialise in miniaturisation and reproduce "passively" through replacement of deaths. As indicated in the figures, extinction events may also occur. The "miniature" species seem to have much more variance in volume than the active species, which is expected given their smaller population sizes.

The division between active (killing), large, numerous species and passive (non-killing), small species was consistently observed. Furthermore, the system seems to favour the presence of only one active species. After extinction events, when the largest species (which is invariably the active one) is split into two groups, only one of these groups remains active, while the other is forced into a passive lifestyle.

The behaviour of the active species generally relies on the use of one or more undulating limbs to achieve locomotion. This can be used in several ways. One option is to emphasise locomotion, in order to cover as much ground as possible. Species following this strategy tend to evolve slim morphologies with very efficient motile behaviour. This strategy often culminates in streamlined, high-speed "worms". Another possibility is to emphasise damage dealt on contact. These species develop morphologies that travel less efficiently, but maximise the probability and intensity of damage when colliding with other individuals.

The presence of different strategies, which emerge in response to each other (rather than being immediate adaptations to external parameters of the system) indicate that the system lies in class 2 of our classification.

10.6 Conclusion

We have described an open environment, Evosphere, in which free-living creatures interact freely. Reproduction can occur either by killing and cloning, or simply by staying alive long enough to be chosen as a parent for replacement of dying creatures that were not directly killed.

This environment is successful in generating natural selection and adaptive evolution. This adaptation consists in the discovery and refinement of efficient behaviours. Adaptive evolution is observed both by visual inspection of creatures, and by fitness transmission calculations. The correlation between the onset of morphological evolution and of positive fitness transmission

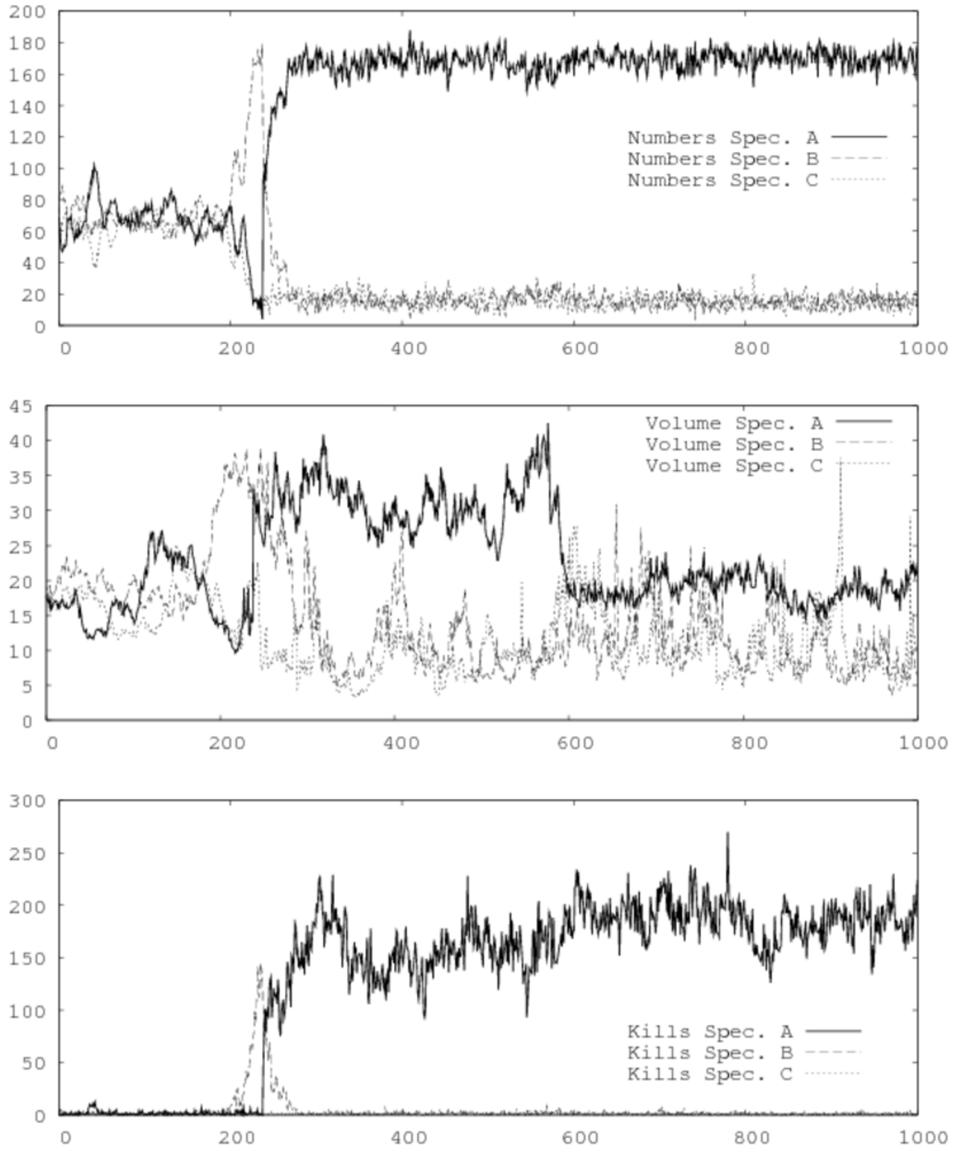

Figure 10.6: From top to bottom: population sizes, average volume, and number of kills inflicted for three species A, B and C. The gradation and general features are similar to those in Figure 10.5. The apparent exchange in positions between species A and B soon after 200 cycles is due to an extinction event: soon after species A discovers a working killing behaviour, species B goes extinct. Species A is split in two and one of the resulting groups is affected to replace species B. This group proves more successful and confines its sister group (still labelled "species A") to a passive, miniature lifestyle.

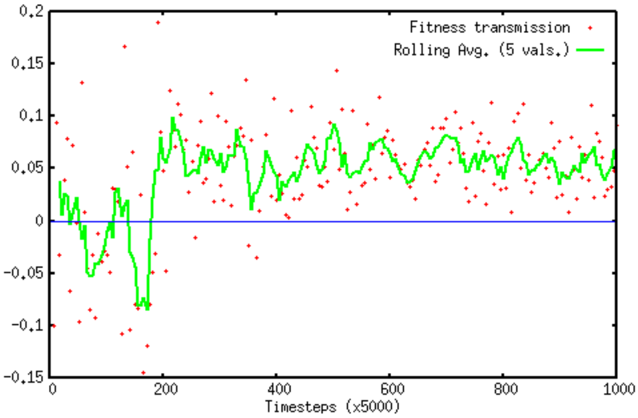

Figure 10.7: Fitness transmission for the run depicted in Figure 10.6. The onset of positive fitness transmission coincides closely with the increases in numbers of kills and population size divergence.

constitutes additional confirmation (on an actual complex system rather than on toy models) that fitness transmission is a usable signature of active natural selection.

Furthermore, when restrictions on mating are imposed through the introduction of species, evolutionary "responses" are observed, in that one species evolves in adaptation to another one's evolutionary changes. Thus evolving species in this system are in principle capable of creating and altering each other's fitness landscape. While the behaviours observed so far remain simple, this is an encouraging indication of this system's potential for future investigations.

Chapter 11

Conclusion: The Road to Everywhere

This work was motivated by a double recognition. First, we acknowledged the prodigious power of evolution in creating complex adaptive designs, and the obvious appeal of using this power for human interests. Second, we took note (as did many authors) of a discrepancy between the expectations placed on artificial evolution, and what could be inferred from the mechanism of evolution itself.

This discrepancy was particularly acute concerning the notion of progress in evolution and coevolution. This applied equally to the concept of progress in an "engineering" sense of performance, or in the more abstract sense of increase in complexity. In both cases, the problem was not so much one of false beliefs or incorrect reasoning, but rather stemmed from confusion about terms and implicit assumptions that were somehow not clearly enunciated.

Accordingly, much of our work consisted in conceptual clarification. By giving full attention to apparently innocuous details, we were able to fend off much unnecessary confusion and tackle certain lingering disputes (e.g. regarding complexity.) Building upon this, we proposed several methods to address the challenges posed by artificial evolution and coevolution. To illustrate these methods, we also produced a complex experimental system, introducing several firsts along the way (first complete replication of Sims' results; first implementation of evolving both morphology and behaviour of 3D creatures for physical combat; first evolutionary environment based on free-living 3D creatures interacting physically.)

The main contributions of this thesis can be summarised as follows:

1. Coevolution, superiority and progress

   In chapter 3 we investigated the notions of performance, superiority and progress in coevolution. We pointed out that the notions of superiority and progress in coevolution were necessarily based on a given superiority criterion *and* on a given set of reference opponents. Depending on the implicit set of opponents considered, several notions of superiority and progress can be identified, namely: local, historical and global. Local superiority implies superiority when comparing respective performances against current opponents; historical superiority implies superiority when comparing performances against all previously encountered opponents; and global superiority implies superiority when comparing performances against all possible opponents. Each of these notions of superiority, when applied between newer and older individuals within an evolutionary process, result in a corresponding notion of progress. We contended that much confusion about coevolution stems from a confusion between local, historical and global progress. Local progress is immediately expected under natural selection; historical progress can be enforced by various algorithmic means, such as the use of an archive; global progress corresponds to the implicit objective of practitioners of coevolution, but cannot be reduced to any of the previous ones.

Comparing our concepts with the recent work of Ficici [53, 54], we found much common ground. However, we pointed out the difficulties posed by monotonic solution concepts, and that practitioners might not necessarily discard intuitive solution concepts (such as best scoring strategy, i.e. global superiority using the average score criterion) and adopt monotonic solution concepts instead.

2. *Evolution and Complexity:* in chapter 8, we addressed the question of whether or not an "arrow of complexity" [12] exists in evolution. We drew from recent evolutionary research (especially from McShea [117, 119]) to avoid common misunderstandings and identify implicit assumptions that may have hampered previous discussion. We argued at length, based on arguments from logic and from evidence, that the clear increase in the maximum of complexity in the course of biological evolution was most adequately explained as a passive trend, whereby a complexity-neutral evolutionary process constantly branches out into many directions and thus mechanically pushes the upper envelope of complexity. Noticing that such a passive trend was in no way trivial (especially considering how much artificial evolution seems to be biased *against* a continuous increase in complexity), we tried to formulate a set of abstract conditions necessary for such a continuous trend to occur.

3. *Analysing coevolution:* in chapter 6 we described several tools and methods to track the performance of coevolutionary algorithms over time, building upon previous work, and making use of our improved understanding of coevolutionary progress. Coarse-grained Master Tournaments (expanding on work by Cliff & Miller [32] and Nolfi & Floreano [133]) were introduced as a method to observe historical progress. Cross-validation of runs was proposed as a way to introduce a component of global progress in the evaluation, suggesting that full cross-validation of many runs (if computationally feasible) would offer a reliable estimation of global progress. Finally, we described how equal-effort comparisons could be performed between coevolutionary runs from various algorithms, in order to compare their relative performance over time.

   In Chapter 7, we illustrated the use of equal-effort comparisons to give a firm answer to a simple question: in the (modified) LEO algorithm, how many interactions per evaluation are optimal, and what is the advantage of using a sliding archive? Our answer was that the advantage of adding more competitions per evaluation plateaued quickly, with a maximum between 3 and 4; and that sliding archives provided a small, but noticeable advantage in the long run, though this was at the cost of a performance hit in the short term.

4. Fitness transmission, a signature of Darwinian evolution: in chapter 9, we introduced

Fitness Transmission as a statistical signature of Darwinian evolution that can be computed from genealogical records. We illustrated the use of this measure on several simple evolutionary experiments, showing that it reliably detected the presence of Darwinian evolution. Perhaps equally importantly, we pointed out restrictions in the usability of Bedau & Packard's evolutionary activity statistics [11] for this purpose.

5. Fighting creatures and Evosphere: our experimental system, as described in chapter 4, was based on an extension of Sims' "virtual creatures" system (apparently the first complete reimplementation and extension of Sims' model). In chapter 5, we expanded our experimental system, by introducing physical combat between creatures. We pointed out the difficulties posed by Newtonian physics, and addressed these by a small non-Newtonian alteration of physical behaviour. To our knowledge, this is the first example of the evolution of actual physical (i.e. contact-based) combat between autonomous 3D agents, where both morphology and behaviour are under full evolutionary control.

   In chapter 10 we built upon this system to create an open environment, called Evosphere, in which free-living creatures interact, fight, and evolve freely, without any explicit fitness function. We observed that this system not only gave rise to actual evolution, but could also accommodate reciprocal evolution between species. While creatures remained conceptually simple, this is an encouraging sign about the potential of this platform for further investigation.

   We wrapped up our discussion by applying fitness transmission to detect the onset of natural selection and adaptive evolution within this virtual world. Fitness transmission measures were in accordance with several other indicators of evolutionary activity, providing further confirmation of the validity of this statistic to detect the presence of Darwinian evolution.

More generally, if this thesis has a central message, it is that the first and foremost outcome of evolution is *diversity*. While "arrows" may emerge on the course of evolution, evolution itself follows no overarching arrow. Evolution does not predictably move along from point A to point B, even less so to some apocalyptic "omega point". Rather, it constantly flourishes into a profusion of branches, opening a spectacular multitude of paths, continuously exploring its immediate vicinity. Every living creature is a trailblazer; every birth is a crossroad. Most branches either die out or immediately fold back into their stem; but some endure, which is to say that they themselves become new stems in their own right, and from them many new crossroads in turn emerge, slowly and patiently projecting the unbounded space of the *possible* onto the almost equally enormous space of the *useful*. Deeply pluralistic, implacable to those who fail, but infinitely tolerant as to the means to succeed, evolution is, in effect, a road to everywhere - or at least, everywhere worth being at.

Appendix A

Details of the software platform

A.1 Creature morphology

In the following sections we provide a broad description of our system, stressing both similarities and differences with Sims' model. In order to facilitate comparisons, our description deliberately follows the same organisation as Sims [163], section by section. Note that many of the basic features described in these sections can also be found in our previous paper [123].

As in Sims' model, the creatures are branching structures composed of rigid 3D blocks. The blocks (or "limbs") are connected to their parent limb by a hinge joint - except for the first, "root" limb which has no parent. The genetic specification of a creature is given as a graph of nodes. Each of these nodes contain morphologic and neural information about one limb. Each node is responsible for storing the description of its limb's physical connection with its parent node's limb, removing the need for connections to carry their own information, as is the case in Sims' model.

The morphologic information in each genetic node specifies the *dimensions* of the limb (i.e. width, length and height), the *orientation* of this limb with regard to its parent (in the form of two parameters indicating polar angles with the xz and the xy planes, i.e. longitude and latitude, in the frame of reference of the parent limb; these two parameters are discrete multiples of $\pi/8$), the *axial direction* of the hinge joint which may be either horizontal or vertical (i.e. aligned either with the y or with the z axis of the limb), and a boolean flag for *reflection* which governs symmetric replication along the xz plane of its parent (see section on Genome Expression). A limb also contains *neural information*, as described in the next section.

A.2 Creature control and neural organisation

Our creatures are controlled by neural networks. As in Sims' model, each limb contains a set of neurons. Genetic information about a given neuron specifies the *activation function* for this neuron, a threshold/bias parameter θ taken in the $[-1, 1]$ range, and connection information. The activation function may be either a sigmoid ($\frac{1}{1+\exp^{-(\sigma+\theta)}}$) or the hyperbolic tangent $\tanh(\sigma + \theta)$ where σ is the weighted sum of inputs (the difference between sigmoid and tanh is that the first has values in $[0, 1]$ while the latter has values in $[-1, 1]$). Connection information specifies, for each connection, the source of this connection (i.e. the neuron whose output is received through

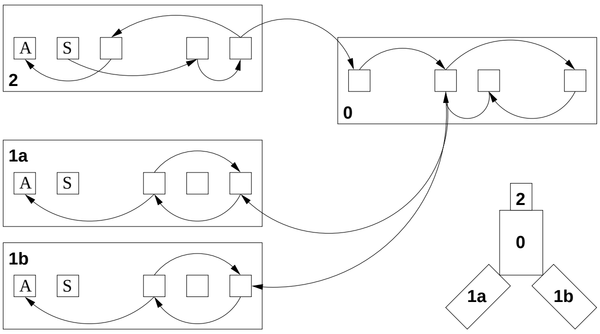

Figure A.1: Organisation of a fictional creature pictured in the bottom-right corner. Limb 0 has no sensor (S) or actuator (A). Limb 1 is reflected into two symmetric limbs 1a and 1b, which share the same morphologic and neural information.

this connection) and a weight in the $[-1, 1]$ range. Neurons can only be connected with other neurons from adjacent limbs, or from the root limb. Each neuron may receive a variable number of connections, up to a maximum value (3 in the present experiments).

The most important difference with Sims' model lies in the choice of standard neurons with traditional activation functions, in contrast to Sims' large set of functions (including arithmetic operations and oscillators). An important consequence of this simpler set of functions is that there is no trivial way for evolution to generate oscillators or other cyclic forms of behaviour, which are necessary for any sustained locomotion to take place. Such behaviours have to emerge out of the interaction between several neurons, assembled together under the guidance of evolution. A more practical consequence is that in our model, each neuron may have an arbitrary number of inputs (up to a maximum value), by contrast to Sims' neurons which had a fixed number of inputs, dependent on their function.

A.2.1 Sensors and actuators

Sensor neurons and actuator neurons are handled specially. In these experiments two types of sensors are used: proprioceptors and external sensors. Proprioceptive neurons measure the current angle formed by the hinge joint to which this neuron's limb is attached, scaled within the $[-1, 1]$ range. External sensors come in two types, which measure the x-distance of the limb containing the sensor to either the trunk of an opponent, or to an inert box. An x-distance is the distance between the centre of mass of the limb and the centre of mass of the object, along the x-axis of the frame of reference of the limb in which the sensor exists (i.e. the x component, in the frame of reference of this limb, of the vector joining the centre of this limb to that of the detected object). The outputs of external sensors are squashed through a $tanh$ function. Actuator neurons command the movement of limbs, that is, the desired angular velocity around their joint. Their inputs are defined similarly as other neurons, but their activation function is always a scaled hyperbolic tangent of the form $MaxSpeed * \tanh(\sigma)$, where $MaxSpeed$ is a system constant. In our model, sensor neurons do not receive any connection from any other neuron, and no neuron may receive a connection from an actuator neuron. Sensors and actuators are, respectively, pure sources and pure sinks of data.

A difference with Sims' model is that an actuator does not specify a force or a torque, but a *desired angular velocity.* The physics simulator implements a motor at each joint, which will constantly attempt to reach the desired speed, with the constraint that the total torque it exerts cannot be larger than a specified maximum. This maximum is a system constant. This mechanism corresponds to a very simple model of servomotors.

Each limb has exactly one actuator and one proprioceptor. It may have other neurons, including external sensors, within a maximum number (in the current experiments the maximum number of neurons for each limb in addition to the actuator and the proprioceptor is 2). Note that while each limb has a sensor and an actuator, there is no requirement that they should receive or send connections from or to other neurons: connections are established in a random manner and no connection toward these special neurons is explicitly enforced. Thus each limb is free to use its sensor and actuator, or not, depending on how its network evolves. This is equivalent to Sims' model.

A.3 Expression of the Genome: the Developmental System

When a creature is to be generated from its genotype, a simple developmental system translates the genotype into a corresponding phenotype, and may introduce additional complexity if the genetic information dictates it. Our system uses developmental features similar to those introduced by Sims (bilateral symmetry and segmental replication). To fully exploit these features, we introduce control flags which enforce fine-grained control of the neural connections in replicated limbs. We also introduce a new mutation operator, recursion unrolling, which allows developmental replications to be transcribed back into the genome.

A developmental system corresponds to the introduction of a bias in the search space: while the total number of different individuals accessible with a given N-bits genome is obviously bounded by 2^N in any case, using a developmental system will allow the system to obtain different (usually more complex) creatures, at the expense of making others impossible. Quite often raw, declarative genomes describe a compact search space of very similar creatures, and the developmental system has the effect of "stretching" the search space, creating a globally wider, but sparser, range of possibilities. Developmental systems must of course be chosen with care so that the newly obtainable individuals can be expected to perform globally better than the "sacrificed" ones. A common feature of developmental systems is a stress on modularity and replication, with the expectation that non-linear interactions between similar elements may bring interesting behaviours with minimal additional information; Hornby and Pollack [83], for example, provide noticeable examples of modularity introduced by a developmental system[1].

A.3.1 Reflection

Symmetry in our model is implemented somewhat differently than in Sims'. In our model, each genetic node (corresponding to a limb) may possess a "reflection" flag, which means that when this node is read and the corresponding limb attached to its parent, a symmetric copy of this limb will also be created. Any further sub-limbs will similarly be duplicated in a symmetric

---

[1] Developmental systems in artificial evolution have been used for a relatively long time, but rigorous analysis of their effects and implications has only recently started to attract interest. Stanley and Miikkulainen [172] provide an enlightening review and discussion of developmental systems in artificial evolution.

fashion, which leads to the appearance of bilaterally symmetric branches. Our present design allows for only one type of symmetry, namely symmetry along the parent's xz plane. When a given limb is randomly generated, its reflection flag is set with probability P_{ref}.

