

Evolution in Natural and Artificial Systems
Thesis Proposal
RSMG Report 3

Thomas Miconi
Supervisor: Alastair Channon
Thesis group members: Aaron Sloman, John Bullinaria

August 24, 2004

Abstract

The goal of this research is to study the conditions in which evolution may lead to the sustained emergence of novel behaviours, and how this may be applied to the automatic design of complex entities. We argue that with regard to artificial evolution, this field of study has been rather overlooked, with a preference given to mathematical or experimental results based on abstract selection models. We argue further that when practitioners of artificial evolution have tried to harness the dynamics of more natural selective processes, this relative lack of connection with traditional evolutionary theory has often led to implicit assumptions and “common-sense” intuitions, which may have prevented a better understanding (and more reasonable expectations) about these systems.

Based on a study of evolutionary theory, we expect to provide reasonably clear definitions of what we may expect to observe in evolutionary experiments. We will put these ideas to test through the realisation of a software platform which will allow us to simulate the evolution of three-dimensional creatures interacting in a physically realistic world. This software platform will place both morphology and control of the creatures under evolutionary control, thus allowing them to construct their own behaviours in response to their environment. By choosing different settings for competition, selection, reproduction and representation of the creatures, we will try to discern which conditions allow interesting behaviours or evolutionary dynamics to emerge.

1 Outline of the Problem

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 would baffle many 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 shaping 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 inventivity of life. Explaining that 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 use of computer to study concepts related to life and evolution is almost as old as computers themselves. John Von Neumann applied Turing's universal computing automaton 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 [36]). Soon afterwards, in an apparently unrelated effort, engineers attempted to use 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 flourished into closely related lines of research in artificial life and evolutionary computation.

1.2 Discrepancies between natural and artificial evolution

The development of evolutionary computation was either empirical (as for Evolutionary Strategies [52], developed by engineers), or based on abstract mathematics (the prime example of which is Holland's Genetic Algorithm, as described below). While both approaches were natural in settings for which the selective process was based on external fitness functions, they may be difficult to translate to the study of artificial evolution in more complex settings which come closer to natural reality. In the latter case, classical evolutionary theory, and more generally the field of biology, might provide a valuable source of insight. Unfortunately, there seems to be 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 (see e.g. Pattee [44]).

A notable example of this lack of grounding is related to the notion of progress in evolution. The idea that evolution, and more precisely natural selection, leads to a globally monotonous progress, and that this progress is linked

to a general increase in complexity, pervades much of the literature. Let us take a few examples:

Humans have been practicing applied evolution since the dawn of agriculture. However, our management of evolution has taken place at the “micro” level, the alteration of existing species. We have never been able to harness and manage the more creative properties of evolution: the origin of new species, and the emergence of complexity itself (Ray [47]).

... the coevolution of competing populations may produce increasingly complex evolving challenges. As discussed by Dawkins and Krebs [15], competing populations may reciprocally drive one another to increasing level of complexity by producing an evolutionary “arms race”... (Nolfi and Floreano [42]).

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” [12]. 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 [50]).

Intuitively, the distinctive mark of 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 [4]).

These passages express an implicit intuition that seems to be deeply entrenched within the educated public; in the (ironic) words of McShea [37], they state “what everybody knows”.

Yet 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 a global overall progress. Are we absolutely certain that modern lifeforms are in any sense better than their ancestors? How could we compare a modern bird with a cretaceous pterosaur? Even less clear is the necessity for any such global improvement to be related to an overall drive towards complexification.

Indeed, as we will see shortly, the notion of evolutionary progress is a highly controversial one in natural science, and has been so from the beginning of evolutionary theory. Let us take two different statements from major evolutionary biologists:

During the past billion years, animals as a whole have evolved upwards in body size, feeding and defensive techniques, brain and behavioral complexity, social organisation and precision of environmental control (...) Progress, then, is a property of the evolution of life as a whole by almost any conceivable intuitive standard, including the acquisition of goals and intentions in the behavior of animals (...) Attentive to the adoration of C.S. Peirce, let us not pretend to deny in our philosophy what we know in our hearts to be true. (E.O. Wilson, quoted by Gould in [21])

and

The notion of progress has a bad name among evolutionary biologists (...) The fossil record shows that many organisms - horseshoe crabs, the coelacanth, crocodiles, for example, have undergone little change, progressive or otherwise, for hundreds of millions of years. On a shorter timescale, sibling species tell the same story. The fruit flies *Drosophila melanogaster* and *D. simulans* are hard to distinguish morphologically, but molecular data indicate that they are separated by several million years of evolution. Hence, either morphological evolution in the two species has been almost exactly parallel, which is implausible, or neither species has changed.

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 Szathmary [55]).

The consensus in biology seems to lean much more on Maynard-Smith and Szathmary's side. At the very least, the term "progress" is regarded as a dangerously imprecise word which should be avoided altogether. Even the more discreet term "complexity" is seen with suspicion. Such caution does not seem to prevail in the artificial evolution community.

Similarly, the distinction between selective processes based on pre-defined fitness function (equated by some to "artificial selection") on the one hand, and

selective processes determined by interaction between coevolving individuals on the other, has led to conjectures about the difference in “power” that might be expected from these two types of selection. For example, Ray [48] states:

Humans have been practicing applied evolution since the dawn of agriculture. However, our management of evolution has taken place at the “micro” level, the alteration of existing species. We have never been able to harness and manage the more creative properties of evolution: the origin of new species, and the emergence of complexity itself. We are able to guide the evolution of poor quality wild corn into high quality domestic corn, however, we can not guide the evolution of algae into corn.

The inoculation of the process of evolution into media created by humans opens a new chapter in our relationship to evolution. Work with the new synthetic evolutions may allow us to enter into a new relationship with evolution, in which we can manage these more creative aspects of the process. However, this will require new approaches to working with evolution. These higher objectives can not be achieved through the traditional approach of breeding captive populations (applying fitness functions) (...)

Similarly, Channon [6] states:

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. (Emphasis original)

Again, the conclusion is not exactly obvious. Experiments based on external fitness functions may, and do, lead to remarkable levels of complexity and to the emergence of interesting, complex behaviours that were not explicitly specified in the fitness function. Again, both argumentations seem to rest on the assumption that by evolving against each other, organisms build a kind of mutual scaffolding that allow them to reach ever higher levels of complexity. This idea is based on the concept of evolutionary “arms race” evoked by Dawkins and Krebs [14] and described in the next section. It was bolstered by the initial successes of Hillis [26] and Axelrod [2]. However the conceptual justification for this transformation of local selection into a global improvement process has been slim for quite some time, despite recent studies (such as e.g. Nolfi and Floreano [42]).

The question is made more difficult by the lack of a precise definition of artificial vs. natural selection in computer experiments. In nature, the defining factor of artificial selection is simply human intervention. But in computer experiments, where humans define the totality of the environment, the distinction may not be so obvious. Packard [43] provides a reasonable distinction between

extrinsic adaptation (governed by external factors such as a fitness function) and *intrinsic* adaptation (governed by the interactions between individuals). However, as we will see, this distinction may prove ambiguous in some situations. It is possible that the distinction between artificial and natural selection is less straightforward in computers than in Nature.

More generally, 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 [59]:

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:

Although [our system] Cosmos was designed to study evolution, and in particular (...) the evolution of multicellular organisms from unicellular ones, it was not built around any particular theory of what the important features of this transition might have been. Features such as the 2D environment, energy tokens and so on were included for the reasons discussed in Chapter 4, but there were no coherent theoretical reasons for deciding which features should be modelled, and which should be left out. This weakness is not specific to Cosmos, but is shared by all the other Tierra-like systems I have come across. (...)

I am certainly not the first person to criticise artificial life models on these grounds (...) For example, Howard Patee 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” [44]. (...)

The ad hoc feel of Tierra-like 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.

1.3 Lack of structural freedom in artificial evolution

Another restriction of most artificial evolution experiments is the severe constraints imposed on how much of the individuals’ capacity to interact with their environment is under evolutionary control. Even when agents do have a body, i.e. when they are not represented as purely abstract entities, but as situated actors which can have an influence over (and be influenced by) 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 [28]) , reproduction or fight (as in Polyworld [62]).

Few experiments allow individual to construct their own actions through evolution, which requires that both morphology and control architecture are under evolutionary control. Bongard [5] attempts at doing this with a strongly biologically-inspired system that is practically a model of biological embryology including complex gene regulation. However the most influential experiments in that domain are undoubtedly Karl Sims’s “block” creatures [54], which despite ad hoc features in their machinery allowed for an unprecedented freedom in the evolutionary control over both morphology and behaviours of agents.

The possibility for organism to construct their own behaviours by themselves, instead of choosing between a few actions among a closed pre-defined repertoire, may allow for an increase in the openness of the evolutionary process that may be on par with the distinction between artificial and natural selection. At any rate, the question deserves to be studied and experimentally tested.

2 Importance of the problem

Understanding the dynamics of evolutionary processes is important if we want to put these dynamics to use. In particular, as can be seen from excerpts presented above, evolutionary techniques that use more natural methods, in which pre-defined selection criteria are replaced with natural interaction between individuals, have induced many expectations, some of which may be difficult to justify formally. Clearing up these concepts, and understanding what evolution can be expected to do (and how it can do it), is clearly a major objective of evolutionary research. This research may be seen as a step on this path.

It is also important that we use a study model that minimises constraints on the evolutionary process (or more precisely, allow us to make these constraint vary in laxity as much as possible). First, because it offers a larger field for evolutionary dynamics to develop in; second, because constraints imposed by our environment may lead to a bias in the outcome of our experiments. The type of model we have in mind (to be described in more detail below) has only been implemented once, by Karl Sims: an overwhelming majority of research in artificial evolution focuses on much simpler systems. It is also notable that attempts at reproducing these works have not been fully successful so far. The results presented by Sims certainly leave a lot of space for further investigation.

We still have a lot to learn from evolutionary experiments in complex, open worlds; our project will try to shed more light on these systems, with the hope of gaining some insight about the possibility to use evolutionary techniques for the design of complex entities.

3 The concept of Progress in Evolution

3.1 “What everybody knows”

The elusive problem of evolutionary progress finds its root in the oldest mystery of biology. A striking aspect of nature is the abundance of integrated functional complexity, that is, the amount of inter-related, inter-dependent parts at each level of the living world. Nature’s capacity to build complex structures out of interdependent parts, each of which often seems uncannily adapted to its role, and in such a way that the removal of one part often prevents the proper functioning of the whole (either at the level of individual organs, organisms, or whole populations) has been a defining character of life ever since Aristotle.

Furthermore, all entities do not exhibit the same amount of complexity. From sponges to human, it is quite obvious that some organisms are more complex than others, whatever the precise meaning of “complex” could be. The idea that there should be a hierarchy between lifeforms according to their complexity, leading to the notion of “higher” and “lower” animals, is not *a priori* contemptible. We are naturally inclined to attribute more value to complex things, if only because we feel that they would be more difficult to make. Thus if an organism is more complex than another, it seems to denote a superior amount of skill and craftsmanship from Nature, or from whichever Creator should be credited. In a pre-evolutionary world, in which species were supposed to be fixed, defining a “great chain of being” that would extend from the “lower”, simpler entities to the “higher”, more complex ones, thus defining complexity as an arbitrary notion of value, seemed rather innocuous.