Symmetric replication introduces information flow issues. When a limb is duplicated by reflection, all genetic information is duplicated in the process, including neural information. A consequence of this duplication is that *a given limb cannot distinguish information it sends to, or receives from, either of its symmetric sub-limbs.* Because neurons from both symmetric sub-limbs share the same connection information, they will receive identical connections (and information) from the same neurons in the parent. Similarly, any connection that the sub-limbs send to the parent will point to the same neuron in the parent, and information from both sub-limbs will be merged at that point. Thus, although both limbs may behave in different manners due to their separate inner neural networks (which may react independently to different sensor information), they will not be able to send distinct information to the parent, or to receive distinct information from it. Sims does not mention this problem, or document his solution to it, in his papers.

We address this problem in the following way: every connection has a special $Reftype$ flag, which can take one of three values: Original, Symmetric, or Both. When a connection (as specified by the genome) originates from a neuron that exists in a reflected limb, then the actual connection in the resulting creature will be connected either to the original version of the limb, or to its symmetric copy, or to both, depending on the value of its $Reftype$ flag. If $Reftype$ is 'Both', then this connection will carry the average of the outputs of the two neurons.

Similarly, if neurons from the two instances of a reflected limb carry a connection originating from their common ancestral limb, the $Reftype$ flag is used to determine the actual wiring, that is, whether only the original instance, or the symmetric copy, or both, will receive input from the parent limb.

A.3.2 Segmentation

In Sims' model, a loop in the genetic graph corresponds to a set of limbs which is repeated a certain number of times. Connections between limbs specify a recursive limit, which is the maximum number of times this connection should be followed when in a recursive cycle. Some connections may also be marked as "terminal", meaning that they will only be applied when the recursive limit is reached, thus allowing for specific "tailing" structures at the end of repeated sequences. This, in essence, is a simple and effective model of *segmentation*, that is, the repetition of homologous modules arranged sequentially, as apparent in many animals (vertebrates, arthropods, anellidae, etc.).

We import this feature in our model, with the restriction that only self-loops are allowed: a loop can only exist between a node and itself. No other loops within the graph can exist. This allows for bio-inspired segmentation (repetition of similar segments), while preventing the appearance of extravagant body plans (such as, say, a human body in which the thumb would contain a "loop" to the thorax).

Segmentation brings in the same information flow issues as were discussed above about symmetry. Imagine that a certain node has a recursive loop to itself, inducing its replication into similar segments. How can we allow for communication between segments ? In the genome, the information about a connection specifies the genetic node (and the neuron within this node) from which this connection originates. But when a connection refers to the same node as the one within which it exists, and the node is recursively replicated, we must decide which instance of the node is actually referred to (so that connections could occur within the same limb, between one limb and its recursive predecessor, or between one limb and its recursive successor).

We address this issue in the same manner as with symmetry. Each connection also carries a *Rectype* flag, commanding its behaviour under recursion, which can take any of three values: Dad, Son, or Same. If a connection for a given neuron originates from the same genetic node as that in which the neuron exists, and this neuron is recursively duplicated, then the value of *Rectype* determine the actual wiring of this connection: the value 'Dad' indicates that this connection should go from one instance to its predecessor in the recursion (so obviously it will not exist in the first, 'original' instance); the value 'Son' indicates that this connection should go from one instance to its successor in the recursion (so it will not exist in the last, 'terminal' instance). The value 'Same' indicates that this connection should be understood to originate from the same instance and will therefore be present in all instances.

A.3.3 Recursion Unrolling

A common source of novelty in Nature is the duplication-exaptation process: one part is duplicated in two, originally identical elements, then the features of these elements diverge and assume different roles. The versatility of arthropod appendages is a striking example of these mechanisms. Our system allows for similar duplication-exaptation patterns indirectly, through a mutation operator called recursion unrolling: a recursive cycle (one genetic node which specifies its own recursive replication during development) is *unrolled* in the genome, that is, it is developed as it would be in a final body, and the resulting limbs are transcribed back into as many new, independent genetic nodes. Originally these new nodes are almost identical (except for neural connections within themselves and between each other, which depend on the flags described in the previous section), but from now on they can evolve independently. This method may be in opposition with the central dogma of genetics (information flows from the genotype to the phenotype), but is not Lamarckian: it is a macro-mutation which occurs randomly, without using selective information, thus respecting the Darwinian mechanism of 'blind' mutations. This mechanism is not present in Sims' system.

A.4 Creature evolution

A.4.1 Genetic operators

We use three genetic operators, similar to those used by Sims, plus one addition.

Crossover is performed by simply aligning the genetic nodes of both parents in two rows, then building a new list of genetic nodes by concatenating the left part of one parent with the right part of the other.

Grafting corresponds to the removal of a branch (i.e. a limb and all its sub-limbs), and its replacement by a branch taken from another individual. Connectivity information is adapted and maintained: the neurons of the trunk establish the same connections with the new branch as they had with the old one, and similarly the new branch has the same connection with its new trunk as it had with its previous trunk.

Picking corresponds to simply taking a branch from a given individual and copying it in the genome of another individual, without removing any material. This is our only addition to Sims' genetic operators, but it seems to have an impact on the continuous appearance of complex features.

Mutation occurs by modifying each parameter of the genome with a certain probability P_{mut}. The mutation operator proceeds in a sequential manner. First, with probability P_{mut}, a random limb may be deleted from the genome (with the restriction that no creature may have less than two limbs). Then a new randomly generated limb (with randomly generated neural information) may be created. Then, a recursive cycle may be unrolled (if there is a recursion in the genome). Then each slot in the neural array may be "flipped" (i.e. empty slots are filled with a new neuron with randomly assigned connections, existing neurons are deleted). Each interneuron may be turned into a sensor, and vice versa (proprioceptors and actuators are fixed for all creatures). Each sensor neuron can change types (i.e. from sensing the box to sensing the opponent and vice versa). Then, the threshold value of each existing neuron may be modified through Gaussian perturbation (standard deviation 0.4) within the $[-1, 1]$ torus. The output function may be changed. Each connection of every existing neuron may be "flipped", i.e. created (and randomly assigned) if it is unassigned, or deleted otherwise. The weight of each existing connection may then be modified through Gaussian perturbation (standard deviation 0.4) within the $[-1, 1]$ torus. The source of each connection (i.e. the neuron from which it originates) may be randomly reassigned. Finally, with probability P_{mut}, morphological information for each node is mutated. Morphological mutation performs one randomly selected operation out of seven possibilities: reassigning a given limb to a different "ancestor" limb (which amounts to moving a whole branch along the organism), randomly assigning a new length, width or height to the limb, modifying either of its orientation angles (possible orientations are discrete multiples of $\pi/4$; mutation occurs by choosing a new value within the range $[-\pi/2, \pi/2]$ around the current value), switching the orientation of its joint (horizontal or vertical), and flipping its "reflection" flag. Again, each of these modifications is applied with probability P_{mut} for each parameter.

A.5 Stability of physics simulation

The basic purpose of a physics simulator is to integrate the differential equations of Newtonian mechanics. Occasionally the system will find itself in a situation when integration errors will accumulate rather than vanish, leading to a dramatic disruption of realistic behaviour (bodies gaining energy, unbounded acceleration, explosion of joints, etc.) As noticed by many authors, including Sims [163], evolution will gladly exploit any such non-physical behaviour to its profit. More recently Chaumont and colleagues [29] describe how the inaccuracies of the simulator (the same as ours) require extensive tweaking of the simulation and fitness function.

However, after a lengthy trial-and-error process incurring much parameter-tuning, code-tweaking, and teeth-gnashing, we eventually stumbled upon non-obvious values for certain simulation parameters that result into satisfying stability. Since then, we did not observe any troublesome case of non-physical behaviour exploited by evolution to its advantage: the system is consistently stable and realistic. The values of these parameters (namely the CFM and ERP parameters of the ODE simulator) are given in Table A.1.

A.6 Simulation parameters

Any complex simulation involves many parameters, some of which have a critical impact on the results of the simulation. Table A.1 summarises the values we chose for various parameters of our program. Most of these parameters were educated guesses which proved satisfactory. In particular, the scaling factors shown in activation function result from an attempt at obtaining well-behaved function shapes (e.g. with regard to steepness of the non-linearity), considering that outputs and weights were all within the $[-1;1]$ range. Some of these parameters, however, (especially the parameters related to the ODE physics engine) were arrived at through a long process of parameter tuning. This is particularly the case for the ERP and CFM parameters. We hope that these figures may prove useful for other researchers, especially those interested in using the ODE simulator.

Simulation constants	
Max. number of neurons per limb	6
Max. number of connections per neuron	3
Initial proportion of assigned connections	75 %
Max. number of genes	8
Max. number of limbs (after development)	17
Joint angle range	$[-3\pi/4, 3\pi/4]$
Range of connection weights	$[-1, 1]$
Range of neuron thresholds	$[-3/2, 3/2]$
P_{ref} (see text, section A.3)	0.1-0.25
P_{mut} (see text, section A.4.1)	0.02-0.06
Max. force applicable by each motor	4 N
Max. angular speed at each joint	4 rad / s
Actual mass of a limb of dimensions (x, y, z)	$0.8 * \tanh(x * y * z)$
Physics engine parameters	
Step size	0.01 s
ERP	0.015
CFM	0.01

Table A.1: Simulation constants

Appendix B

Artificial evolution - from bits to worlds

> We are like dwarfs sitting on the shoulders of giants, so that we can see more than they, and things at a greater distance, not by virtue of any sharpness of sight on our part, or any physical distinction, but because we are carried high and raised up by their giant size.
>
> Bernard of Chartres, quoted by John of Salisbury, *Metalogicon*

B.1 Evolving software objects

Evolutionary computation, that is, the idea that software objects could undergo evolution just as biological objects do, is not a recent idea. Von Neumann's self-replicating organism, based on a cellular automaton, is the first explicit instance of this concept. What differentiates this complex organism from more trivial self-replicating automata (such as a 2-state automaton that would simply propagate a given state) is precisely the capacity to evolve. In Von Neumann's view, the difference between trivial and non-trivial replication was the capacity to produce an offspring more complex than its parent : a non-trivial reproducing machine should be able to construct any arbitrarily complex machine, and self-replication is only a particular case of this ability. This property is what makes evolution possible, and what ultimately differentiates trivial, non-living reproduction - which occurs frequently in mineral objects - from actual life (see [114] for an enlightening discussion of this fact, and how it was somewhat overlooked by subsequent efforts in self-replicating automata).

Besides the works of Von Neumann, the idea of using artificial evolution to solve actual problems can be traced back to the 50s. A rich and fascinating account of early experiments in evolutionary computation has been compiled by [59].

B.1.1 Genetic Algorithms

We begin our survey with the particular model that has more or less become the standard model of evolutionary computation, namely the Genetic Algorithm introduced by Holland [82] (a much more readable introduction is provided by Mitchell [128]).

The genetic algorithm operates on strings of bits that can be evaluated by a given "fitness function". From the viewpoint of the algorithm, this fitness function is the complete, self-contained definition of the problem. Starting with random strings, each string is evaluated according to this function. Some of the fitter strings are then selected for survival and reproduction. Genetic operators such as crossover (recombination of two parents into one offspring) and mutation (random modification of some bits in a string) are applied. Then the cycle starts again with the new population.

The genetic algorithm was more than a simple engineering trick. It came with a thorough mathematical analysis which broke with several preconceptions of evolutionary dynamics, mostly by asserting the preeminence of recombination over mutation. In the genetic algorithm, or at least in the original analysis found in [82] and [65], crossover is the major evolutionary force, and mutation is regarded as a secondary (although essential) tool whose main role is to maintain some degree of diversity within the population. This affirmation is formally expressed by the Schema Theorem, which relies on the notion of schemas.

A schema is a sequence of characters which can be either 1, 0 or a "don't care" symbol (#). Thus the schema 100# represents both 1001 and 1000. 1001 and 1000 are two instances of the 100# schema. What the Schema theorem says is that at each timestep, the number of instances of a given schema increases geometrically with the relative fitness of its current instances ("relative" with regard to the average fitness of the population). In other words, schemas which appear to provide above-average fitness will grow exponentially in the population (minus two terms that account for the disruption of a schema by mutation and recombination). [82, 65].

The Schema theorem provides insight into the mechanism of evolution under recombination. It shows that the genetic algorithm exponentially allocates more tries to schemas which, according to their observed instances, seem to be beneficial. "Evolution mixes things up, and things that look good get mixed (geometrically) more often" - this is the message of the Schema theorem[1]. Holland was able to show that such a policy was, in a sense, optimal in terms of profit maximisation.

From this analysis, it follows that crossover-based artificial evolution, just like its natural counterpart, is not a global optimiser. It is a rather greedy algorithm, that constantly tries to perform a very specific kind of local optimisation: by exploiting current information in order to maximise the expected profit at each step, the overall result is that genetic algorithm really optimise the expectancy of *accumulated* fitness over the whole history of the process. This was precisely the objective of Holland, who was more interested in lifelong adaptation than in global function optimisation.

However it must be kept in mind that evolution does not physically operates on abstract schemas : it operates on individuals, based on full genotypes. In full genotypes, several schemas are present and interfere with each other. Together with the fact that populations are finite, this introduces a strong sampling effect on the evaluation of any schema: the *observed* fitness of a schema (as observed over all its currently available instances) may differ strongly from its *real*, "static" fitness (as defined over the whole set of all possible instances). Of course, this is

---

[1] Note that the Schema theorem really provides a lower bound : it does not describe how new schemas are created by the system, only how these schemas are exploited

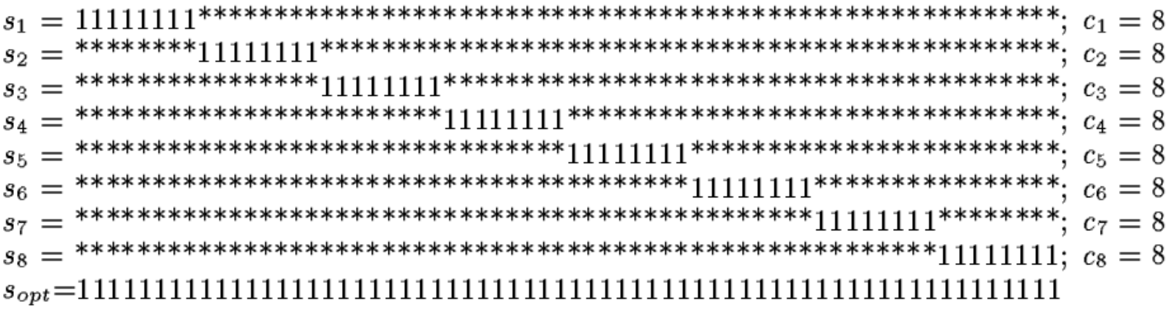

$s_1 =$ 11111111**; $c_1 = 8$
$s_2 =$ ********11111111**; $c_2 = 8$
$s_3 =$ ****************11111111**; $c_3 = 8$
$s_4 =$ ************************11111111********************************; $c_4 = 8$
$s_5 =$ ********************************11111111************************; $c_5 = 8$
$s_6 =$ **11111111****************; $c_6 = 8$
$s_7 =$ **11111111********; $c_7 = 8$
$s_8 =$ **11111111; $c_8 = 8$
s_{opt}=1111111111111111111111111111111111111111111111111111111111111111

Figure B.1: The R1 function, one of the Royal Road functions described in [62]. The fitness of a string is the total sum of the fitness of all building blocks ($s1$ to $s8$) that are represented in this string. For example, R1(1111111100..0) = 8, and R1(1111111100..0011111111) = 16. The maximum possible fitness is achieved by the string s_{opt} = 1111. . . 1 (image reproduced from [62]).

even more true if the fitness of a schema has a strong variance over all its particular instances. Another problem is that of collateral convergence: instances of a good schema may quickly swamp the population, thus biasing the estimation of other schemas. The problems of high variance and collateral convergence have been described and explained in [74].

The latter problem has been made more evident by the Royal Road functions (see figure B.1) introduced by Forrest and Mitchell [62]. At first sight, these functions seem to be well-suited for GA optimisation, because explicit building blocks exist, and all of them contribute to reaching the global optimum. However, the GA was found to perform poorly on these functions: as soon as strings containing a superior schema were discovered, they proliferated and virtually eliminated their competitors, even though these competitors contained other useful schemas which were necessary to obtain a global optimum. This led the algorithm to premature convergence toward suboptimal solutions. This phenomenon of premature convergence, common in greedy algorithms, was described as a *hitch-hiking* phenomenon: “bad” components were “hitch-hiking” on good ones, and relied on the performance of these good components to propagate themselves through the population, eliminating better alleles that belonged to less-favoured individuals.

The problem of fitness variance within a schema is directly linked to the problem of *epistasis*, that is, the fact that the effect of a given gene may depend significantly on other genes. For almost any non-trivial problem (and even for quite a lot of trivial ones), some degree of epistatic interaction between genes is expected. As an example, in a function optimisation process, having a 1 or a 0 in a given locus may not bring the same benefit depending on the value of the rest of the string (e.g. if the current string is just below the optimum, changing the lowest-order bit from 0 to 1 will increase fitness; but if it is just over the optimum, the same transformation will decrease it). However, by averaging schema evaluation over many instances, the genetic algorithm manages to overcome most situations of moderate epistasis.

But when epistasis comes to dominate the evaluation function, it becomes impossible to evaluate the fitness contribution of a given schema; in fact, in such a situation, the notion of schema fitness becomes meaningless. The presence or the absence of a given schema may bring no recognisable information in terms of fitness contribution, because epistatic interactions garble any correlation that might be observed : completely epistatic fitness landscapes are essentially random. In such a situation, the genetic algorithm, being constantly misled, has very limited chances of success (Kauffman [95] provides a systematic discussion of “tunable” epistatic fitness landscapes known as the NK-landscapes, described by genomes of N genes in which the fitness contribution of a given gene is randomly modified by the value of K other genes; Kauffman showed that in such landscapes, as soon as $K > 2$, the fitness landscape becomes highly irregular).

On the other hand, it should be noted that no search algorithm can cope much better with such conditions. The “No Free Lunch” theorems for searching algorithms [193], which state that no search algorithm performs better than any other over the set of all possible functions in a given space (including random exhaustive search), is precisely based on the fact that the vast majority of possible functions are uncorrelated. The No Free Lunch theorems simply formalise the intuitive idea that a searching method can only perform better than random search by exploiting certain kinds of regularities. The Schema theorem indicates that recombination-based evolution exploits regularities based on schemas.

This discussion does not diminish the fact that genetic algorithms have been successfully used as optimisers, and that their domain of application can still grow (e.g. [30], or a panorama in [128]). Their attractiveness owes much to their conceptual simplicity and to their agnosticism regarding domain knowledge : as long as the problem can be expressed in a fitness function, the genetic algorithm can be applied with reasonable chances of success. Then again, this means that the design of the fitness function (and of the bitstring encoding method) becomes the real problem to solve for human engineers.

B.1.2 Genetic Programming

Genetic programming [99] essentially uses the same ideas as the genetic algorithm, but generates programs instead of simple bit strings. More precisely, it usually operates on tree-shaped programs written in a functional language. This idea has two immediate advantages : First, it eliminates issues related to the necessarily variable length of individuals (programs of fixed length are usually not suitable). Second, by their very structure, functional programs are naturally decomposed into autonomous elements; branches of program trees can be swapped as whole, self-contained functions.

This immediate modularity would be more difficult to obtain by crossing sequential, imperative programs. Of course, one can note that in nature the DNA program is purely linear, but the problem is different : in nature, crossover usually occurs among very similar programs. Furthermore, the natural “computer” through which DNA is interpreted has little in common with artificial computers. In artificial evolution, GP seems to have no major competitor for the generation of programs.

An interesting side-effect of free code growth in GP is the fact that over time, average program size tends to grow almost continuously throughout the evolutionary process. This phenomenon is known as “bloat”, and is seen as an annoyance which must be fought, usually by adding a size penalty in the fitness function. This bloat can be seen as a natural consequence of the fact that while code size has a lower bound (i.e. the empty tree), it has no upper bound; more precisely, for most problems, given any program of size N, there is an infinity of programs of size superior to N which have exactly the same effect, and thus the same fitness. [102][2]

B.1.3 Evolutionary Strategies

Other evolutionary methods simply do not rely on crossover at all. Indeed, historically, mutation-based methods were the first algorithms to be thought of. This has long been the case of evolutionary strategies [159] although recombination operators have since been included in these as well. In contrast to genetic algorithm, evolutionary strategies operate directly on vectors of real values, and create new individuals by mutating some or all values within the ancestor by a random (usually Gaussian) amount. The variance of this mutation is actually a genetic parameter, so each individual has its own mutation rate. Thus mutation rates are expected to evolve together with the population.

Classical ES use either the $(M + L)$ method (populations consist of M individuals; after L offspring are created, the M best of the $L + M$ individuals available will be chosen for the next population) or the (M, L) scheme (only the L offspring are considered for choosing the M parents within the next population - of course, $L > M$). If one includes recombination, which may come in many flavours too (especially if it applies to parameters of the algorithm, such as variances, etc.), evolutionary strategies offer a wide range of possibilities for evolving real-valued parameters.

B.2 Neural Networks

The idea to create an artificial intelligence by directly simulating cerebral processes actually predates digital computers : McCulloch & Pitts' model of an artificial neuron was published in 1943. This "formal neuron" is essentially a 0/1 threshold function operating on a linear combination of 0/1 inputs : inputs are combined into a weighted sum, and if that sum exceeds a threshold, the neuron's output is 1 - otherwise it is 0. This model is of course an extreme simplification of real neurons, if only because no time considerations are included (real neurons seem to communicate by frequency of impulses rather than by amplitude of signal). Yet this very simplicity made it extremely attractive for simulation processes. A readable introduction to neural networks (centred on a pattern recognition viewpoint) is provided by Bishop. [13]

The first neural network was actually not really a network : Rosenblatt's Perceptron was simply a collection of McCulloch & Pitts neurons. For a classification problem involving N possible categories, the network consisted of N such formal neurons, each of them assessing whether the input fell into the category it represented. However, in their book "Perceptrons" [127], Minsky and Papert, describing the system and its underlying theory in (rather excruciating) detail, exposed important limitations. In particular, they showed that because a perceptron can only classify linearly separable problem (a linear equation of variables is the equation of a

---

[2] This idea can be put in parallel with the controversy about the inevitability of complexification in natural evolution [71, 194].

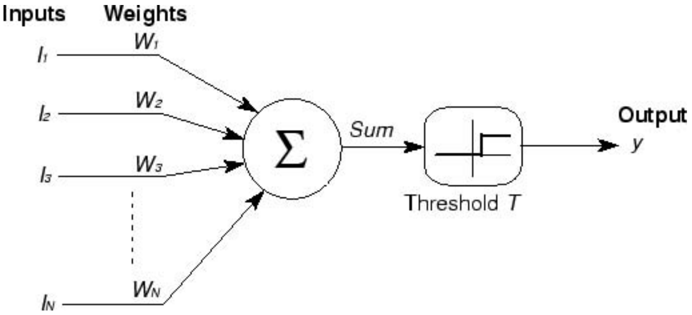

Figure B.2: The formal neuron described by McCulloch and Pitts. Input values are multiplied by weights and summed. A threshold on the resulting value provides the output of the neuron.

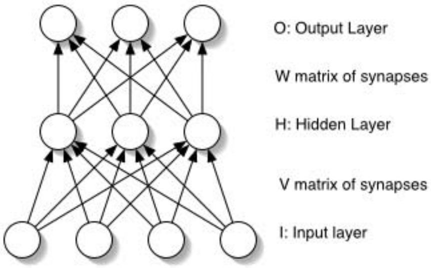

Figure B.3: The multi-layer perceptron (MLP) model. Note that there may be more than one intermediary ("hidden") layer.

hyperplane), it cannot solve the XOR problem or the parity problem. They considered the possibility of multi-layer perceptrons, but did not seem to think that they would bring any major advantage.

Proving that perceptrons are simply unable to solve even the XOR problem had a devastating effect on neural network research altogether. Neural networks only re-emerged when a practical way to use multi-layer perceptrons (MLP) was found. Contrarily to their single-layer predecessors, multi-layer perceptrons are not limited to linearly separable problems. In fact, neural networks with one or more hidden layer are universal function approximators. [176]

The most widely used algorithm for MLP learning is the gradient backpropagation method. Like Rosenblatt's method, it essentially follows the idea of slightly modifying weights that contribute to incorrect behaviour, in a direction inversely proportional to their perceived error. The idea is to calculate, for every training pattern and every weight $w_{i,j}$ in the network, the quantity dE / d$w_{i,j}$, where E is a quadratic error. The method was apparently introduced independently by several authors, but is mostly associated with Rumelhart and McClelland [156].

However, this algorithm can oscillate and even diverge if the error-correction rate is too large. On the other hand, if this rate is too low, learning will be painfully slow. In rugged error landscapes, choosing the right parameter may be a problem. Some adaptive methods that modify the learning rate can be used. Other problems include the fact that backpropagation may get stuck into local minima of the error function and that it is at any rate a long learning process, which requires many presentations of the training set. Another possible inconvenience is that it is usually quite difficult to understand a neural network, to extract a model of the data from it, as opposed to, say, decision trees.

B.3 Evolving controllers for autonomous agents

Evolutionary methods seem well-suited to the generation of controllers for autonomous agents[3]: after all, biological agents were produced by an evolutionary algorithm. It is clear that for such a task, the core of the problem lies not only in the choice of a particular evolutionary method, but also in the choice of an adequate representation for the controller.

B.3.1 Control Pattern Generators

A simple way to control a limb in an agent's body is to specify an activation pattern in the form of oscillating values. These values are then applied to the actuators that they control, resulting in oscillating movements. The role of the controller is to provide an adequate pattern with regard to the task at hand (usually a locomotive task). This method, which may seem a bit artificial at first glance, is in fact deeply rooted in animal neurobiology, from which the term Central Pattern Generator (CPG) originates.

For example, Ijspeert and colleagues [88] have used biological models to evolve such CPG for a swimming lamprey. The lamprey is composed of many segments, each containing a few intercoupled neurons. These neurons can produce an individual oscillation pattern in a single segment. Interconnected segments are capable of synchronising their oscillation in a coordinated wave movement. Proprioceptive information (i.e. input from the body state) allows the lamprey to maintain its movement in the face of a disturbed environment (contrary currents, etc.) The authors successfully evolved neural connections for an artificial lamprey. Using the same framework, the same author took a bold evolutionary leap over the whole Pisces group, and stepped directly from the lamprey to the salamander; while keeping lamprey-like swimming abilities in water, they evolved controllers for quadruped locomotion on land. [87]

Complex oscillating patterns may also be generated more directly by simply specifying a set of Fourier components. In such a model, evolving the CPG amounts to evolving sets of numeric parameters for a Fourier series which describes a periodic activation pattern. A bias ensures that low-frequency components are favoured. This elegant approach, described by Arnold [7], was applied to simple locomotive tasks for 3D articulated organisms (quite similar to Sims [163], see below) and generated a wide range of interesting behaviours, from clumsy creeping to frog-like leaps-and-bounds.

B.3.2 Classifier Systems

Classifier systems were introduced by Holland [82, 81] as a framework to apply the adaptive properties of the genetic algorithm to the production of autonomous behaviours. The Schema theorem emphasised the adaptive aspect of the GA, especially with regard to the maximisation of *cumulated* reward, which corresponds to what one might expect from a living organism

---

[3] Here the word "agent" is to be taken in the most general sense of "something that acts", which corresponds to its etymology. It only implies that the considered agent must perform some kind of action in a given environment.

throughout its life. Therefore it seemed natural enough to exploit these properties in an adaptive system, instead of simple function optimisation.

A classifier system is essentially a set of rules, each rule being composed of a condition part and an action part. They are encoded in binary strings, with the addition of a special "don't care" character (#). External conditions (the environment) are described by binary strings. The internal state of the agent is a set of "messages" (also encoded as binary strings) resulting from previous rule executions. When the condition part of a given rule matches (parts of) this input, the action part of this rule is executed. The action part of a rule may lead to an effective action from the animat, but it may also trigger the release of a "message", which is added to the message list; thus a rule may influence the execution of other rules.

The rules constantly evolve under a genetic algorithm, which generate new rules and eliminates old, unsuccessful ones. The fitness of a given rule is a measure of its expected pay-off, as estimated from the correlation between rewards and the activity pattern of this rule. When the agent receives a reward, the expected pay-off of rules that have been activated are modified. Equipped with such a controller, an adaptive system is expected to adapt to a given environment and to learn from its interactions with it.

The first problem is to know how each rule should be rewarded when chains of rules lead to the obtention of a reward. Clearly if we only reward rules that were active when the reward was received, we will not adequately estimate the value of rules that were activated in the past and played a role in the obtention of this reward. Such a system would not be able to make sequential plans (first do this, then do that, and finally do this): only rules that lead to the immediate acquisition of a reward would be elected by the system. Obviously such a system would easily be confused in any moderately complex environment. This is a typical *credit sharing* problem, as occurs frequently in collective action models.

The first model maintained quite an elaborate set of statistics for each rule, mainly concerned with the frequency of use and the delay between application and reward obtention. Expected pay-off were only modified when an external reward was issued, and this modification occurred for each rule that had been active since the last reward obtention, according to these statistics.

A simpler system called "bucket brigade" algorithm was later introduced, in which the expected pay-off is replaced by a strength value. When a rule is activated it must transfer a part of its current strength to all the rules that led to its executions. When a reward is issued, the strengths of rules that are currently active are increased. Thus, over a large number of iterations, efficient rules will be reinforced, and rules that led to their execution will also be indirectly reinforced by the bucket brigade system of strength transfer: chains of efficient rules, indicating sequential behaviours, will spontaneously emerge (Wilson [190] provides a clear description of the standard classifier system).

A problem with this method is that establishing reliable chains of rules may require a lot of evaluations [192]. Another problem is that if a given action provides very high pay-off, the population might be overtaken by similar classifiers all exploiting this high pay-off. A solution to this problem is to share the reward between all active classifiers, instead of attributing the whole of it to each classifier. But then the strength value, produced by aggregating received pay-offs over all cases where the rule was active (which may correspond to very different conditions) looses representativity. Other solutions restrict the application of the genetic algorithm to groups of rules that match the same set of conditions; thus no takeover may occur.

It is also possible to change the evaluation process by determining the genetic fitness, not by the expected pay-off, but by the accuracy of this predicted pay-off : unreliable, overgeneral classifiers would then be penalised. This is the rationale behind XCS, introduced by Wilson [191]. XCS aims at obtaining a complete, accurate mapping from inputs and actions to payoff, while using maximally general classifiers (thus offering a minimal, complete description of the environment).

B.3.3 Evolving neural networks

Applying evolutionary algorithms to the generation of neural networks seems a straightforward idea. In theory, the flexibility of evolutionary algorithms (which require only a fitness function a an indicator of progress) might make it possible to overcome the limitations of backpropagation training, especially when it comes to complex, recurrent architectures.

However applying evolutionary algorithms to neural networks is not without problems - especially with regard to the crossover operator. Parameters in a neural networks are highly intricate and interdependent, so crossing two efficient networks is prone to produce a poorer offspring (this interdependence of parameters in a genetic representation is known as epistasis and is the plague of evolutionary algorithms).

On the other hand, two neural networks that correspond to equivalent architectures can be very dissimilar, just because similar neurons occur in different orders in the two networks; this may cause problems for recombination operators (e.g. applying crossover between two differently-ordered networks loses some of the nodes and duplicates some of the others). This is essentially the permutation problem, also known as the “competing convention” problem. Admittedly, this problem has been found to be less severe than expected [77]. Still, it has prompted a number of authors to abandon crossover altogether, and to rely exclusively on mutation-based operators. An impressive survey of the application of evolutionary techniques to neural networks can be found in Yao [198] (who also favour mutation-only evolution).

Another way to proceed is to cooperatively coevolve neurons in a network. This method has been applied in several ways by collaborators of Miikkulainen, e.g. [66], [131].

A way to preserve structures in the evolution of neural network, especially when the architectures are to be evolved together with the weight, is to allow recombination only between equivalent nodes. One can note that if all networks start from similar, elementary networks, then any differing structure must come from a mutation (either the addition, or the removal of one of several nodes). In this context, equivalent nodes are nodes that correspond to different version of the same original node. One way to detect this is that, every time a node is created by a mutation, it should be tagged with a unique ID (e.g. a simple sequential number), and all its descendants that carry this node should keep this tag. This way it will be possible to recombine only similar nodes. This is, in essence, the idea behind NEAT, or Neuro-Evolution of Augmenting Topologies. [171]

This tagging idea does not only provide a way to control crossover; it also allows one to compare two networks and to determine their degree of “parenthood”, i.e. their genealogical closeness, thus providing an objective way to classify individuals into species (which NEAT does). It might also make it possible to trace evolutionary progress by determining which additions proved more beneficial (those that are frequent in the current gene pool) and which weren't (the "holes" in the distribution of ID numbers).

B.4 Generative Encodings

A neural network (or indeed any complex system) of significant size can depend on a very large number of parameters. If each parameter is to be encoded directly within the genotype, the resulting search space can quickly become intractable for evolutionary methods. As an example, in the case of a fully recurrent neural network of N nodes, in theory N*N connections should be stored. This makes evolution impractical for more than a few tens of nodes.

One way to overcome this problem is to use genetic encodings, that is, genotypes that do not directly and explicitly describe each and every parameter of the system, but rather provide a "recipe" that is to be interpreted according to pre-defined rules. Thus, depending on the encoding method, systems of arbitrary complexity can be evolved with genotypes of reasonable size.

It is important, however, to keep in mind the trade-off that such methods imply. Whatever the way we use them, N bits of information can only describe 2^N different configurations. Using a complex developmental scheme can allow us to reach complex structures, but will not increase the number of different structures that can be reach. This means that, while generative encodings push the envelope of possible solutions, they also reduce the density of possible solutions within that envelope : the subset of solutions that can be reach is wider, but more sparse.

This is very reasonable if one considers that not all possible configurations are desirable, or that only some types of configurations should be sought after. For example, generative encodings allow one to introduce constraints and symmetries into a system. However, if the set of reachable solutions becomes to sparse, the encoding may become "brittle" : a small modification in the genotype leads to unpredictably large modifications of the phenotype, thus making evolution more difficult.

B.4.1 Cellular Encoding

A popular generative encoding for neural networks is cellular encoding (CE, see [75]). In CE, networks are represented as graphs of cells, each cell having a number of input and output weighted links. Cells develop and divide following instructions contained in grammar trees. These trees are the genotype of an individual, and can be evolved by standard GP methods.

The developmental process is represented in Figure B.4. Each cell has a complete copy of the instruction trees. In these trees, division instructions are branching points from which two sub-trees stem. When a cell encounters a dividing instruction, it performs the specified division (several division operators exist, which differ in the repartition of input and output links from the mother cell to the child cells) and each of its child cells follow a different sub-tree. Other instructions may involve modification of weights or links. It has been shown that CE can theoretically produce any graph of nodes. A parametric "recursion" instruction allows a cell

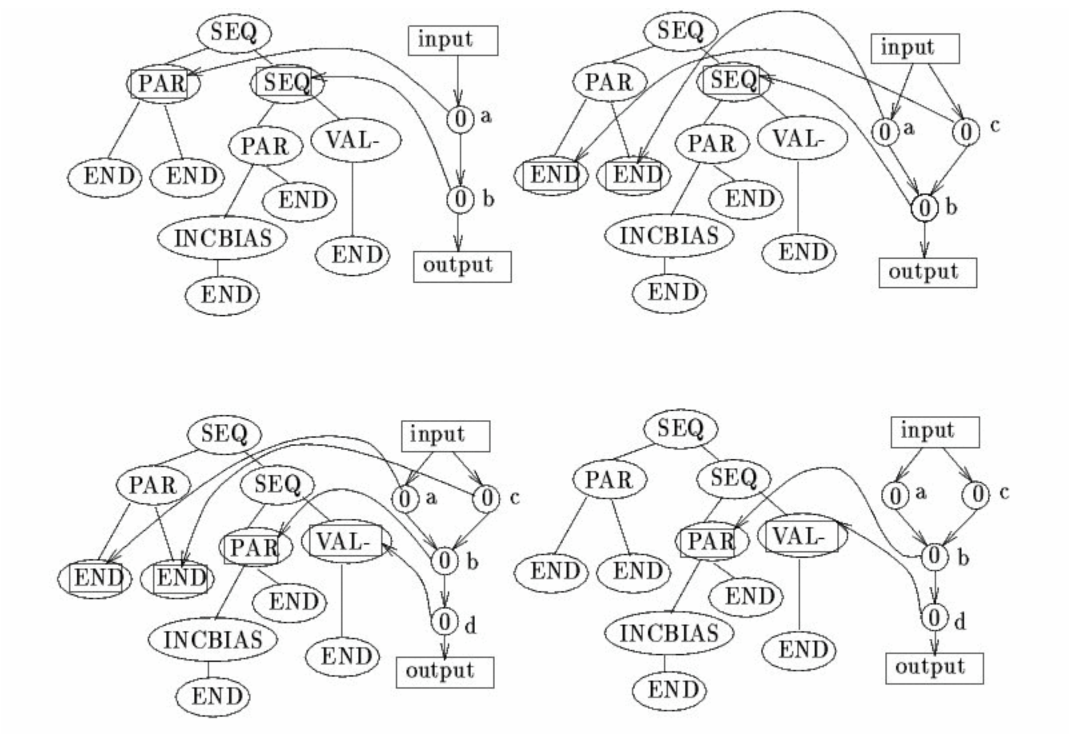

Figure B.4: Application of cellular encoding (image from [75]). The instruction tree is on the left, the corresponding network is on the right. In the first picture, the SEQ instruction has already been executed, leading to a sequential division (two cells connected in sequence). Each new cell follows a different path on the rest of the tree. Cell *a* encounters a PAR instruction, leading to a parallel division (two unconnected cells), leading to the division of *a* into *a* and *c*. Each cell follows a different sub-tree, and both of them read an END instruction, which stops their division process. Cell *b* first finds a SEQ instruction, prompting a sequential division. This results in the creation of cell *d*, sequentially connected to *b*. From there on, *b* will undergo a parallel division; its child will find an END instruction, while *b* will read the INCBIAS instruction (which will increase the bias of its output function) before ending development. Cell *d* will read a VAL- instruction which will decrease the value of the current link (the link currently pointed to by the link pointer; this pointer can be modified by instructions in the tree) before ending its development.

to start back to the top of the tree a given number of times, thus creating similar sub-graphs. This instruction, among others, allow CE to produce modular networks in a natural way. This orientation toward modularity is a distinctive feature of CE. The emblematic experiment for CE was its success in evolving hexapod locomotion from scratch ([75]), in which the modularity of the system allowed it to discover (and exploit) useful symmetries. CE has also been reported to evolve successful networks for difficult non-markovian (i.e. history-dependent) problems such as double pole balancing ([76]).

CE has been used and adapted a number of times. A possible modification is to apply operators to edges instead of nodes, thus correcting the (somewhat clumsy) CE mechanism for edge handling. This is the main idea behind Edge Encoding ([108]), which also modifies the order of execution of a tree (depth-first instead of breadth-first) to ensure a greater repeatability of a sub-tree's results, thus making the process more resistant to genetic recombination.

B.4.2 L-Systems

L-Systems [144] are grammatical substitution systems whose original aim was to model the development of trees and plants as simply as possible. An L-system is composed of two-parts substitution rules, in which the left part (the “predecessor”) specifies what should be replaced, and the right part (the “successor”) specifies what it should be replaced with.

L-Systems may provide a “branching” mechanism. In a standard grammar, a string is replaced by a sequential set of strings. But in L-Systems, replacing elements need not be sequential. Some (or all) of them can “branch off” the replaced section. In practise, this means that after “branching” parts of the replacement have been added, subsequent part of the replacement will be added immediately after the point from which the branching part stemmed. This geometric addition allows L-Systems to describe complex tree-like shapes with very little information.

L-Systems can also be made contextual (the left part of each rule contains fragments that are not meant to be replaced, but simply indicate that the replacement should only happen where these fragments occur) or parametric (elements in the system are given a parameter, replacement rule depend on this parameter and may change it in their replacing part).

L-Systems offer a simple and elegant method for generating complex, recursive structures out of very compact rules. While they were not originally designed to be evolved, these properties make them highly interesting for evolutionary experiments (e.g. [90]). However, as we will see below in a specific application, their high expressivity may come at the price of a excessive brittleness and irregularity in the fitness landscape.

Kitano [96] was first to use L-systems for evolving neural network (a critical description can be found in [161]). Boers and Kuiper [15] used contextual L-Systems to generate complex feed-forward neural networks, which were then trained with backpropagation. The alphabet used for the encoding described groups of nodes as well as “skip” values, that specified to which group the output of a given group should be “projected”. This alphabet was translated to bit strings in order to apply evolutionary operators, then translated back as rewriting rules for evaluation. This method allowed them to evolve networks that solve simple problems (e.g. XOR), and apparently more complex separation problems. [16]

In this perspective, the evolutionary process (together with the developmental method) is seen as a support for backpropagation : it is supposed to find better architectures for a class of problems, that is, architectures that allow backpropagation to learn better and faster. Architectures described in these experiments are usually fairly simple, but simplification was actually one of the goals.

L-Systems and CE are two methods that rely on the application of rewriting rules to generate complex neural networks. It is not difficult to see where their main difference resides : in L-Systems, at each timestep, all rule are applied to all nodes in the graph. Contextualisation may refine the target for a given L-System rule, but then finding the right context for rules becomes an additional problem for evolution to solve. In CE, by contrast, each instruction is to be executed by one single cell - the one that reads it. CE only allows for a rather crude control of global structure (pure recursion) at the benefit of a finer control of local modifications. L-Systems do exactly the opposite. A mutation in CE will tend to have more local effects than a mutation in L-Systems.

B.4.3 Biological generation

Besides grammar-based approaches, other types of generative encoding rely on more biological metaphors, oriented toward cell chemistry and chemical neuron growth methods. The neurons is such approaches are physical, situated objects that interact with their environment. The final network emerges out of interactive process between various parts of the system. Such models can get to a very low level of description, including the growth of axons and the geographical location of the neurons.

In a series of experiments, Nolfi and Parisi [132] developed a descriptive model in which neurons are specified by their position in a 2D map, together with the length and the angle of their branching patterns. After the genotype has been read, neurons are located on the plane and their axons begin to grow, branching off sub-axons at angles and lengths specified in the genotype. Axons that reach another neuron establish a connection. After the process is complete, non-functional axons and unconnected nodes are removed and the resulting network is evaluated. Cangelosi [25] extends this model by allowing cells to divide : rewriting rules are applied to an original egg-cell. The rules also incorporate information about the location and the growth parameters of the daughter cells. Those models were applied to the evolution of a simple animat. Learning and development can also be incorporated. [132]

An even lower-level experiment can be found in the experiments performed by Jakobi [91]. In this rather uncompromising model, genomes code for actual "proteins" that affect not only the behaviour of the cell, but also the expression of other genes, thus giving rise to a complete genetic regulatory network.Some proteins can navigate between cell, thus providing inter-cellular communication. Depending on the quantities of protein present in a cell at a given time, this cell can divide, move, or differentiate into a neuron. When this latter stage is reach, axons begin to grow from the cell.