The problems begin when evolution is added in the model. If evolution, that is, the change of species and their derivation from each other, can explain the variety of living species, the complex arrangement of their parts and their adaptation to their environment, there is a strong temptation to use the same mechanism to explain the apparent progression from lower organisms to higher one. This is exemplified in E.O. Wilson’s comment cited above.

3.2 Progress in pre-darwinian evolution

Curiously enough the first evolutionist, the arch-naturalist Buffon¹, does not assert any evolutionary progress: even though he does use the verb “degenerate” to describe the branching of new species from others, this mechanism does not involve a scale of value. Zebras are not superior to horses, horses are not superior to donkeys (Buffon’s vivid, moving eulogy of the donkey is a famous passage of his Natural History). Their differences simply reflect the influences of the external environment upon them, slowly integrated into the species through inheritance of acquired characters. However it should be noted that Buffon

¹Inspection of Buffon’s works confirms that he did embrace the concept of evolution, seen as the derivation of species from and into each other, and was probably the first to do so. Mayr [35] contends that Buffon only evoked evolution to dismiss it; this is true in the passage quoted by Mayr and others, but other passages leave no doubt about the changes in Buffon’s ideas.

only envisaged evolution and branching of similar species within a genus, not common descent of all living organisms.

Lamarck's theory of evolution extends and emphasises the mechanism that Buffon suggested, but does so with a clearly progressive stance. Lamarck explicitly posited that evolution was first and foremost an immanent tendency of living forms to complexify. Just as the physical influenced the moral, the moral could influence the physical. Species evolve because the appearance of new needs lead to the emergence of new organs, or the development of existing ones:

First law: In every animal which has not exceeded the limit of its development, the more frequent and sustained use of any organ gradually strengthens this organ, develops it, makes it larger, and gives it a power proportional to the duration of this use; whereas, the constant lack of use of such an organ imperceptibly weakens it, makes it deteriorate, progressively diminishes its faculties, and ends by making it disappear.

Second Law: Everything which nature has made individuals acquire or lose through the influence of conditions to which their race has been exposed for a long time and, consequently, through the influence of the predominant use of some organ or by the influence of the constant disuse of this organ, nature preserves by reproduction in the new individuals arising from them, provided that the acquired changes are common to the two sexes or to those who have produced these new individuals.

All along his work, Lamarck explicitly and relentlessly stresses not only the notion of a ladder from lower to higher lifeforms (an idea which, as he acknowledged, had a long tradition), but also how his own theory is essentially an explanation of this ladder. The words "more perfect" and "more imperfect" are used extensively throughout, especially in his description of the animal kingdom in 14 classes (from unicellulars to mammals), all arranged within a continuous progression of complexity. This gradation in complexity does correspond to a common descent: species derive from each other, and elementary forms of life are constantly being created².

Man is the explicit point of reference: "it is evident that the organic structure of human beings is the most perfect and must be looked upon as the standard against which we must judge the improvement or the degradation of other organic structures in animals". Viviparity, internal skeletons, lungs supported by a diaphragm, are seen as intrinsically and evidently "more perfect" than their non-human alternatives.

²Indeed Lamarck was the first to explicitly propose (with a remarkable amount of oratory precautions) that men descend from apes, and provided a hypothesis for this transition (a particularly clever species of "quadrumanes" lost the habit of climbing up trees, thereby freeing their hands for other purposes, and managed to spread over every inhabitable region of the world) which vaguely resembles current theories.

3.3 Darwin's progress

Darwin, who posited natural selection among variants as the central mechanism of evolution, has had notoriously fleeting opinions about the theme of evolutionary progress. It must first of all be said that in Darwin's mind, natural selection is undoubtedly progressive in a local, adaptive way: because living organisms are constantly engaged in a struggle for existence, those who survive are in some way better than those who don't. But how much of this local improvement may be said to be cumulative, and thus result in a global progress - and a progress towards what?

The following passage of the *Origins of Species* [11] contains a valuable summary:

As natural selection acts by competition, it adapts the inhabitants of each country only in relation to the degree of perfection of their associates; so that we need feel no surprise at the inhabitants of any one country, although on the ordinary view supposed to have been specially created and adapted for that country, being beaten and supplanted by the naturalised productions from another land. Nor ought we to marvel if all the contrivances in nature be not, as far as we can judge, absolutely perfect; and if some of them be abhorrent to our ideas of fitness. We need not marvel at the sting of the bee causing the bee's own death; at drones being produced in such vast numbers for one single act, and being then slaughtered by their sterile sisters; at the astonishing waste of pollen by our fir-trees; at the instinctive hatred of the queen bee for her own fertile daughters; at ichneumonidae feeding within the live bodies of caterpillars; and at other such cases. The wonder indeed is, on the theory of natural selection, that more cases of the want of absolute perfection have not been observed.

It is obvious that the bee would be "more perfect" if it could sting and live, or that the fir-tree would be "more perfect" if every grain of pollen had a chance to germ. Similarly the vertebrate eye is "more perfect" than the spider's, not because it is more complex, but because this complexity allows us to see more things, more precisely. A spider with the eyes of a sparrow might be expected to perform much better in the struggle for existence (and reproduction). In our modern, gene-centered view, if the genome of a given spider allowed it to have the eyes of a sparrow, we may suspect that these genes would propagate quickly throughout the population. This is what Darwin means in this famous passage:

As all the living forms of life are the lineal descendants of those which lived long before the Silurian epoch, we may feel certain that the ordinary succession by generation has never once been broken, and that no cataclysm has desolated the whole world. Hence we may look with some confidence to a secure future of equally inappreciable length. And as natural selection works solely by and for the good

of each being, all corporeal and mental endowments will tend to progress towards perfection.

The point is made even more clearly in a specific section “On the state of Development of Ancient Forms” (Chap. X):

There has been much discussion whether recent forms are more highly developed than ancient. I will not here enter on this subject, for naturalists have not as yet defined to each other’s satisfaction what is meant by high and low forms. But in one particular sense the more recent forms must, on my theory, be higher than the more ancient; for each new species is formed by having had some advantage in the struggle for life over other and preceding forms. If under a nearly similar climate, the eocene inhabitants of one quarter of the world were put into competition with the existing inhabitants of the same or some other quarter, the eocene fauna or flora would certainly be beaten and exterminated; as would a secondary fauna by an eocene, and a palaeozoic fauna by a secondary fauna. (...) but I can see no way of testing this sort of progress. Crustaceans, for instance, not the highest in their own class, may have beaten the highest molluscs. From the extraordinary manner in which European productions have recently spread over New Zealand, and have seized on places which must have been previously occupied, we may believe, if all the animals and plants of Great Britain were set free in New Zealand, that in the course of time a multitude of British forms would become thoroughly naturalized there, and would exterminate many of the natives. On the other hand, from what we see now occurring in New Zealand, and from hardly a single inhabitant of the southern hemisphere having become wild in any part of Europe, we may doubt, if all the productions of New Zealand were set free in Great Britain, whether any considerable number would be enabled to seize on places now occupied by our native plants and animals. Under this point of view, the productions of Great Britain may be said to be higher than those of New Zealand.

Darwin’s notion of global, historic progress does not involve a metaphysical tendency of Life to complexify. It relies on the idea that local improvement, defined as a better performance in the struggle for existence, is *cumulative*, at least in part, over time and space (that is, in relation to the amount of competition to which a given lineage has been exposed). In itself, this notion makes sense and may be argued about. But even in this very passage the Great Chain of Being insinuates its pervasive grip: more recent, “higher” groups (crustaceans) are supposed to have “beaten”, i.e. to be better than, mollusks.

We may discuss the idea that species which are alive now have in a sense “beaten” those which have disappeared (or at least, those with which they were in competition). But if crustaceans have “beaten” molluscs, why do we still see molluscs in our modern world, even where crustaceans abound? In which

sense have lobsters “beaten” mussels? Darwin, even though he provides the first objective measure of progress (better chances of survival and spreading), still cannot completely escape the idea of a correlation between complexity and perfection. This tension is particularly perceptible in the last paragraph of the *Origins*:

It is interesting to contemplate an entangled bank, clothed with many plants of many kinds, with birds singing on the bushes, with various insects flitting about, and with worms crawling through the damp earth, and to reflect that *these elaborately constructed forms*, so different from each other, and dependent on each other in so complex a manner, have all been produced by laws acting around us. These laws, taken in the largest sense, being Growth with Reproduction; Inheritance which is almost implied by reproduction; Variability from the indirect and direct action of the external conditions of life, and from use and disuse; a Ratio of Increase so high as to lead to a Struggle for Life, and as a consequence to Natural Selection, entailing Divergence of Character and the Extinction of less-improved forms. Thus, from the war of nature, from famine and death, the most exalted object which we are capable of conceiving, namely, *the production of the higher animals*, directly follows. (Emphasis added)

At the beginning of this passage, the worm and the bird both deserve the title of “elaborately constructed form”; their very presence around us indicate that they have been able to survive millenia of competition, and thus should be regarded as equals. But in the end of the passage, the idea that evolution through natural selection naturally and mechanically leads to “the production of the higher animals” is asserted.

3.4 Coevolution: Red queens and arms races

In 1973 Leigh Van Valen tried to determine the distribution of survival times for different taxa [61]. He found this distribution to be 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 explanation for this fact was coevolution. Coevolution expresses the idea that in the natural world, species do not merely evolve against a fixed or slowly moving environment: they *coevolve* against each other, because their fitness is based in part on their interactions. Therefore, if one species A develops a given, specific feature, this feature may have an impact on the fitness landscape of another species B. This change in the fitness landscape may in turn prompt evolutionary changes in species B. Thus, as different species shape each other’s fitness landscapes, they may drive each other into series of *echoing* genetic changes.

While this notion may look like a mere re-statement of Darwin's ideas about natural selection, it brings about several interesting observations. In particular, a consequence of coevolution is that, as a given species evolve and improve its fitness, other species evolve too; thus each species is confronted to a constantly deteriorating environment, i.e. an environment that constantly changes so that current state is not as optimised as it was³. Thus, like the Red Queen in Lewis Carroll's *Through the Looking Glass*, species must constantly run (i.e. adapt) just in order to survive; those who fail, become extinct.

When this constant coevolution occurs in a directional fashion, that is, leads to the constant improvement in certain particular characteristics of coevolving groups (improvement in the sense that these characteristics become better at performing whichever function they support), the metaphor of choice is that of an "arms race", as introduced by Dawkins and Krebs [15]. The rabbit and the fox both have to find ways to run faster. The cuckoo must find a way to imitate the eggs of his hosts, while the hosts must find a way to identify the cuckoo's eggs. In this directional perspective, evolution is cumulative and leads to a constant "improvement" with regard to the characteristics selected for in this particular arms race.

This leads Dawkins and Krebs to the following suggestion:

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.

This offers an intuitive, reasonable notion of progress: progress is said to have occurred if it can be said that modern organisms would out-perform their ancestors. The obvious caveat being that, since species coevolve, it is quite possible that the global environment of a given group changes so much over time that the particular improvements developed by a late descendent would be of no particular help in the environment of its distant ancestor. More generally, the arms race metaphor begs the question of *intransitivity* in superiority: if in a given organism A can be said to be superior to B, and B superior to C, it is not always the case that A is necessarily superior to C.