Both axon growth and cell movement respond to particular gradients in the combination of signal proteins in their environment. The original breaking of symmetry is provided by a extra-cellular sources that emit specific signal proteins. The model is applied to the evolution of a simple robot (following Jakobi's own "minimal simulation" method [92]) that moves forward in a curve pathway without hitting the walls. The author makes interesting remarks on how this model, despite its complexity, can evolve symmetries quite effortlessly, simply by not responding to specific, symmetry-breaking signals.

These growing processes may be modelled more abstractly. For example, in GasNets [86], each neuron specifies its connection by defining "branching zones", which correspond to portions of a circular area in a 2D plane. Any other neuron that falls within this area (as defined by the radius, the opening angle and the orientation of the area) is connected to the originating neuron. Alternatively, direct position coordinates may be specified, such that any neuron within a given radius from the specified position will be connected to. Neurons can then alter each other's transition functions through the diffusion of a "gaseous" transmitter which diffuses from each neuron (in analogy to a similar process in animal brains).

Many more biologically-inspired models are possible, differing mainly by their degree of fidelity in the reproduction of currently known biological processes. For example, gene regulation and genomic networks, through which some genes influence the expression of others, are an important feature of biological organisms. Because these effects determine the process of development, they are an important component of evolution. Several possibilities exist for modelling these mechanisms, some of them will be described later in the following sections.

B.5 Evolving embodied agents

The challenge of evolving the morphology of an agent together with its control architecture has attracted significant research; yet it seems that virtually all of them share at least a few common points:

- The control architecture of choice is usually a neural network. Rule-based controls (e.g. classifier systems) or oscillating patterns (such as the CPG evolved by Ijspeert and colleagues [88] for the control of a salamander) do not fare well with unpredictable morphologies.
- A generative, indirect developmental system is almost always chosen. Specifying each and every part of the morphology, together with each of its parameter, is not common. A first reason for this is that generative encodings lead to immediate symmetries that result in more interesting or "natural" appearances. Another reason is that since Nature itself uses generative encodings anyway, doing the same may lead to interesting models of biological development. The proportion of each motivation vary according to the kind of experiment considered.

Stanley and Miikkulainen provide a remarkable survey of developmental systems for evolving controllers and morphologies simultaneously. [172]

B.5.1 Dellaert and Beer's abstract morphogenesis

Dellaert and Beer [44, 45] describe a developmental model which idealises biological processes in a very straightforward way. Organisms are represented by groups of single 2D squares or vertical rectangles. When a cell divides, it is simply cut in width (for rectangles) or in height (for squares). Thus after 2 divisions a square cell gives rise to 4 square cells, with each grandchild cell's side length being half that of the original cell. At each division, the complete genome is transmitted to child cells.

In this model, the genome encodes a Random Boolean Network (RBN - see section B.8.1 for details). At each timestep, the RBN is allowed to run until it reaches a stable state. Then, depending on the value of a special "division" node in the network, the cell decides whether or not to divide. If it does, the current state of its RBN is transmitted to its child cells along with the genome. The RBN can also use information from neighbouring cells (presented as an OR combination of the state vectors of their RBNs, plus some limited spatial information in the form of special bits indicating whether or not the cell lies on the median axis, or on a border). At the end of development, cells can be tagged as neurons, actuators or interneurons depending on their state, thus giving rise to small neural networks within the organism.

Figure B.5: Division and differentiation of rectangular cells according to the state of their RBN in Dellaert and Beer's model (from [45]).

The model has first been tested for simple geometric features, in order to demonstrate its capacity to produce organisms that exhibit some degree of physical complexity. It has also been successfully applied to the evolution of a curve-following animat. Its suitability for more complex tasks would depend pretty much on the possibility to add actual physical properties to the organisms (e.g. muscle cells).

B.5.2 Bongard's "Artificial Ontogeny"

Although they were not introduced as such, the experiments made by Bongard and colleagues [18] consist essentially in taking the previous model a big step further toward both physical embodiment and biological plausibility. The purpose of these experiments was clearly to investigate the mechanisms of natural embryogeny and their application to the generation of complex, modular agents. Instead of a 2D, fixed-size organisms, this model works with 3D organisms composed of spheres, each having its own size and position. More importantly, instead of modelling genetic regulatory dynamics by RBNs, it uses a quantitative, asynchronous, spatial genetic regulation model called "Artificial Ontogeny".

In this model, genomes are strings of real values, and the beginning of each gene is indicated by a value corresponding to a "promotor site" (since all values below a given threshold are regarded as promotor sites, it seems to be quite possible for genes on a genome to overlap). The values that follow the promotor site indicate the type of gene product that this gene produces if it is expressed, the gene product that regulate this gene's expression (and whether this regulation is positive or negative), and the range of this product's concentration to which the gene responds. It also indicate from which "diffusion site" its own gene product should diffuse within the organism.

A "diffusion site" of a unit is one of its 6 geometrical poles (N, S, E, W, plus "forward" and "backward"). If a gene is to be expressed, it is injected into the unit at the specified diffusion site, with a genetically specified concentration, as long as the gene is active. At each timestep, a decay parameter reduces the concentration of every product by a fixed amount. Gene product can diffuse between neighbouring diffusion sites at a given rate, and between units at a significantly lower rate.

Two specific gene products regulate the growth of each unit, respectively activating and repressing it (no size diminution is possible). When a unit exceeds a given size, it simply splits, and the new unit is attached to it at the diffusion site that has the maximal concentration of growth enhancer product. The two units are joined centre-to-centre by a rigid link, and a rotational joint is created at the centre of the new unit.

By contrast to this degree of biological realism in the generation of the organism's morphol-

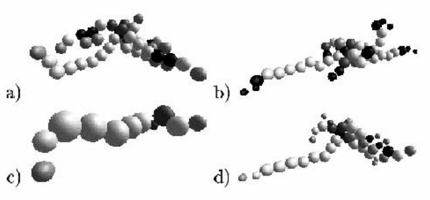

Figure B.6: Four organism created by "artificial ontogeny" (from [18]). Colours indicate the presence or absence of motor / sensor neurons (dark cells have no neurons at all). Note the presence of relatively structured appendages.

ogy, a much more abstract method is used for the generation of the neural system : it is simply Cellular Encoding, in which the various instructions to be executed by the current neuron and synapse are indicated by several specific gene products. When a unit splits, all the neurons at the corresponding diffusion site are moved to the opposing diffusion site of the new unit, but keep their connections intact, which allows for the migration of neurons and the maintaining of long-range connections in the organism. Neurons can be sensor or motor. At each timestep, within each unit, the activities of all motor neurons are averaged and the result is seen as a target angle (in the $[-\pi/2; \pi/2]$ range) for the cell's rotational joint.

The model was applied to simple physical tasks such as block-pushing. The analysis of the organisms produced indicates that some degree of modularity was achieved by the system, in the form of genes that had similar activation patterns (thus probably responding to the same "master" gene) and giving rise, under some conditions, to similar appendages.

More recently, Bongard and colleagues have explored new directions in the joint evolution of morphology and behaviour. They created a system in which real robots engaged in continuous self-modelling and simulation (generating models of themselves, testing these in internal simulation, and evaluating the results by comparing with reality in an evolutionary algorithm among models). This allows the robot to recover from random damage: "when a leg part is removed, [the robot] adapts the self-models, leading to the generation of alternative gaits." [17]

B.5.3 The Brandeis Lab

This remarkable degree of fidelity to biological processes gives a striking contrast with the experiments described in [83], in which the use of neural networks seem to be the only concession to biological plausibility. Creatures in this model are composed of rigid bars, possibly equipped with motor joints. The specification of the animat is given by a series of instruction that are followed by a LOGO-like "turtle", such as "add a bar, turn left N degrees, add a joint", plus the ability to "branch off" and come back to branching point later (note that this is very similar to the kind of descriptive language used in original L-Systems).

The controlling neural network is built in parallel with the morphology : the genomes consists of both morphological and neural instruction. Each time a new joint is created, a new output neuron is attached to is. The neural construction language is based on Edge Encoding. The resulting creature is then evaluated on a simple locomotive task in a simple 3D physics simulator.

After testing a standard genetic algorithm based on direct encoding, the authors decided to use a generative encoding - namely, a parametric L-System. In this case, the additional complexity allowed by the generative led to a striking increase in efficiency. This may be related to the fact that with such simple structural elements, no efficient behaviour can emerge unless a large number of elements are present, which is only possible with a generative encoding. The organisms created with parametric L-Systems exposed a high degree of modularity and inner symmetry.

A related project, coming from the same laboratory, is the GOLEM system described in [107]. In this experiment, simple animats composed of extensible and joints are evolved in a 3D simulator. The model in itself is quite simple: the controller is a fully connected neural network, some bars may be associated with an actuator neuron (in which case they are seen as retractable) and the whole organism is directly represented by a string of parameters. No input to the neural network is possible.

The originality of the experiment comes from the fact that the evolved creatures were actually built, not by hand, but by a “3D-printing” system. The neural network was implemented by a micro-controller and actuators were constructed as sliding motors in the specified bars.

B.6 Sims’ virtual creatures

B.6.1 The system

We will now introduce the works of Karl Sims, which we intend to discuss more thoroughly than other systems. These experiments have had a deep influence on the field of artificial life as a whole, and many of the works that we have related so far (in fact, all those that came after it) mention it as an inspirational source.

Sims’ experiments [163, 162] deal with articulated creatures composed of 3D rigid block. These blocks are arranged in a tree-like hierarchy, that is, each block (or “limb”) stem from another “parent” limb (except of course for the first, “root” limb). Thus all creatures are organised as trees of limbs. These creatures are controlled by a special kind of neural network in which neurons compute functions of their weighted inputs (e.g. summation, product, trigonometric... or even a simple sigmoid function, in which case the neuron behaves like a standard McCulloch-Pitts neuron).

Each limb has its own local network, which may be linked to other networks from adjacent limbs. In addition, creatures have a “central” network which is not associated to any limb, and can be linked to any network in the creature. Joints between two limbs are associated with actuators. An actuator takes one weighted input from a given neuron or sensor, and applies a corresponding torque to the limbs it controls. This torque is scaled by a maximum value proportional to the maximal cross section area of these limbs, thus scaling differently from their mass (which scales with volumes). Similarly, various kinds of sensors are implemented, which include contact sensors, proprioceptive sensors returning the current value of each degree of freedom of each joint within the organism, or “photo-sensors” which indicate the direction of a given global light source.

The genotype of these organisms is not a direct tree-like representation, but a graph of nodes (see Figure B.7). Each node correspond to a limb. The difference between graphs and trees is the presence of loops. When a node has a recursive link to itself, similar copies of these nodes

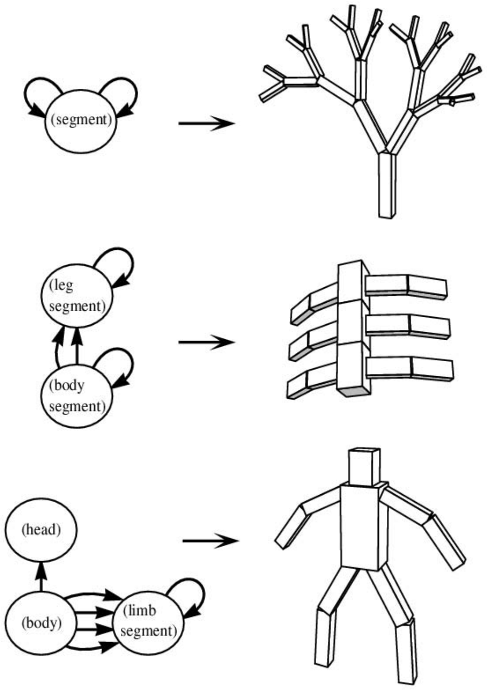

Figure B.7: Sims' Creatures are represented by directional graphs. Each vertex in the graph corresponds to a limb. Edges correspond to articulations between a limb and a sub-limb, and contain information about the placement, resizing and orientation of the sub-limb. A cycle in the graph can be followed several times, giving rise to recursive structures (from [163]).

branch off each other. Because connections also contain information in the form of parameters (scaling, angle, etc..), these recursive sub-parts need not be exactly identical. Similarly, if the recursive link does not point toward the node itself, but toward an ancestor nodes, similar (but usually not identical) sub-trees will be created, branching off each other. Two nodes may be joined by more than one link, each with possibly different information, which means that one limb can have many similar sub-trees branching from it.

This mapping between genotype and phenotype is quite different from what we have seen so far. It naturally produces modularity and symmetry, as generative encodings do; but it also keeps a fine-grained control of the phenotype, by ensuring that each part of the organism is fully specified, just as in direct encoding. The language of the system still deals with actual body parts, not abstract terms in a developmental grammar, thus keeping a link between genetic and morphological complexity. The balance between complexity and modularity on one hand, and control and efficiency on the other, is undoubtedly one of the strong points of the model.

The genetic algorithm used consists, at each generation, in selecting a given proportion of the organisms and mating them in order to replace the discarded ones with offspring. Mutation can affect either the structure of the graph or the value of some parameters. Recombination my occur either by grafting (simply taking connecting a node from one parent to a node in another and removing all newly unconnected nodes) or by graph crossover : graphs are aligned and the offspring is copied, node by node, from one of the parents, switching between both parent once or more in the process. One can easily imagine applications of historical marking (see NEAT, above) to this model.

The model was first applied to simple locomotive tasks [163], either on a plane (walking / creeping), or in a liquid space (swimming), as well as on a source following experiment. Efficient organisms were discovered, exhibiting a wide range of behaviours (e.g. snakes, bipeds, paddlers) and complexities. The strikingly natural appearance of these organisms owes much to the natural symmetries of the generative system.

Interesting difficulties were encountered. For example, choosing a "good" fitness function, even for simple locomotive tasks, was not as trivial as one might think. Simply evaluating organisms after the distance they covered led to the emergence of "opportunistic" behaviour e.g. producing a sharp initial thrust, then stop moving altogether. Although clearly sub-optimal, these behaviours quickly dominated the population and led to premature convergence. A turn-around was to give a stronger weight to the distance covered near the end of evaluation time.

Another difficulty came from the fact that any physical simulator is bound to be imperfect. In some situations, the simulation may become unstable : "fictional" forces begin to appear, and the body parts seem to undergo incorrect and incoherent movements. These instabilities were more often than not exploited by evolution, resulting in creatures that exhibited fast, but completely non-realistic motion.

B.6.2 Coevolution and the LEO algorithm

In another set of experiments, creatures were evaluated in a competitive setting, in which two animats competed for the "control" of an inert block. At the end of the evaluation period, the creature closest to the block was the winner. Since individuals are evaluated after the outcome of their interactions with other, evolving individuals, this is a case of coevolution.

As we saw in chapter 3, coevolution immediately poses the problem of how to evaluate individuals reliably and efficiently. Sims discusses several methods, among which the division of the population in two species (such that competition occurs only between individuals from another species), or the possibility of pitting all individuals against the best individual from the previous generation. Sims eventually settled on a combination of these two.

The resulting algorithm consists in dividing the population between two species A and B. Each creature from each species is then evaluated by competing against the current best from the other species. More precisely, at every generation, each individual from population A is pitted against the current champion of population B. The resulting score is used as a fitness value for selection and reproduction among individuals from population A; it is also used to choose a new champion for population A (the individual from A which obtained the best score against the current champion of B). Then the same process is applied to population B, using the new champion of population A for evaluation, etc.

This algorithm has since been called the Last Elite Opponent (LEO) algorithm by Cliff and Miller [32].

The system generated a wide variety of behaviours, some of which are shown in Figure B.8. Some creatures jab their opponents away from the cube before dragging it closer (b, d, e), some pin down their opponent (b). Some were capable of following the cube after they (or their opponent) displaced it (f). Some two-armed creatures flap the cube away with an arm and

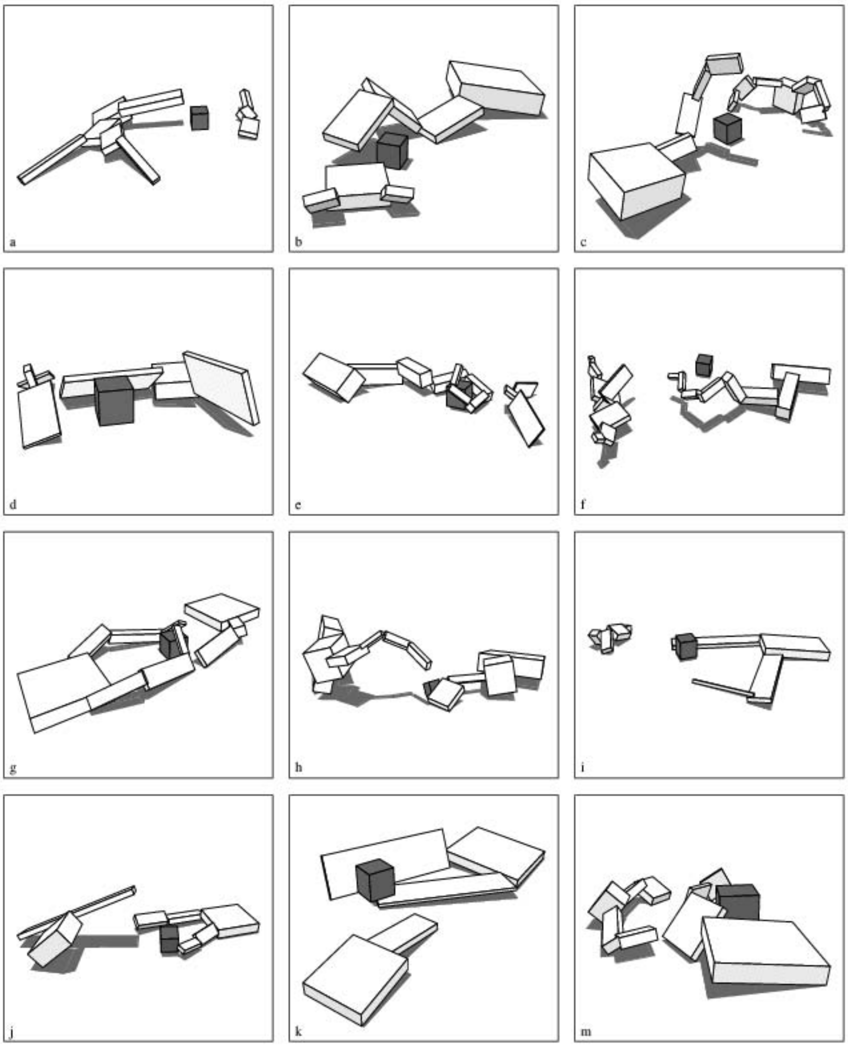

Figure B.8: Creatures evolved by Sims in a coevolutionary experiment (from [162]). See text for details.

catch it from the other. Some simply caught the cube and away with them (m). The population dynamics in the two species of each run exhibited various behaviours as well.

These experiments have had a rich legacy (see Section 4.1.2 in Chap. 4 of this thesis). Of all the projects discussed so far, they are probably the one that come closest to the intuitive notion of “artificial life”. The evocative morphology of these creatures, and the immediate realism of their behaviour, bears an uncanny resemblance to biological organisms that can hardly be found in any other experiment in the field. Once again, it is clear that even though the realism of the physical simulator plays an important role in making the organisms “look real”, the genetic language is what really makes them “look *alive*”. The division in trees or chains of sub-modules comes directly from the remarkably efficient use of generative encoding.

B.7 Artificial Worlds

Most of the experiments discussed so far essentially rely on purely artificial selection. Coevolutionary experiments are the simplest way to implement a more natural, possibly open-ended form of evolution. Yet it is reasonable to think that much richer dynamics can emerge from more complex environments in which many individuals (instead of just two) interact, especially if the reproduction process becomes a part of these interactions (instead of being enforced separately in the algorithm).

However, the simulation of full populations is computationally more demanding than simple individual-based experiments. Not only the sheer number of individuals, but the potentially complex nature of their interactions can impose significant computational costs that may increase dramatically as the size of the population grows (e.g. if each agent must know the position of each other, or even of the closest one, this imposes $N^2 - N$ calculations).

Thus the first simulations of these kind applied to abstract worlds and abstract organisms, instead of physical, embodied agents. However, a few decades of continuous application of Moore's law[4] have made it possible to build up increasingly complex simulation, including 3D, physically realistic, situated organisms.

B.7.1 Abstract Ecologies

In Conrad and Pattee's system [33] (cited by Fogel [59]), a one-dimensional discrete world harbours abstract organisms which consist of extremely simple instructions. Organisms must collect resources (under the form of items called "chips") in order to reproduce and repair themselves. At each timestep, an individual executes its current instruction, which may result in trying to seize the resources present at that place, trading resources with other individuals, or mating. Asexual reproduction occurs as soon as an individual has accumulated enough resources. Some results were produced, indicating adaptation of some sort, but the crudeness of the system (and, more particularly, the very limited control that organisms had over their destiny) made further evolution difficult. However, as the authors noticed, this study "indicated the feasibility and usefulness of this synthetic approach as a source of new ideas and information about fundamental problems in evolutionary biology."