The particular traits selected for in a given arms race may be at odds with other necessities in the environment, or particular constraints imposed by the

³This description of the Red Queen effect owes much to David Rand's online course.

existing body plan dictated by the taxon. Dawkins and Krebs mention hypothetical cases of sexual selection, in which the environment dictates a maximum viable size, but sexual selection favours individuals larger than average. The population will be driven towards a fragile equilibrium, close to the maximum size, at the risk of being swiftly wiped out in case a minor catastrophe (i.e. change in the environment) occurs. It is difficult not to think of the Irish elk and its gigantic woods [20].

Arms race may also end up in cycles: when species A has co-adapted to a change in species B, the incentive for that change no longer exists and may prompt a return to the previous type from species B. They may also reach a stable point, where interacting species find themselves at a local optimum from which they have no incentive to deviate. The latter point is an example of the concept of Evolutionary Stable Strategies (ESS), developed mostly by John Maynard Smith [30]: an ESS is a strategy, or a mix of strategies, that is *robust* against small perturbations or invasion by a different strategy. In terms of game theory, such a situation is reminiscent of a Nash equilibrium, in which none of the parties has any interest in changing its strategy in the current global state.

3.5 Lower and higher animals

Thus the notion of progress that seems to underlie the “arms race” metaphor must be taken with caution. At any rate, it certainly does not lead to any hierarchy between modern lifeforms. In a particularly acidic paper, Dawkins [13] criticises the notion of “higher” and “lower” animals, arguing that these vague terms actually cover “a phantasmagoric muddle-up of the following relations:

- A is ancestral to B.
- A resembles their common ancestor more than B does.
- A is simpler than B.
- A is stupider than B.
- A is less well adapted than B.
- A is less adaptable (or versatile) than B.
- A resembles “man” less than B does.”

While the three first notions may have a definable meaning (actually, when related to modern lifeforms, the first one can only be understood as a variant of the second), they are not necessarily correlated. Evolution does not necessarily imply complexification, and the common ancestor may actually prove more complex than some of its descendents.

Furthermore, each of these notions may be difficult to use in the general case, especially with regard to common assumptions regarding “lower” and “higher” lifeforms. For example, just because A is more primitive (i.e. resembles the

last common ancestor more) than B in some respect, does not mean that it is less primitive in *all* respects. Similarly, just because one member of taxon X is more primitive than a member of taxon Y, does not mean that all members of taxon X are more primitive than all members of taxon Y. The complexity of a given organism does not necessarily indicate at which point it branched from its ancestral line with regard to simpler cousins.

As for the other notions, they are quickly dismissed. In particular, the idea that lower creatures should be “less well adapted” than higher creatures is sternly addressed by noticing that “rat, cockroaches and bacteria are very successful species while animals that we like to think of as ‘high’, such as whales and gorillas, are threatened with extinction.”

3.6 The drunkard’s walk

Stephen J. Gould’s book “Full House” [21] can be seen as a deconstruction of the notion progress in evolution. The attack may actually be a bit too efficient: in particular, the language in which Cope’s rule ⁴ is attacked may lead the reader to think that Gould rejects the very idea of evolutionary trends in general (which is not the case, as explicated in another paper [23]). The main point of the book is that the undeniable generation of ever more complex species over time can, and should, be regarded as a statistical effect: because evolution constantly and, Gould asserts, randomly branches “up” as well as “down” in complexity space, thus following essentially a “drunkard’s walk” (which a mathematician would more dryly call a random walk), and because there is a lower limit, but no upper limit on the complexity of organism, then mechanically the maximum and (to a lesser degree) average complexity among all organisms are bound to increase.

Among others, the book features a powerful section on bacteria, in which the author makes the point that the living world is, and has always been, dominated by bacteria:

Not only does the Earth contain more bacterial organisms than all others combined (...) not only do bacteria live in more places and work in a greater variety of metabolic ways; not only did bacteria alone constitute the first half of life’s history, with no slackening in diversity thereafter; but also, and most surprisingly, total bacterial biomass (even at such minimal weight per cell) may exceed all the rest of life combined, even forest trees, once we include the subterranean populations as well. Need any more be said in making a case for the modal bacter as life’s constant center of maximal influence and importance?

Gould also uses McShea’s work on the subject. McShea has been studying the fossil record to determine whether or not evolution actually followed a progressive trend. He defined several types of complexity, globally related to either the number of different parts or the number of hierarchical levels. The notion

⁴Cope’s rule states that animals tend to grow in size over their evolutionary history.

of complexity as number of different parts echoes that of Pringle, as cited by Dawkins [13, 14]:

Pringle called complexity an epistemological concept, meaning a measure applied to our *description* of something rather than to that something itself. 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 for each segment and would therefore have a higher information content.

McShea also makes the fundamental distinction between *passive* trends and *driven* trends. A passive trend is just a random walk (or more precisely a random branching process) with a lower limit and no upper limit: it is bound to produce an increase in both maximum and average value. In a driven trend, however, the branching process is not symmetrical, and branches that go “up” are favoured over branches that go “down”. McShea precisely attempted at finding a driven trend in evolution towards higher complexity - and failed [37]. McShea insist that his result so far only support “emphatic agnosticism” on the question.

The review of Gould’s book by Dawkins [13], as well as Gould’s review of another book by Dawkins in the same volume [22], provide an enlightening debate about the notion of progress. In particular, Dawkins argues that Gould uses a tailor-made definition of progress which in his opinion no modern evolutionary theorist would defend: “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.” At least on this point, Dawkins overlooks Gould’s quotation of E.O. Wilson mentioned earlier.

But Dawkins goes further:

...my adaptationist definition of progress (...) takes progress to mean an increase, not in complexity, intelligence or some other anthropocentric value, but in the accumulating number of features contributing towards whatever adaptation the lineage in question exemplifies. By this definition, adaptive evolution is not just incidentally progressive, it is deeply, dyed-in-the wool, indispensably progressive. It is fundamentally necessary that it should be progressive if Darwinian natural selection is to perform the explanatory role in our world view that we require of it.

This idea that evolution is progressive in the (local) sense that species become better in the niche that they occupy, which may (and often does) lead

to improvement and /or complexification of parts and characteristics that allow these species to perform whichever functions this niche specifies, is possibly our best candidate for a definition of what evolutionary progress might mean; when coupled with the notion of coevolution, which entails that these changes (progressive or otherwise) modify the environment of other species, and may lead to changes in these other species, it provides a coherent explanation for the sustained emergence of complex adaptive innovations - including behaviours - when such emergence occurs.

3.7 Summary

The notion of evolutionary progress has long been a difficulty in evolutionary theory. Indeed, from the very beginning of evolutionary thought, we see that several very different notions of progress are intimately mingled:

1. Local: The appearance of specific features that allow a given creature to cope better with its environment, e.g. how did the giraffe get its long neck?
2. Global: The idea that these improvements are cumulative over time and space (that is, over the amount of competition), resulting into organisms of ever-increasing “perfection”, e.g. how come the indigenous species of small islands are often displaced by invaders ? How did the vertebrates get their eye?
3. Global and oriented: A trend towards ever more rich, complex, intricate, intelligent forms, e.g. why did bacteria turn into Man?

When exposed in this (non exhaustive) way, the question becomes relatively tractable and hypotheses can at least be formulated - and possibly even tested. The problem is precisely that this distinction has not always been obvious. Unfortunately much of this confusion has been inherited by the field of artificial evolution.

4 Artificial evolution

The apparent power of Darwinian mechanisms in the shaping of Nature inevitably led scientists to wonder if they might harness this power and apply it to their own constructions. John Von Neumann’s thoughts on life, reproduction and evolution offered not only the first formal inquiry into this line of research, but also provided an enlightening description of evolutionary and hereditary mechanisms in general.

4.1 Evolution as a workman: the genetic algorithm

Besides (and shortly after) the works of Von Neumann, a distinct line of research, traced back as far as the 50s, attempted to use evolutionary techniques to solve

actual problems related to optimisation and design. Fogel [17] has compiled a fascinating account of early experiments in artificial evolution.

These methods usually work in the following way: At first a number of possible solutions to a particular problem are randomly generated. Then a two-phase cycle begins: First, the performance of these solutions for this problem is evaluated. Then various operator (mutation, recombination, selection) are applied to the candidate solutions in order to generate new ones, and the process starts again.

The most well-known model of this type is Holland's genetic algorithm [27] (Mitchell provides a much more readable introduction [38]).

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: solving the problem is equivalent to finding the string of bits for which the fitness function is optimised. Starting with random strings, each string is evaluated according to this function. Some of the fitter strings, that is, those for which the fitness function returns a better value, 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.

A pseudocode implementation of the simple genetic algorithm might be written as such⁵:

```
begin GA
  g:=0 { generation counter }
  Initialize population P(g)
  Evaluate population P(g) { i.e., compute fitness values }
  while not done do
    g:=g+1
    Select P(g) from P(g-1)
    Crossover P(g)
    Mutate P(g)
    Evaluate P(g)
  end while
end GA
```

As an engineering tool, the genetic algorithm is interesting enough. But the really important part of the genetic algorithm was the thorough mathematical analysis provided by Holland⁶. Until then, most of the work in artificial evolution had concentrated on mutation as the central genetic operator to create new

⁵This pseudocode is taken from a web page by Moshe Sipper.

⁶Indeed it has been said that the genetic algorithm was devised without any computer simulation, just out of pen-and-paper mathematical demonstrations[34]. Holland's book [27] certainly supports this view.

solutions; but in the genetic algorithm, or at least in the original analysis found in [27] and [18], recombination 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). The Schema theorem is written as follows:

$$M(H, t + 1) \geq M(H, t) * \frac{f(H, t)}{\bar{f}(t)} * (1 - p_d(H, t))$$

where $M(H, t)$ is the expected number of instances of schema H at time t , $f(H, t)/\bar{f}(t)$ is the observed relative fitness of H at time t (that is, the average fitness of this schema as observed among its current instances, divided by the average fitness of all current individuals) and $p_d(H, t)$ is the probability that this schema should be disrupted by such operators as crossover and mutation.

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). It means 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 theorem⁷.

Holland was able to show that such a policy was, in a sense, optimal in terms of profit maximisation. From his analysis, it follows that crossover-based artificial evolution 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.

4.2 Artificial and natural selection

Artificial selection is the selection of living organism by humans, based on subjective human criteria. in the *Origins*, Darwin used artificial selection as an illustration of the power of selective processes, and an example through which

⁷Note 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. Indeed, one of the benefits of Exact Schema Theorems (such as the ones provided by Stephens and Walbroeck [58]) is a better understanding of the role of crossover.

he introduced natural selection - the selection of living organism by interaction with the external environment:

As man can produce and certainly has produced a great result by his methodical and unconscious means of selection, what may not nature effect? Man can act only on external and visible characters: nature cares nothing for appearances, except in so far as they may be useful to any being. She can act on every internal organ, on every shade of constitutional difference, on the whole machinery of life. *Man selects only for his own good; Nature only for that of the being which she tends.* (...) Under nature, the slightest difference of structure or constitution may well turn the nicely-balanced scale in the struggle for life, and so be preserved. How fleeting are the wishes and efforts of man! how short his time! and consequently how poor will his products be, compared with those accumulated by nature during whole geological periods. Can we wonder, then, that nature's productions should be far 'truer' in character than man's productions; that they should be infinitely better adapted to the most complex conditions of life, and should plainly bear the stamp of far higher workmanship? (Emphasis added)

Thus for Darwin natural selection is seen as a better, more efficient breeder, because it can both detect and affect more characters than human selection.

As we have already mentioned, for some authors, the genetic algorithm, genetic programming, and other methods based on externally defined fitness functions, are equivalent to artificial selection. The human programmer defines a set of criteria, and selection occurs in accordance to these criteria.

We have also seen that some authors regard this form of selection as fundamentally more limited than real natural selection. The argument is, in part, similar to Darwin's: humans can only select after a few characteristics, while natural selection (supposedly) puts to test all the features of a given individual.

Another part of the argument rests on the implicit notion of evolutionary progress caused by coevolution. Thus, Ray writes [47]:

How might we work with digital evolution to produce useful products? Although we have a well established practice of plant and animal breeding, our relationship to digital evolution is quite different. Our ancestors were able to go out into nature and observe many highly evolved and complex organisms. They found uses for some of these, such as the ancestors of, rice, corn, wheat, chickens, pigs, dogs, etc. They then bred them to produce the much improved domesticated plants and animals that we know today.

However, in the case of digital evolution, we are starting with very simple organisms that have not yet achieved the complexity to be useful, so our first objective is to evolve complexity. Yet we have no prior experience with managing the evolution of complexity.

Probably any attempt to guide the evolution of algae to become corn, through artificial selection in the context of a breeding program, would prevent such a transition from occurring. I believe that we will never be able to guide the evolution of complexity by the use of artificial selection. To facilitate complexity increase we need a new and different approach. (...)

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”.

This passage can be interpreted in two ways: the first interpretation essentially relies on the idea of arms races, as described in the previous sections, with the assumption that these arms race will naturally lead to an increase in complexity. Mutual scaffolding will not only allow, but force creatures to become more complex. Again, we encounter the Lamarckian idea of evolution as a tendency to complexify, using Darwin’s natural selection as a motor.

It is also possible to read this passage in a more neutral way, as a statement that natural selection allows for the “drunkard’s walk” (which mathematically leads to a passive increase in complexity, unless something prevents it), while artificial selection prevents it by “railroading” the evolutionary process into a dead end towards which the population inexorably converges.

These two readings are somewhat mingled throughout the literature. The first one is based on the arms race concept that we have already mentioned. The second one underlies the term “open-ended evolution”, that is, perpetual evolution towards unpredictable forms⁸.

However, discussing the relative properties of artificial vs. natural selection requires that we can distinguish between them in the first place, that is, that we have a working definition (or, equivalently, separation criteria) between artificial and natural selection. In Nature, the distinction is obvious enough: artificial selection is operated through human intervention, while natural selection occurs in the wild. But in a computer experiment, human intervention is everywhere. The rules of an artificial world are set by the programmer; the programmer is

⁸Note that even in “artificial selection” models, some authors have attempted at implementing special genetic algorithms in which the selection process and the genetic operators would allow for open-ended evolution. A prominent example is Harvey’s Species Adaptation Genetic Algorithm [24, 25], in which the population acts as a semi-converged quasispecies, constantly sampling the fitness landscape in the immediate genotypic vicinity of its current location by single-bit mutations and small increases or decreases in the genome size. SAGA’s message is that we should embrace convergence (or more precisely quasi-convergence) instead of fearing it. Local optima are escaped by following neutral networks, that is, networks of fitness-neutral (or even slightly deleterious) mutations that eventually lead to a new, improved optimum. SAGA applies Kauffman’s results on NK landscapes [32] and his assimilation of long jumps (i.e. multiple changes in the genome) with blind, random jumps. While this is correct on NK landscapes, the idea is certainly very much alien to the standard GA culture, which relies on the schema theorem to justify the exchange of long pieces of genome.

bound to set the criteria according to which a given individual lives, dies, or reproduce. Thus artificial evolution poses a question which biology did not feel the need to address, yet happens to have significant importance for artificial evolution.

Packard [43] makes a distinction between what he calls extrinsic selection (driven by an external fitness function) and intrinsic selection (driven by the interactions between creatures). But from the viewpoint of each individual agent, would it be possible to see all other agents as one large fitness function? What is the real, defining difference between these two? Clearly it is not a matter of fixity vs. change: a fitness function is not necessarily fixed. The abiotic environment (e.g. climate) provides an external fitness function that may change much faster than any evolutionary process would allow for - let us think of Sahara, or of the once aptly-named "Greenland". Similarly, human tastes and requirements may change over centuries. Artificial / extrinsic selection is anything but fixed.

The main difference between artificial and natural selection seems to be the possibility of coevolution, and thus of evolutionary feedback: in natural selection, because of coevolutionary dynamics, each organism may have a direct influence on the fitness landscape of others, so changes in one group may prompt a change in another, which in turn modifies the fitness landscape of the former group. The idea of coupled fitness landscapes seems to be the main difference between artificial and natural selection.

However consider a situation in which simple algae grow in a small pond. They may evolve all kinds of processes to extract resources from the pond, which is sufficiently small that particularly performant individuals will actually deplete resources for their less performant neighbours. However, because of their body plan and organisation, they cannot physically interact with each other. Their only interaction is through their relative efficiency in exploiting the resources of the pond; this is however a strong interaction, and clearly an efficient organism will have a strong negative influence over the fitness of other organisms, by depriving them of a significant amount of resources.

Certainly this example seems reasonable, and one might be tempted to see it as an example of natural selection in the wild. Yet this situation happens to be a close biological equivalent to the standard genetic algorithm. The algae are the individuals. The pond is the finite size population. The fitness function is their efficiency in exploiting resources. The difference between this situation and real coevolution is clearly that here, different organism only alter the *level* of each other's fitness landscape, and not its *shape*.

Another defining factor in natural selection is the presence of ecological cycles (food chains, etc.) between many different groups. In Chapter III of the *Origins*, Darwin writes about a complex cycle involving birds who prey on insects, who prey on young herbivorous animals, who browse young trees, in which the birds live, "and so onwards in ever-increasing circles of complexity." Clearly if one of these links were altered, the whole chain would be affected, possibly causing avalanches of cascading effects. The fact that these avalanches do occur in the fossil record, and that they seem to obey a power law (the

frequency of events of a given magnitude M is inversely proportional to a power of M^9), have led some authors to state that life exhibits self-organised criticality (see in particular Bak’s disturbingly simple model of coevolutionary dynamics among many species [3]).

5 Artificial Coevolution

As we have seen, coevolution occurs when the fitness of each given individual is the result of its interaction with other individuals. This leads to a situation in which each individual’s fitness landscape is coupled to others’. In Nature, such selective interaction during lifetime is the rule; but in artificial evolution, evaluation methods based on external fitness functions dominate. They are simple to implement, and are readily suitable to practical tasks. Thus the field of artificial coevolution has been comparatively less studied than, say, the dynamics of the genetic algorithm.

It may be noted that some problems are competitive by nature, e.g. two-player games such as checkers, chess or go. Obviously for such problem the idea of evaluating individuals by pitting them against each other seems natural enough. For example, Chellapilla and Fogel [8] successfully evolved neural networks to estimate checkers positions, which were then used into a standard minimax algorithm for generating expert-level checkers players.

However, the case for coevolution embraces a wider scope. It is expected that, thanks to the “arms race” concept, coevolution will enrich the dynamics of evolution, by shifting the selective process from a fixed, pre-defined fitness landscape, to an ever-changing fitness landscape with constantly moving optima. In other words, coevolution is a first step in the replacement of artificial selection with something closer to natural selection. Because humans define the rules of competitive or cooperative interaction, it can still be constrained for engineering purposes; but by relinquishing the direct definition of what a “fit” organism is, we open the possibilities of evolution.

In artificial evolution, coevolution may also prove useful in its own right when the problem is not easily defined, or when evaluation requires a complex test suite. It also has the added advantage of being an adaptive evaluation method: a constant, pre-defined fitness function may run the risk of being too harsh in the beginning (weeding out potentially promising, but yet ineffective strategies from the gene pool) or not enough in the end (which lowers the selective pressure and slows down evolution). But in coevolutionary models, the first, unadapted organisms compete against equally unadapted opponents. The “harshness” of evaluation for each organism thus grows slowly and incrementally over time as opponents refine their strategies, which allows for a balanced evaluation.

⁹Power laws have several implications. First, large events are much more likely to happen than they would be under a Gaussian hypothesis. Also, they lead to an apparent clustering effect: in the words of Bak [3], “the longer you have waited since a large [event], the longer you can expect still to have to wait. . . [Events] are clustered in time, not periodic.”

| | | |
|--------------|--------------------|--------------------|
| | B cooperates | B defects |
| A cooperates | R = 3 (Reward) | S = 0 (Sucker) |
| A defects | T = 5 (Temptation) | P = 1 (Punishment) |

Table 1: The reward matrix of the prisoner’s dilemma, from the viewpoint of player A. In Axelrod’s experiment, this matrix is identical for both players. Note that $2R > T > P > S$.

5.1 The introduction of artificial coevolution

Reed’s poker player ([31], in [17]) seems to be the oldest known example of artificial coevolution. The competition was based on a simplified game of Poker, in which each player received only one card (either “low” or ”high”) and placed a bet (pass, low or high, costing 2, 3 and 7 pennies). The highest bet wins; in case both bets are equal, the highest hand wins. If bets and hands are equal, each player keeps its money. Genotypes encoded the relative probabilities of each bet depending when receiving a given hand.

An interesting fact was that sexual recombination did not provide any benefit to the population, which seemed to contradict the widely held opinion that sex could accelerate evolution dramatically. The authors argued that strong epistasis in the genetic parameters (the performance of an allele depended strongly on the value of other alleles) was the cause of this absence of result: in a modified version of the system with no epistatic parameters, recombination did provide an evolutionary speedup¹⁰.

Despite these results, coevolution (or more precisely competitive evaluation in evolutionary processes) did not gain any significant attention until twenty years later, with Axelrod’s results about the iterated prisoner’s dilemma (IPD, [2]). The prisoner’s dilemma is a simple game in which each player can either cooperate, or defect. Mutual cooperation provides the highest benefit for both players, but defecting against a cooperating opponent favours the defector, and mutual defection is less costly than cooperation against a defector. The game becomes interesting when played for several rounds (i.e., iterated), and when each player uses the history of his opponent’s behaviour to choose his next action.

Axelrod had organised an IPD contest, in which the most efficient strategy happened to be the surprisingly simple Tit-For-Tat (cooperate on first round, then do whatever the opponent did at the previous round). He then decided to evolve strategies from scratch, using the eight most representative strategies collected from the previous tournament as evaluators¹¹. This provided a fixed fitness landscape, in which efficient strategies did emerge.

Yet, as the author remarked, ”while they are better in the particular environment consisting of fitness proportions of the eight representatives (...) they

¹⁰It should also be noted that in this model, recombination led to the destruction of both parents, thus making comparisons with natural or GA-like crossover somewhat difficult.

¹¹The most representative strategies were actually those which explained the majority of the variance in the scores of all strategies.

are probably not very robust in other environments”. A coevolutionary experiment was then set up, in which each strategy played against each other. The dynamics of the population were reported to lead to a first phase of general defection, then to an increase in cooperating behaviours (which are more efficient, in terms of general payoff).

Hillis [26] applied a coevolutionary model to an optimisation problem. He considered the task of evolving sorting networks (sequences of instructions that could sort any sequence of number in a minimal number of operations). At first, the potential solutions were evaluated over a fixed test suite, leading to acceptable but suboptimal results. By allowing the population of test problems to coevolve with the sorting networks, a host/parasite cycle was initiated (the solutions being the hosts, and the testing problems being the parasites), leading to significant improvement in the quality of solutions found. The remarkable result obtained demonstrated the potential of coevolution in optimisation tasks.