It is difficult not to notice the similarities of this early model with Echo, the more recent artificial ecology model designed by Holland (as described in [84, 128]). In Echo, the world is a two-dimensional grid. Each agent is defined by a genome, which encodes an external appearance (composed of alphabetic strings called "tags") and a set of internal conditional rules (which determine the activation of pre-defined behaviours, according to the tags of other agents). Agents also have a reservoir, in which they can accumulate resources; these resources are represented as alphabetic characters, just like the rules and the tags of agents.

As in Conrad and Pattee's model, asexual reproduction occurs when an agent has acquired enough resources to replicate itself; resources can be acquired either from the environment, or from trading with other agents; and agents can reproduce sexually if they choose to. However, they can also fight each other, and move randomly on the grid when they fail to acquire any re-

---

[4] "The number of transistors in an integrated chip doubles every year." [130]

source. The system can be tailored in various way to study its ecological dynamics. Comparisons with natural systems are also possible, although somewhat risky (see [84] for an example).

B.7.2 Ackley and Littman: Evolutionary Reinforcement Learning

The model described by Ackley and Littman [1] is an attempt at modelling the interactions of between lifetime learning and evolution. The world (called AL) is a 2D grid inhabited by agents, trees, food, and predators. Food grows geometrically. Carnivores reproduce when sufficiently nourished, and may die of starvation. They are controlled by hand-coded rules that do not change throughout the simulation. Trees provide shelter from predators for only 1 agent at a time. Agents can eat food and dead agents, be eaten by predators, and die of starvation. Their inputs indicate the quantity and proximity of each other entity (agent, predator, food...) in each of the cardinal direction, plus internal information (health). Both evolutionary and reinforcement learning are used for the agents. The interplay between lifelong and evolutionary learning was prominently featured in the authors' discussion, illustrating both advantages and drawbacks of the combination.

B.7.3 Evolving competing programs - Tierra

Competing programs in virtual computers have been implemented for decades, usually as games. An example is Core Wars [46], for which evolutionary experiments have been run. However, the most famous evolutionary experiments on machine code organisms is indisputably the Tierra system [146].

Tierra is essentially a virtual computer, in which small self-replicating programs are executed by the virtual processor. Programs are stored in a sequential access memory and are composed of various machine-code instructions. The system operates in multitasking mode between all programs. The processor executes the current instruction (as indicated by an instruction pointer IP), then increases the local program's IP, so that it will read and execute the next instruction when its execution is resumed; Note, however, that some branching and jumping instructions can modify the IP directly, allowing the execution flow to jump to other sequences of instruction (possibly out of the program's code). After each operation, the CPU switches to the next program in the queue.

While branching operations such as CALL and RET operate as their classical assembly language counterparts, the JMP instruction works in a specific manner : each JMP instruction is followed by a series of NOP0 and NOP1 (no-operation) instructions, that work as a pattern to be matched by the jumping process. The processor will then look forward (or backward) in memory until it finds the complement of that series. For example, if the JMP instruction is followed by, say, NOP0 NOP0 NOP1 NOP0, the processor will look forward until it finds the sequence NOP1 NOP1 NOP0 NOP1. This "*addressing by template*" mode was designed to mimic the interactions of cells through proteins that match specific receptors on the cells' surface.

Another feature of Tierra is that it avoids numeric arguments altogether: a numeric register (the equivalent of the CX register) holds the results of arithmetic operations and can be operated upon by bit-flipping and shifting. If a numeric value is needed, the same instructions can create any given integer in this register. Thus each bit in the genome actually codes for an instruction, and the total alphabet of the system is limited to the instruction set.

Programs with which the Tierra system is seeded are delineated by specific sequences of NOPs. So are the reproducing loop, and the copying procedure it uses. Programs can write either inside their own memory zone, or on a section of memory that they explicitly allocated with a specific ALLOC instruction. After they have copied their own code in such a memory zone, a DIVIDE instruction removes their writing privileges to that zone and creates a new IP for the newly created program. In order to prevent memory overflow, a destructive mechanism (the "reaper") stores programs in a LILO queue (new programs enter the bottom of the queue) and destroys the program at the top of the queue every time memory occupation becomes too high.

Genetic change occurs by "cosmic ray" mutation - that is, about 1 bit in the memory is flipped for every 10000 instruction executions. Mutation can also occur while copying instructions, at a higher rate (1/25000 to 1/1000). The result of operations can also be altered by +/- 1 at low frequency.

The results are, at first sight, impressive. The first distinct species to appear are parasites, smaller pieces of code that cannot reproduce by themselves, but use the copying code from other organisms to do so. Being smaller, they enjoy a comparative advantage because they can be copied more quickly. However, because they need hosts to replicate them, they cannot eliminate them completely, and the population enters a Lotka-Volterra cycle. Some hosts can defeat these parasites, which leads to their disappearance.

Later on, cleverer hosts become hyper-parasites, by ensuring that, when a parasite's IP points to their replicating code, it ends up copying not the parasite itself, but the hyper-parasite - and never returns. After the parasites have gone extinct, these hyper-parasites (which are able to reproduce on their own) evolve social characteristics - i.e. their reproduction begins to depend on the presence of similar organisms next to them. Ray argues this by the fact that high genetic uniformity in neighbouring organisms supports such "social" behaviours, especially in the presence of a selective advantage brought by their shorter size. Ineluctably, such behaviour (passing instruction pointers around for replicating each other) is eventually exploited by "cheaters", which never return the IP that get through them. Other results in Tierra include the appearance of various optimisers, along with cases of code mixing that indicate a crude, spontaneous (and unfortunately unstable) form of recombination.

Unfortunately, Channon [28] downplayed these results somewhat, by showing that some of them could be obtained quite easily. While Ray had documented how the first parasites evolved out of a single bit-flip mutation, the fact that hyper-parasites could appear with just one more mutation made their complexity more doubtful (at least from an evolutionary viewpoint).

Still, Tierra remains an interesting experiment which is still under development. Multi-threading has been introduced in an attempt to model multi-cellular organisms [148]. Similarly, a network adaptation of Tierra has been developed in order to examine the system's adaptive capacities in heterogeneous environments. Unfortunately, published results have been scarce. The main source of documentation for Tierra and the ongoing results (if any) is Ray's website.

Figure B.9: One of the more efficient swimmers evolved by Ventrella (from [179]).

B.7.4 Ventrella's Swimmers

Ventrella's swimmers [179] is among the first examples of an artificial world in which organisms are physically modelled, and interact physically with their environment. These swimmers are 2D coloured figures composed of interconnected line segments. The morphology is constrained (one to four limbs emanating from a single "head" node, two to four segments in each limb, all segments have same lengths). Organisms are controlled by oscillating motors, governed by a set of parameters in the genotypes. These parameters include amplitude, frequency, shifts in phase or amplitude between successive limbs and segments, plus a few parameters that are used when changing direction. Thus, while the decision part of the controller is fixed, the response (locomotion and direction change) depends entirely on genetic parameters.

Movements and reproduction consume energy; additionally, a small amount of energy is lost at every time step, even if no action is taken. Energy can be restored by eating food bits, which multiply at a fixed rate in the environment. A swimmer's field of view covers 360 degrees, allowing it to detect other swimmers or food from any direction (within a maximal radius). A swimmer can detect the relative orientation of a food bit, or of another swimmer, as well as colour information on other swimmers. At any time, the behaviour of a given swimmer is dictated by a fixed, non-evolving finite states automaton. The possible states correspond to the search and pursuit of food and mates, as well as eating and mating if food or mates are reached.

Mating occurs when a mating organism is sufficiently close to the partner it sought. The offspring's genotype (a string of 17 floating point values within the $[0.0; 1.0]$ range) is created by n-point crossover and mutation of the parents' genomes. Repetitive reproduction between two swimmers is possible as long as they have enough energy. Both parents (even if one of them did not seek reproduction) contribute a fixed amount of their energies to their offspring. If an agent's energy fall under a "hunger" threshold, it will stop seeking mates and start looking for food instead. If it falls under 0, or if the agent has reach its maximal life span, the agent dies and disappears.

Sexual preference plays an important role in this simulation. Each swimmer has a head colour, a colour shift parameter (which controls colour gradient between successive segments), and a "favourite colour" for potential mates. A swimmer will only choose mates that have their favourite colour, or a closely related one. The amount (*not* the percentage) of this colour (and of its close derivatives) will be the criterion of choice among several potential mates.

Experiments with this model lead to the emergence of several genetically-related groups (as indicated by their colour), developing increasingly efficient locomotion skills. Sexual selection may also lead to a "peacock tail" effect, in that some organisms prefer to invest in sexual attractiveness (by displaying as many colours as possible, which entails having many limbs) at the detriment of efficiency [180].

This model can undoubtedly be seen as a successful application of natural selection in an artificial world. The Darwinian metaphor (survival of the fittest in an environment with strictly limited resources) applies in full. Even though only a few behaviour were possible, the actual realisation of these behaviours by each organism was completely dependent on evolution, which came up with various solutions. A limitation of this work is that the control architecture was fixed, thus restricting the range of possible strategies. Nevertheless, this represents a choice in the balance between simplicity and completeness which, in this model, provided highly interesting dynamics.

B.7.5 Polyworld

Polyworld [197] is a flat environment in which simulated organisms reproduce sexually, fight and kill and eat (either food items, or other dead organisms' bodies). Organisms have a fixed repertoire of behaviours (move, turn, attack, eat, mate, plus the possibility of controlling the brightness of a part of their bodies) which are controlled by neural networks with Hebbian learning synapses. Neural networks are divided into connected groups in order to reduce the number of genetic parameters (which still remains very large). Vision is implemented in a one-dimensional way : the visual input is provided by extracting the strip of pixel just above the ground in front of the organism.

Mating simply occurs when two spatially overlapping individuals both have their "mating" neuron activated. An offspring is generated (through crossover and mutation) if the maximum number of organisms has not been reached, and (depending on a simulation parameter) if the genetic difference between both parents is not too high. Parents transmit a part of their energy to their offspring. An organism can fight another one if they overlap. The power of the attack is determined by the activity of the "fighting" neuron and the organism's strength, size and current energy level, and is deduced from the attacked organism's energy. The level of excitement of the "fighting" and "mating" neurons are indicated in the blue and red components of its current colour (the green component is equivalent to the organism's ID number).

Note that all actions, including locomotion, induce a consumption of energy proportional to their intensity. Energy is also consumed continuously in proportion to an organism's size, strength and neural activity. Thus being big and strong is a competitive advantage in terms of fighting, but leads to a higher energy consumption. Energy can be replenished by consuming patches of food that grow naturally, or by eating carcasses of dead animals, allowing for various possible strategies. A distinction exists between health and food energy (the former being related to fights, both being replenished by food).An organism dies when it outlives its life span, or when any of its energy levels drops under 0.

Simulations began with an artificial selection method, based on an ad hoc fitness function, until self-sustaining populations occurred. Experiments led to the emergence of various strategies. The simplest one was simply to run forward, always wanting to eat and to mate, which worked well in open worlds with sufficient food supply. Organisms turning around in a tight radius, following edges and barriers (when they were present), or simply staying very close to each other (thus providing abundant resources for both mating and eating to each other) also appeared, depending on the global parameters of the world. The most complex behaviours include examples of swarms, i.e. groups of organisms that follow each other.

The fact that any strategy did emerge at all can be seen as a small miracle when one considers the size of the search space that was imposed upon these creatures. The number of (8-bits) genes went literally into the thousands. This comes from the size of the neural networks considered, as well as their genetic. According to the author, "A network of 185 neurons can be fully specified by just *[sic]* 2146 genes". The "just" is a reference to a previous method based on a full connectivity matrix, implying N^2 genes for N neurons, i.e. 200 neurons led to 40000 genes (simulations were run with this model, although no results are mentioned - perhaps unsurprisingly).

Such a large search space is difficult to scan in a reasonable amount of time, even for evolution. This may explain the relative simplicity of resulting strategies (indeed, the capacity of reacting to food or to visual stimuli is described as a "complex, emergent behaviours"). Additionally, the fact that Hebbian learning was implemented for all synapses makes it difficult to distinguish the role of evolution in the appearance of these behaviours.

B.7.6 Geb

Imagine that we take the Polyworld experiment and simplify it as drastically as possible, while still retaining enough features to allow for natural selection. We would probably obtain something that looks very much like Geb.

Although created independently of Polyworld, Geb [28, 27] also features abstract organisms with a limited repertoire of behaviours, namely moving (and turning), mating and killing. They are also controlled by neural networks. However, in contrast to Polyworld, the world has been thoroughly simplified. First, Geb is a discrete world, both in space in time : the environment is a 2D grid, and time occurs in discrete timesteps. At each timestep, an organism acts accordingly to the outputs of its neural network. No energy quantities are considered in the system (essentially, Geb organisms do not eat). The population is naturally limited by the fact that no more organisms can exists than there are squares in the grid (the standard being a $20 * 20 = 400$ squares grid).

Mating and fighting occur in a rather straightforward way. If an organism enters "fighting" behaviour, then the organism right ahead (Geb organisms have an orientation) is simply killed. If an organism enters "mating" behaviour, then it reproduces with the organism that stands right ahead. The offspring is placed on the square that lies just beyond the mating partner, erasing any organism that might be there. Organisms can also turn (by an angle proportional to the excitement of corresponding output neurons) and move forward, if the square ahead is empty. Note that, since the behaviours are abstract and pre-defined, the goal of the organisms is not to evolve behaviours, but to know how and when to use them.

Genotypes are simply binary-encoded L-Systems [144], that is, recursive grammars suitable for the generation of trees and graphs. Note that since Geb organisms have no physical property, the genotype only describes a neural network. Neural networks in Geb are graphs of nodes, in which each node has a character (i.e. a bitstring). An organism's genotype is thus a binary string

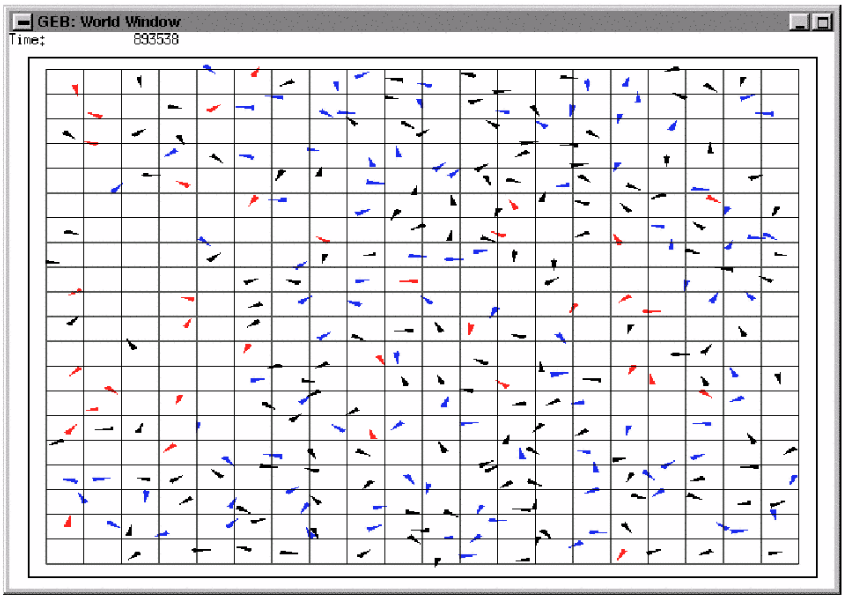

Figure B.10: The experimental world in Geb.

that is decomposed into a set of rewriting rules; these rules operate on the nodes' characters and connections. A rule matches a node if the predecessor of that rule matches the beginning of a node's character, and if no better-matching rule for that character is found. The rules are encoded according to a somewhat intricate binary language that identifies meaningful segments by checking whether the next odd-rank bit is 0 (otherwise, a separator has been found).

The resulting networks may be arbitrarily complex, and the character of the final nodes determines their status : input or output, and in the latter case, what kind of action they encode (this is determined from the first few bits on their character). Inputs are sensitive to the excitation of output nodes of other organisms that match their character. This may provide the basis for some communication between organisms. Reproduction occurs by crossover and mutation. Crossover offsets the cutting point in the second organism by one (in a randomly chosen direction) with regard to the first organism. Thus genotypes' length is expected to vary smoothly between generations.

Identifiable behaviours reliably emerge. Running forward and killing everything, or doing all possible actions at a time, initially dominate the population. Consistently, a simple behaviour consisting in turning in a tight circle while always trying to kill and to reproduce emerges rapidly. Some basic level of mutual detection occurs by the appearance of a non-action output character (000) that also appears among the inputs of (some of) these organisms. Some other organisms use this as an input to turn against these "turners" and kill them.

Identifying any further behaviour by visual inspection is difficult, however. In fact, the complexity of the neural networks often makes it difficult to analyse the exact behaviour of any given agent.

B.7.7 Framsticks

Framsticks is an artificial 3D world which defines itself as "a three-dimensional life simulation project". Framsticks organisms are composed of articulated, deformable sticks controlled by neural networks. These neural networks have sensors for contact, proprioception, spatial orientation and presence of food; actuators may control either "bending" or "rotational" motors (which are actually two orthogonal axes of rotation). The neurons themselves are standard McCulloch-Pitts neurons, but with several additional parameters : intensity, sigmoid control and especially an inertia parameter, which may allow for interesting dynamics at neuron level.

Several encoding systems have been implemented, including direct representation, CE-like encoding, and an intermediate kind of generative encoding which is highly evocative of Sims' [98]. The 3D simulation is physically realistic except in the fact that an organism can self-penetrate (two of its stick may go through each other); however, collisions and contacts between different agents are taken into account.

The simulation environment is based on an artificial selection model : organisms are chosen from a pool, placed in the simulator, evaluated according to some human-defined criterion and then mate according to this fitness. However the system allows for the simultaneous simulation of several agents within the environment. Together with the simulation of food and energy consumption by the agents, this leads the author to propose a way to emulate natural selection : simply evaluate agents according to their life time [97]. An agent that lives longer is allowed to produce more offspring.

Over time Framsticks has been enriched with various features, including sine wave neurons and even neurons encoding fuzzy rules [78]. Yet for all the innovations in Framsticks, the results appear somewhat scarce. Beyond extremely simple crawling creatures, no significant behaviours have been evolved from scratch by the system. However it should be kept in mind that Framsticks is a recent work in constant progress. An active community has developed, and intensive development continues together with further experimentation, which means that, in contrast to previous experiments, new results can reasonably be expected.

B.8 Appendix

B.8.1 Random Boolean Networks

A RBN is a graph of N nodes (a, b, c... n), each having one boolean output and K boolean inputs coming from other nodes. At each timestep, the state of a node is determined by a boolean function of its inputs, that is, of the state of its input nodes at the previous timestep RBNs were introduced by Kauffman [95] as a tool for modelling genetic regulatory networks, that is, the interaction of genes in a genome that can influence the expression of each other.

Obviously a synchronous, discrete model such as RBNs may be an extreme simplification of the asynchronous, quantitative process of genetic regulation. On the other hand, RBNs offer a wide range of possibilities : N nodes having each K inputs can provide $(N^{K})^{N} * (2^{(}2^{k}))^{N}$ different RBNs; for $N = 3, K = 2$, this gives about 3.10^{6} networks. Besides, RBN dynamics have been analysed quite in depth: whatever its initial state, an RBN is bound to fall into a state that was already encountered, thus resulting into a cycle. These cycles are attractors in the state space, and can be thought of as the global "state" of the whole system after a stabilisation phase.

In a biological metaphor, this system would be a cell, and each node would correspond to a given protein, that would be produced or not depending on the presence or absence of other proteins. Eventually each cell falls into a given attractor, that corresponds to a particular cycle of protein production. This particular cycle corresponds to the type of that cell. RBNs seem to exhibit a surprising stability with regard to mutation, which makes them interesting for evolutionary purposes. Wuensche [195] provides a clear description of RBNs and their application to gene regulation modelling.

Appendix C

Is evolution "just a tautology?"

> When it comes to proving such obvious things, one will invariably fail to convince.
>
> Montesquieu, *L'esprit des lois*, Book XXV, Chapter XIII.

The claim that evolution is "just a tautology" is one illustration of how subtle the Darwinian mechanism is, and how easy it is to thoroughly misunderstand it.

At the suggestion of several correspondents, including Wallace, Darwin attempted to make his argument more intuitive by replacing the term "natural selection" with Herbert Spencer's vivid expression, "survival of the fittest". Unfortunately, this has led to even more confusion. It is sometimes argued that the expression "survival of the fittest" (and therefore, it is claimed, Darwinian evolution in general) is nothing more than a tautology. The reasoning is that if one takes the term "fit" to mean "endowed with phenotypic characteristics which improve chances of survival and reproduction" (which is roughly how Spencer understood it), then "survival of the fittest" can simply be rewritten as "survival of those who are better equipped for surviving and reproducing."

Note that this is not *exactly* a tautology: we might imagine a playful deity that would consistently favour the poorly adapted, and destroy well-adapted creatures, so that "survival of the fittest" might actually *not* occur. However, this is not a very informative statement: it simply reduces to the statement that the game of Life is not rigged in favour of the poorly adapted, which is not controversial. Furthermore, the expression does become a tautology if one uses the most widely accepted definition of "fitness" in modern biology, namely reproductive success itself (rather than any set of characters conducive to this reproductive success). This reasoning is sometimes used to claim that Darwin's entire theory of evolution by natural selection is fundamentally tautological, and therefore devoid of any explanatory power.