5.2 Coevolution as a (better) workman

Thus, from the beginning, coevolution was seen as a way to improve the performance of evolutionary algorithms. It often occurs that only one of the population is under consideration, the other being simply used as a “testing” population. This host-parasite distinction, as exemplified by Hillis, is also used in Rosin and Belew’s experiments [51].

Nolfi and Floreano [42] summarise the potential advantages and problems of coevolutionary settings from an optimisation viewpoint:

First, the coevolution 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 coevolution is not involved (...)

Finally, competing coevolutionary systems are appealing because the ever-changing fitness landscape, due to changes in the coevolving species, is potentially useful in preventing stagnation in local minima. (...)

Unfortunately a continuous increase in complexity is not guaranteed. In fact, coevolving populations may cycle between alternative classes of strategies that, although they do not produce advantages in the long run, may produce a temporary improvement over the coevolving population.

The tacit assumption that quality is correlated with complexity, and that evolution is supposed to be progressive in both senses (if it weren't for those annoying cycles) is quite common in the literature [1, 51, 41]. However, in artificial evolution, it may be easier to defend (at least in some cases) than in Nature, if only because artificial evolution experiments usually start with completely random (and thus very poorly adapted) individuals.

The first argument proposed by Nolfi and Floreano, based on the arms race concept, has two symmetric aspects: first, it is expected that arms race dynamics and mutual scaffolding may lead to ever-increasing progress and complexity. But the second aspect of that argument, which is often important in artificial simulations, is that the coevolution may provide a gradually increasing degree of difficulty which allows the selective process to be adapted, at any time, to the current level of the population. As Hillis stressed, a fixed selective environment may prove too difficult in the first stages of evolution, or too lax in the later stages. In a coevolutionary environment, as individuals progress (provided they do), so does the difficulty of the task they impose on others. Of course this relies on a regular, incremental evolutionary progress.

The second advantage expresses the idea that, because a given species has to evolve against many different successive competitors, the more diverse fitness landscape will prompt the appearance of more complex behaviours. This, it seems, is at best a confusing formulation. Evolution only prompts a given species to adapt to its *current* competitors. If no special care in the competitive process (that is, in the *evaluation* phase) is taken, there seems to be no reason why evolution would care about previous individuals in any way. If a modification that destroys previous adaptations happens to prove beneficial against the current competitors, it will be adopted.

In Nature, a certain inertia is caused by the fact that very few macromutations are viable, so previous adaptations cannot be destroyed so easily. It may be so in some artificial evolution experiments, but it is not obvious that simple neural networks have this limitation.

An evolutionary process has no memory beyond the current state of the populations at hand. This reliance on evolutionary “memory” may not be justified unless the evaluation method enforces it in some way, which is precisely why many authors (including these ones) use methods in which recent individuals compete against ancient ones (see below).

The third advantage, which relies on the idea that ever-changing fitness landscapes should help preventing local minima, seems to disregard the fact that in coevolutionary situations, cycles and stable points *are* the local minima. However this idea is expressed in the conclusion of the paper:

These results point to the conclusion that in certain tasks it is always

possible to find a simple strategy that is able to defeat another single, albeit complex and general, strategy (although such simple strategy is a specialized strategy, i.e. it is able to defeat only that individual complex and general strategy) (...) From this point of view, the fact of a coevolutionary dynamics leading to a limit cycle (...) should not be considered as a failure but as an optimal solution.

This consequence of the Red Queen hypothesis is developed below.

5.3 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, [49], or Gomez and Miikkulainen's ESP method for the cooperative coevolution of neurons in a neural network [19]). 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 used by Sims [53], see below). In cases where only one species is considered, i.e. when all individuals are functionally equivalent, a direct elimination tournament can also be held [1].

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, only $N - 1$ matches are needed, yet each individual undergoes (in average) $\log_2(N)$ 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 (competing against a very good individual in the first round will lead to an unfairly low score). See [49] 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 and

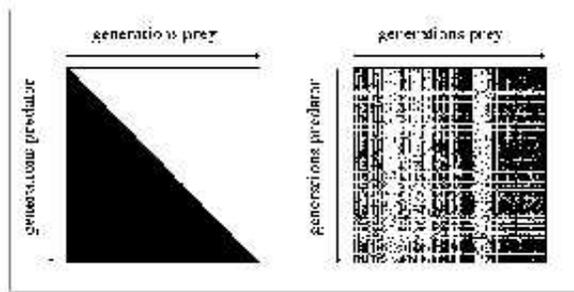


Figure 1: Master Tournament: Performance of the best individual of each generation, tested against all the best opponents of each generation. The black dots represent predator victories, while the white dots represent prey victories. The picture on the left represents an ideal situation of coevolutionary progress, in which the best individual of each species at generation N is able to defeat all opponents of previous generations. The picture of the right shows a typical run of a specific experiment (figure taken from [42]).

Belew [51]. Note that this process can be applied to all competing populations simultaneously.

5.4 Beware of the Red Queen

Because coevolution was introduced as a way to enhance the performance of artificial evolution, it was necessary to find a way to measure this performance. But when the fitness of an individual is evaluated after its interaction with other individuals, because of the Red Queen effect, the observed fitness of an individual at a given time provides little information about the global quality of the individual - provided such a global quality may be defined.

Artificial coevolution is subject to the same ambiguities as its natural counterpart, but these are made more acute by the necessary simplification of the environment. The effects of intransitive superiority (B can defeat A, C can defeat B, but A can defeat C), in particular, may become more acute. Thus it is theoretically possible that the constant relative “progress” of individuals with regard to each other does not lead to any global progress at all.

In a series of experiments, Cliff and Miller coevolved predators and preys controlled by neural networks (generated according to a biologically-inspired development model reminiscent of Nolfi’s [41], described below). They used this model to test several ideas about “tracking the Red Queen” [9]. They measured the performance of each individual at generation n against opponents taken from previous generations. If new individuals can consistently defeat old ones, then some kind of progress occurs.

This idea of Current Individual against Ancestral Opponent (CIAO) is conceptually simple enough and corresponds to the idea of progress in arms race

evoked by Dawkins, in which modern preys and predators might be expected to outperform their Eocene competitors significantly.

Rosin and Belew [51] provide a similar method which Nolfi and Floreano called “Master Tournament”: the best individual of each generation is pitted against the best individual of *each* generation in the competing population, which provides an idea of its “global” quality. By displaying the results of each competition in a square grid, it is possible to have an immediate indication of coevolutionary progress in the arms race sense - or lack of it (see Figure 1). This idea may also be used in the selective process, by keeping a “Hall of Fame” of the best previous individuals.

This illustrates one of the advantages of artificial evolution: a complete fossil record. We have seen that one of the more tractable notions of progress implied the idea that modern organisms might be expected to out-perform their ancestors. “Jurassic Park” scenarios aside, this idea is obviously difficult to test with the natural fossil record. In artificial evolution, however, a complete record can be kept, allowing for arbitrary experiments.

While these method help leading the process towards ever more efficient behaviours (in the sense of “efficient against a broad range of opponents”), they are not as straightforward than they seems. The biological implausibility of such methods is not really a problem. A more significant objection is that, in the words of Nolfi and Floreano,

the effect of the coevolutionary dynamic will be progressively reduced (...) with the increase in number of previous opponents. In fact, as the process goes on, there is less and less pressure to discover strategies that are effective against the opponent of the current generation and mgreater and greater pressure to develop solutions capable of improving performance against opponents of previous generations.

In other words, as more and more ancient opponents are added in the evaluation process, the coevolutionary aspect of the algorithm becomes smaller and smaller: the fitness landscape slowly solidifies.

5.5 Observing Arms Races

If the quality of a given candidate is measured after its performance against all (or a randomly selected subset of) previous opponents, we might expect that putting such a measurement in the selective process would allow us to obtain better individuals, according to this measurement. In other words, if evaluation occurs against previous opponents (as in the “Hall of Fame” method mentioned above), then final quality as observed against previous opponents should be expected to increase.

However, Nolfi and Floreano found that in an experiment in which both the predator and the prey had a sufficiently refined sensory system, simple coevolution *without* Hall of Fame actually performed better than *with* it. In other

words, coevolution against opponents of the current generation only allowed better performance against individuals of all generations, than coevolution against individuals of all previous generations.

This is a strong confirmation of the arms race concept: coevolution, by itself, may be able to find increasingly better, more general solutions in some cases; however the circumstances in which such situations arise were not studied in much detail. In this case, this successful coevolutionary run occurred when the prey was endowed with enhanced visual capacities.

6 Artificial Worlds

Coevolutionary experiments are the simplest way to implement a more natural, possibly open-ended form of evolution. Yet it is reasonable to think that 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 law¹² have made it possible to build up increasingly complex simulation, including 3D, physically realistic, situated organisms.

6.1 Abstract Ecologies

Once again, it is in Fogel [17] that we find trace of the first artificial ecologies. Conrad and Pattee [10] introduces 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

¹²"The number of transistors in an integrated chip doubles every year" [40]

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 in Hraber [28] and Mitchell [39]). 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 resource. The system can be tailored in various way to study its ecological dynamics. Comparisons with natural systems are also possible, although somewhat risky (see Hraber [28] for an example).

6.2 Evolving competing programs - Tierra

Competing programs in virtual computers have been implemented for decades, usually as games. An example is Core Wars [16], for which evolutionary experiments have been run. However, the most famous evolutionary experiments on machine code organisms is indisputably the Tierra system [46].

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 are 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 LIFO 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 instable) form of recombination.

Unfortunately, [6] 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 devel-

opment. Multi-threading has been introduced in an attempt to model multicellular organisms ([48]). 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.

6.3 Geb

Geb [6] is a simple environment designed for the study of evolutionary dynamics among many agents. It may be seen as a drastic simplification of a much more complex experiment known as Polyworld [62], although it was conceived independently.

Geb 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 severely 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 exist 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.

Genotypes are simply binary-encoded L-Systems [45]. 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 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. (note that an empty predecessor means that the rule may potentially apply to *all* nodes). 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

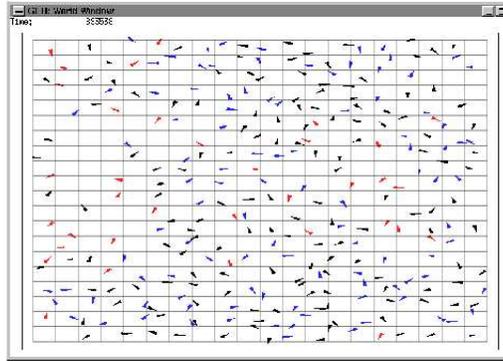


Figure 2: The experimental world in Geb.

their character). Inputs are sensitive to the excitation of output nodes of other organisms that match their character. This might 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.

Some behaviours do 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 other behaviour by visual inspection is difficult. In fact, the dynamics of the system as a whole appear somewhat unclear. This is not surprising when one considers the characteristics of Geb: An organism can kill any other, just by wanting it (indeed, when the world reaches saturation, simply wanting to reproduce causes the death of another organism). Genetic encoding relies on a rather fragile bit-per-bit description to encode L-systems, in which rules with no predecessor may be applied to the whole network. The brittleness of the encoding, together with the high unreliability of the environment, make the system rather inhospitable to evolution ¹³. It might be interesting to see what could emerge in a slightly less chaos-prone version of Geb.

¹³Actually, such a system comes dangerously close to the chaotic side of the order-to-chaos phase transition, which may prevent the emergence of just about any recognisable property in the system, as explained by Kauffman [32].

7 Karl Sims' experiments

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 [29] 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 remarkable survey of developmental systems for evolving controllers and morphologies simultaneously [57]. This study is centred on description and conceptual analysis of development processes and generative encodings.

Karl Sims' experiments [54, 53] 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 3). Each node correspond to a limb. The difference between graphs and trees is the presence of loops. When a node

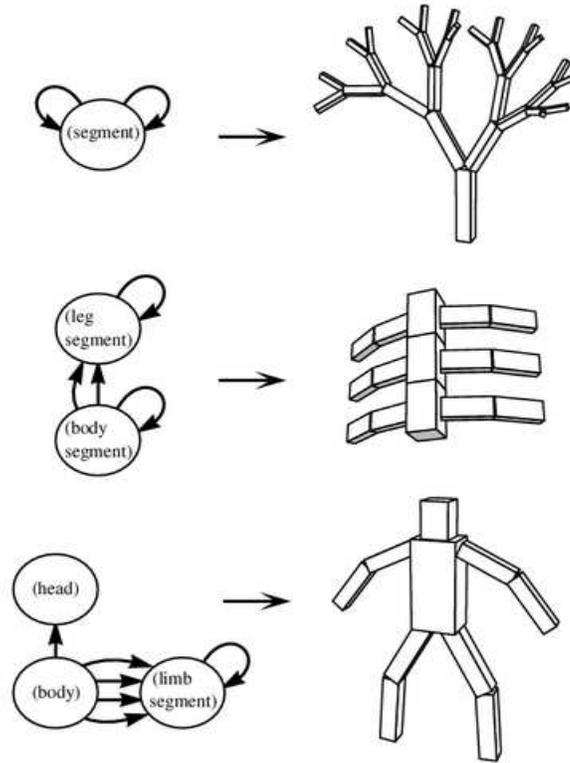


Figure 3: Sims' Creatures are represented by directional graphs. Each vertex in the graph corresponds to a limb. Edges correspond to articulations between a limb and a sub-limb, and contain information about the placement, resizing and orientation of the sub-limb. A cycle in the graph can be followed several times, giving rise to recursive structures.

has a recursive link to itself, similar copies of these nodes 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 rein on “junk 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 may 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 [54], 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.

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. Note that in this experiment, different sensors were available for detecting the block and the opponent.

When evaluation is determined by competition between individuals, it may be difficult to estimate - or even to define - the overall quality of a given creature. As we saw before, having each creature compete against just one opponent will hardly provide a meaningful measure of fitness, but pitting each creatures against all possible opponents entails dissuasive computing costs. 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. A combination of these two (each creature from a species competing against the best from the other species) was used for most runs. The result showed a wide variety of behaviours, some of which are shown in Figure 4. 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 catch it from the other. Some simply caught the cube and moved away with it (m). The population dynamics in the two species of each run exhibited various behaviours as well.

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.

8 Summary and potential inadequacies of previous work

In summary of this review, our opinion is that the mechanisms through which artificial evolution techniques could be employed for open-ended evolution are still very much a field of exploration.

An example of this is the notion of evolutionary progress. We point out that although this notion often seems to be taken as self-evident fact in artificial evolution, it is held in high suspicion or even outright defiance by biologists. What kind of progress can be expected from evolution, if any, should be studied carefully if evolutionary dynamics are to be better understood.

Similarly, it has been claimed that the emergence of novel behaviours (or even the emergence of complexity) is favoured by the introduction of natural selection, by opposition to more abstract models based on pre-defined fitness functions (e.g. the genetic algorithm) which are equated to “artificial” selection. However it turns out that we have no clear definition of the distinction between artificial evolution and natural selection, because the obvious criterion that exists in nature (human intervention) is inoperant in artificial evolution, where the human programmer necessarily controls the parameters of the environment.

In another direction, we observe that experiments in artificial evolution so far have mainly been concerned with abstract entities (such as numerical parameters or program code) or with specific parameters in simulated agents (neural networks, etc.) Even in those experiments that studied embodied agents, that is, agents that were able to interact in more complex ways with their environment, the possible behaviours were often pre-defined as a fixed repertoire of actions,

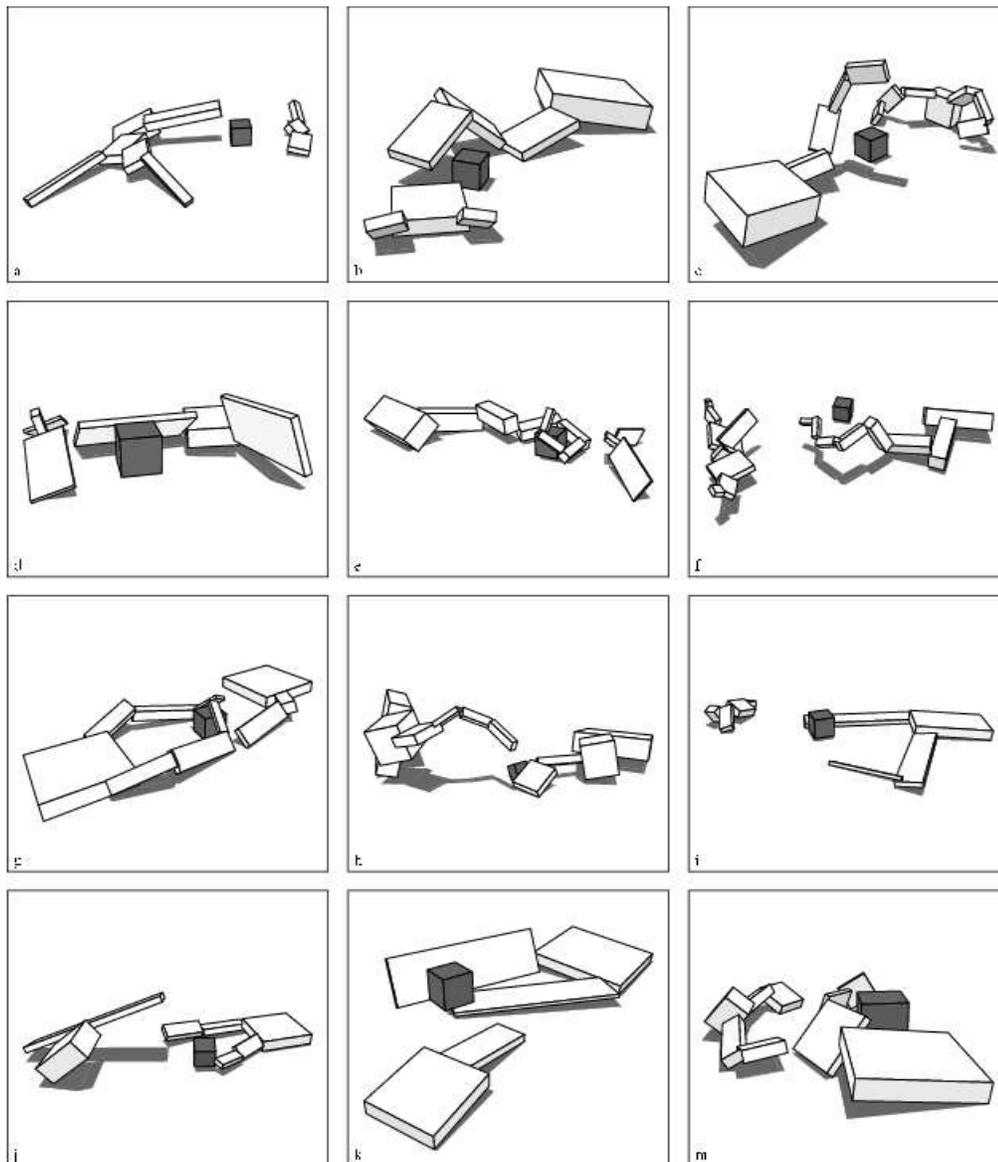


Figure 4: Creatures evolved by Karl Sims in a coevolutionary experiment. See text for details.

and the role of evolution was mainly to build an action selection architecture. Yet it is possible that evolutionary creativity might benefit if agents were able not only to choose among pre-defined actions, but to construct their own actions and behaviours in accordance to their environment. This can only be obtained if both morphology and control undergo evolution. While few experiments on this subject have been reported, one specific model (Karl Sims' creatures) seems to offer an interesting example in that regard.

It is felt that our knowledge of the mechanisms of evolution could benefit from the study of a complex evolutionary environment in which different types of tasks could be performed by creatures that would have evolutionary control over their own behaviours. To avoid the seemingly common pitfall of relying too much on "what everybody knows", such a study ought to be based on a sound conceptual grounding borrowed from the biological literature.

9 Proposed work

9.1 Objectives

Our objective is to provide a broad panorama of evolution in natural and artificial systems, and, based on this conceptual grounding, to study evolutionary dynamics with computer experiments based on semi-realistic artificial animals. The centre of our study concerns the conditions under which behaviours can appear in response to each other, and the consequences of these situations with regard to the performance, complexity or other characteristics of the evolving organisms.

More precisely, our objectives are the following:

- To produce a comprehensive conceptual framework of life and evolution, from definitions to causal descriptions, in order to add to the understanding of these concepts and of what can be expected in different conditions in artificial evolution experiments.
- Based on this framework, to create a software platform in which the evolution of artificial creatures can be adequately studied under a broad range of conditions, including at least artificial selection and interactive coevolution, in order to observe the emergence of novel behaviours.
- As an additional objective, we will consider the feasibility of using this platform to create a natural selection system in which evolution among populations of significant size can be simulated and analysed.

Let us stress the obvious: We do not seek to produce an all-encompassing model of evolution. Rather, our simulation concentrates on a small part of the larger picture, namely that which relates to coevolutionary dynamics. Our interest is to use our conceptual framework and our software simulation to pursue on research objectives stated by Channon [7]: "To add to the list of known requirements for the generation of evolutionary systems in which increasingly

complex advantageous behaviours can emerge”, with a little addition: “after determining whether, and how, such emergence actually occurs in evolutionary systems”.

In summary, there are really three components in our proposal:

1. A theoretical study
2. The realisation of a software platform
3. The design and analysis of experiments with this platform.

10 The software platform

10.1 Requirements

Let us precise our expectations for the software system we intend to build. If we want to use computer experiments to study evolution, we should come up with a system that offers the widest possible field for evolution to explore, while still keeping evolution manageable. In line with Von Neumann’s intuition, our system should allow for virtually unlimited growth in complexity (practical considerations aside, of course). It should not be limited to the adjustment of weights in a neural network, or even to the construction of neural networks in abstract agents with a limited repertoire of actions; ideally, it should allow agents to evolve their own actions as well as the behaviours that regulate these actions. 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 that were not explicitly defined a priori.

The difficulty with such objectives is that if the search space is too large and uncorrelated, it may become difficult for evolution to produce any meaningful search at all. Fortunately, Karl Sims’ experiments seem to offer a workable framework that fulfills our requirements: both control and morphology are under evolutionary control, and the behaviours that arise, as identified by human observation, were not defined a priori by the programmer, but emerged by themselves through evolutionary “tinkering”.

Therefore, we propose to build an environment similar to Sims’, in which creatures (defined as morphologic trees of rigid blacks) may interact physically with their environment, featuring a control system based on neural networks. We think that this model is rich and flexible enough to allow for the appearance of novel behaviours through the joint evolution of morphology and control.