However, the expression "survival of the fittest" (taken on its own and out of context) gives a very incomplete account of the mechanism of natural selection. The reason is that it does not mention a key requirement for natural selection, namely the requirement of *heritability.* It is true that the phrase "survival of the fittest", by itself and out of context, is a tautology if fitness is defined by survival and reproduction. However, natural selection is *not* just survival of the fittest. Natural selection is the portion of variation in reproductive success, that is caused by *heritable* (phenotypic) characters (see chapter 2.)

If certain heritable characters increase or decrease the chances of survival and reproduction of their bearers, then it follows mechanically (by definition of "heritable") that those characters that improve survival and reproduction will increase in frequency over generations. This is precisely what is called "evolution by natural selection." On the other hand, if the characters which lead to differential reproductive success are not heritable, then no meaningful evolution will occur, "survival of the fittest" or not: if improvement in reproductive success is caused by traits that are not heritable, then there is no reason why these traits should increase in frequency over generations. In other words, natural selection does not simply state that "survivors survive" or "reproducers reproduce"; rather, it states that "survivors survive, reproduce and *therefore* differentially propagate any *heritable* characters that have affected their survival and reproductive success." This statement is not tautological: it hinges on the testable hypothesis that such fitness-impacting heritable variations actually exist (a hypothesis that has been amply confirmed.)

Appendix D

Levels of selection and the gene's eye view of evolution

D.1 Levels of selection

D.1.1 Evolution and natural selection occur at many levels

We saw that any population of entities that exhibit variation, multiplication and heredity will evolve. Organisms are an obvious example of such entities. But quite clearly, entities of this sort can be found at many levels of observation. For example, it is possible to see entire species as evolving entities: when a speciation event occurs, "daughter" species inherit many of the characteristics of their "parent" species. But they may also differ in other respects, and the differences (which accumulate after the separation) will in turn be transmitted to later species. So species do exhibit heritable features which can change over time.

Natural selection may also occur at many levels. Any entity that can see its replication rate affected by inherited features will undergo natural selection. This is clearly not limited to individuals: tightly bound groups, such as certain insect colonies, may behave in this fashion [189]. On the other hand, within multi-cellular organisms, natural selection can certainly happen among individual cells - though it is of course severely curtailed by the machinery of the organism. In fact, when a group of cells manage to break free from the control process and replicate "selfishly", without any regard for the well-being of the organism, they often gain a massive selective advantage - which is to say, their reproductive success is vastly increased with that of their more urbane (and therefore less successful) neighbours. The result, of course, is cancer. At an even lower level, the cases of meiotic drive mentioned above show that selection among genes can occur quite independently of any organismal interaction.

Natural selection does not just occur at many levels: it occurs at many levels *concurrently.* In fact, almost any entity which undergoes variation, multiplication and heredity (that is, any entity that evolves) will encounter various forms of spontaneously occurring selective pressure. New entities will find ways to out-replicate each other, simply because the natural environment is so rich and complex as to offer many opportunities for natural selection to act. The consequence is that, when a selective force acts at a certain level, we must remember that other selective forces keep acting at all lower levels.

D.1.2 Natural selection “at” a certain level

Note that there can be some confusion about whether natural selection occurs at a certain “level” or another. For example, let us consider again the example of species: if an individual of a certain species is born with a highly advantageous feature, this feature will spread through the population, greatly increasing its success; as a result, the species may thrive and branch off into many successful species; these daughter species, which will inherit the beneficial feature, will also prove successful - at least as long as the feature itself is beneficial. On the other hand, related species that do not possess this feature will not share in this success, and may even suffer from the increased competition. This can certainly be seen as a process of evolution by natural selection: change in the inherited characteristics of a population of entities (species) caused by the effect of these characteristics on reproductive success.

However, it is clear that in this case the selective effect on species is entirely *indirect*: the increased success of species is merely an arithmetic consequence of the success of individuals, endowed as they are with a beneficial feature. Crucially, the spread of the differentiating feature can entirely be explained by interactions and selective advantage at the level of individuals alone. It does not seem useful to say that selection has occurred “at the level of the species”, when any selective effect on species is a purely arithmetic artifact of lower-level selection. When can we say that selection occurs “at” a certain level?

Importantly, the level of selection should *not* be equated with the level at which the feature of interest, or the interaction, exists. In this sense, we cannot state that selection occurs (solely) at the “proximal” level. A feature can have independent selective effects at other levels than that of its bearer - for example, the sting of a stinger bee has a (positive) effect on the success of their hives, which is different from the (dramatically negative) effect on its bearer. The former cannot be reduced to mere arithmetic aggregation of the latter. The widespread occurrence of the stinger apparatus, so harmful to those that make use of it, can only be explained by considering the positive effect at the higher level of the hive (and, as we will see, the indirect effect at the level of the gene that codes for it). Neither should it be simply reduced to a question of which entity “benefits” from a certain feature - if only because natural selection can have both beneficial and harmful effects.

Therefore, it seems reasonable to say that natural selection occurs “at” a certain level, if there is variance in contribution to future generations (caused by heritable phenotypic features) between entities at that level, that is not merely the aggregation of effects at another level. As shown by the example of the stinger bee, a single feature can create many selective pressures, sometimes contradictory, at different levels. The net balance of these contradicting forces will determine the fate of the feature: propagation or elimination.[1]

D.1.3 The basis of heredity: the “gene”

In summary, both evolution and natural selection take place at many levels concurrently. However, in nature, these many levels share one crucial aspect, namely the physical basis of heredity. In almost all cases of interest, and independently of the level being considered, this basis happens to be genes, chemically encoded in nucleic acids: genes are the medium through which

---

[1] See also Endler. [49, p.23-24]

heritable features are being transmitted. Whether we consider populations, individuals or (obviously) alleles, we find that the ultimate support of the transmission almost invariably consists in genetic material. For example, at the population level, when populations or groups split, the long-term hereditary message is not transmitted through some diffuse, 'zeitgeist' mechanism: it is embedded within the individuals that set out in each group, and in each of them it can be located in the tangible form of nucleic acids. For almost any evolutionary event, we see that in last analysis, the differential transmission of inherited features (that is, evolution itself) is essentially a differential transmission of genes.

Importantly, this account is only valid for a wide acceptation of the word "gene." In particular, it is difficult to equate "genes" in this sense with individual alleles, simply because some characters have a genetic basis which cannot be reduced to alleles. An obvious example (mentioned with this intent by Sober and Lewontin [166]) is that of heterozygote advantage, in which having two different alleles at a given locus brings an advantage over homozygotes. While heterozygosity certainly has a genetic basis, by definition, no single allele can be the genetic basis for heterozygosity. In fact, the notion of "gene" in this sense is probably most closely approximated by the term "group of alleles." [2]

Evolutionary explanations seek to find out the circumstances which lead certain particular features to propagate in a population, and as we saw this phenomenon can occur at many levels. But if heredity is (almost) always physically embedded in genes, then in order to answer this question, we must ensure that our answer also includes an answer to its gene-level equivalent, namely: why, or when, do this particular gene propagate within a population? In other words, in almost all situations of interest, selection at any level must somehow result in a coherent force of genic selection. Importantly, this does not mean that natural selection only occurs at the level of the gene: it is clear that natural selection can exert its primary force at any level. But to have any real effect, this force must (indirectly) translate into a coherent impulse at the level of the gene.

It follows that any selective force at a certain level can only be have an effect if it can be translated into an indirect, but coherent force at lower levels. Genes are only special in that they happen to occupy the bottom rung of the ladder: since (almost) all heredity is ultimately based on genes, any form of evolution (and natural selection) must use genes as a ultimate medium.

D.1.4 Gene-level selection: *primus inter pares*

This simple fact has profound consequences regarding the validity of various selective explanations of evolution. The point is made concisely and efficiently by Dawkins [37] (emphasis added):

> Before the gene-centred view of natural selection became fashionable, people used to say that if something was good it would happen. This has led some to believe that the adaptationist approach is an easy game. It's been said that you can easily come up with some Darwinian idea to explain anything. As against that, the proper

---

[2] In short, a gene "for" a certain character is any piece of genetic material which, when present in an individual's genotype, will cause or favour the presence of a certain characteristic, assuming all other relevant factors are in their "normal" state (see section D.1.5.)

> understanding of Darwinism at the gene level severely limits you to a certain kind of explanation. It's not good enough just to say that if something is vaguely advantageous it will evolve. You have to say that it's good *for the genes that made it.* That automatically wipes out great swathes of possible facile explanations.

The most prominent of these "facile explanations" is probably the concept of adaptations "for the good of the species." This once-popular notion professed that natural selection could drive individuals to deliberately restrict their own success for the benefit of the species as a whole. A refined version of this idea, centring on durable congregations rather than species in the strict sense, was put forward by Wynne-Edwards under the name "group selection." For example, drawing a comparison with human overfishing and the resulting depletion of fish stocks, Wynne-Edwards contended that animals deliberately limited their own numbers below what immediate resources could sustain, in order to avoid over-exploitation and future starvation [196, pp. 5-9].

The problem of such explanations, of course, is that widespread altruism is vulnerable to the invasion of cheaters, who selfishly enjoy the results of others' altruism without repaying the favour themselves. Williams' critique [188] showed not only that many altruistic behaviours could be explained simply by direct costs and benefits to the bearers of altruistic genes, but that the vulnerability of altruistic groups to invasion by cheaters posed enormous problems to the onset of group selection in general. This illustrates how the gene's eye view of evolution helps avoiding the heed of intuitively appealing, but ultimately misleading explanations.

Again, this does not at all imply that group selection cannot happen; rather, it means that explanations based on group selection (that is, natural selection acting between groups rather than individuals or gene) must ensure their compatibility with the gene's eye view: saying that a feature is "good for the group" is insufficient, it must also be shown to be (indirectly) "good" for the gene that encodes it. An explanation cannot be accepted until it has be translated into gene-selection terms, and the balance of all forces acting at the level of the corresponding genes has been shown to be favourable. This is exactly the path that has been taken recently by authors such as Sober and D. S. Wilson [167] in their thought-provoking efforts to revive interest in group-selection approaches.

It is in this sense that the gene's eye view of evolution can be called *primary*: because genes are the ultimate medium of almost all heredity in Nature, it follows that all levels of natural selection must somehow *include* gene-level selection. A selective process at a certain level will only have an evolutionary effect if it somehow creates a selective force at the gene level. Whether the selective process occurs between groups, or individuals, or genes themselves, if it cannot be translated into a selective advantage for genes, then it can not influence evolution. This is a simple fact; but keeping this simple fact in mind may open vast fields of explanatory mechanisms - some of which are famously described in *The Selfish Gene* [36]. Forgetting it may lead to gross errors of logic.

D.1.5 Limitations of the gene's eye view of evolution

While the gene's eye view of evolution is in a sense primary, it should not be considered as absolute: there are limitations to the centrality of the gene in evolution.

Epistasis and pleiotropy

A widespread, though questionable, objection is that the gene's eye view is inoperant because most phenotypic traits result from the complex interactions of many genes, rather than the mere addition of well-defined contributions from each single gene. The effect of a particular gene on the phenotype will typically be affected by the effects of other genes, a phenomenon known as *epistasis.* This is expressed quite forcefully by Gould [70]:

> Richard [Dawkins] has taken [Darwin's] posture of trying to beat the level of explanation down, and has carried it to its ultimate extreme: it's not even the organisms that are struggling, it's only the genes. The organisms are "vehicles..." The only active agents in Richard's worldview are genes. He's wrong... Richard is basically wrong, because organisms are doing the struggling out there. If organisms could be described as the additive accumulation of what their genes do, then you could say that organisms are representing the genes, but they're not. Organisms have hosts of emergent characteristics. In other words, genes interact in a nonlinear way. It is the interaction that defines the organism, and if those interactions, in a technical sense, are non-additive that is, if you can't just say that it's this percent of this gene plus that percent of that gene then you cannot reduce the interaction to the gene. This is a technical philosophical point. As soon as you have emergent characteristics due to non-additive interaction among lower-level entities, then you can't reduce to the lower-level entities, because the non-additive features have emerged. These features don't exist until you get into the higher level. His argument is wrong.

This objection requires a closer look at the relationship between genes and heritable features. That genes interact non-linearly is obvious - nobody ever argued that epistasis doesn't exist. However, this does not prevent the fact that specific alleles can have specific effects on the phenotype. The reason is that, when we say that a certain gene "controls" a given phenotypic trait, what we really mean is that a variation in that particular gene will cause a (predictable) variation in the trait - *assuming* that other genes remain in their usual, "wild" state. A certain allele at this particular locus will cause the appearance of a certain trait, while other alleles will lead to different traits, with the implicit assumption of a common "genetic background" against which this variation occurs. Any statement that a certain gene "codes for" a certain trait implies a silent "all other things being equal" (or at least, "normal") clause.

Thus, despite non-linear interactions, single-allele differences can have important, predictable effects over the configuration of the organism. Of course, if they did not, then there could not be any heritable characteristic, and Darwinian evolution would be impossible in the first place. If, and only if, these phenotypic effects have an impact on the reproductive success of the organism (or another entity which contains the same allele), then this allele will be favoured by natural selection, and the resulting characteristic will propagate throughout the population. The related argument of pleiotropy, which points out that a single gene can have an effect on many characters at once, seems equally relative.

The limits of the gene-centric view

The true limitation of the gene's eye view of evolution lies elsewhere; in fact, it results from the very same cause which ensures its power in the first place. As we saw, the importance of the gene's eye view derives from the fact that genes are the support of heredity. It immediately follows that, in those case where genes are *not* the support of heredity (that is, when important changes in inherited characteristics occur in a way that is not mediated by genes), then the gene's eye view looses relevance.

There are important examples of momentous evolutionary events in the history of Life, which were not based on genetic changes. Perhaps the most well-known is the endosymbiotic origin of eukaryotic cells, which resulted from the fusion of free-living prokaryotes[112] into single organisms. While it is clear that the fusion events themselves led to numerous later genetic changes, both in the hosts and in the symbionts, and even if we admit that certain genetic predispositions were necessary for this fusion to take place, this does not alter the fact that the fusion *itself* was clearly not a genetic event. Other examples of endosymbiosis (most familiarly the bacterial "gut flora" in humans) also qualify as examples of evolutionary event not directly based on genetic change.

Since the present work is very much oriented towards artificial evolution and artificial life, we must also point out that the primacy of the gene is very much a feature of "life as it is" - that is, biological life as it exists today in nature. Artificial models need not be constrained by the mechanisms of natural life. For example, we may imagine an artificial environment in which individuals reproduce by self- (or mutual-) inspection, rather than relying on an explicit genetic message.[3] In such a world, there is simply no "genes" in any meaningful sense, but evolution and natural selection can still very well occur.

A different problem arises when the gene-centric view of evolution becomes a philosophical description, not just of evolution, but of life itself. In this view, life *is* the process through which replicators (genes) replicate, and the construction of entities (Dawkins' "lumbering robots" [36]) is essentially an interesting side effect. However, these concepts cannot suffice to capture the essence of life: although they address efficiently one of the fundamental aspects of life (reproduction and evolution), they do not consider (or rather, explicitly see as irrelevant) the other fundamental aspect of life, that is, autonomous self-construction and self-organisation. This latter aspect is at least as important as the former to our concept of "life" - indeed, approaches to the problem of defining life are commonly classified according to the relative importance that they attach to each of these two aspects.[155, 31]

This discussion must not be taken as a minimisation of the gene-centric view of evolution. Rather, it suggests that the gene's eye view should be seen as an emerging description, deriving from the organisation of nature, rather than as the philosophical basis of evolution and natural selection. The gene's eye view of evolution is a logically sound, enormously useful way to look at evolution, but reducing the very concept of evolution itself to genes is not desirable. Again, evolution is about changes in heritable characteristics, which are not synonymous with genes. The "selfish gene" is the most fruitful way to assess evolutionary problems (and explanations)

---

[3] Taylor [173, Chap. 7.2.3] discusses these two modes of reproduction, as well as their respective implications and impact over the evolutionary process.

in general, but it should not be seen as a self-sufficient philosophical concept of evolution or (a fortiori) of life.

Appendix E

Solution concepts, monotonicity and progress

In this section we introduce some of Ficici's results on the question of progress in coevolution, as summarised in an important recent paper [54] and in his Ph.D. thesis [53]. While the concepts discussed in the previous sections were developed independently from Ficici, they certainly benefited from comparison with Ficici's results.

E.1 Solution concepts

The principal notion introduced by Ficici is that of a *solution concept* [54]:

> Fundamental to all search problems is the notion of a *solution concept*. Whatever properties our problem domain may possess, and however we embed that domain into a search space, we require a solution concept to indicate which locations in the search space - if any - constitute *solutions* to our problem. A solution concept thus partitions a search space into two classes: Solutions and non-solutions. Typically, the two classes are distinguished in a systematic way - by some number of measurable properties that are present or absent in class members. . . When we apply a particular solution concept to a search space, we obtain a particular *search problem*.

Thus defining a solution concept is equivalent to stating what exactly we are looking for, before starting to look for it. To quote Ficici, "while this point may seem obvious, years of coevolutionary practice indicate otherwise." [53, Conclusion]. Solution concepts discussed by Ficici include the conventional concept of the Best Scoring Strategy (BSS), Pareto-dominance and the Nash equilibrium, among others.

E.2 Preference relations: a set-theoretic definition of superiority

Once a solution concept is formulated, it is possible to define a *preference relation* between two individuals A and B in the following way: essentially, we prefer an individual A to an individual

B if any set of known opponents in which B is a solution (i.e. corresponds to the solution concept when only this particular set is considered) is a sub-set of a set in which A is a solution. For example, for the simple solution concept of "numerical maximum," B=10 is a solution in all sets of integers that include 10 and numbers lower than 10. A=12 is a solution in all sets of integers that include 12 and numbers lower than 12. Every set in which B is a solution is a subset of a set in which A is a solution - as can be trivially shown by simply adding A to any such set. Therefore, we prefer A over B.[1]

One hardly needs to emphasise the elegance of the system: instead of defining a superiority relation, we just "plug in" a given solution concept, and a (global) superiority relation is automatically provided by elementary set theory.

E.3 Monotonicity: when historical progress is equivalent to global progress

An important point of Ficici's formalism is that set theoretic considerations allow for the derivation of powerful results. In particular, Ficici investigates the capacities of search processes which do not discard any information. Let us imagine a search process endowed with an archive, such that every individual encountered by the process so far is kept in the archive indefinitely. At every step, we pick a new individual from the search space (according to some unspecified heuristic), incorporate it into our population, and determine the "champions" of the new population - i.e. those that correspond to the solution concept for this new population - if any. In our terminology it is clear that such a process follows a pattern of perfect historical progress: the new champion is, by definition, superior to its predecessor against the entire history of previous opponents.

At this point Ficici introduces the notion of *monotonicity.* A solution concept can be monotonic or non-monotonic. Monotonicity, for a given solution concept, means that for any three sets $S_1 \subset S_2 \subset S_3$, if A is a solution in S_1 and S_3, then A is also a solution in S_2. If it is possible that an individual A is solution of two sets S_1 and S_3 such that $S_1 \subset S_3$, but not of a set S_2 where $S_1 \subset S_2 \subset S_3$, then the solution concept is called non-monotonic.

Monotonicity means that if no information is discarded, the new solutions found by the evolutionary process will never contradict the preference relation described above: old solutions will never be preferred to new ones. Once we have found a new solution, we can be certain that no previously discarded solution will end up being preferred to this new solution, even after new information is added: no matter how much new information we gain, no previously replaced solution will ever be "rehabilitated". Thus a monotonic solution concept implies that "the quality of the result returned by a search heuristic (assuming that it does not discard information) will also increase monotonically over time." [54]

Using our terminology, a monotonic solution concept is superiority criterion for which historical progress *implies* global progress: whenever a current champion is replaced by a new one (based on historical superiority), we can be certain that the new champion is also *glob-*

---

[1] The terminology here is a simplification of Ficici's. We believe the meaning is appropriately conveyed for the purpose of our discussion. Also note that Ficici's discussion apply to single-population coevolution, but is easily extended to inter-species coevolution.

ally superior to the previous one: we have a guarantee that the current champion is not just historically superior, but *globally* superior to all previous champions. Thus, with a monotonic solution concept, a process that follows a strict historical progress concept will also mechanically observe strict global progress. In fact, this property (of equivalence between historical progress and global progress) can be seen as an alternate definition of monotonicity.[2]

But what if the solution concept is non-monotonic ? In that case, even with perfect information, a constant historical progress (solutions which are consistently superior to their predecessors in competing with the set of all previously encountered opponents) will *not* mechanically lead to global progress; even after thousands of iterations, we will have no guarantee that our current champion is globally superior to previously encountered champions. With a non-monotonic solution concept, no amount of information or algorithmic cunning can make absolutely sure that evolution will not need to backtrack.

E.4 Monotonic solution concepts

Ficici's major achievement is not only to have put together a consistent theoretical framework for coevolution, but also to have demonstrated the monotonicity (or lack thereof) of several solution concepts. For example, Ficici demonstrates (with a simple counter-example) that the "conventional" solution concept of maximum expected score, or "best scoring strategy," is non-monotonic. However, the solution concept of Nash equilibria (as described in section 3.4.3) *is* monotonic. These results, as well as others obtained by Ficici, are undoubtedly of paramount importance for the understanding of coevolution.

While Ficici does not explicitly distinguish between global and historical superiority, these notions can be applied to his work: the preference relation is global in nature (because it is to be estimated over the whole search space), while the complete-information process is clearly historical; the concept of monotonicity provides a link between the two. However this leads Ficici to assert that monotonicity is a pre-requisite for the existence of an objective scale of value between individuals:

> There exists a common belief that objective metrics of goodness are an ill-conceived proposition in coevolution (...) We show that this belief does not transfer to algorithms designed around appropriate solution concepts. [53, Chap. I]
>
> The non-monotonicity of [several solution concepts such as best-scoring strategy] is certainly consistent with the red-queen effect and with the common sentiment that objective metrics of goodness are difficult to obtain for coevolutionary domains (...) The work we present here shifts the focus of attention away from the game and to the solution concept; "monotonicity", which must exist in one sense or another if we are to have an objective metric of goodness, reveals that difficulty in obtaining an

---

[2] Importantly, this only applies to the sequence of *solutions*. Even with a monotonic solution concept, it is perfectly possible that a known individual, which is not currently considered a solution, may be regarded as a solution in the future, as new information is added. However, once it has become a solution, then if it is ever discarded and replaced with a new solution, it will never become a solution again: monotonicity means that demoted solutions will never be promoted back to solution status. (Ficici, personal communication)

> objective metric of performance can be due, in no small part, to the solution concept itself. [53, Chap. IX]

Using our terminology, we can alter this statement. As we have seen, non-monotonous solution concepts can, and do, exhibit a consistent, global scale of "goodness": global superiority provides precisely such a global, intransitive scale of goodness through which any two individuals can (in theory!) be unambiguously compared, based on any solution concept, and on any fitness landscape, that we might think of. The problem, of course, is that this global scale, even though it does exists (and corresponds to the result that we intuitively desire from coevolutionary optimisation) is not directly *observable* in the general case: to evaluate global progress reliably, we would need to have complete knowledge of the results of all possible interactions within the entire search space, which is clearly unfeasible. Ficici's major contribution, based on the concept of monotonicity, is precisely to have determined in which conditions historical progress (which is readily achievable through the use of an archive) does, or does not, automatically create global progress.

E.5 Should we use monotonic solution concepts?

We have seen that in coevolution, progress and superiority can only be meaningfully discussed with regard to a certain set of opponents, and with a well-defined superiority criterion, or solution concept. Ficici has shown that not all solution concepts are equal: some are monotonic, while others are not. "Monotonic" means that we will never have to backtrack and rehabilitate a previously discarded champion: if we enforce perfect historical progress, then we have a guarantee that newer champions are always, not only historically, but globally superior to previous champions. Clearly monotonicity is a desirable property. Does it mean that we should discard non-monotonic solution concepts (including the intuitive, "conventional" Best Scoring Strategy solution concept) and use monotonic solution concepts by default?

The problem with this proposition is that monotonicity comes at a cost. Ficici demonstrates the monotonicity of two solution concepts: Nash equilibria and Pareto dominance (under conditions[3]). While Nash equilibria appear conceptually simple, they imply the use of mixed strategies, that is, sets of strategies from which the individual chooses randomly (according to a certain distribution) at every interaction. This requirement is rather unwieldy from an implementation viewpoint, especially considering that neither the number of strategies in the set nor the probability distribution, are fixed. Note that monotonicity is only guaranteed in the case of exhaustive enumeration of the search space. Exhaustive enumeration of all possible mixed strategies is clearly more costly than exhaustive enumeration of all possible pure strategies. As for Pareto dominance, we have already highlighted its specific problems in section 3.4.3.

In addition to these difficulties, it is worth pointing out the intuitive appeal of the Best Scoring Strategy concept (maximisation of expected score against a random opponent), especially from an optimisation viewpoint. This concept seems to embody the common notion of superiority and progress that coevolutionary optimisation is expected to produce. This is precisely

---

[3] The condition for monotonicity of Pareto dominance is that new candidates which appear identical (that is, obtain identical scores in all currently known dimensions) to those in the current dominance front, should not be included into the front.[53]

why Ficici calls it the “conventional” concept. Pareto dominance is intuitive and easily understood, but widespread mutual non-dominance in complex search spaces limits its appeal. Nash equilibria may not be quite as intuitive, especially considering the need to use mixed strategies. For example, it seems counter-intuitive for a checkers player to randomly pick a new playing method at the start of each new game, independently of its opponent. And how much would we trust a robot that picks a random strategy from its repertoire before performing its task?

Finally, we must keep in mind that even with monotonic solution concepts, steady historical (and global) progress is only guaranteed if we do not discard any information. This may not be realistic in practice.

In summary, the drawbacks of the conventional BSS criterion (namely its non-monotonicity) must be balanced against its advantages. Monotonicity, while desirable, may come at a high costs, either in terms of computation and intuitiveness (for Nash equilibria) or discriminative power (for Pareto dominance). In practice, the intuitive appeal of the BSS criterion may overcome the theoretical cost of non-monotonicity. From a pragmatic perspective, in many cases, it may well be that the best way to deal with non-monotonicity is simply to accept it as a fact of life, and try to find algorithms which are as resistant to it as possible, rather than to use non-conventional solution concepts.

E.6 Summary and conclusion

In coevolution, the concepts of superiority and progress are not as self-evident as in fitness-based evolution and must be carefully defined. To evaluate the superiority of an individual A over another individual B, we must define two things:

1. A superiority criterion (maximum expected utility, Pareto dominance, etc.)
2. A common set of opponents over which this criterion will be assessed and compared between A and B.

In the course of coevolution, different sets of “reference” opponents characterise different notions of progress, namely local progress (superiority of newer individuals over their predecessors, evaluated against current competitors), historical progress (superiority of newer individuals, evaluated against all previously encountered competitors) and global progress (superiority of newer individuals, evaluated against all possible competitors). The first is all that natural selection is concerned with. The second can be brought about by means of an archive or similar algorithmic devices. The third is the most desirable form of progress in coevolution, but is difficult to achieve reliably. In particular, it does *not* mechanically follow from any of the previous two.

Ficici has shown that, for some solution concepts (a notion closely related to that of superiority criteria), historical progress is demonstrably equivalent to global progress: newer champions are better than previous ones, not just against previously encountered opponents, but against the entire search space. However, monotonic solution concepts (such as Nash equilibria and Pareto non-dominated fronts) carry additional requirements and limitations which may make them rather unwieldy in practice. The conventional superiority criterion (Best Scoring Strategy, or maximum expected score against a random opponent), while non-monotonic, may be seen as preferable due to its intuitive appeal and its conceptual and technical simplicity.

Appendix F

Physical complexity

Adami and colleagues [3, 2] have introduced a specific concept called “physical complexity”. They suggest [2] that this quantity measures “the amount of information that is stored in [a genome] about [its] environment”. Using this notion of complexity, they claim [3] that “because natural selection forces genomes to behave as a natural 'Maxwell Demon', within a fixed environment, genomic complexity is forced to increase.” However, physical complexity is quite distinct from the everyday concept of complexity.

Physical complexity (as calculated by equation (4) in [3]) is the total length of the genome (or some section of interest), minus the sum of the entropies of each locus of this genome or section within the population. Entropies are calculated simply by observing the distribution of alleles at this locus, and calculating the standard Shannon entropy of this distribution. The consequence is that physical complexity is, very precisely, a quantitative measure of *convergence* among a population of genomes: it explicitly measures how genetically similar individuals are within a given population.

One consequence is that it is independent of the specific features of the individuals or the environment: physical complexity is only concerned with the degree of convergence within the population, regardless of the environment it has evolved in, or the state it has converged to. Thus, a fully converged population of 1000-loci genomes will have a physical complexity of 1000, whether it has converged on the optimum of the OneMax function (a function that returns the number of ‘ones’ in a sequence of bits) or on an optimal neural network to solve the double-pole balancing problem. Thus, the assumption that “an organism’s complexity is a reflection of the physical complexity of its genome” [3] requires caution.

In practice, however, the environment may have an influence over the dynamics of physical complexity: in a simple environment, with one single, easily reach optimum (say, the OneMax function), the population is more likely to converge totally than in a more complex environment in which many optima are present, so physical complexity is likely to be higher in trivial environments. What’s more, if we consider a complex environment with many optima, then several dynamics can be observed: the evolving population can cover the entire environment and simultaneously explore several optima; or it may narrowly concentrate on one single local optimum and utterly ignore the rest of the environment. By definition, in the latter case (concentration on one single optimum) physical complexity will be higher than in the former case (simultaneous exploration of many optima), precisely because the population will be more concentrated. A naive understanding of the term “information about the environment” might have led to opposite conclusions.

Furthermore, genetic drift, being a cause of diversity loss, will also increase physical complexity. This would imply that random fluctuations qualify as a ‘Maxwell’s demon’. So do population bottlenecks (i.e. episodes of severe random selection), for the same reason. Thus, the statement that physical complexity measures how much “information” the population has accumulated about its environment [3, 2] should not be understood in a standard information-theoretic sense.

We believe that the expression ‘arrow of complexity’ usually refers to the more common notions of morphological and functional complexity. However, those quantities of functional and morphological complexity were explicitly *not* considered (“skirt[ed]” [3]) by Adami and colleagues. Accordingly, we do not include this concept in our discussion.

Bibliography

[1] D. Ackley and M. Littman. Interactions between learning and evolution. In Langton et al. [103], pages 478–507.

[2] C. Adami. What is complexity? *BioEssays*, 24(12):1085–1094, 2002.

[3] C. Adami, C. Ofria, and T.C. Collier. Evolution of biological complexity. *Proceedings of the National Academy of Sciences*, 97(9):4463, 2000.

[4] Christoph Adami, Richard K. Belew, Hiroaki Kitano, and Charles E. Taylor, editors. *Artificial Life VI: Proceedings of the Sixth International Workshop on the Synthesis and Simulation of Living Systems*, Cambridge, MA, 1998. MIT Press.

[5] J. Alroy. Cope's Rule and the Dynamics of Body Mass Evolution in North American Fossil Mammals. *Science*, 280(5364):731–734, 1998.

[6] Peter J. Angeline and Jordan B. Pollack. Competitive environments evolve better solutions for complex tasks. In S. Forrest, editor, *Procs 5th Intl Conf on Genetic Algorithms*, pages 264–270. Morgan Kaufmann, 1993.

[7] Dirk Arnold. Evolution of legged locomotion. Master's thesis, School of Computing Science, 1997.

[8] R. Axelrod. The evolution of strategies in the iterated prisoner's dilemma. In Lawrence Davis, editor, *Genetic algorithms and simulated annealing*, pages 32–41. Morgan Kaufman, 1987.

[9] J.D. Barrow and F.J. Tipler. *The anthropic cosmological principle.* Oxford University Press, 1988.

[10] M. A. Bedau and N. H. Packard. Measurement of evolutionary activity, teleology, and life. In C. Langton, C. Taylor, D. Farmer, and S. Rasmussen, editors, *Artificial Life II: Proceedings of the Workshop on Artificial Life*, pages 431–461, Redwood City, CA, 1992. Addison-Wesley.

[11] M. A. Bedau, E. Snyder, and N. H. Packard. A classification of long-term evolutionary dynamics. In Adami et al. [4], pages 228–237.

[12] M.A. Bedau. Four puzzles about life. *Artificial Life*, 4(2):125–140, 1998.

[13] Christopher M. Bishop. *Neural networks for pattern recognition.* Oxford University Press, 1996.

[14] S.H. Black, M. Shepard, and M.M. Allen. Endangered invertebrates: the case for greater attention to invertebrate conservation. *Endangered Species UPDATE*, 18:42–50, 2001.

[15] E.J.W. Boers and H. Kuiper. Biological Metaphors and the Design of Artificial Neural Networks. Master's thesis, Leiden University, Niels Bohrweg 1, 2333 CA, Leiden, The Netherlands, aug 1992.

[16] E.J.W. Boers, H. Kuiper, B.L.M. Happel, and I.G. Sprinkhuizen-Kuyper. Designing modular artificial neural networks. In H.A. Wijshoff, editor, *Procs of Computing Science in The Netherlands*, pages 87–96, SION, Stichting Mathematisch Centrum, 1993.

[17] J. Bongard, V. Zykov, and H. Lipson. Resilient Machines Through Continuous Self-Modeling. *Science*, 314(5802):1118, 2006.

[18] Josh C. Bongard and Rolf Pfeifer. Repeated structure and dissociation of genotypic and phenotypic complexity in artificial ontogeny. In Spector et al. [168], pages 829–836.

[19] J.T. Bonner. *The Evolution of Complexity by Means of Natural Selection.* Princeton University Press, 1988.

[20] R. Boyd and J.B. Silk. *How humans evolved.* W. W. Norton, New York, 2000.

[21] R. Brooks and P. Maes, editors. *Procs 4th Intl Works on Synthesis and Simulation of Living Systems (ALIFE IV).* MIT Press, 1994.

[22] A. Bucci and J. B. Pollack. A mathematical framework for the study of coevolution. In *Foundations of Genetic Algorithms 7 (FOGA VII)*, 2002.

[23] Seth Bullock and Mark Bedau. Exploring the dynamics of adaptation with evolutionary activity plots. *Artificial Life*, 12:1–5, 2006.

[24] Mikhaïl Burtsev. Measuring the dynamics of artificial evolution. In W. Banzhaf et al, editor, *Proceedings of the Seventh European Conference on Advances in Artificial Life (ECAL 2003)*, volume 2801 of *LNCS*, pages 580–587. Springer, 2003.

[25] A. Cangelosi, S. Nolfi, and D. Parisi. Cell division and migration in a 'genotype' for neural networks. *Network*, 5:497–515, 1994.

[26] J.P. Cartlidge. *Rules of Engagement: Competitive Coevolutionary Dynamics in Computational Systems.* PhD thesis, The University of Leeds, 2004.

[27] A. D. Channon. Unbounded evolutionary dynamics in a system of agents that actively process and transform their environment. *Genetic Programming and Evolvable Machines*, 7(3):253–281, 2006.

[28] A. D. Channon and R. I. Damper. Perpetuating evolutionary emergence. In Pfeifer et al. [140], pages 534–539.

[29] N. Chaumont, R. Egli, and C. Adami. Evolving Virtual Creatures and Catapults. *Artificial Life*, 13(2):139–157, 2007.

[30] J. Clark, J. J. Dolado, M. Harman, R. Hierons, B. Jones, M. Lumkin, B. S. Mitchell, S. Mancoridis, K. Rees, M. Roper, and M. Shepperd. Reformulating software engineering as a search problem. *Journal of IEE Procs - Software*, 150(3):161–175, 2003.

[31] C.E. Cleland and C.F. Chyba. Defining life. *Origins of Life and Evolution of the Biosphere*, 32(4):387–93, 2002.

[32] Dave Cliff and Geoffrey F. Miller. Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In *Proceedings of the Third European Conference on Artificial Life (ECAL-95)*, volume 929 of *Lecture Notes in Artificial Intelligence*, pages 200–218. Springer, 1995.

[33] Michael Conrad and H.H. Pattee. Evolution experiments with an artificial ecosystem. *Journal of Theoretical Biology*, 28:393–409, 1970.

[34] Charles Darwin. *On the Origin of Species.* John Murray, London, 1859.

[35] R. Dawkins. *The extended phenotype: the gene as the unit of selection.* Freeman, 1982.

[36] R. Dawkins. *The selfish gene.* Oxford University Press, 1989.

[37] R. Dawkins. A survival machine. In J.D. Brockman, editor, *The Third Culture: Beyond the Industrial Revolution*, pages 74–95. Simon & Schuster, 1995.

[38] R. Dawkins. *Climbing Mount Improbable.* Penguin, 1996.

[39] Richard Dawkins. *The Blind Watchmaker.* Penguin, London, 1986.

[40] Richard Dawkins. Human chauvinism. *Evolution*, 51(3), 1997.

[41] Richard Dawkins and J. R. Krebs. Arms races between and within species. *Procs of the Royal Society of London, Series* B, 205:489–511, 1979.

[42] E.D. De Jong. The Incremental Pareto-Coevolution Archive. *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-04*, pages 525–536, 2004.

[43] Edwin D. De Jong and Jordan B. Pollack. Ideal evaluation from coevolution. *Evolutionary Computation*, 12(2), 2004.

[44] F. Dellaert and R.D. Beer. Toward an evolvable model of development for autonomous agent synthesis. In R. Maes and P. Maes, editors, *Artificial Life IV.* MIT Press, 1994.

[45] F. Dellaert and R.D. Beer. A developmental model for the evolution of complete autonomous agents. In Maes et al. [109], pages 393–401.

[46] A. K. Dewdney. Computer recreations: In the game called core war hostile programs engage in a battle of bits. *Scientific American*, 250(5):14–22, may 1984.

[47] David Dingli and Martin A. Nowak. Cancer biology: Infectious tumour cells. *Nature*, 443(7107):35–36, 2006.

[48] R.I.M. Dunbar. The social brain hypothesis. *Evolutionary Anthropology*, 6(5):178–190, 1998.

[49] J.A. Endler. *Natural Selection in the Wild*. Princeton University Press, 1986.

[50] D.P. Feldman and J.P. Crutchfield. Measures of statistical complexity: Why? *Physics Letters A*, 238(4):244–252, 1998.

[51] C. Fernando and J. Rowe. Natural selection in chemical evolution. *Journal of Theoretical Biology*, 247(1):152–167, 2007.

[52] C. Fernando and J. Rowe. The origin of autonomous agents. *Biosystems*, 2007. in press.

[53] Sevan G. Ficici. *Solution Concepts in Coevolutionary Algorithms*. PhD thesis, Brandeis University, May 2004.

[54] Sevan G. Ficici. Monotonic solution concepts in coevolution. In *GECCO '05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 499–506, New York, NY, USA, 2005. ACM Press.

[55] Sevan G. Ficici and Jordan B. Pollack. Challenges in coevolutionary learning: arms-race dynamics, open-endedness, and medicocre stable states. In *Procs 6th Intl Conf on Simulation and Synthesis of Living Systems (ALIFE VI)*, pages 238–247. MIT Press, 1998.

[56] Sevan G. Ficici and Jordan B. Pollack. Pareto optimality in coevolutionary learning. In J. Kelemen and P. Sosik, editors, *Advances in Artificial Life: 6th European Conference (ECAL 2001)*. Springer, 2001.

[57] Sevan G. Ficici and Jordan B. Pollack. A game-theoretic memory mechanism for coevolution. In E. Cant-Paz, James A. Foster, Kalyanmoy Deb, Lawrence Davis, Rajkumar Roy, Una-May O'Reilly, Hans-Georg Beyer, Russell K. Standish, Graham Kendall, Stewart W. Wilson, Mark Harman, Joachim Wegener, Dipankar Dasgupta, Mitchell A. Potter, Alan C. Schultz, Kathryn A. Dowsland, Natasa Jonoska, and Julian F. Miller, editors, *Proc. GECCO 2003*. Springer, 2003.

[58] RA Fisher. The genetical theory of natural selection, 1930.

[59] David B. Fogel, editor. *Evolutionary Computation: The Fossil Record*. IEEE Press, 1998.

[60] R.A. Foley. Speciation, extinction, and climatic change in hominid evolution. *Journal of Human Evolution*, 26:277–289, 1994.

[61] R.A. Foley. Pattern and process in hominid evolution. In J. Bintliff, editor, *Structure and Contingency: Evolutionary Processes in Life and Human Society*. Leicester University Press, London, 1998.

[62] Stephanie Forrest and Melanie Mitchell. Relative building-block fitness and the building-block hypothesis. In L. Darrell Whitley, editor, *Foundations of Genetic Algorithms 2*, pages 109–126. Morgan Kaufmann, 1993.

[63] P.D. Gingerich. Stratigraphic record of Early Eocene Hyopsodus and the geometry of mammalian phylogeny. *Nature*, 248:107–109, 1974.

[64] P.D. Gingerich. Rates of Evolution: Effects of Time and Temporal Scaling. *Science*, 222(4620):159–161, 1983.

[65] David E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning.* Addison-Wesley, 1989.

[66] Faustino Gomez and Risto Miikkulainen. Incremental evolution of complex general behavior. *Adaptive Behavior*, 5:317–342, 1997.

[67] G.C. Gould and B.J. MacFadden. Chapter 17: Gigantism, Dwarfism, and Cope's Rule:Nothing in Evolution Makes Sense without a Phylogeny. *Bulletin of the American Museum of Natural History*, 285(1):219–237, 2004.

[68] S.J. Gould. *The Panda's Thumb: More Reflections in Natural History.* WW Norton & Company, 1992.

[69] S.J. Gould. The evolution of life on earth. *Scientific American*, 271(4):84–91, 1994.

[70] S.J. Gould. The Pattern of Life's History. In J.D. Brockman, editor, *The Third Culture: Beyond the Industrial Revolution*, pages 51–73. Simon & Schuster, 1995.

[71] Stephen J. Gould. *Full House / Life's Grandeur.* Jonathan Cape, London, 1996.

[72] Stephen J. Gould. Self-help for a hedgehog stuck on a molehill. *evolution*, 51(3), 1997.

[73] A. Grafen. Biological signals as handicaps. *Journal of Theoretical Biology*, 144(4):517–46, 1990.

[74] John J. Grefenstette. Deception considered harmful. In L. Darrell Whitley, editor, *Foundations of Genetic Algorithms 2*, pages 75–91. Morgan Kaufmann, 1992.

[75] Frederic Gruau. Automatic definition of modular neural networks. *Adaptive Behavior*, 3(2):151–183, 1995.