10.2 Experiments

Using this platform, we intend to perform experiments related to the dynamics of evolution and coevolution among these creatures. As we have already mentioned, a major goal of these experiments would be to study the emergence and

consequences of coevolutionary dynamics in a minimally constrained environment.

The first type of experiments we intend to perform are simple genetic algorithms based on a very simple criterion such as locomotion (i.e. distance travelled), similar to the experiments described by Sims [54]. The goal of these experiments would be to test the suitability of the software platform, as we explain in the next section. However this task should not be underestimated: as we said, replication of Sims' results has proven elusive so far [60].

A second set of experiments will be based on the "box grabbing" task described in Sims's second paper [53]. This competitive task provides an interesting support for studying coevolutionary dynamics: it is conceptually simple, relatively open (there are many different ways to succeed) and is tolerant to simple, unrefined strategies (even poor strategies may obtain a better-than-zero score), which is important in the earlier stages of evolution. Sims' paper, which was rather short and predated the studies of Rosin and Belew or Nolfi and Floreano, has certainly not exhausted the scientific study of this experiment. Again, simply reproducing results similar to Sims would not be a trivial achievement and, contrarily to the previous experiment, could provide valuable insight.

The next step is to devise a more complex task, that could possibly be used in a system involving many individuals. Box-grabbing can only be implemented with two (or little more than two) individuals. We said that natural selection involves coevolutionary coupling between the fitness landscape of various creatures. In Geb and Polyworld, this coupling rests on the fact that creatures can kill each other. In Echo and similar systems, it rests more on the exchange of resources.

Physical fight between 3D creatures seems to be an interesting task to study (we are not aware of any such research in evolutionary computing). It may seem prohibitively complex at first, but it may simply be equated with repeated contact (i.e. hitting an opponent). The main problem with such a direct fighting method is Newton's law of equal action and reaction: when an individual hits another, both undergo the same shock - and potentially the same damage. It might be possible to arbitrarily give an advantage to the individual that actually performs the blow (possibly detecting this by checking which of the colliding limbs had the highest velocity with regard to the external environment). However this might simply favor simplistic strategies such as "jerking" all limbs as fast as possible. This, in turn, might be addressed by giving more resistance to heavier bodies (or simply dividing the damage of any blow by the mass of the hit limb). All this could involve subtle trade-offs and parameter tuning. Another solution is to state that only the trunk of each creature (the "zero" limb from which all other limbs stem) is sensible to shocks and blows.

All of these possibilities, however, face a difficulty in the initial stages of an evolutionary process: it is clear that poorly adapted, randomly generated individuals, would find it difficult to obtain any non-zero score, thus preventing any evolution at all. This can be solved by adding another component to the fitness, for example based on locomotion: a small amount of fitness could be attributed to creatures based on the distance they covered (in a coevolutionary

setting based on competition between two individuals, it might be desirable to take into account distance covered in the direction of the opponent). Of course this should be much lower than any fitness gained or lost by actual fight.

10.2.1 Analysis

Implementing these tasks would be an interesting engineering project, and obtaining meaningful behaviours might have significant scientific interest; however it is important that we use these experiments to support our study of evolutionary dynamics.

Several factors and settings may influence the results of these experiments. We are not talking about “technical” parameters such as mutation rate or friction coefficients, but about choices that influence the evolutionary dynamics of the system. They mostly fall in two broad categories:

- The number of opponents against which each individual is evaluated, and the way these opponents are chosen (tournaments, leagues, random sampling, competing populations, etc.)
- The way genetic information flows between individuals (subdivision of populations, niching, etc.)

By modifying these parameters, we may obtain various results which may be identifiable to the human eye. However, qualitative observation, although important, is not sufficient to analyse an evolutionary system. Fortunately, several tools in the literature allow us to dissect both the results and the internal dynamics of evolutionary experiments. For coevolution, we have already mentioned the “Master Tournament” method, which graphically summarises the results of competitions between the best individual of each generation and the best individuals of all other generations. As we noted, this provides a simple way to assess the presence of arms race dynamics in the system.

Stanley and Miikkulainen [56] provide a method for tracing genetic changes over lineages. Each time a new gene is created, be it by duplication, random generation or any method, it is given a unique identification number. This identification number is then passed on to children who inherit the gene. This method was introduced to enforce coherent crossover between genomes representing neural networks: crossover takes identification numbers into account to determine analogous genes, thus allowing to preserve structure information while performing crossover. It is certainly possible to use the same technique for our creatures, which might serve both for evolutionary guidance and a posteriori analysis of the results.

In short, we have experiments based on various parameters and settings. We want to modify these parameters and settings and see what happens. We want to use objective measures that allow us to classify the results, in order to say “these conditions lead to these effects”. We want to observe whether coevolutionary “echoing” dynamics occur, and whether they lead to mutual scaffolding towards anything visible (be it complexity, performance, or anything else). More

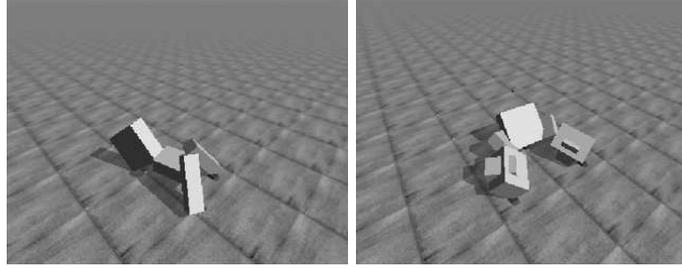


Figure 5: Some of the creatures obtained during our preliminary experiments.

precisely, we definitely want to observe coevolutionary echoing dynamics, and thus find out which conditions favour the emergence of arms races. This, more than anything, is the centre of our research.

10.2.2 Optional development towards larger populations

We would like to assess the feasibility of a system involving many concurrently evolving individuals. Computing costs may make such a system unfeasible in practice, but this can only be decided after previously mentioned experiments are performed.

If such a system turns out to be feasible in terms of computing time, then we might explore the possibility to set up an experiment of this type and to observe the joint evolution of many different creatures. This objective amounts to using our software platform to produce small artificial ecosystems comparable to Geb, but with creatures that are both more complex and less constrained in their interactions.

10.2.3 Results obtained so far

As mentioned in our previous report, we already built an environment corresponding to these requirements (see figure 5 for screenshots of the system). In accordance to Karl Sims’ design, our creatures are essentially articulated trees of 3D blocks, controlled by neural networks. Each 3D block (which we may call a “limb”) contains a small neural network, with two input neurons and two output neurons connected to the joint between this limb and its parent limb (i.e. the limb is attached to). These two inputs (r. outputs) receive (r. send) information about one of the two degrees of freedom of the joint. The neural network of a given limb can also receive input from other limbs attached to it.

In contrast to Sims’ work, in which the output of the neuron specified the desired *torque* to be applied, in our simulation the output of this neuron specifies a desired *angular speed*. A maximum possible torque is specified as a global simulation parameter. The joints will automatically apply the necessary torque to reach the desired speed, within the limit set by the maximum torque

parameter (this may be seen as a grossly simplistic model of a speed-regulating servo-controller).

Another difference with Sims' work is that instead of using special ad hoc neurons which perform complex operations on their inputs, we chose to use standard sigmoid-based neurons, with bias and weighted inputs. Our preliminary results indicate that these simple neurons allow for the emergence of at least midly complex behaviours, sufficient to emulate the first part of Sims' experiments (evolving locomotion).

The first stage of the project was a very simple system in which only morphology was under evolutionary control. In this preliminary model, each joint rotates at constant speed and without limits. Body parts can interpenetrate freely. Locomotive behaviours were obtained quite easily.

The second stage was to implement neural controllers in the creatures. Joint limits were imposed, but body parts could still interpenetrate at will. After some time, a variety of locomotive behaviours emerged again.

These results, for all the creativity they showed, did not look very natural. We decided to prevent interpenetration between body parts that were not linked by a joint. Together with joint limits, this should provide a good approximation of Sims' more complex model. Unfortunately, this proved to be more difficult than expected. Only after a long process of parameter tuning did locomotive behaviours appear.

We did spend some time on the parameters of the physics simulator in order to obtain a more stable environment. Perhaps more interestingly, it seems that other parameters related to the structure of the neural graph (e.g. the maximum number of input connections for each neuron) had a strong influence on the result - or absence of results. If this is confirmed, this might be related to results on the appearance of order in random networks [32].

We also designed coevolutionary experiments corresponding to Sims' second experiment (box grabbing). These entailed the addition of a sensory system, which simply provided one new input to one of the limbs (more precisely limb zero, the "trunk" of the creature): this input was simply the angle between the current direction of this limb and the line between this limb and the box.

Overall results are mitigated. Creatures did manage to produce meaningful behaviours, and did exploit sensory information (as could be seen by "lesion" experiments in which this information was suppressed, or experiments in which the box was displaced and followed by some creature). However the locomotive systems, if present, were often unimpressive ("jittering" movements which allowed the creature to move in a way that was easy to direct through movements of very small amplitude).

10.3 Contingencies

Considering the nature of the project, two types of risks can be identified:

- The underlying physical simulation system proves unstable to the point of preventing any experiment altogether, or slow enough to become unusable.

We feel that experiments performed so far allow us to regard this risk as minimal. We managed to find simulation parameters that make our simulations stable, and the current speed of the system allowed us to perform experiments similar in computational costs to those described by Sims.

- Our experiments fail to bring any results. The difficulty of the task should not be underestimated; we may find ourselves into a situation where our experiment refuse to produce any meaningful behaviour. However, previous experiments indicate that evolution does manage to find solution to these problems.
- Initial results are produced, but turn out to be rather uninteresting, and cannot be extended. Only simplistic behaviours emerge, preventing any exploration of multiple solutions and nullifying the principal object of this work, that is, coevolutionary dialogue between evolving entities. The point of this project is precisely to overcome this risk. Should such a situation occur, even though it is difficult to predict exactly which specific problems may arise, several possible routes can be followed
 - Adding knowledge into the system: this include any higher-level facility that we may provide to the system to facilitate evolution. Note that Sims does just this by allowing his creatures to use complex functions as neurons. However, these facilities should always be components from which evolution has to choose in order to build behaviours, and should not directly define these behaviours, otherwise any self-constructing aspect of our project would be lost.
 - Seeding the system, that is, starting with initial populations composed of hand-designed, simple creatures which are known to have non-zero fitness. The limit there is that these creatures should be simple enough to allow for significant further evolution, which is the central point of this project.
 - Simplifying the system, that is, adding constraints that limit the search space in some way. This might be done by specifying patterns of connection between nodes, imposing a constrained developmental process, or any other means; the important thing is that we still provide enough room for evolution and coevolution to occur and construct behaviours of their own.

Those options certainly do no exhaust the possible answers to any problem that might occur, especially given the notorious unpredictability of evolutionary experiments.

11 Criteria for success or failure

The success / failure criterion for the software platform is simple enough to define: we expect to obtain an environment in which meaningful evolution is possible. This can be tested with experiments based on genetic algorithms for simple tasks (or at least tasks which can be simply defined) such as locomotion. If our software platform does not allow for the emergence of such straightforward behaviours as “moving forward”, it will probably not be suitable for the study of more intricate behaviours. A success in this regard is a necessary prerequisite for the success of the entire project.

If we succeed to build such a platform, the experiments we are to perform on it constitute the centre of the thesis. The basic success / failure criterion, regarding our objectives, is to reach at least a point where evolutionary feedback is possible, in such a way that some behaviours emerge in response to others. Studying the circumstances in which such echoing dynamics occur, as well as their overall direction or lack thereof, is the main point of this work.

An additional success might be related to the ability of our systems to exploit the capacities of the platform. By estimating the possibilities that the platform offer to evolution, and comparing these possibilities with what our evolutionary systems actually come up with, we can further evaluate the efficiency of our settings.

The success / failure criterion is necessarily more subjective for the conceptual part of the work. We set ourselves to provide definitions and descriptions for concepts and mechanisms that constitute whole fields of research. The point is that if we can provide a framework that is both comprehensive and useful (in that it allows to understand, classify and possibly predict), we may be reasonably satisfied with our work. Essentially, we do not commit ourselves to provide definitive answers to questions that have puzzled Aristotle, Darwin and Von Neumann. What we seek are satisfying concepts and justifications for the evolutionary mechanism we expect to simulate.

12 Timetable for remainder of the research

Here is a tentative timetable for the production of the final document.

- September 04 - October 04 Reconstruction of the software platform. Several waves of additions, improvements and structural changes have made the code rather unwieldy. We expect to freshen it by cleaning it up or possibly starting a new codebase, with the constraint that the code should be easily maintainable.
- November 04 - December 04 Design, implementation and analysis of experiments based on Sims’ first paper (standard genetic algorithm with locomotion as a fitness criterion). Most of this timeframe will probably be spent fixing and tweaking the codebase.

- January 05 - February 05 Design, implementation and analysis of experiments based on Sims' second paper (coevolution for a box-grabbing task). Our objective here is to reproduce Sims' results in a publishable form as soon as possible. Writing a paper summarising experiments, results and discussion.
- March 05 - May 05 Design, implementation and analysis of further experiments based on Sims' second paper (coevolution for a box-grabbing task). Our objective is to obtain a working reproduction of Sims' results, then use the software for a study of coevolution. By contrast to the previous stage, much of this timeframe will be spent in research and experimental design. We expect to produce much of our conceptual research during that time.
- June 05 - August 05 Improvement of the software platform in order to allow or a new set of experiments based on a more complex task (e.g. physical fight). Writing a paper summarising experiments, results and discussion.
- September 05 - December 05 Experiments going beyond single-task evolution. Writing a paper.
- January 06 - March 06 Detailed analysis (results,successes,failures) of experiments and conversion of the conference paper to a journal article.
- April 06 - Jul 06 Writing the thesis.

References

- [1] Peter J. Angeline and Jordan B. Pollack. Competitive environments evolve better solutions for complex tasks. In S. Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 264–270. Morgan Kaufmann, 1993.
- [2] 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.
- [3] Per Bak. *How Nature Works*. Oxford University Press, 1997.
- [4] 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, Santa Fe Institute Studies in the Sciences of Complexity, Vol. X*, pages 431–461. Addison-Wesley, 1992.
- [5] Josh C. Bongard and Rolf Pfeifer. Repeated structure and dissociation of genotypic and phenotypic complexity in artificial ontogeny. In Lee Spector, Erik D. Goodman, Annie Wu, and W. B. Langdon, editors, *Proceedings of*

- the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 829–836. Morgan Kaufmann, 2001.
- [6] A. D. Channon and R. I. Damper. Perpetuating evolutionary emergence. In R. Pfeifer, B. Blumberg, J.-A. Meyer, and S. Wilson, editors, *Proceedings of the Fifth Conference on the Simulation of Adaptive Behavior (SAB '98)*, pages 534–539. MIT Press, 1998.
 - [7] Alastair Channon. *Evolutionary Emergence - The Struggle for Existence in Artificial Biota*. PhD thesis, University of Southampton, 2001.
 - [8] K. Chellapilla and D. B. Fogel. Evolution, neural networks, games, and intelligence. *Proceedings of the IEEE*, 87(9):1471–1496, 1999.
 - [9] Dave Cliff and Geoffrey F. Miller. Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In *Proceeding of the European Conference on Artificial Life (ECAL-95)*, pages 200–218, 1995.
 - [10] Michael Conrad and H.H. Pattee. Evolution experiments with an artificial ecosystem. *Journal of Theoretical Biology*, 28:393–409, 1970.
 - [11] Charles Darwin. *On the Origin of Species*. John Murray, London, 1859.
 - [12] Richard Dawkins. *The Blind Watchmaker*. Penguin, London, 1986.
 - [13] Richard Dawkins. Progress. In E. Fox Keller and E. Lloyd, editors, *Keywords in Evolutionary Biology*, pages 263–272. Harvard University Press, 1992.
 - [14] Richard Dawkins. Human chauvinism. *evolution*, 51(3), 1997.
 - [15] Richard Dawkins and J. R. Krebs. Arms races between and within species. *Proceedings of the Royal Society of London, Series B*, 205:489–511, 1979.
 - [16] 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.
 - [17] David B. Fogel, editor. *Evolutionary Computation: The Fossil Record*. IEEE Press, 1998.
 - [18] David E. Goldberg. *Genetic Algorithms in Search Optimization and Machine Learning*. Addison-Wesley, 1989.
 - [19] Faustino Gomez and Risto Miikkulainen. Incremental evolution of complex general behavior. *Adaptive Behavior*, 5:317–342, 1997.
 - [20] S.J. Gould. *Ever Since Darwin*. W.W. Norton, New York, 1977.
 - [21] Stephen J. Gould. *Full House / Life's Grandeur - The Spread of Excellence from Plato to Darwin*. Jonathan Cape, London, 1996.

- [22] Stephen J. Gould. Self-help for a hedgehog stuck on a molehill. *evolution*, 51(3), 1997.
- [23] Stephen J. Gould. On replacing the idea of progress with an operational notion of directionality. In David L. Hull and Michael Ruse, editors, *The Philosophy of Biology*. 1998.
- [24] Inman Harvey. Species adaptation genetic algorithms: a basis for a continuing saga. In F. J. Varela and P. Bouguine, editors, *Proceedings of the First European Conference on Artificial Life (ECAL 92)*. MIT Press, 1992.
- [25] Inman Harvey. Artificial evolution: A continuing saga. In Takashi Gomi, editor, *Proc. of 8th Intl. Symposium on Evolutionary Robotics (ER2001)*. Springer Verlag Lecture Notes on Computer Science LNCS 2217, 2001.
- [26] W.D. Hillis. Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D*, 42:228–234, 1990. NOT READ.
- [27] John H. Holland. *Adaptation in natural artificial systems*. University of Michigan Press, 1975.
- [28] Peter T. Hraber, Terry Jones, and Stephanie Forrest. The ecology of echo. *Artificial Life*, 3(3):165–190, 1997.
- [29] 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.
- [30] Maynard Smith J and Price G R. Logic of animal conflict. *Nature*, 246(5427):15–18, 1973.
- [31] Reed J., Toombs R., and Barricelli N. A. Simulation of biological evolution and machine learning. *Journal of Theoretical Biology*, 17:319–342, 1967.
- [32] SA Kauffman. *The Origins of Order*. Oxford University Press, 1993.
- [33] C. G. Langton, C. Taloy, J. D. Farmer, and S. Rasmussen, editors. *Artificial Life II: Proceedings of the workshop on the Synthesis and Simulation of Living Systems*. Addison Wesley, 1991.
- [34] Steven Levy. *Artificial Life: The Quest for a New Creation*. Pantheon, 1992.
- [35] Ersnt Mayr. *The Growth of Biological Thought*. Harvard University Press, 1982.
- [36] Barry McMullin. John von neumann and the evolutionary growth of complexity. *Artificial Life*, 6(4):347–361, 2000.
- [37] D. W. McShea. Metazoan complexity and evolution: is there a trend? *Evolution*, 50:477–492, 1996.

- [38] Melanie Mitchell. *An introduction to genetic algorithms*. The MIT Press, 1996.
- [39] Melanie Mitchell and Stephanie Forrest. Genetic algorithms and artificial life. *Artificial Life*, 1(3):267–289, 1994.
- [40] Gordon D. Moore. Cramming more components onto integrated circuits. *Electronics*, 38(8), 1965.
- [41] 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.
- [42] Stefano Nolfi and Dario Floreano. Coevolving predator and prey robots: Do “arms races” arise in artificial evolution? *Artificial Life*, 4(4):311–335, 1998.
- [43] N. H Packard. Intrinsic adaptation in a simple model for evolution. In Langton et al. [33].
- [44] H. H. Pattee. Simulations, realizations, and theories of life. In Christopher Langton, editor, *Artificial life, volume VI of Santa Fe Institute Studies in the Sciences of Complexity*, pages 63–77, 1988.
- [45] P. Prusinkiewicz and A. Lindenmayer. *The Algorithmic Beauty of Plants*. Springer, 1990.
- [46] Thomas S. Ray. An approach to the synthesis of life. In Langton et al. [33], pages 371–408.
- [47] 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.
- [48] Thomas S. Ray and Joseph Hart. Evolution of differentiated multi-threaded digital organisms. In Christoph Adami, Richard K. Belew, Hiroaki Kitano, and Charles E. Taylor, editors, *Proceedings of the 6th Conference on Artificial Life (ALIFE VI)*. MIT Press, 1998.
- [49] Craig W. Reynolds. Competition, coevolution and the game of tag. In Rodney Brooks and Pattie Maes, editors, *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. MIT Press, 1994.
- [50] Christopher D. Rosin and Richard K. Belew. Methods for competitive coevolution: Finding opponents worth beating. In Larry Eshelman, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 373–380. Morgan Kaufmann, 1995.
- [51] Christopher D. Rosin and Richard K. Belew. New methods for competitive coevolution. *Evolutionary Computation*, 5(1):1–29, 1997.

- [52] Hans-Paul Schwefel and Gunter Rudolph. Contemporary evolution strategies. In *European Conference on Artificial Life*, pages 893–907, 1995.
- [53] Karl Sims. Evolving 3d morphology and behavior by competition. In *ALife IV : Proceedings of the 4th Conference on Artificial Life*, pages 28–39. MIT Press, 1994.
- [54] Karl Sims. Evolving virtual creatures. In *SIGGRAPH 94*, pages 15–22. ACM Press, 1994.
- [55] John Maynard Smith and Eors Szathmary. *The Major Transitions in Evolution*. Oxford University Press, 1995.
- [56] Kenneth O. Stanley and Risto Miikkulainen. Efficient reinforcement learning through evolving neural network topologies. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2002)*. Morgan-Kaufmann, 2002.
- [57] Kenneth O. Stanley and Risto Miikkulainen. A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130, 2003.
- [58] Christopher R. Stephens and Henri Waelbroeck. Schemata evolution and building blocks. *Evolutionary Computation*, 7(2):109–124, 1999.
- [59] Tim Taylor. *From artificial evolution to artificial life*. PhD thesis, University of Edinburgh, 1999.
- [60] 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.
- [61] Leigh Van Valen. A new evolutionary law. *Evolutionary Theory*, 1:1–30, 1973.
- [62] L. Yaeger. Computational genetics, physiology, metabolism, neural systems, learning, vision and behaviour or polyworld: Life in a new context. In C. G. Langton (ed.), editor, *Artificial Life III, Vol. XVII of SFI Studies in the Sciences of Complexity*, pages 263–298. Addison-Wesley, 1994.