[76] Frederic Gruau, Darrell Whitley, and Larry Pyeatt. A comparison between cellular encoding and direct encoding for genetic neural networks. In John R. Koza, David E. Goldberg, David B. Fogel, and Rick L. Riolo, editors, *Genetic Programming 1996: Procs of the First Annual Conference*, pages 81–89, Stanford University, CA, USA, 28–31 1996. MIT Press.

[77] P. J. B. Hancock. Genetic algorithms and permutation problems: a comparison of recombination operators for neural net structure specification. In D. Whitley, editor, *Procs of COGANN workshop, IJCNN, Baltimore.* IEEE, 1992. Permutation problem moins severe que prevu.

[78] M. Hapke, M. Komosinski, and D. Waclawski. Application of evolutionarily optimized fuzzy controllers for virtual robots. In *Procs of the 7th Joint Conference on Information Sciences (JCIS7)*, pages 1605–1608. Association for Intelligent Machinery, 2003.

[79] F. Heylighen. The growth of structural and functional complexity during evolution. In F. Heylighen and D. Aerts, editors, *The Evolution of Complexity*, pages 17–44. Kluwer, 1999.

[80] W.D. Hillis. Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D*, 42:228–234, 1990.

[81] J. H. Holland and J. S. Reitman. Cognitive systems based on adaptive algorithms. In D. A. Waterman and F. HayesRoth, editors, *Pattern directed inference systems*. Academic Press, 1978.

[82] John H. Holland. *Adaptation in natural and artificial systems.* University of Michigan Press, Ann Arbor, MI, 1975.

[83] Gregory S. Hornby and Jordan B. Pollack. Body-brain co-evolution using L-systems as a generative encoding. In Spector et al. [168], pages 868–875.

[84] Peter T. Hraber, Terry Jones, and Stephanie Forrest. The ecology of Echo. *Artificial Life*, 3(3):165–190, 1997.

[85] J. Hunt, L.F. Bussière, M.D. Jennions, and R. Brooks. What is genetic quality? *Trends in Ecology & Evolution*, 19(6):329–333, 2004.

[86] P. Husbands, T. Smith, N. Jakobi, and M. O'Shea. Better living through chemistry: Evolving gasnets for robot control. *Connection Science*, 10(4):185–210, 1998.

[87] A. J. Ijspeert, J. Hallam, and D. Willshaw. From lampreys to salamanders: evolving neural controllers for swimming and walking. In Pfeifer et al. [140], pages 390–399.

[88] A.J. Ijspeert, J. Hallam, and D. Willshaw. Evolving swimming controllers for a simulated lamprey with inspiration from neurobiology. *Adaptive Behavior*, 7(2):151–172, 1999.

[89] D. Jablonski. Body-size evolution in Cretaceous molluscs and the status of Cope's rule. *Nature*, 385:250–252, 1997.

[90] Christian Jacob. Evolution programs evolved. In Hans-Michael Voigt, Werner Ebeling, Ingo Rechenberg, and Hans-Paul Schwefel, editors, *Parallel Problem Solving from Nature IV, Procs of the International Conference on Evolutionary Computation*, volume 1141, pages 42–51. Springer Verlag, 22-26 1996.

[91] N. Jakobi. Harnessing morphogenesis. In *Procs of the International Conference on Information Processing in Cells and Tissues*, pages 29–41, 1995.

[92] N. Jakobi. The minimal simulation approach to evolutionary robotics, 1998.

[93] W.R. Jeffery and D.P. Martasian. Evolution of Eye Regression in the Cavefish Astyanax: Apoptosis and the Pax-6 Gene. *Integrative and Comparative Biology*, 38(4):685, 1998.

[94] Edwin De Jong. The MaxSolve algorithm for coevolution. In *GECCO '05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 483–489, New York, NY, USA, 2005. ACM Press.

[95] SA Kauffman. *The Origins of Order.* Oxford University Press, 1993.

[96] H. Kitano. Designing neural networks using genetic algorithms with graph generation system. *Complex Systems*, 4(4):431–476, 1990.

[97] M. Komosinski. The world of framsticks: simulation, evolution, interaction. In *Procs of 2nd International Conference on Virtual Worlds (VW2000), Paris*, pages 214–224. Springer-Verlag (LNAI 1834), 2000.

[98] M. Komosinski and A. Rotaru-Varga. Comparison of different genotype encodings for simulated 3d agents. *Artificial Life*, 7(4):395–418, 2001.

[99] John R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection.* MIT Press, 1992.

[100] N. Krasnogor and D.A. Pelta. Fuzzy Memes in Multimeme Algorithms: a Fuzzy-Evolutionary Hybrid. In José L. Verdegay, editor, *Fuzzy Sets Based Heuristics for Optimization.* Springer, 2003.

[101] M.M. Lahr and R.A. Foley. Towards a theory of modern human origins: geography, demography, and diversity in recent human evolution. *Yearbook of Physical Anthropology*, 41:137–176, 1998.

[102] W. B. Langdon and R. Poli. Fitness causes bloat. In P. K. Chawdhry, R. Roy, and R. K. Pan, editors, *Second On-line World Conference on Soft Computing in Engineering Design and Manufacturing*, pages 13–22. Springer-Verlag London, 23-27 1997.

[103] C. G. Langton, C. Taloy, J. D. Farmer, and S. Rasmussen, editors. *Artificial Life II: Procs of the workshop on the Synthesis and Simulation of Living Systems.* Addison Wesley, 1991.

[104] N. Lassabe, H. Luga, and Y. Duthen. A new step for artificial creatures. In *Procs 1st IEEE Conference on Artificial Life (IEEE-ALife '07)*, page 243. IEEE Press, 2007.

[105] R. Lewin and R.A. Foley. *Principles of Human Evolution.* Blackwell Publishing, 2004.

[106] RC Lewontin. The Units of Selection. *Annual Review of Ecology and Systematics*, 1:1–18, 1970.

[107] H. Lipson and J. Pollack. Automatic design and manufacture of artificial lifeforms. *Nature*, 406:974–978, 2000.

[108] Sean Luke and Lee Spector. Evolving graphs and networks with edge encoding: Preliminary report. In John R. Koza, editor, *Late Breaking Papers at the Genetic Programming 1996 Conference Stanford University July 28-31, 1996*, pages 117–124. Stanford Bookstore, 1996.

[109] Pattie Maes, Maja J. Mataric, Jean-Arcady Meyer, Jordan Pollack, , and Stewart W. Wilson, editors. *From Animals To Animats 4: Procs 4th Intl Conf on the Simulation of Adaptive Behaviour (SAB 96)*. MIT Press, 1996.

[110] C. C. Maley. Four steps toward open-ended evolution. In Wolfgang Banzhaf, Jason Daida, Agoston E. Eiben, Max H. Garzon, Vasant Honavar, Mark Jakiela, and Robert E. Smith, editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 99)*, pages 1336–1343. Morgan Kaufmann, 1999.

[111] D. Marbach and A.J. Ijspeert. Co-evolution of configuration and control for homogenous modular robots. In F. Groen, editor, *Procs of the Eighth Conference on Intelligent Autonomous Systems (IAS8)*, pages 712–719. IOS Press, 2004.

[112] L. Margulis. *Symbiotic planet: a new look at evolution*. Basic Books, 1998.

[113] E. Mayr. Answer to "The abundance of life-bearing planets". *Bioastronomy News*, 7(4), 1995.

[114] Barry McMullin. John von neumann and the evolutionary growth of complexity. *Artificial Life*, 6(4):347–361, 2000.

[115] D. W. McShea. Complexity and evolution: what everybody knows. In David L. Hull and Michael Ruse, editors, *The Philosophy of Biology*. Oxford University Press, 1998.

[116] D.W. McShea. Mechanisms of Large-Scale Evolutionary Trends. *Evolution*, 48(6):1747–1763, 1994.

[117] D.W. McShea. Perspective: Metazoan Complexity and Evolution: Is There a Trend? *Evolution*, 50(2):477–492, 1996.

[118] D.W. McShea. The minor transitions in hierarchical evolution and the question of a directional bias. *Journal of Evolutionary Biology*, 14(3):502–518, 2001.

[119] D.W. McShea. The evolution of complexity without natural selection, a possible large-scale trend of the fourth kind. *Paleobiology*, 31(2):146–156, 2005.

[120] Betrand Mesot. Self-organisation of locomotion in modular robots: A case study. Master's thesis, EPFL, Lausanne, feb 2004.

[121] T. Miconi. When evolving populations is better than coevolving individuals. In Georg Gottlob and Toby Walsh, editors, *IJCAI-03, Procs of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico, August 2003*. Morgan Kaufmann, 2003.

[122] Thomas Miconi and Alastair Channon. Analysing coevolution among artificial creatures. In E. G. Talbi, editor, *Procs Evolution Artificielle 2005 (EA 05)*. Springer-Verlag, 2005.

[123] Thomas Miconi and Alastair Channon. A virtual creatures model for studies in artificial evolution. In *IEEE Congress on Evolutionary Computation (CEC 2005)*, 2005.

[124] Thomas Miconi and Alastair Channon. An improved system for artificial creatures evolution. In Luis Rocha, Mark Bedau, Dario Floreano, Robert Goldstone, Alessandro Vespignani, and Larry Yaeger, editors, *Procs 10th Intl Conf on Simulation and Synthesis of Living Systems (ALIFE X)*. MIT Press, 2006.

[125] Thomas Miconi and Alastair Channon. The N-strikes-out algorithm: A steady-state algorithm for coevolution. In Gary Yen, editor, *Procs IEEE Congress on Evolutionary Computation (CEC 2006)*. IEEE Press, 2006.

[126] G. Miller. *The Mating Mind: How Sexual Choice Shaped the Evolution of Human Nature*. Vintage, 2001.

[127] M. Minsky and S. Papert. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, 1969.

[128] Melanie Mitchell. *An introduction to genetic algorithms*. The MIT Press, 1996.

[129] Jerome Monod. *Chance and Necessity*. William Collins Sons & Co Ltd, 1972.

[130] Gordon D. Moore. Cramming more components onto integrated circuits. *Electronics*, 38(8), 1965.

[131] David E. Moriarty and Risto Miikkulainen. Efficient reinforcement learning through symbiotic evolution. *Machine Learning*, 22:11–32, 1996.

[132] S. Nolfi and D. Parisi. Genotypes for neural networks. In M. A. Arbib, editor, *The Handbook of Brain Theory and Neural Networks*. MIT Press, 1995.

[133] Stefano Nolfi and Dario Floreano. Coevolving predator and prey robots: Do "arms races" arise in artificial evolution? *Artificial Life*, 4(4):311–335, 1998.

[134] R.G. Northcutt. Understanding Vertebrate Brain Evolution. *Integrative and Comparative Biology*, 42(4):743, 2002.

[135] Michael J. T. O'Kelly and Kaihen Hsiao. Evolving mutually perceptive creatures for combat. In P. Vogt, editor, *Procs 9th Intl Conf on Simulation and Synthesis of Living Systems (ALIFE IX)*. MIT Press, 2004.

[136] E. Pachepsky, T. Taylor, and S. Jones. Mutualism Promotes Diversity and Stability in a Simple Artificial Ecosystem. *Artificial Life*, 8(1):5–24, 2002.

[137] N. H Packard. Intrinsic adaptation in a simple model for evolution. In Langton et al. [103].

[138] L. Pagie and M. Mitchell. A comparison of evolutionary and coevolutionary search. *International Journal of Computational Intelligence and Applications*, 2(1):53–69, 2002.

[139] H. H. Pattee. Simulations, realizations, and theories of life. In C. G. Langton, editor, *Artificial life, volume VI of Santa Fe Institute Studies in the Sciences of Complexity*, pages 63–77, 1989.

[140] R. Pfeifer, B. Blumberg, J.-A. Meyer, and S. Wilson, editors. *From Animals To Animats 5: Proceedings 5th International Conference on the Simulation of Adaptive Behavior (SAB '98)*, Cambridge, MA, 1998. MIT Press.

[141] T. Pfeiffer, S. Schuster, and S. Bonhoeffer. Cooperation and Competition in the Evolution of ATP-Producing Pathways. *Science*, 292(5516):504, 2001.

[142] Jordan B. Pollack, Alan D. Blair, and Mark Land. Coevolution of a backgammon player. In Christopher G. Langton and Katsunori Shimohara, editors, *Artificial Life V: Procs. 5th Int. Workshop on the Synthesis and Simulation of Living Systems.* MIT Press, 1997.

[143] Mitchell A. Potter and Kenneth De Jong. A cooperative coevolutionary approach to function optimization. In Yuval Davidor, Hans-Paul Schwefel, and Reinhard Männer, editors, *Parallel Problem Solving from Nature – PPSN III*, pages 249–257. Springer, 1994.

[144] P. Prusinkiewicz and A. Lindenmayer. *The Algorithmic Beauty of Plants.* Springer, 1990.

[145] T. S. Ray. Aesthetically evolved virtual pets. In Carlo C. Maley and Eilis Boudreau, editors, *Artificial Life 7 Workshop Procs*, pages 158–161, 2000.

[146] Thomas S. Ray. An approach to the synthesis of life. In Langton et al. [103], pages 371–408.

[147] Thomas S. Ray. Artificial life. In Renato Dulbecco, David Baltimore, Francois Jacob, and Rita Levi-Montalicini, editors, *Frontiers of Life, Volume One: The Origins of Life*, pages 107–124. Academic Press, 2001.

[148] Thomas S. Ray and Joseph Hart. Evolution of differentiated multi-threaded digital organisms. In Adami et al. [4].

[149] Craig W. Reynolds. Competition, coevolution and the game of tag. In Brooks and Maes [21].

[150] Mark Ridley. *Evolution.* Blackwell Scientific, 1993.

[151] Christopher D. Rosin and Richard K. Belew. Methods for competitive co-evolution: Finding opponents worth beating. In Larry Eshelman, editor, *Procs of the Sixth International Conference on Genetic Algorithms*, pages 373–380. Morgan Kaufmann, 1995.

[152] Christopher D. Rosin and Richard K. Belew. New methods for competitive coevolution. *Evolutionary Computation*, 5(1):1–29, 1997.

[153] B. Rosslenbroich. The Notion of Progress in Evolutionary Biology–The Unresolved Problem and an Empirical Suggestion. *Biology and Philosophy*, 21(1):41–70, 2006.

[154] G. Roth, K. C. Nishikawa, and D. B. Wake. Genome size, secondary simplification, and the evolution of the brain in salamander. *Brain behavior and evolution*, 50(1):50–59, 1997.

[155] Kepa Ruiz-Mirazo, Juli Peretó, and Alvaro Moreno. A universal definition of life: Autonomy and open-ended evolution. *Origins of Life and Evolution of the Biosphere*, 34(3), 2004.

[156] D.E. Rumelhart, G.E. Hinton, and R.J. Williams. Learning representations by back-propagating errors. *Nature*, 323:533–536, 1986.

[157] PC Sabeti, SF Schaffner, B. Fry, J. Lohmueller, P. Varilly, O. Shamovsky, A. Palma, TS Mikkelsen, D. Altshuler, and ES Lander. Positive Natural Selection in the Human Lineage. *Science*, 312(5780):1614–1620, 2006.

[158] K. Safi, M.A. Seid, and D.K.N. Dechmann. Bigger is not always better: when brains get smaller. *Biology Letters*, 1(3):283–286, 2005.

[159] Hans-Paul Schwefel and Gunter Rudolph. Contemporary evolution strategies. In *European Conference on Artificial Life*, pages 893–907, 1995.

[160] Y.S. Shim and C.H. Kim. Evolving Physically Simulated Flying Creatures for Efficient Cruising. *Artificial Life*, 12(4):561–591, 2006.

[161] A. Siddiqi and S. Lucas. A comparison of matrix rewriting versus direct encoding for evolving neural networks. In *Procs of the 1998 IEEE International Conference on Evolutionary Computation*, pages 392–397. IEEE Press, 1998.

[162] Karl Sims. Evolving 3d morphology and behavior by competition. In Brooks and Maes [21], pages 28–39.

[163] Karl Sims. Evolving virtual creatures. In *SIGGRAPH 94*, pages 15–22. ACM Press, 1994.

[164] B. Sinervo and C.M. Lively. The rock-paper-scissors game and the evolution of alternative male strategies. *Nature*, 380(6571):240–243, 1996.

[165] John Maynard Smith and Eors Szathmary. *The Major Transitions in Evolution.* Oxford University Press, 1995.

[166] E. Sober and R.C. Lewontin. Artifact, Cause and Genic Selection. *Philosophy of Science*, 49(2):157–180, 1982.

[167] E. Sober and D.S. Wilson. *Unto Others: The Evolution and Psychology of Unselfish Behavior.* Harvard University Press, 1999.

[168] Lee Spector, Erik D. Goodman, Annie Wu, and W. B. Langdon, editors. *Proceedings of the GECCO 2001 conference*, San Francisco, 2001. Morgan Kaufmann.

[169] Russell K. Standish. Diversity evolution. In Russel K. Standish, Mark Bedau, and Hussein A. Abbas, editors, *Artificial Life VIII: Proceedings of the Eighth International Workshop on the Synthesis and Simulation of Living Systems*, pages 131–137. MIT Press, 2002.

[170] Kenneth O. Stanley and Risto Miikkulainen. The dominance tournament method of monitoring progress in coevolution. In *Procs GECCO 2002 Workshop*. Morgan Kaufman, 2002.

[171] Kenneth O. Stanley and Risto Miikkulainen. Efficient reinforcement learning through evolving neural network topologies. In *Procs GECCO 2002*. Morgan-Kaufmann, 2002.

[172] Kenneth O. Stanley and Risto Miikkulainen. A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130, 2003.

[173] Tim Taylor. *From artificial evolution to artificial life*. PhD thesis, University of Edinburgh, 1999.

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

[175] J. G. Thewissen, M.J. Cohn, L.S. Stevens, S. Bajpai, J. Heyning, and W.E. Horton Jr. Developmental basis for hind-limb loss in dolphins and origin of the cetacean bodyplan. *Proc Natl Acad Sci U A*, 2006.

[176] D. Tikk, L. T. Koczy, and T. D. Gedeon. A survey on the universal approximation and its limits in soft computing techniques. *Int. J. of Approx. Reasoning*, 33(2):185–202, June 2003.

[177] Leigh Van Valen. A new evolutionary law. *Evolutionary Theory*, 1:1–30, 1973.

[178] J.W. Valentine, A.G. Collins, and C.P. Meyer. Morphological Complexity Increase in Metazoans. *Paleobiology*, 20(2):131–142, 1994.

[179] Jeffrey Ventrella. sexual swimmers: Emergent morphology and locomotion without a fitness function. In Maes et al. [109].

[180] Jeffrey Ventrella. Attractiveness vs efficiency. In Adami et al. [4], pages 484–492.

[181] G.P. Wagner and L. Altenberg. Complex adaptations and the evolution of evolvability. *Evolution*, 50(3):967–976, 1996.

[182] Alfred Wallace. The limits of natural selection as applied to man. In *Contributions to the Theory of Natural Selection*, chapter 10. Macmillan, 1875.

[183] R. A. Watson and J. B. Pollack. Coevolutionary dynamics in a minimal substrate. In Spector et al. [168].

[184] J. Weiner. *The Beak of the Finch*. Vintage, 1994.

[185] W.B. Whitman, D.C. Coleman, and W.J. Wiebe. Prokaryotes: The unseen majority. *Proc. Natl. Acad. Sci. USA*, 95:6578–6583, 1998.

[186] R. Paul Wiegand. *An Analysis of Cooperative Coevolutionary Algorithms.* PhD thesis, George Mason University, 2003.

[187] R. Paul Wiegand, William C. Liles, and Kenneth A. De Jong. An empirical analysis of collaboration methods in cooperative coevolutionary algorithms. In Spector et al. [168], pages 1235–1242.

[188] G.C. Williams. *Adaptation and natural selection.* Princeton University Press, Princeton, 1966.

[189] E.O. Wilson. Kin Selection as the Key to Altruism: Its Rise and Fall. *Social Research: An International Quarterly of Social Sciences*, 72(1):1–8, 2005.

[190] Stewart W. Wilson. Classifier systems and the animat problem. *Machine Learning*, 2:199–228, 1987.

[191] Stewart W. Wilson. Classifier fitness based on accuracy. *Evolutionary Computation*, 3(2):149–175, 1995.

[192] Stewart W. Wilson and David E. Goldberg. A critical review of classifier systems. In *In Procs of the Third International Conference on Genetic Algorithms (ICGA-3)*, pages 244–255. Morgan Kaufmann, 1989.

[193] David H. Wolpert and William G. Macready. No free lunch theorems for search. Technical Report SFI-TR-95-02-010, Santa Fe Institute, 1995.

[194] Robert Wright. *Nonzero: The Logic of Human Destiny.* Vintage, 2001.

[195] A. Wuensche. Genomic regulation modeled as a network with basins of attraction. In R.B.Altman, A.K.Dunker, and L.Hunterand T.E.Klien, editors, *Pacific Symposium on Biocomputing*, pages 89–102. World Scientific, 1998.

[196] V.C. Wynne-Edwards. *Animal Dispersion in Relation to Social Behaviour.* Oliver and Boyd, 1962.

[197] L. Yaeger. Computational genetics, physiology, metabolism, neural systems, learning, vision and behaviour or polyworld: Life in a new context. In C. G. Langton, editor, *Artificial Life III, Vol. XVII of SFI Studies in the Sciences of Complexity*, pages 263–298. Addison-Wesley, 1994.

[198] Xin Yao. Evolving artificial neural networks. *Procs of the IEEE*, 87, 1999.